Generative AI Adoption, Labor Restructuring, and Corporate Capital Structure

Iftekhar Hasan, Yunying Huang, and Buhui Qiu*

This Version: March 2026

Abstract

We study the connection between Generative AI and capital structure. We use 350 million LinkedIn job postings from Revelio Labs to construct a firm-level measure of Generative AI adoption based on hiring for "integrator" roles that build or deploy large-language-model systems. Using a staggered difference-in-differences approach, we find that firms adopting Generative AI exhibit reduced leverage and improved debt-service capacity. Results are consistent and robust across alternative leverage measures and modern staggered-adoption estimators, and withstand identification checks including entropy balancing, stacked-cohort matching, parallel-trends tests, synthetic DiD, and two instrumental variable strategies based on the evolution of LLM capabilities and geographic proximity to AI research universities. One mechanism reveals that Generative AI adoption shifts firms toward intangible, less collateralizable capital and raises technology risk, consistent with lower debt capacity. Second, adoption increases mass layoffs, especially among firms with high ex ante generative-AI exposure, and deleveraging is stronger in low-union industries, consistent with reduced reliance on debt as a labor-discipline device.

Keywords: Generative AI adoption, Capital Structure, Intangible Assets, Labor Restructuring
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. Buhui Qiu is affiliated with University of Sydney Business School, The University of Sydney; Email: buhui.qiu@sydney.edu.au.

1 Introduction

This study investigates how Generative AI adoption affects corporate leverage. Corporate financing decisions shape firm value, investment, and, in the aggregate, economic growth. Understanding what drives these decisions has therefore been a persistent focus of financial economics (Rajan & Zingales, 1995; Graham, 2000). A firm's borrowing capacity reflects the pledgeability of its assets, the value of tax shields, and the expected costs of financial distress (Myers, 1977). A large literature further shows that labor market frictions enter this trade-off: when wages are rigid, renegotiation is costly, or bargaining power is concentrated, the effective costs of distress rise and firms adjust financing choices accordingly (e.g., Titman, 1984; Bronars & Deere, 1991; Matsa, 2010; Simintzi et al., 2015; Matsa, 2018). Against this backdrop, the diffusion of generative artificial intelligence (Generative AI) raises a first-order question for capital structure because it is transforming both the composition of firms' workforces and the nature of the productive inputs on which firms rely.

Generative AI is a general-purpose technology whose economic implications differ from earlier automation waves in ways that matter for financing. Previous technology waves, from steam engine and electrification to computing and the internet, reconfigured production and organizational processes, but much of complex language-based reasoning, writing, and analysis remained largely human-intensive (Autor et al., 2003; Acemoglu & Restrepo, 2018). Generative AI crosses this boundary. Large language models (LLMs) and the agentic systems built on them generate text, code, and analysis from natural-language instructions, expanding the scope of automation into cognitive tasks that previously required professional expertise (Eloundou et al., 2024; Brynjolfsson et al., 2025). Beyond improving individual decisions, generative AI crowds out the human cognitive effort that sustains collective knowledge (Acemoglu, Kong, & Ozdaglar, 2026).[1] Because these tasks are performed dis-

[1]ChatGPT reached 100 million monthly active users within two months of its November 2022 launch (UBS Global Research, February 1, 2023, based on Similarweb data; see https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/). According to a McKinsey report, by early 2025, 71 percent of organizations report regularly using Generative AI in at least one business function (Singla et al., 2025), and worldwide spending is projected to exceed $640 billion.

proportionately by non-unionized professionals who bargain individually rather than collectively, Generative AI can directly alter the labor frictions that the financing literature has identified as determinants of capital structure.[2] At the same time, firm-level adoption of Generative AI may shift production toward more intangible, less collateralizable inputs (e.g., software, data systems, and organizational capital) and change firm risk during implementation. These shifts in workforce composition, productivity, and asset pledgeability raise a first-order question for capital structure: how does Generative AI adoption reshape firms' leverage decisions?

We measure Generative AI adoption using job-posting data. Deploying Generative AI at scale requires specialized human capital, and firms' hiring choices therefore capture organizational implementation rather than generic tool usage. Our data comprise over 350 million LinkedIn job postings from U.S. public firms, drawn from a worldwide universe of 1.33 billion postings sourced from Revelio Labs, which compiles workforce records from online professional profiles and job boards (Eisfeldt et al., 2023; Lichtinger & Maasoum, 2025; Kwan et al., 2026). Job postings represent forward-looking resource commitments: a firm that creates positions to build or deploy LLM systems commits salary, infrastructure, and organizational capital to embed Generative AI into its operations. This deliberate investment is conceptually distinct from the diffusion of off-the-shelf tools across the existing workforce. A posting that merely lists ChatGPT proficiency as a desired qualification may reflect individual employees' adoption of readily available tools rather than a firm-level implementation decision. Our adoption measure therefore focuses on postings that recruit for building, fine-tuning, integrating, or deploying LLM-based systems, each of which reflects an explicit decision to embed Generative AI into the production process.

A firm's workforce shapes its capital structure. Workers who accumulate firm-specific skills demand compensation for bearing the risk of financial distress (Hart & Moore, 1994;

[2]Union membership in 2022 was 1.6% in professional, scientific, and technical services and 1.6% in finance, compared with 8.0% in durable goods manufacturing and 20.0% in transportation and warehousing. Source: Hirsch et al. (2026).

Berk et al., 2010; Baghai et al., 2021); organized workers extract rents that firms counter with strategic debt (Bronars & Deere, 1991; Matsa, 2010); and labor adjustment frictions generate operating leverage that limits borrowing (Simintzi et al., 2015; Favilukis et al., 2020). Generative AI directly alters these labor frictions. It substitutes at scale for cognitive and language-intensive professional work (Eloundou et al., 2024), displacing non-unionized workers who bargain individually rather than collectively. When firms can restructure their professional workforce through technology, the strategic value of debt as a labor-discipline device weakens. The technology also creates intangible organizational capital that cannot serve as collateral, reducing debt capacity. We expect Generative AI adoption to lower leverage, and adopters in our sample span industries well beyond the technology sector.

We find that Generative AI adoption reduces corporate leverage and improves debt-service capacity. Using a staggered difference-in-differences design from 2020Q1 through 2025Q2, we estimate that 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 ratios improve by 17 to 20%. These estimates remain stable when we replace year-quarter fixed effects with industry-by-year-quarter fixed effects, indicating that the leverage reduction reflects firm-level financing choices rather than industry-wide trends in AI investment or deleveraging.

Dynamic event-study estimates show no evidence of differential pre-trends and reveal that the effects emerge gradually after adoption, consistent with firms reoptimizing capital structure as implementation progresses. The results hold across alternative leverage definitions, including market leverage, net leverage, and long-term leverage. The effects also scale with adoption intensity: moving from the 25th to the 75th percentile of cumulative Generative AI hiring among adopters reduces leverage by an additional 2.0 percentage points. For comparison, Kim (2020) finds that large manufacturing plant openings in the local labor market, typically involving thousands of new workers, increase leverage by 2.5 to 3.9 percentage points. The interquartile range of Generative AI hiring among adopters, from 2 to

36 hires, generates a leverage reduction of comparable magnitude.

Our identification strategy combines transparent event-time evidence with complementary robustness and validation exercises. We verify the findings using modern DiD estimators designed for staggered treatment timing and heterogeneous treatment effects (Callaway & Sant'Anna, 2021; Sun & Abraham, 2021; Borusyak et al., 2024; De Chaisemartin & d'Haultfoeuille, 2024), all of which yield qualitatively similar patterns. To address the concern that insignificant pre-trends may have low power to detect economically meaningful violations (Roth, 2022), we construct robust confidence intervals following Rambachan and Roth (2023); the estimated effects remain statistically significant even when post-treatment trend deviations reach 40% of the worst pre-treatment violation. We also reweight the control group using entropy balancing to achieve exact covariate balance on pre-treatment characteristics, and implement the synthetic difference-in-differences estimator of Arkhangelsky et al. (2021), which matches each treated firm's pre-adoption leverage trajectory. The gradual emergence of the treatment effects after Generative AI adoption further argues against reverse causality: the leverage reduction follows rather than precedes adoption.

To further strengthen the causal interpretation, we exploit two sources of quasi-experimental variation in adoption incentives. First, we construct a novel Bartik shift-share instrument. LLMs have expanded rapidly from text generation to data analysis, code writing, and autonomous multi-step workflows, and each new capability changes the set of occupations susceptible to augmentation or replacement. We measure this expanding frontier using text-based similarity between occupation task descriptions from O*NET and model capability records from the Epoch AI large-scale models database, which tracks LLM-specific capabilities across model generations. As new models are released, the alignment between LLM capabilities and certain occupations increases, while occupations outside the language and reasoning domain remain comparatively unaffected. The shift is this time-varying, occupation-level alignment; the share is the predetermined occupation composition of each firm's industry from the BLS. Industries that concentrate employment in occupations whose

tasks resemble LLM capabilities receive a greater instrument value as the technology advances. Second, following Babina et al. (2024), we use geographic proximity to AI research universities. Both instruments produce strong first stages, and the corresponding second-stage estimates corroborate our baseline findings.

We focus on two mechanisms. The first is an intangible capital and distress-cost channel. Generative AI adoption requires firms to reorganize production around software, data infrastructure, AI-enabled workflows, and organizational know-how. Relative to tangible assets, these productive inputs are less collateralizable and harder to redeploy in financial distress (Rampini & Viswanathan, 2013; Falato et al., 2022). As Generative AI adoption increases firms' reliance on intangible and organizationally embedded capital, expected distress costs rise and debt capacity falls. The technology itself evolves rapidly, with new model generations and entirely new capabilities emerging within months. Physical equipment generates predictable cash flows over known useful lives and retains residual value in secondary markets; Generative AI capital has no such stability. Under standard trade-off logic (Graham, 2000), this uncertainty raises the expected cost of financial distress and further reduces optimal leverage.

We find evidence consistent with this channel. Generative AI adoption raises idiosyncratic stock return volatility by 3.4% under the Fama-French three-factor model, and excess kurtosis also rises. The deleveraging effect concentrates among firms with higher ex-ante intangible intensity: high-R&D firms reduce leverage by 2.8 percentage points after adoption, with interest coverage rising by 24.5% relative to the sample mean, while low-R&D firms show no effect. The same pattern holds for SGA intensity and patent activity.

The second mechanism is a labor restructuring and debt-discipline substitution channel. A large literature argues that debt serves as a strategic commitment device in labor bargaining by tightening the firm's budget constraint and strengthening management's negotiating position (Bronars & Deere, 1991; Perotti & Spier, 1993; Hennessy & Livdan, 2009; Matsa, 2010). This strategy is costly, and firms bear those costs only when labor bargain-

ing frictions are strong enough to justify them. Generative AI adoption may reduce the need to rely on leverage for this purpose. Because Generative AI substitutes for cognitive and language-intensive tasks in professional occupations (Eloundou et al., 2024), it displaces workers who are largely non-unionized and who negotiate wages individually rather than through collective bargaining (Stole & Zwiebel, 1996). When firms can adjust white-collar employment through technology-based restructuring, the indirect mechanism of threatening financial distress becomes less necessary. Moreover, high leverage is costly to managers because financial distress can destroy firm-specific managerial human capital (Titman, 1984; Berk et al., 2010; Baghai et al., 2021). If Generative AI adoption makes labor adjustment more feasible, managers may optimally choose lower leverage.

We find direct evidence of workforce displacement. Adopting firms are 47% more likely to file a WARN Act mass layoff notice, and the increase reflects mass layoffs within continuing operations rather than plant closings. The layoff response concentrates among firms with high ex-ante Generative AI exposure, consistent with labor substitution in cognitive occupations rather than routine manual tasks. The leverage reduction itself is stronger in industries with lower unionization rates, where workers bargain individually and the strategic motive for debt is weakest. In high-union industries, the leverage effect attenuates and the interest coverage effect disappears.

This study contributes to the emerging literature on Generative AI and firm outcomes (e.g., Eisfeldt et al., 2023; Noy & Zhang, 2023; Babina et al., 2024; Eloundou et al., 2024; Brynjolfsson et al., 2025). We study a core corporate finance outcome, capital structure, and focus on realized firm-level adoption rather than ex-ante task exposure. This distinction is important. Eisfeldt et al. (2023) measure workforce exposure to large language models; their measure captures the susceptibility of a firm's task mix to AI automation. Our adoption measure captures the firm's deliberate choice to commit resources to Generative AI integration, observed through hiring. Exposure is fixed before ChatGPT's release; adoption varies across firms and over time. A firm may have high exposure but delay implementation, or

adopt aggressively despite moderate exposure. The two constructs can produce opposite coefficient signs without contradiction. Babina et al. (2024) study how firm-level AI investment affects sales growth, employment, and market valuation. Hampole et al. (2025) examine how AI exposure reshapes labor demand across skill levels, and Lichtinger and Maasoum (2025) study how Generative AI adoption shifts the seniority composition of hiring. We contribute to the literature by documenting a negative effect of Generative AI adoption on corporate leverage.

