

Generative AI and Firm Values*

Andrea L. Eisfeldt[†] Gregor Schubert[‡] Bledi Taska[§] Miao Ben Zhang[¶]

November 18, 2023

Abstract

What are the effects of recent advances in Generative AI on the value of firms? Our study offers the first quantitative answer to this question for U.S. publicly traded companies based on the exposures of their workforce to Generative AI. Using Artificial Minus Human portfolios that are long firms with higher exposures and short firms with lower exposures, we show that higher-exposure firms earned excess returns that are 0.4% higher on a daily basis than returns of firms with lower exposures following the release of ChatGPT. Our novel firm-level measure of workforce exposure to Generative AI is validated by data from earnings calls, has intuitive relationships with firm and industry-level characteristics, and is robust to considering product-market exposures. Consistent with a labor-substitution effect, we show that hiring at more exposed firm shifts away from exposed occupations, and these occupations experienced relative wage declines. Consistent with a data-complementarity effect, we document a positive interaction effect on firm values for firms that have more exposed workforces as well as evidence of in-house data corpuses.

Keywords: Artificial Intelligence, Generative AI, Chat GPT, Corporate Valuations, Equity Returns, Technological Change

*We thank participants in the UCLA Anderson Finance Bag Lunch, KDD Finance Day, NYU Stern Finance Seminar, and the 2023 “Artificial Intelligence and the Economy” conference. for helpful feedback.

[†]UCLA Anderson School of Management and NBER, Email: andrea.eisfeldt@anderson.ucla.edu

[‡]UCLA Anderson School of Management, Email: gregor.schubert@anderson.ucla.edu.

[§]Lightcast, Email: bt540@nyu.edu.

[¶]USC Marshall School of Business, Email: Miao.Zhang@marshall.usc.edu

We construct the first firm-level measure of the exposure of firms’ workforces to being disrupted by Generative AI and show that US public firms in the highest quintile of exposed workforces increased in value by 4.8% relative to firms in the lowest quintile of labor-force exposures in the two weeks following the release of ChatGPT 3.5. Our results are not driven by industries with larger product market exposures, and occur during a period in which the overall market declined by over 4%. Indeed, there is substantial heterogeneity in returns *within* industries. The relative return differential is 3% within industries, and industry effects explain only 18% of the total firm-level variation in Generative AI exposures.

Thus, we argue that the release of ChatGPT 3.5 generated a substantial divergence in firm values across firms with higher and lower workforce exposures to Generative AI disruption. We show that the increase in firm values for firms with more exposed