We also contribute to the literature on labor market frictions, bargaining, and capital structure. A large body of work shows that labor protections and labor bargaining affect corporate 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; Chava et al., 2020; Favilukis et al., 2020; Kim, 2020). A common feature of these studies is that the labor adjustments involve organized workers and tangible capital investments.

We examine a setting in which technology reshapes the workforce segment that prior studies have not reached. The labor adjustments in the existing literature involve organized production workers protected by collective agreements, minimum wages, or employment regulation. Generative AI operates on a different margin: it displaces non-unionized professionals who perform cognitive and language-intensive tasks and who bargain individually rather than collectively (Eloundou et al., 2024; Acemoglu, Autor, & Johnson, 2026). Acemoglu, Autor, and Johnson (2026) argue that firms' private incentives tilt AI deployment toward automation that commodifies this professional expertise rather than toward applications that extend it. When firms can restructure their professional workforce through technology, the strategic value of debt as a labor-discipline device weakens; at the same time, the capital that generative AI creates is intangible and cannot be pledged as collateral. Our findings document that this combination of workforce displacement and intangible capital formation produces a leverage reduction, contributing new evidence on how the nature of labor market

frictions shapes corporate financing decisions.

Finally, we contribute to research on intangible capital and debt capacity. A broad literature emphasizes that debt capacity depends on the collateralizability and composition of firms' assets (e.g., Myers, 1977; Rampini & Viswanathan, 2013; Falato et al., 2022; Li, 2025). Falato et al. (2022) document that the secular rise of intangible capital over the past four decades has shrunk corporate debt capacity and driven a broad-based increase in cash holdings among U.S. public firms. Li (2025) formalizes this mechanism, showing that limited pledgeability of intangible capital drives firms to substitute liquidity for debt capacity. Our setting provides a specific, measurable source of new intangible investment with a well-defined adoption date and clear firm-level variation. Rather than relying on secular trends or cross-sectional differences in intangible intensity, we study the within-firm transition as individual firms shift their productive base toward intangible capital through Generative AI investment. More broadly, our results suggest that technology adoption is not a monolithic determinant of capital structure: the sign and magnitude of the leverage response depend on how the technology changes asset composition, operating flexibility, and expected distress costs. Generative AI adoption is a financing-policy shock whose consequences for debt capacity, distress risk, and the strategic role of leverage extend beyond the productivity gains that the existing literature emphasizes.

The rest of the paper proceeds as follows. Section 2 describes the data and measurement of Generative AI adoption. Section 3 presents the main results and robustness tests. Section 4 examines the intangible capital and labor bargaining channels. Section 5 concludes.

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.[3] 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 location, 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. Industry-level union coverage and membership rates come from the Current Population Survey as compiled by Hirsch et al. (2026).

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

[3]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.

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 Lichtinger and Maasoum (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).[4] The classifier assigns each posting to one of four categories: *integrator* (the role involves developing, deploying, or operating generative AI systems and applications 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. Only integrator postings enter our adoption measure. The full keyword list, classification prompt, model specifications, and robustness to alternative models and prompts are in Online Appendix OA1.

[Insert Figure 1 about here]

[4]Results are robust to classification using Meta's Llama-3.1-8b and Google's Gemini-2.5-flash. See Online Appendix OA1.

Our adoption measure specifically includes firms whose first generative AI integrator posting predates ChatGPT's November 2022 release.[5] 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. Online Appendix Table OA1 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 Online Appendix Table OA3.

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

[5]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).

& Maasoum, 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.[6] We merge Revelio workforce data with Compustat quarterly fundamentals and exclude financial firms (SIC 60–67) and utilities (SIC 49), and drop firm-quarter observations with zero total debt, because book leverage is mechanically zero and interest coverage is undefined for these observations. The final sample contains 41,007 firm-quarter observations for 2,650 unique firms, of which 642 adopt generative AI during the sample period. 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 quarterly salary exceeding $200 million by early 2025. Salary figures include Revelio-predicted compensation where reported salary is unavailable. Panel A of Figure 3 shows that software engineers and data analysts account for

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

most integrator postings, followed by application engineers and scientists. 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 Figure 2 and 3 about here]

3 Main Results

3.1 Generative AI Adoption and Leverage

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, the natural logarithm of one plus the ratio of operating income before depreciation to interest expense.

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 intensity. 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, and the relationship is both economically and statistically significant at the 1% level. Column (2) adds standard capital structure controls, and the statistical and economic significance of the coefficient remains consistent. To address the possibility that industry-level trends in AI investment or deleveraging confound the estimate, Column (3) replaces year-quarter fixed effects with industry-by-year-quarter fixed effects, and the coefficient remains at 1.6 percentage points, confirming that the leverage reduction reflects firm-level decisions rather than broader industry dynamics. The interest coverage estimates further confirm this finding. Column (4) shows that adoption increases log interest coverage by 0.178, approximately 11% of the sample mean, implying roughly a 20% increase in the interest coverage ratio. Columns (5) and (6) yield similar magnitudes of 0.157 and 0.162, approximately a 17% increase in the coverage ratio, after adding controls and industry-by-year-quarter fixed effects.

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 Figure 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 two-way fixed effects (TWFE) baseline. 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.[7]

Figure A2 in the 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

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

differential pre-existing trends.

C. Matched sample estimates

We construct two matched samples that pair adopting firms with non-adopting firms that are observably similar. Panel A of Table A4 uses entropy balancing. At 2022Q3, the quarter before ChatGPT's public release, we reweight control firms to match treated firms on size, tangibility, market-to-book, profitability, R&D intensity, and book leverage. 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. Table A5 in the Appendix confirms near-exact balance: all standardized differences fall below 1%. The reweighted estimates are, if anything, larger than the unmatched baseline. Column (1) yields 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; adding the leverage trend in Column (2) increases the estimate to 2.0 percentage points. Interest coverage rises by 12.4% without trend balancing in Column (5) and by 13.7% with it in 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) through (4) 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-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. 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 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 18%. 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. These 2 to 36 integrator hires produce a leverage reduction comparable 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]

[8]Panel A coefficients under industry-by-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: $e^{0.066 \times 1.18} - 1 = 0.081$ and $e^{0.066 \times 2.50} - 1 = 0.179$.

[9]Panel B coefficients under industry-by-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.

B. Alternative leverage measures

Our findings are robust to alternative leverage definitions. Table A3 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-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 after 2022Q4, to align with the economic shock of ChatGPT's public release used in Eisfeldt et al. (2023). Figure A3 in the 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 A4 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-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 A6 in the 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.

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. Two-way fixed effects remove level differences across firms and quarters but cannot account for firm-specific trends, while the SDID weighting scheme absorbs 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 A6 in the Appendix reports the results. All five leverage measures yield statistically significant coefficients, directionally consistent with the TWFE baseline, both with and without controls entered via the projected method. That our results survive an estimator designed to match pre-adoption leverage trajectories strengthens the case that the capital structure effects reflect a causal response to generative AI adoption.

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 A7 in the Appendix reports the results. All six placebo coefficients are economically small and statistically insignificant across both outcomes and all three specifications. The placebo leverage effect is 0.1 percentage points without controls in Column (1), compared with 1.7 percentage points in Column (1) of Table 2; the placebo interest coverage effect is 2.4%, compared with 17.8%. Adopting firms were not deleveraging before adoption.

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

[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. Appendix OA2.1 provides full construction details.

strong in both panels, with Kleibergen-Paap F-statistics of 36.5 and 33.1. The second-stage leverage coefficient is negative and the interest coverage coefficient is positive, both significant at the 5% level or better, consistent with the OLS estimates. Table A8 in the 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.[11] 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-quarter fixed effects absorb time-varying

[11] 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 institutions within a fixed radius of each firm's headquarters. Results are robust to alternative ranking cutoffs and distance thresholds. Appendix OA2.2 provides the full construction details, university rankings, and excluded corporate laboratories.

industry shocks, so identification comes from within-industry geographic variation in proximity to AI research.

Table A9 in the 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

Generative AI is a general-purpose technology that, unlike prior waves of technological change, substitutes directly for cognitive labor (Trajtenberg, 2018; Eloundou et al., 2024). From legal drafting to data analysis to software engineering, it performs cognitive tasks previously requiring professional judgment across a wide range of occupations, reducing firms' reliance on knowledge workers. This fundamentally restructures the firm's production function: human capital gives way to software and algorithms. As a factor of production, this new input has two properties relevant for capital structure. First, it is intangible and non-collateralizable; it erodes the asset base that supports debt. The technology also evolves rapidly, and the resulting uncertainty raises the expected cost of financial distress. Second, generative AI displaces individually contracted professionals who lack collective bargaining power. The strategic motive for holding leverage as a disciplining tool against labor vanishes. Both channels operate on any firm with positive debt: the intangible capital channel tightens the collateral constraint whenever the asset base shifts from pledgeable to non-pledgeable inputs, and the labor channel reduces the strategic motive for debt wherever firms can replace individually bargaining professionals with technology. We trace each channel to a specific set of empirical tests below.

4.1 Intangible Capital

A. Intangible investment intensity

Falato et al. (2022) show that intangible capital cannot be easily verified or liquidated, and thus cannot be pledged as collateral to raise debt financing. As the intangible share of firms' productive assets rises, debt capacity shrinks (Rampini & Viswanathan, 2013), a constraint that binds even for the substantial share of corporate borrowing that is formally cash-flow-based rather than asset-based (Lian & Ma, 2021). Generative AI accelerates this shift. When firms invest in generative AI integration, the resulting capital takes the form of trained models, reorganized workflows, and specialized data infrastructure, all of which are embedded in the organization and have little value outside it. The firm can pledge less collateral, which constrains its ability to borrow. Financial distress also becomes more costly when the firm's assets cannot be liquidated, which lowers optimal leverage. The leverage reduction should therefore concentrate among firms whose operations are already intangible-intensive, where generative AI adoption amplifies the existing collateral deficit.

We test this prediction using three measures of intangible investment intensity determined prior to ChatGPT's release. In Panel A of Table 5, Column (1) shows that high-R&D firms reduce leverage by 3.7 percentage points after generative AI adoption. The effect remains economically and statistically significant at 2.8 percentage points with controls in Column (2), with a 24.5% increase in interest coverage in Column (4), both at the 1% level. Low-R&D firms show no effect in Columns (5) through (8). Panel B examines SGA intensity, which captures spending on organization capital, brand equity, and employee training. Column (1) shows high-SGA firms reduce leverage by 3.0 percentage points; the effect remains at 2.4 percentage points with controls in Column (2), with a 34.0% increase in interest coverage in Column (4). Low-SGA firms show no response. In Panel C, Column (1) shows patenting firms reduce leverage by 2.6 percentage points, holding at 1.7 percentage points with controls in Column (2), with a 19.1% increase in interest coverage in Column (4). Non-patenting firms show no leverage reduction in Column (6). When generative AI shifts the

productive base toward intangible capital, the collateral constraint tightens most for firms already reliant on unpledgeable assets.

[Insert Table 5 about here]

B. Technology risk

The returns to generative AI investment are potentially large but inherently uncertain. Large language models not only advance through new generations within months but also acquire entirely new capabilities, from web search to autonomous reasoning to agentic workflows, and each generation reshapes how firms integrate the technology.[12] Firms stand to gain large productivity improvements, but the magnitude and durability of those gains are difficult to forecast when both the models and their capabilities evolve this quickly. Physical equipment generates relatively predictable cash flows over known useful lives and retains residual value in secondary markets. Generative AI capital operates in a different environment: capabilities improve rapidly, costs decline unpredictably, and the direction of technological progress is difficult to foresee. Under the trade-off theory, this higher business risk raises the expected cost of financial distress and lowers optimal leverage.

We measure this risk using two properties of stock returns. Idiosyncratic volatility captures firm-specific uncertainty after removing systematic risk factors; if generative AI increases the firm's own business risk, idiosyncratic volatility should rise. Excess kurtosis captures tail risk, the frequency of extreme return realizations. In Table 6, Column (1) reports that generative AI adoption increases idiosyncratic volatility by 2.7% under the Fama-French three-factor model, and Column (3) shows a 2.8% increase under the five-factor model, both significant at the 5% level. The results hold with controls in Columns (2) and (4), at 3.4% and 3.5%. Columns (5) and (6) show that excess kurtosis also increases, so adopting firms experience both higher volatility and more frequent extreme outcomes. Both results are

[12]Key milestones include GPT-3 (June 2020), ChatGPT (November 2022), GPT-4 (March 2023), GPT-4o (May 2024), chain-of-thought reasoning models (OpenAI o1, September 2024; DeepSeek-R1, January 2025), agentic AI coding tools (Claude Code, ChatGPT Codex, 2025), and GPT-5 (2025).

consistent with generative AI introducing a technology whose payoffs are genuinely uncertain, and this elevated risk provides a second force that lowers optimal leverage beyond the collateral channel.

[Insert Table 6 about here]

4.2 Labor

The second channel operates through the labor market. A long-standing theoretical literature, beginning with Bronars and Deere (1991), argues that firms hold debt to strengthen their hand against workers. The strategy works because debt imposes hard financial constraints (Hart, 1995): interest and principal payments are contractual obligations that cannot be renegotiated away. By raising the probability of financial distress, leverage commits the firm to resist wage demands, and workers, recognizing the constraint as binding, accept concessions to preserve their jobs. Matsa (2010) confirms this prediction: firms facing stronger union protections carry higher leverage. Chava et al. (2020) document the same pattern, showing that right-to-work laws weaken union bargaining power and firms reduce leverage in response. But high leverage is itself costly, and firms bear these costs only so long as the strategic benefit justifies them. When labor is organized, the incentive to maintain leverage as a credible threat is strongest.

Generative AI decreases the strategic value of debt because it targets a segment of the workforce that prior waves of automation largely bypassed. Over the past four decades, computerization and industrial robots displaced workers performing routine manual and cognitive tasks, contributing to a broad decline in the labor share of income (Autor et al., 2003; Karabarbounis & Neiman, 2014; Acemoglu & Restrepo, 2018, 2020, 2022). The affected occupations were concentrated in manufacturing and clerical work, where collective bargaining has historically been strongest. Generative AI operates on a different margin. It substitutes for non-routine cognitive and language-intensive tasks across professional occupations (Noy

& Zhang, 2023; Eloundou et al., 2024; Brynjolfsson et al., 2025), and white-collar professionals have unionization rates far below those of production workers.[13] When firms can replace these workers directly through technology, the indirect mechanism of threatening financial distress becomes unnecessary. [14] The leverage reduction should therefore concentrate among firms whose workers lack collective bargaining power.

A. WARN layoff notices

We first test whether 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 7 shows that adopting firms are 1.7 percentage points more likely to file a layoff notice, and Column (4) estimates an 8.6% 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 1.8 percentage points, a 47% increase relative to the sample mean of 3.8%.

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 not laying off workers because of financial distress. They are substituting generative AI for labor within continuing operations.

[Insert Tables 7 about here]

[13] In 2022, union membership in professional, scientific, and technical services was 1.6% and in finance 1.6%, compared with 8.0% in durable goods manufacturing and 20.0% in transportation and warehousing. These rates are largely unchanged in 2025 (1.7%, 1.2%, 7.6%, and 20.3%, respectively). Source: Hirsch et al. (2026), based on Bureau of Labor Statistics Current Population Survey data (https://unionstats.com/).

[14] Serfling (2016) documents a countervailing channel in which firing costs raise operating leverage, crowding out financial leverage. By reducing effective firing costs, generative AI should free capacity for more debt. That we instead observe deleveraging indicates the strategic debt and intangible capital channels dominate.

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. 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). Table 8 shows that the layoff response concentrates among high-exposure firms. Column (1) shows that adoption increases the probability of layoff filing by 2.1 percentage points among high-exposure firms, statistically significant at the 5% level. Column (3) estimates an 8.6% increase in the number of affected workers, significant at the 1% level. Low-exposure firms show no layoff response in Columns (4) through (6). Generative AI displaces knowledge workers, not production workers.

[Insert Table 8 about here]

Generative AI displaces workers in high-exposure occupations. If this substitution also drives the leverage reduction, high-exposure firms should deleverage more. Table 9 confirms this prediction. In Panel A, Column (1) shows that firms in high-exposure industries reduce leverage by 1.7 percentage points after generative AI adoption. The effect remains economically and statistically significant at 1.3 percentage points with controls in Column (2), a 3.9% decline relative to the sample mean, with a 21.4% increase in interest coverage in Column (4). Low-exposure industries show no leverage response in Columns (5) through (8). Panel B uses firm-level exposure, computed from each firm's occupation mix. Column (1)

shows a 3.3 percentage point reduction in leverage; the effect holds at 2.3 percentage points with controls in Column (2), 6.9% of the 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 in Column (6). Firms with greater labor substitution potential deleverage more after adopting generative AI.

[Insert Table 9 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, and sales per employee captures revenue per knowledge worker. All three measures are averaged over firm-quarters through 2022Q3, before ChatGPT's release made generative AI commercially available.

Table 10 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 following Belo et al. (2014). 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 10 about here]

D. Union bargaining power

Generative AI targets cognitive and language-intensive occupations that are, by their nature, largely non-unionized. In high-union industries, the workforce is concentrated in production and operational roles where collective bargaining preserves worker power and the strategic motive for holding debt persists. The leverage reduction should attenuate in these industries.

Table A10 tests this prediction using Current Population Survey union data fixed at ex-ante 2022 values. In Panel A, Column (3) shows that firms in low-coverage industries reduce leverage by 2.1 percentage points, and Column (4) shows a 23.6% increase in interest coverage, both significant at the 1% level. In high-coverage industries, Column (1) shows the leverage effect falls to 1.2 percentage points with marginal significance, and Column (2) shows the interest coverage effect is indistinguishable from zero. Panel B examines union membership and produces the same pattern. The leverage reduction does not disappear entirely in high-union industries, consistent with the intangible capital channel operating independently of labor bargaining power. But the attenuation is sharp. Generative AI targets workers who bargain individually; where collective representation persists, the strategic motive for debt survives.

5 Conclusion

This paper examines how firm-level Generative AI adoption affects corporate capital structure. Using LinkedIn job-posting data from Revelio Labs, we construct a novel measure of Generative AI adoption based on firms' hiring for Generative AI "integrator" roles (i.e.,

positions that build, deploy, or maintain Generative AI systems within the organization). Combining this measure with a staggered-DiD design, we show that Generative AI adoption is followed by economically meaningful deleveraging and improved debt-service capacity. Adopting firms reduce leverage and increase interest coverage, with stronger effects at higher levels of adoption intensity. The results are robust across alternative leverage measures, modern staggered-DiD estimators, and additional identification checks, and event-study evidence supports parallel trends.

We provide evidence consistent with two complementary mechanisms. First, Generative AI adoption shifts firms toward more intangible, less collateralizable capital and greater technology-related uncertainty, consistent with lower debt capacity and higher expected distress costs. This mechanism helps explain why deleveraging is stronger among firms with higher technology intensity. Second, Generative AI adoption is associated with greater workforce reductions, particularly in firms with higher ex ante Generative AI exposure, and deleveraging is stronger in low-union industries. These patterns are consistent with a labor-restructuring mechanism in which firms rely less on leverage as a strategic bargaining or disciplinary device when technology-enabled labor adjustment becomes more feasible.

Our findings contribute to the emerging literature on Generative AI and firm outcomes by documenting a financing-policy consequence of adoption. They also contribute to the capital structure literature by showing that the financing implications of technology adoption depend on the nature of the technology shock. In contrast to traditional automation settings, where leverage may rise as labor-cost rigidity declines, Generative AI adoption can induce deleveraging because it increases reliance on intangible capital and alters labor-adjustment incentives.

References

Acemoglu, D., Autor, D., & Johnson, S. (2026). *Building pro-worker artificial intelligence* (Working Paper). National Bureau of Economic Research.

Acemoglu, D., Kong, D., & Ozdaglar, A. (2026). *Ai, human cognition and knowledge collapse* (Working Paper). National Bureau of Economic Research.

Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, *108*(6), 1488–1542.

Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, *128*(6), 2188–2244.

Acemoglu, D., & Restrepo, P. (2022). Tasks, automation, and the rise in us wage inequality. *Econometrica*, *90*(5), 1973–2016.

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.

Baghai, R. P., Silva, R. C., Thell, V., & Vig, V. (2021). Talent in distressed firms: Investigating the labor costs of financial distress. *The Journal of Finance*, *76*(6), 2907–2961.

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

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.

Chava, S., Danis, A., & Hsu, A. (2020). The economic impact of right-to-work laws: Evidence from collective bargaining agreements and corporate policies. *Journal of Financial Economics*, *137*(2), 451–469.

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

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.

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.

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

Graham, J. R. (2000). How big are the tax benefits of debt? *The Journal of Finance*, *55*(5), 1901–1941.

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

Hart, O. (1995). *Firms, contracts, and financial structure*. Clarendon press.

Hart, O., & Moore, J. (1994). A theory of debt based on the inalienability of human capital. *The Quarterly Journal of Economics*, *109*(4), 841–879.

Hennessy, C. A., & Livdan, D. (2009). Debt, bargaining, and credibility in firm–supplier relationships. *Journal of Financial Economics*, *93*(3), 382–399.

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

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.

Karabarbounis, L., & Neiman, B. (2014). The global decline of the labor share. *The Quarterly Journal of Economics*, *129*(1), 61–103.

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.

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.

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

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.

Perotti, E. C., & Spier, K. E. (1993). Capital structure as a bargaining tool: The role of leverage in contract renegotiation. *American Economic Review*, 1131–1141.

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.

Singla, A., Sukharevsky, A., & Yee, L. (2025). *The state of AI: How organizations are rewiring to capture value* (tech. rep.). McKinsey & Company.

Stole, L. A., & Zwiebel, J. (1996). Intra-firm bargaining under non-binding contracts. *The Review of Economic Studies*, *63*(3), 375–410.

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.

Trajtenberg, M. (2018). *Ai as the next gpt: A political-economy perspective* (Working Paper). 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.

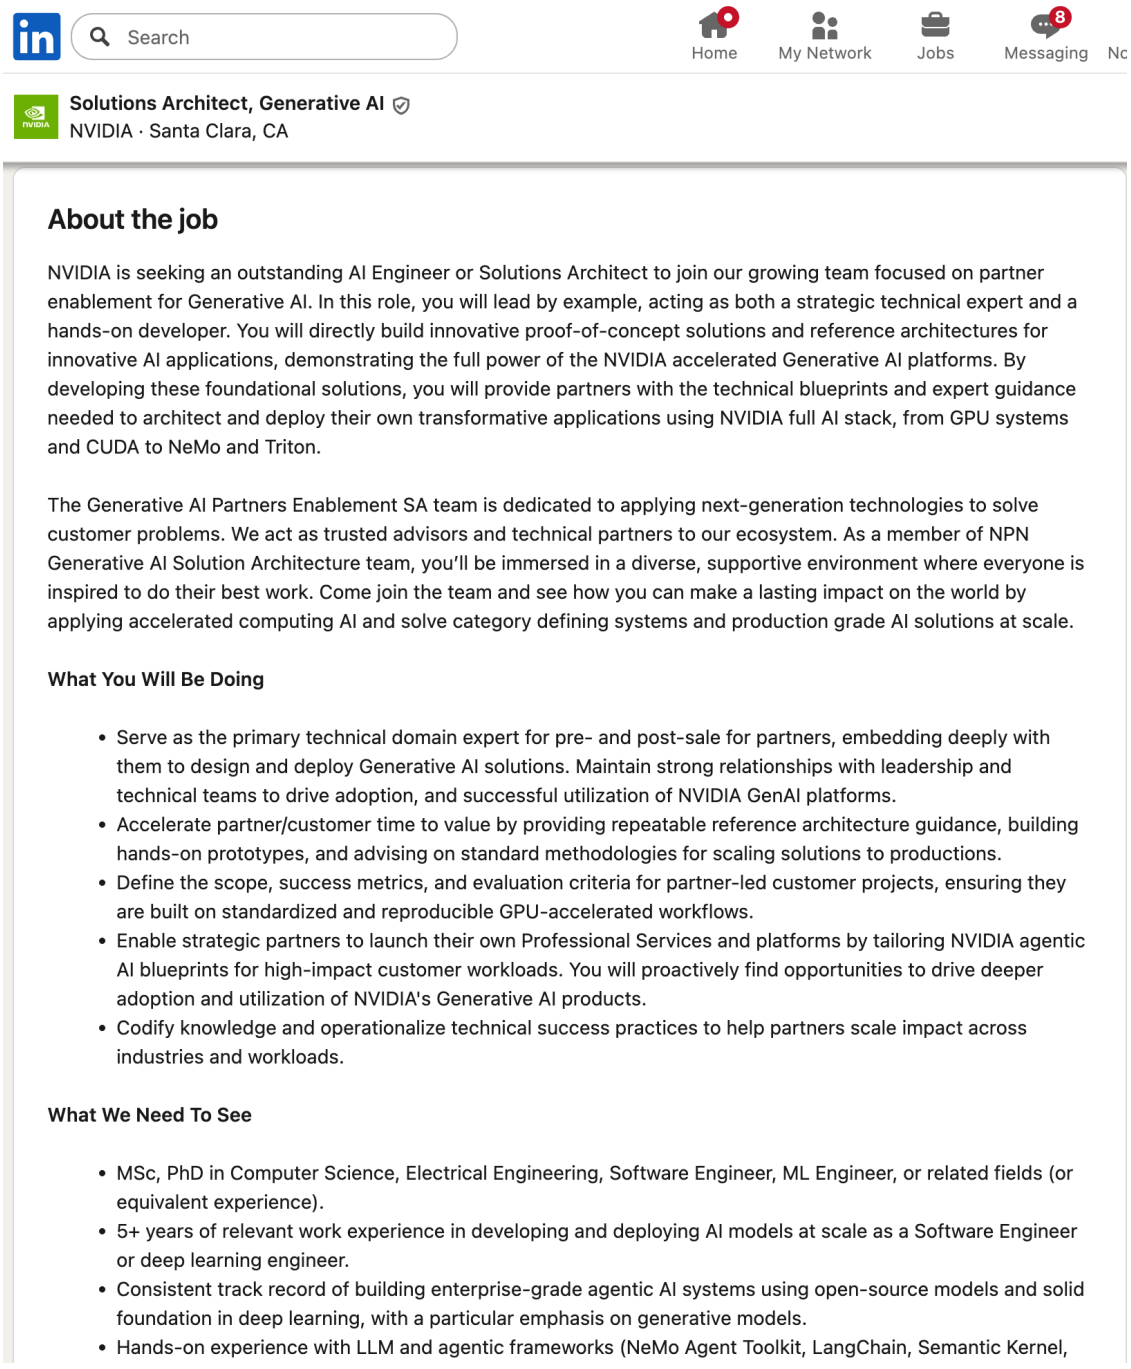


Figure 2: **Human Capital Investment in Generative AI and Job Posting Trends**

Figure 2 plots the quarterly human capital investment in generative AI and the share of generative AI postings among all technology job postings for Compustat firms from 2020Q1 through 2025Q2. Human capital investment is measured as the aggregate annualized salary allocated to generative AI job postings in a given quarter, expressed in millions of dollars. The Generative AI posting share is the number of generative AI postings divided by total technology postings within the matched Compustat sample. The vertical dashed line marks the first public release of ChatGPT (GPT-3.5) in November 2022. Generative AI job postings are restricted to integrator roles, defined as positions that involve building, developing, or deploying generative AI systems, identified through keyword filtering of job descriptions followed by LLM classification. Postings that mention generative AI solely as a general skill or tool for everyday use are excluded.

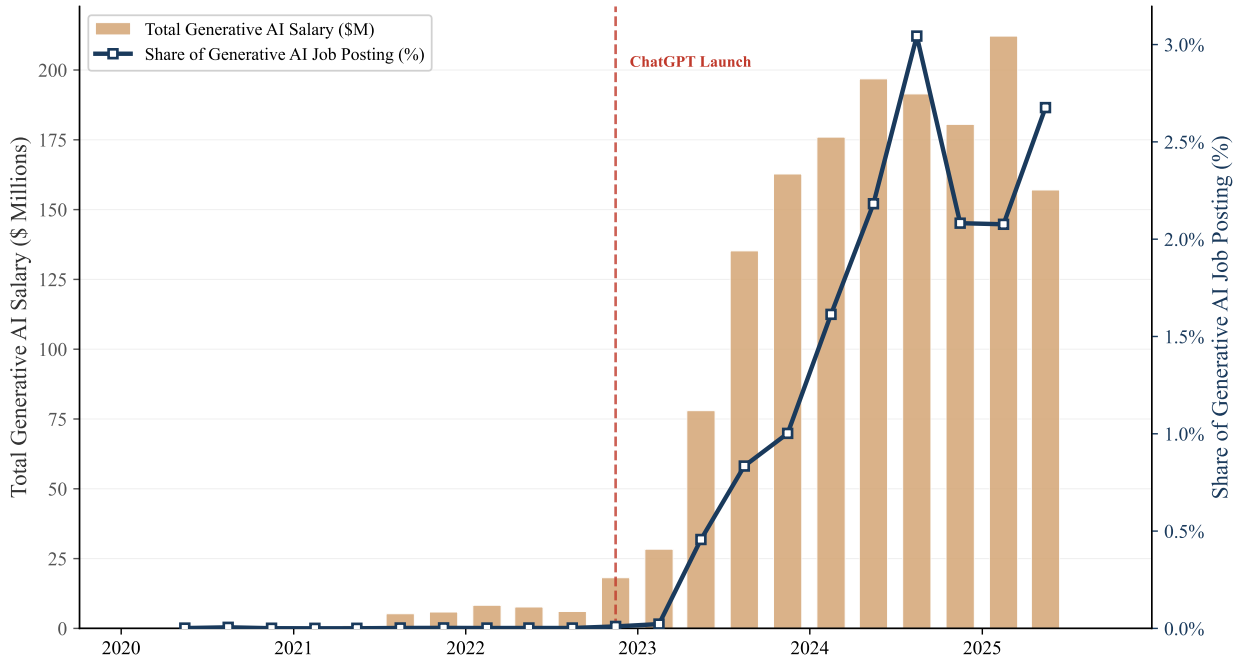
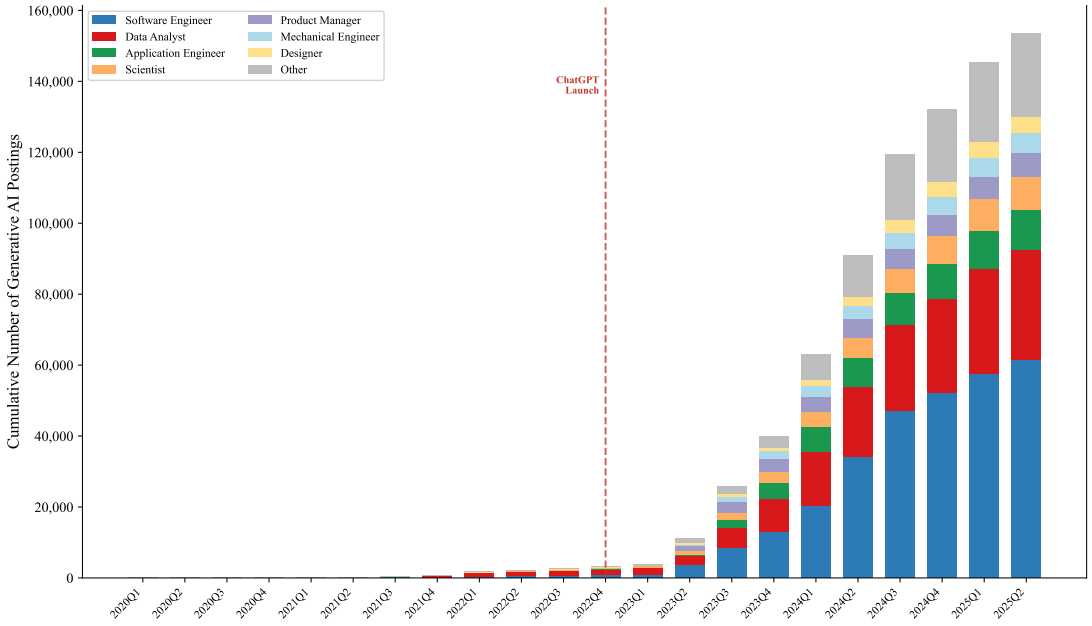
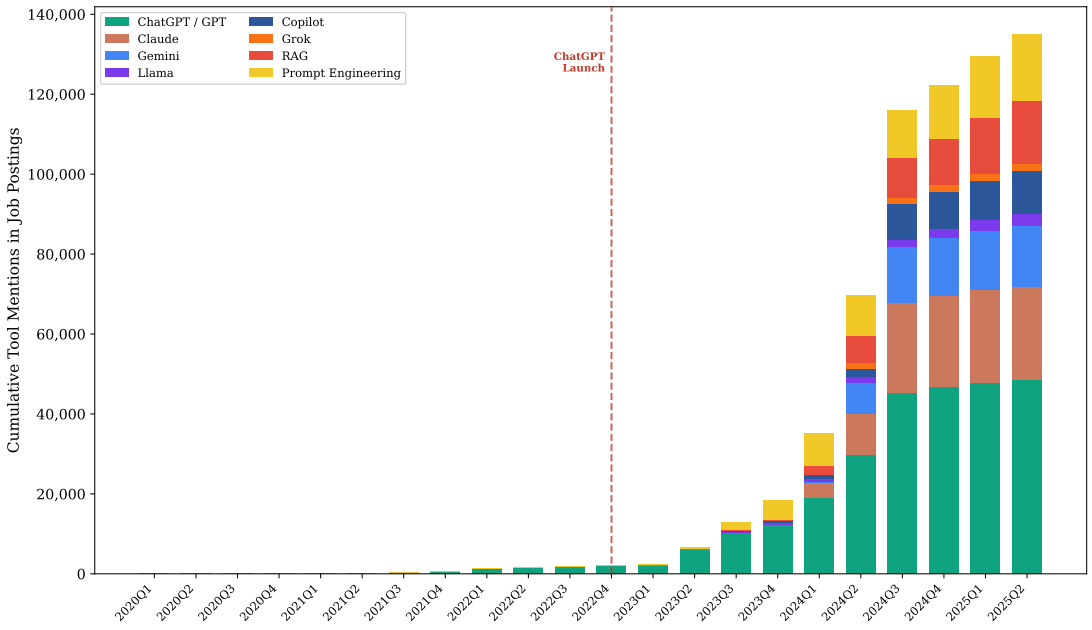


Figure 3: **Generative AI Human Capital Stock by Role and Tool**

Figure 3 plots the cumulative number of generative AI job postings by quarter from 2020Q1 through 2025Q2 for Compustat firms. 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 five difference-in-differences estimators robust to heterogeneous treatment effects and staggered adoption: traditional TWFE, 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 quarter fixed effects. Standard errors are clustered at the firm level.

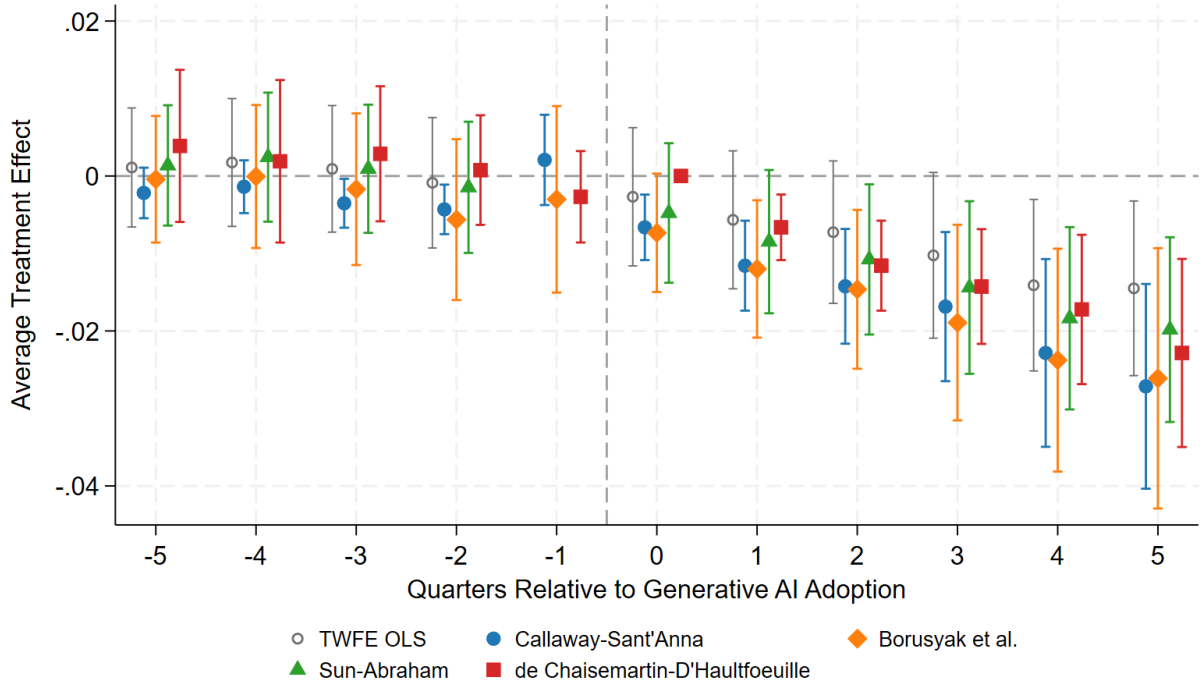


Figure 5: **Event Study: Generative AI Adoption and Interest Coverage**

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 five difference-in-differences estimators robust to heterogeneous treatment effects and staggered adoption: traditional TWFE, 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 quarter fixed effects. Standard errors are clustered at the firm level.

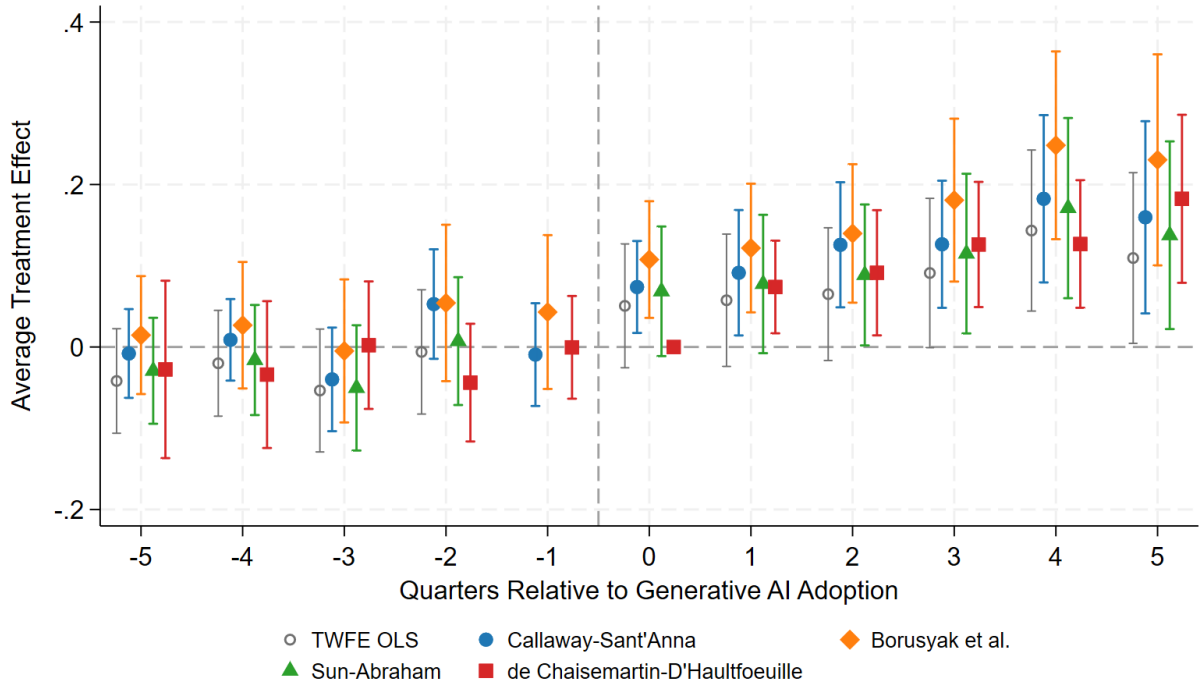
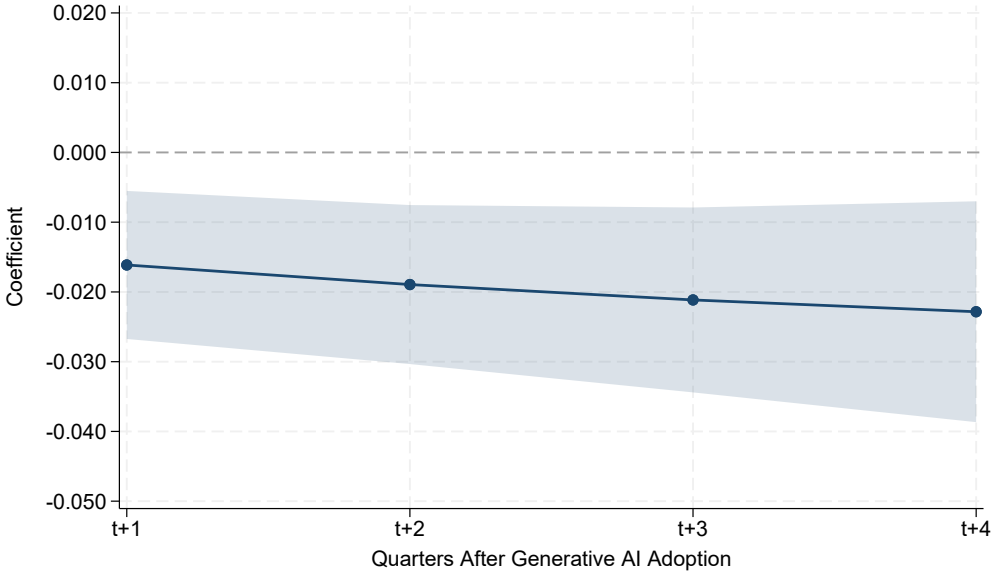
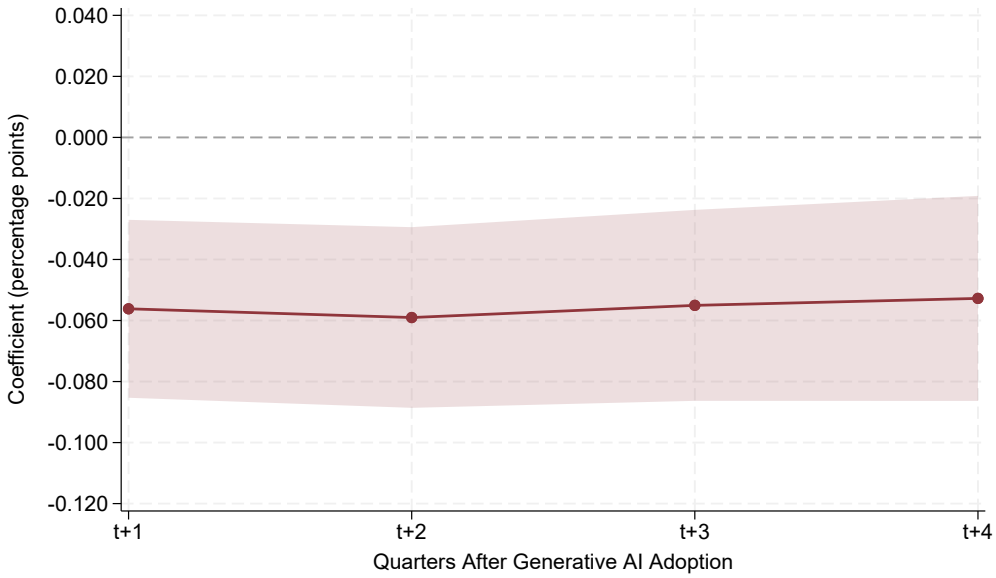


Figure 6: **Leverage Reduction: Debt Levels and Interest Expense**

Figure 6 plots coefficient estimates and 95% confidence intervals from regressions of future debt and interest expense on the Generative AI adoption indicator 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 following the first Generative AI-related job posting. All specifications include firm and year-quarter fixed effects. The dashed line marks zero. Standard errors are clustered at the firm level.



(a) Total Debt at $t+h$ / Total Assets at t



(b) Interest Expense at $t+h$ / Total Assets at t

Tables

Table 1: **Summary Statistics**

Table 1 reports summary statistics for the main analysis sample. The sample covers firm-quarters from 2020Q2 to 2024Q4, 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
ROA	41,007	−0.023	0.084	−0.034	0.003	0.017
Cash Ratio	41,007	0.225	0.242	0.046	0.122	0.330
R&D/Assets	41,007	0.018	0.032	0.000	0.001	0.022
Capex/Assets	40,934	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
IVOL (FF3)	40,872	0.554	0.451	0.298	0.445	0.668
IVOL (FF5)	40,872	0.547	0.449	0.293	0.438	0.660
Kurtosis	40,872	10.411	20.060	2.186	4.790	10.377
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	0.449	0.085	0.399	0.448	0.511
Generative AI Exposure (Industry)	41,007	0.060	0.023	0.045	0.049	0.069
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 TWFE 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 following the first Generative AI-related job posting. Columns (1) and (4) include firm and quarter fixed effects without controls. Columns (2) and (5) add firm-level controls. Columns (3) and (6) include industry-by-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 × Qtr 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 quarter fixed effects without controls. Columns (2) and (5) add firm-level controls. Columns (3) and (6) include industry-by-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 × Qtr 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. 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 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.

	First Stage	2SLS	
	Generative AI Adoption	*Book Leverage*	*Log(Interest Coverage)*
	(1)	(2)	(3)
*Panel A: O*NET Work Activities / LLM Task-Labels*			
LLM Capability Bartik	0.047***		
	(0.008)		
$\widehat{\text{Generative AI}}$ Adoption		−0.135**	2.376***
		(0.066)	(0.611)
Kleibergen-Paap F	36.5		
Adjusted R^2	0.466	−0.005	−0.251
*Panel B: O*NET Work Activities / LLM (Task + Domain)*			
LLM Capability Bartik	0.044***		
	(0.008)		
$\widehat{\text{Generative AI}}$ Adoption		−0.168**	2.559***
		(0.071)	(0.660)
Kleibergen-Paap F	33.1		
Adjusted R^2	0.466	−0.043	−0.302
N	40,832	40,832	40,832
Firm FE	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Table 5: **Channel Tests: Intangible Investment Intensity**

Table 5 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 following the first Generative AI-related 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
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 6: **Stock Return Risk**

Table 6 reports TWFE difference-in-differences estimates of the effect of Generative AI adoption on stock return risk measures. All measures are computed from CRSP daily returns over a trailing 252-trading-day window, winsorized at the 1st and 99th percentiles; the dependent variable is the one-quarter-ahead value. Log IVOL is the natural logarithm of annualized idiosyncratic volatility, computed as the standard deviation of residuals from Fama-French factor model regressions of daily excess returns, scaled by $\sqrt{252}$. Columns (1)–(2) use the three-factor model and columns (3)–(4) use the five-factor model. Kurtosis is the excess kurtosis of daily returns (normal distribution = 0), where higher values indicate fatter tails and more frequent extreme returns. Generative AI Adoption is an indicator equal to one for firm-quarters following the first Generative AI-related job posting. Odd-numbered columns exclude firm-level controls; even-numbered columns include Log(Assets), Tangibility, Market-to-Book, Profitability, and R&D. All specifications include firm and 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.

	Log IVOL (FF3)		Log IVOL (FF5)		Kurtosis	
	(1)	(2)	(3)	(4)	(5)	(6)
Generative AI Adoption	0.027**	0.034**	0.028**	0.035**	1.371**	1.492**
	(0.014)	(0.014)	(0.014)	(0.014)	(0.622)	(0.625)
Observations	40,861	40,861	40,861	40,861	40,861	40,861
Adjusted R^2	0.831	0.834	0.831	0.835	0.293	0.295
Firm-level Controls	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes

Table 7: **Generative AI Adoption and Layoffs**

Table 7 reports TWFE 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 following the first Generative AI-related job posting. Columns (1) and (4) include firm and quarter fixed effects without controls. Columns (2) and (5) add firm-level controls. Columns (3) and (6) replace quarter fixed effects with industry-by-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 × Qtr FE	No	No	Yes	No	No	Yes
Firm-level Controls	No	Yes	Yes	No	Yes	Yes

Table 8: **Generative AI Adoption, Layoffs, and Ex-Ante Generative AI Exposure**

Table 8 reports subsample TWFE difference-in-differences estimates of the effect of Generative AI adoption on layoffs, split by firm-level ex-ante Generative AI exposure from Eisfeldt et al. (2023). This exposure measure captures the share of a firm's pre-ChatGPT workforce in occupations whose tasks are susceptible to automation by large language models, based on O*NET task descriptions classified at the occupation level. Firms are split at the median of this cross-sectional exposure index, measured as of 2022Q3. Has WARN is an indicator for any WARN notice filing in quarter t. Mass Layoff (500+) is an indicator for layoffs affecting 500 or more workers. Log(1+Workers) is the log of one plus the number of workers named in WARN filings. Generative AI Adoption is an indicator equal to one for firm-quarters following the first Generative AI-related job posting. 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.

	High Generative AI Exposure			*Low Generative AI Exposure*		
	Has WARN	*Mass (500+)*	*Log(1+Workers)*	*Has WARN*	*Mass (500+)*	*Log(1+Workers)*
	(1)	(2)	(3)	(4)	(5)	(6)
Generative AI Adoption	0.021**	0.005***	0.086***	0.005	0.010	0.018
	(0.008)	(0.002)	(0.033)	(0.011)	(0.007)	(0.042)
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	20,505	20,505	20,505	20,502	20,502	20,502
Adjusted R^2	0.146	0.056	0.146	0.135	0.068	0.137

Table 9: **Ex-Ante Generative AI Exposure**

Table 9 reports subsample difference-in-differences estimates of the effect of Generative AI adoption on capital structure outcomes, split by ex-ante Generative AI exposure measured as of 2022Q3 (pre-ChatGPT). Panel A splits firms at the median of industry-level exposure. To construct this measure, we first score each O*NET task for exposure to large language models following Eisfeldt et al. (2023) and average task scores to the occupation level. We then compute each three-digit NAICS industry's exposure as the employment-weighted average of its constituent occupation scores, where employment weights come from the BLS Occupational Employment and Wage Statistics (OEWS) survey for May 2022. Panel B splits firms at the median of firm-level exposure, computed as the position-weighted average of occupation-level exposure scores within each firm's workforce using Revelio Labs position data as of 2022Q3. 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 following the first Generative AI-related 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: Industry-Level Exposure								
	High Industry Exposure				Low Industry Exposure			
Generative AI Adoption	−0.021***	−0.016**	0.225***	0.204***	−0.013	−0.008	0.110*	0.089
	(0.007)	(0.007)	(0.055)	(0.052)	(0.008)	(0.008)	(0.066)	(0.060)
Observations	24,814	24,814	24,814	24,814	15,777	15,777	15,777	15,777
Adjusted R^2	0.832	0.841	0.717	0.728	0.874	0.882	0.713	0.734
Panel B: Firm-Level Exposure								
	High Firm Exposure				Low Firm Exposure			
Generative AI Adoption	−0.033***	−0.023***	0.171***	0.166***	−0.011*	−0.008	0.138**	0.098*
	(0.008)	(0.008)	(0.061)	(0.058)	(0.006)	(0.006)	(0.055)	(0.051)
Observations	18,599	18,599	18,599	18,599	22,408	22,408	22,408	22,408
Adjusted R^2	0.829	0.840	0.746	0.753	0.866	0.874	0.673	0.697
Firm-level Controls	No	Yes	No	Yes	No	Yes	No	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 10: **Channel Tests: Labor Intensity**

Table 10 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 following the first Generative AI-related 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
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Appendix

Figure A1: **Monthly Generative AI Job Postings by Role**

Figure A1 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.

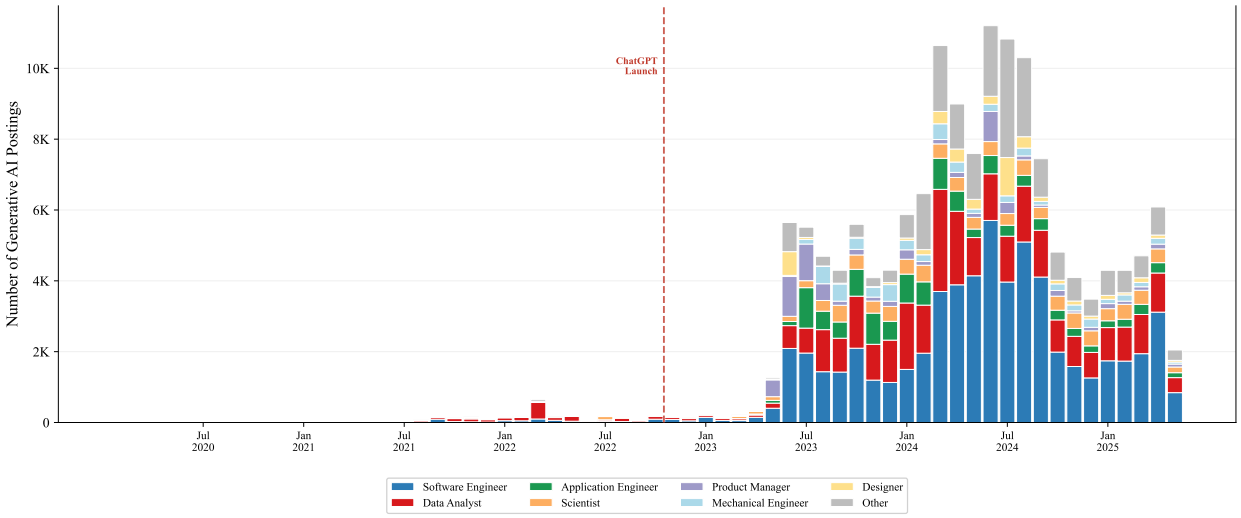
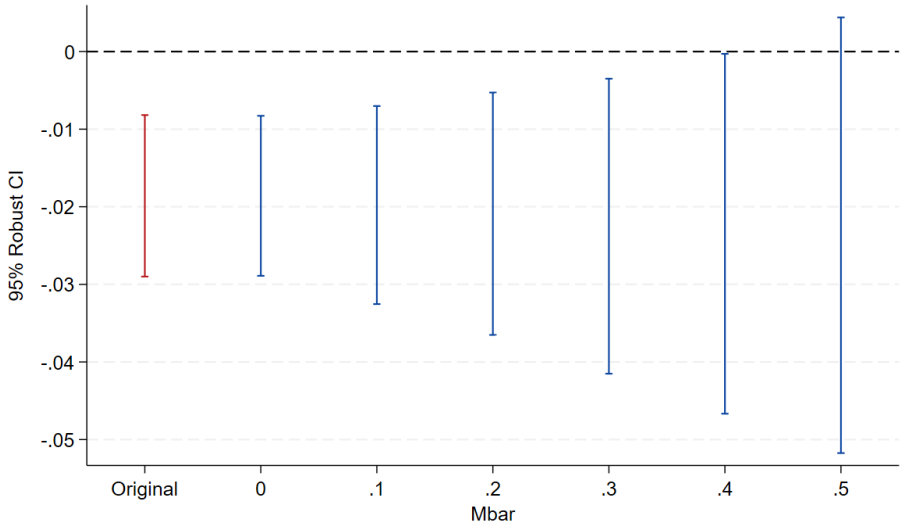
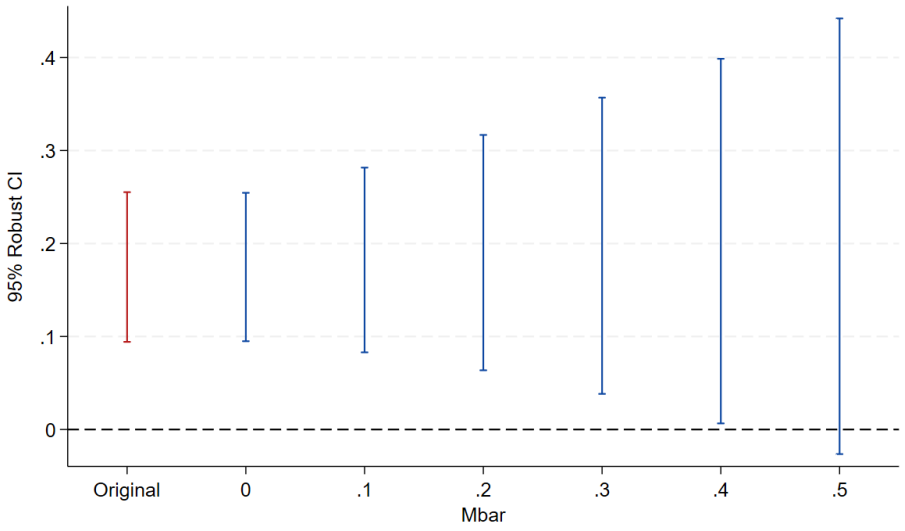


Figure A2: **Sensitivity to Violations of Parallel Trends**

Figure A2 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 A3: **Alternative Classification Models: Coefficient Estimates**

Figure A3 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-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.

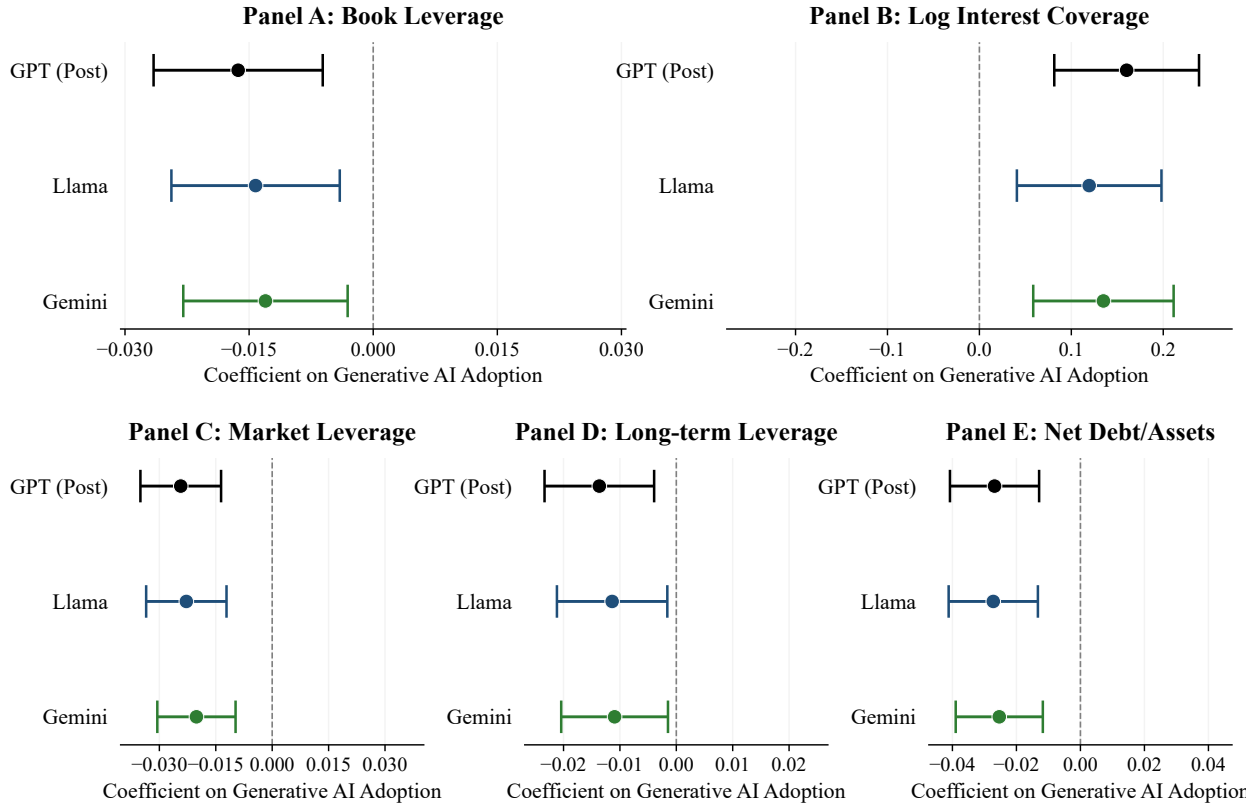
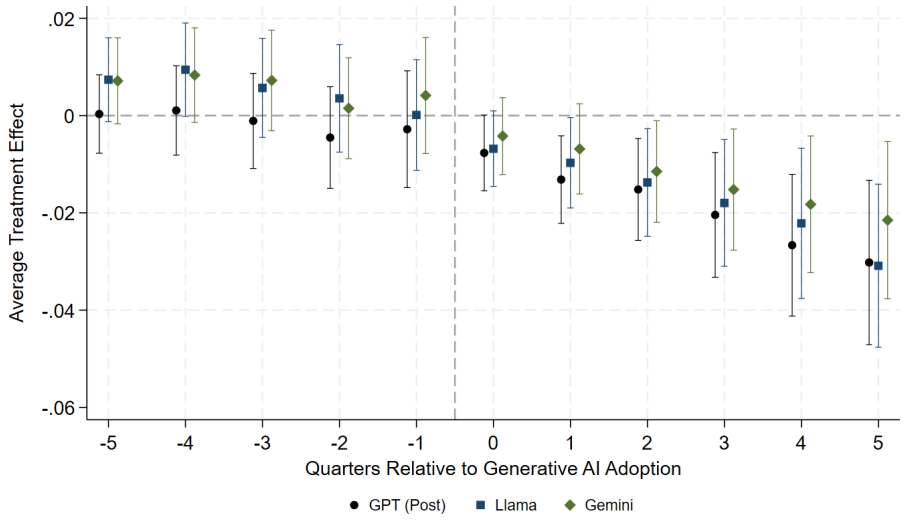
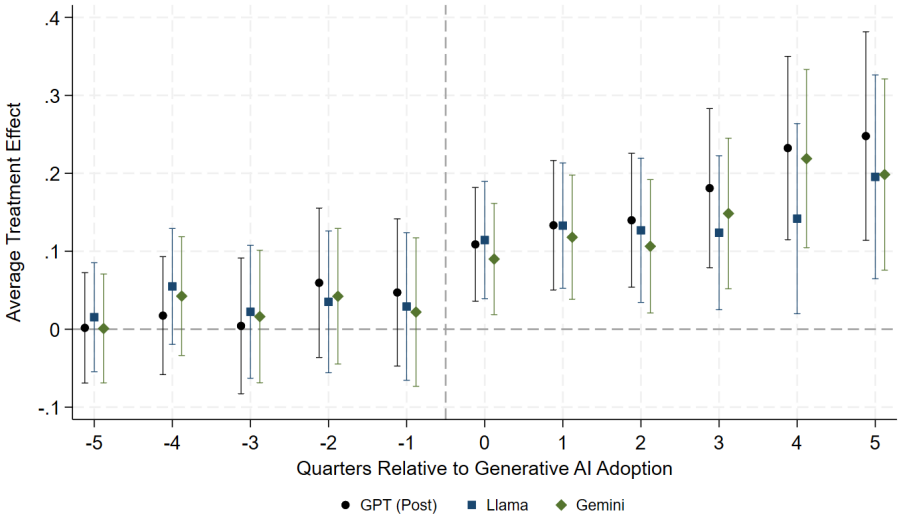


Figure A4: **Alternative Classification Models: Event Study**

Figure A4 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-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.



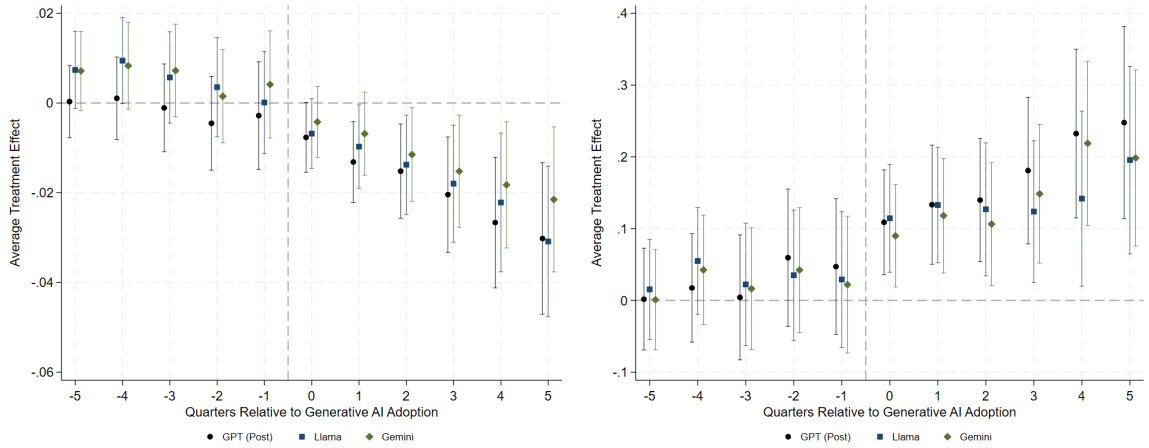
(a) Book Leverage



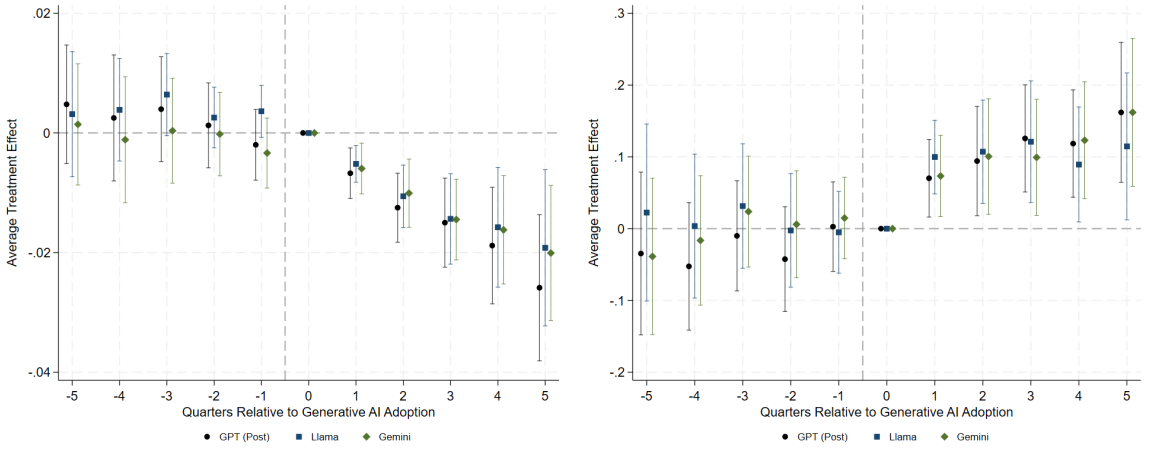
(b) Log Interest Coverage

Figure A5: **Alternative Classification Models: Event Study**

Figure A5 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-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)

Figure A6: **Robustness to Alternative Industry Definitions**

Figure A6 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 intensity), and standard errors clustered at the firm level. The dashed vertical line marks zero.

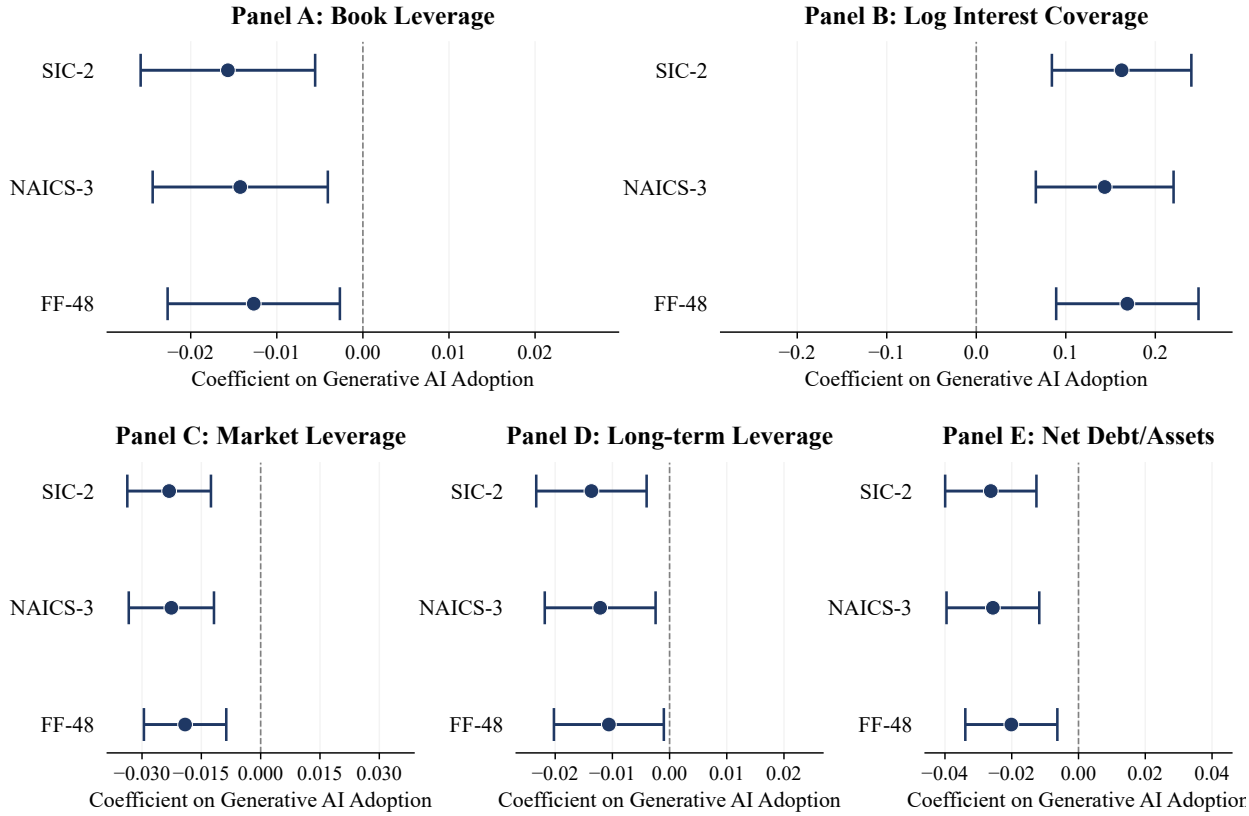


Table A1: **Variable Definitions**

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
Log IVOL	Natural logarithm of annualized idiosyncratic volatility, computed as the standard deviation of residuals from Fama-French factor model regressions of daily excess returns over a trailing 252-day window, scaled by $\sqrt{252}$.	CRSP
Kurtosis	Excess kurtosis of daily stock returns over a trailing 252-day window.	CRSP
Sales Growth	Quarter-over-quarter percentage change in net sales (saleq).	Compustat
ROA	Operating income before depreciation (oibdpq) divided by total assets.	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
R&D Intensity	Average R&D/Assets over all firm-quarters before 2022Q4.	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 (Industry)	Employment-weighted average of occupation-level exposure scores within each three-digit NAICS industry. Employment weights from BLS OEWS May 2022.	BLS OEWS, O*NET
GenAI Exposure (Firm)	Position-weighted average of 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)

Table A3: **Generative AI Adoption and Alternative Leverage Measures**

Table A3 reports TWFE 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 following the first Generative AI-related 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: Two-Way 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: Absorbing Industry-Time Variation						
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 × Qtr 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 A4: **Matching Estimators**

Table A4 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×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×quarter fixed effects, which absorb time-invariant firm heterogeneity within each cohort and cohort-specific aggregate time shocks. "Max Std. Diff." reports the largest absolute standardized difference across the six balancing covariates in the matched sample at event time −1. 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 × Qtr 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 × Qtr 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 A5: **Entropy Balancing: Covariate Balance**

Table A5 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-related job at any point during the sample period. Control firms never post a Generative AI-related 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 A6: **Synthetic Difference-in-Differences**

Table A6 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 following the first Generative AI-related 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 A7: **Pre-Treatment Placebo Test**

Table A7 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
Year-Quarter FE	Yes	Yes	No	Yes	Yes	No
Ind × Qtr 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 A8: **Bartik IV Robustness: Excluding Technology Firms**

Table A8 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)–(3) use the O*NET Work Activities / LLM Task-Labels alignment; columns (4)–(6) use the O*NET Work Activities / LLM (Task + Domain) alignment. Column (1) and (4) report first-stage estimates where the dependent variable is Generative AI Adoption. Remaining columns report second-stage 2SLS 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)*		
	First Stage	2SLS		First Stage	2SLS	
	Adoption	*Leverage*	*Int. Cov.*	*Adoption*	*Leverage*	*Int. Cov.*
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Excluding LLM Developers (8 firms)						
LLM Capability Bartik	0.045***			0.042***		
	(0.008)			(0.008)		
$\widehat{\text{Generative AI}}$ 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	0.460	−0.008	−0.271	0.460	−0.048	−0.325
N	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	0.029***			0.027***		
	(0.007)			(0.007)		
$\widehat{\text{Generative AI}}$ 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	0.438	−0.231	−0.226	0.438	−0.363	−0.319
N	37,191	37,191	37,191	37,191	37,191	37,191
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Table A9: **IV Estimates: AI Research Proximity**

Table A9 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 OA4). 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 × Year-Quarter FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Table A10: **Channel Tests: Union Heterogeneity**

Table A10 reports subsample TWFE difference-in-differences estimates of the effect of Generative AI adoption on leverage outcomes, split by industry-level unionization. Union data are from the Current Population Survey via Hirsch et al. (2026), measured at the Census Industry Classification (CIC) level and fixed at 2022 values to reflect pre-existing unionization conditions before ChatGPT's release in 2022Q4. Panel A splits at the median of the industry union coverage rate, defined as the percentage of workers covered by a collective bargaining agreement. Panel B splits at the median of the industry union membership rate, defined as the percentage of workers who are union members. 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 following the first Generative AI-related job posting. 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.

	Book Leverage	*Log(Interest Coverage)*	*Book Leverage*	*Log(Interest Coverage)*
	(1)	(2)	(3)	(4)
Panel A: Union Coverage				
	High Coverage		Low Coverage	
Generative AI Adoption	-0.012^{*}	0.076	-0.021^{***}	0.236^{***}
	(0.007)	(0.055)	(0.008)	(0.060)
Observations	20,115	20,115	20,042	20,042
Adjusted R^2	0.871	0.731	0.847	0.729
Panel B: Union Membership				
	High Membership		Low Membership	
Generative AI Adoption	-0.011^{*}	0.092^{*}	-0.021^{***}	0.219^{***}
	(0.006)	(0.055)	(0.008)	(0.060)
Observations	20,189	20,189	19,968	19,968
Adjusted R^2	0.872	0.726	0.846	0.732
Firm-level Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes

Table A11: **Leverage Reduction: Debt Levels and Interest Expense**

Table A11 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 following the first Generative AI-related 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
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table A12: **Profitability Mediation Test**

Table A12 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 (4) replicate the baseline from Table 2. Columns (2) and (5) add the one-quarter change in profitability. Columns (3) and (6) 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, 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	-0.013^{**}	-0.013^{***}	-0.012^{***}	0.157^{***}	0.158^{***}	0.151^{***}
	(0.005)	(0.005)	(0.005)	(0.041)	(0.041)	(0.040)
Δ Profitability (1Q)		0.264^{***}			-1.269^{***}	
		(0.047)			(0.311)	
Δ Profitability (4Q)			0.241^{***}			0.439^{*}
			(0.046)			(0.258)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-level Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	41,007	37,819	29,775	41,007	37,819	29,775
Adjusted R^2	0.855	0.867	0.894	0.734	0.742	0.766

Table A13: **Balance Sheet Composition and Firm Performance**

Table A13 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 following the first Generative AI-related 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
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes

Online Appendix

This appendix provides details on the construction of our two main variables. Section OA1 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 OA2 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.

OA1 Firm-Level Generative AI Adoption

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 Lichtinger and Maasoum (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.

OA1.1 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 publicly 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 OA1, 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[15] 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
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."
```

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

```
# 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" →often Cap Gemini (consulting company), NOT Google Gemini AI
- "RAG" in project management →Red-Amber-Green status reports, NOT retrieval-augmented
generation
- "Copilot" in aviation/non-tech →co-pilot or other software, NOT GitHub/AI Copilot
- "Claude" as person's name →NOT Anthropic Claude AI
- "GPT" in surveying/construction →Topcon GPT equipment, NOT language model
## Critical Rule: SKILL ≠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.[16] 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.[17] 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.

[16] 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).

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

OA1.2 Robustness

Robustness of Classification. We re-run classification using the original Lichtinger and Maasoum (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 OA1: **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

OA1.3 Examples

Table OA2: **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 list Generative AI only as a preferred skill.

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**.
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 OA3: **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.

OA2 Instrumental Variable Estimates

We report two instrumental variable strategies that exploit distinct sources of exogenous variation in firms' generative AI adoption. Section OA2.1 constructs a Bartik-style shift-share instrument from the evolution of frontier AI model capabilities. Section OA2.2 instruments adoption using geographic proximity to AI research universities.

OA2.1 LLM Capability Bartik

Our first instrument is a Bartik shift-share IV (Goldsmith-Pinkham et al., 2020) that exploits the interaction between sector-level occupation structures and the evolving capability profile of large language models. As new LLMs are released, their capabilities align more closely with certain occupation tasks than others. Sectors that concentrate 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. This requires a text representation for each side and a method to compare them.

On the occupation side, we draw from the O*NET 30.0 database, which provides structured text descriptions for each 6-digit SOC occupation. We use two types of text. Detailed Work Activities (DWA) describe specific workplace behaviors at a finer granularity than task statements; O*NET maps each occupation's tasks to a standardized set of DWA descriptions. Task Statements describe the duties performed in each occupation. Each panel in Table 4 varies which O*NET text source is used.

On the model side, we use the Epoch AI large-scale AI models database, which catalogs AI systems trained with substantial compute resources along with structured metadata including publication date, demonstrated capabilities, and application domain. We restrict to models classified in the Language or Multimodal domains published after 2010, yielding a cumulative set of models that grows each quarter as new systems are released. Each model record contains a Task field listing demonstrated capabilities (e.g., "question answering," "summarization," "code generation") and a Domain field indicating the application area. Each panel in Table 4 varies which Epoch AI text fields enter the alignment computation.

To compute alignment, we represent both occupation and model texts as TF-IDF vectors using a shared vocabulary fitted jointly on all texts, with 10,000 features, unigrams and bigrams, English stop words removed, and sublinear term-frequency scaling. For each quarter t, we average the TF-IDF vectors of all models published up to t to obtain a single vector $\bar{\mathbf{m}}_t$ representing cumulative LLM capabilities. The raw alignment between occupation s and the model frontier at quarter t 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 that alignment can only increase over time. This reflects the cumulative nature of AI capabilities: once a model demonstrates proficiency in a task domain, that capability persists in subsequent quarters. The constraint also smooths transient fluctuations in the similarity measure caused by shifts in the composition of newly released models.

B. Instrument Construction

The instrument for firm i in quarter t is:

$$Z_{it} = \sum_{s \in \mathcal{S}} \omega_s^{g(i)} \cdot A_{st} \tag{3}$$

where $\mathcal{S}$ is the set of 6-digit SOC occupations, $\omega_s^{g(i)}$ is the employment share of occupation s in firm i's NAICS sector $g(i)$ from the 2022 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 a firm's individual hiring decisions do not affect its instrument value. Cross-sectional variation comes from differences in sector occupation structures: a technology firm receives different shares than a manufacturing firm because the two sectors employ different occupation mixes. Time-series variation comes from A_{st} changing as new models are released and the alignment frontier expands differentially across occupations.

Each panel in Table 4 varies the text inputs to the alignment computation. 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 incorporate

the model's application domain alongside its task categories. The two panels test whether results are robust to the particular text representation used to measure alignment.

C. Exclusion Restriction

The exclusion restriction requires that the Bartik instrument affects capital structure only through generative AI adoption. Two features of the construction support this. First, the occupation shares come from the BLS at the three-digit NAICS industry level, not from individual firms. A firm's decision to hire more software engineers does not change its instrument value; 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 that the shares are predetermined relative to the outcome. This severs 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 determined by the global research community, not by any individual firm's adoption decision.

Firm fixed effects absorb permanent differences in the level of the instrument across sectors, and year-quarter fixed effects absorb aggregate movements in alignment common to all firms. The identifying variation is the differential rate at which AI capabilities become aligned with occupations in different sectors as new models are released. One remaining concern is that sectors whose occupations align closely with LLM capabilities may experience correlated trends in leverage from other technology-driven forces. The three text-representation panels in Table 4 provide a partial check: if results hold across different alignment measures, they are less likely to reflect a single spurious correlation embedded in one particular text combination.

OA2.2 Geographic Proximity to AI Research

Our second instrument exploits geographic variation in 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, which increases adoption through a labor supply channel. The instrument is cross-sectional: a firm's headquarters location relative to research universities reflects historical factors unrelated to its current financing decisions. 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.[18]

[18]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),

We count publications across leading AI and computer science venues and journals in the 2017–2022 window, which begins with the introduction of the transformer architecture and ends before the ChatGPT release.

Each publication–author–institution combination constitutes one observation. For each US institution, we observe the GRID identifier, which allows us to classify institutions by type using the Global Research Identifier Database. 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 in the country, but we exclude them because corporate publication activity may directly reflect firms' own AI investments, violating the exclusion restriction. Table OA4 reports the top-75 education institutions ranked by total paper count, along with 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:

$$Z_{it} = \log\left(\sum_{u \in \mathcal{U}} \mathbb{1}\left[d_{iu} \leq \bar{d}\right] \cdot R_u\right) \times \text{Post}_t \tag{4}$$

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 are sourced from SEC filings.

We construct two variants of R_u. The first counts only first-author papers at each institution, which proxies primary research contribution and reduces double-counting from large 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 any time-invariant effect of location, and industry-by-quarter fixed effects absorb time-varying industry shocks, so identification comes from within-industry geographic variation in proximity to AI research.

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.

C. Exclusion Restriction

The exclusion restriction requires that geographic proximity to AI research universities affects capital structure only through generative AI adoption. Two concerns merit discussion. First, proximity to research universities could proxy for local economic conditions or agglomeration effects. Firm fixed effects absorb permanent location characteristics, including local industry composition, access to venture capital, and historical ties to universities. The remaining variation is the differential effect of proximity after ChatGPT's release, which is difficult to attribute to pre-existing agglomeration patterns. Second, firms near AI universities might differ in unobservable ways that independently affect capital structure. Industry-by-quarter fixed effects address this by absorbing any industry-level trends correlated with geography. The identifying variation is therefore across firms in the same industry, headquartered at different distances from AI research centers, after versus before the release of ChatGPT.

Table OA4: **Top AI Research Institutions**

Table OA4 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			

OA3 Ex-ante Generative AI Exposure

OA3.1 Ex-ante Exposure Measurement

We construct two measures of generative AI exposure. First, following Eisfeldt et al. (2023), we build a firm-level measure: 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. Second, we create an industry-level measure by weighting the same occupation-level scores by national employment shares from the Bureau of Labor Statistics (BLS) Occupational Employment and Wage Statistics (OEWS) survey for May 2022.

A key consideration for our analysis is timing. 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 make two choices to ensure our measures are genuinely ex-ante. First, we draw task descriptions from O*NET Database Version 27.1, which was released in November 2022, shortly before ChatGPT became public on November 30. Second, we classify tasks using `gpt-3.5-turbo`, a model whose capabilities roughly match those of the original ChatGPT.

OA3.2 Occupation-level Exposure using Task Classification

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, score = 0.5), E3 (image-dependent, score = 0.0), or E0 (no exposure, score = 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.[19] In pilot runs, the model rarely assigned the E2 label without a worked demonstration, tending to classify ambiguous cases as either E1 or E0. We also randomize the order of few-shot examples across tasks to mitigate ordering effects. The complete prompts are available in the replication code.

Our classification yields replicable results: across two independent runs with identical

[19]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.

random seeds, label distributions are nearly identical.[20]

OA3.3 Firm-level and Industry-level Exposure

We aggregate task scores to occupations by simple averaging:

$$X_O = \frac{1}{|T_O|} \sum_{t \in T_O} X_t \tag{5}$$

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:

$$X_f = \sum_{o \in O_f} w_{fo} \cdot X_o \tag{6}$$

where w_{fo} is occupation o's share of firm f's workforce. Employment shares come from Revelio Labs position records as of 2022Q3.[21]

Industry-level exposure aggregates the same occupation scores using national employment shares:

$$X_k = \sum_o \frac{E_{ko}}{E_k} \cdot X_o \tag{7}$$

where E_{ko} is employment in occupation o within industry k, $E_k = \sum_o E_{ko}$ is total industry employment, and X_o is the occupation-level exposure score from Equation (1). Employment data come from the BLS Occupational Employment and Wage Statistics survey for May 2022.[22] We aggregate to the three-digit NAICS level, which provides 84 subsectors, comparable in granularity to the two-digit SIC codes standard in the finance literature.

OA3.4 Summary Statistics

Table OA5 reports summary statistics for our occupation-level generative AI exposure scores.

[20] Run 1: E0: 62.3%, E1: 22.4%, E2: 14.4%, E3: 0.9%; Run 2: E0: 62.2%, E1: 22.3%, E2: 14.6%, E3: 0.9%. As a robustness check, we classify O*NET 30.0 tasks using GPT-5.1; the resulting distribution (E0: 58.4%, E1: 26.3%, E2: 13.6%, E3: 1.6%) shows substantially higher direct exposure, reflecting the broader capabilities of frontier models.

[21] Revelio Labs, https://www.reveliolabs.com. We use position records as of 2022Q3, the quarter immediately preceding the ChatGPT release.

[22] Bureau of Labor Statistics, Occupational Employment and Wage Statistics, May 2022 National Industry-Specific Estimates, https://www.bls.gov/oes/2022/may/oes_nat.htm.

Table OA5: **Occupation-Level Generative AI Exposure**

Table OA5 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

Table OA6 reports generative AI exposure scores for each three-digit NAICS industry.

Table OA6: **Industry-Level Generative AI Exposure by Three-Digit NAICS**

Table OA6 reports industry-level generative AI exposure scores for each three-digit NAICS subsector. Each industry's score is the employment-weighted average of its constituent occupation-level exposure scores, where employment weights come from the BLS Occupational Employment and Wage Statistics (OEWS) survey for May 2022. Industry names are from the same source.

NAICS	Industry	Exp.	NAICS	Industry	Exp.	NAICS	Industry	Exp.
113	Forestry & Logging	0.031	337	Furniture & Related Product Mfg.	0.038	517	Telecommunications	0.092
115	Support Activities for Agriculture	0.031	339	Miscellaneous Manufacturing	0.054	518	Computing Infrastructure & Data	0.123
211	Oil & Gas Extraction	0.067	423	Wholesale: Durable Goods	0.085	519	Web Search, Libraries & Archives	0.118
212	Mining (except Oil & Gas)	0.030	424	Wholesale: Nondurable Goods	0.075	521	Monetary Authorities	0.112
213	Support Activities for Mining	0.037	425	Wholesale Agents & Brokers	0.109	522	Credit Intermediation	0.124
221	Utilities	0.063	441	Motor Vehicle & Parts Dealers	0.063	523	Securities & Commodity Contracts	0.125
236	Construction of Buildings	0.053	444	Building Material & Garden Equip.	0.066	524	Insurance Carriers	0.130
237	Heavy & Civil Engineering Constr.	0.033	445	Food & Beverage Retailers	0.054	525	Funds, Trusts & Financial Vehicles	0.118
238	Specialty Trade Contractors	0.036	449	Furniture & Electronics Retailers	0.083	531	Real Estate	0.083
311	Food Manufacturing	0.034	455	General Merchandise Retailers	0.063	532	Rental & Leasing Services	0.067
312	Beverage & Tobacco Mfg.	0.044	456	Health & Personal Care Retailers	0.081	533	Lessors of Intangible Assets	0.119
313	Textile Mills	0.046	457	Gasoline Stations & Fuel Dealers	0.059	541	Professional & Technical Services	0.110
314	Textile Product Mills	0.039	458	Clothing & Jewelry Retailers	0.075	551	Management of Companies	0.114
315	Apparel Manufacturing	0.045	459	Sporting Goods & Hobby Retailers	0.074	561	Administrative & Support Services	0.056
316	Leather & Allied Product Mfg.	0.033	481	Air Transportation	0.039	562	Waste Management & Remediation	0.040
321	Wood Product Manufacturing	0.028	482	Rail Transportation	0.025	611	Educational Services	0.075
322	Paper Manufacturing	0.037	483	Water Transportation	0.045	621	Ambulatory Health Care Services	0.069
323	Printing & Related Support	0.063	484	Truck Transportation	0.038	622	Hospitals	0.063
324	Petroleum & Coal Products Mfg.	0.052	485	Transit & Ground Passenger Trans.	0.038	623	Nursing & Residential Care	0.040
325	Chemical Manufacturing	0.060	486	Pipeline Transportation	0.059	624	Social Assistance	0.033
326	Plastics & Rubber Products Mfg.	0.037	487	Scenic & Sightseeing Transport.	0.037	711	Performing Arts & Spectator Sports	0.064
327	Nonmetallic Mineral Product Mfg.	0.036	488	Support Activities for Transport.	0.059	712	Museums & Historical Sites	0.068
331	Primary Metal Manufacturing	0.038	491	Postal Service	0.042	713	Amusement, Gambling & Recreation	0.046
332	Fabricated Metal Product Mfg.	0.042	492	Couriers & Messengers	0.040	721	Accommodation	0.041
333	Machinery Manufacturing	0.051	493	Warehousing & Storage	0.039	722	Food Services & Drinking Places	0.022
334	Computer & Electronic Product Mfg.	0.076	512	Motion Picture & Sound Recording	0.050	811	Repair & Maintenance	0.040
335	Electrical Equipment & Appliance	0.051	513	Publishing Industries	0.117	812	Personal & Laundry Services	0.043
336	Transportation Equipment Mfg.	0.040	516	Broadcasting & Content Providers	0.096	813	Religious & Civic Organizations	0.095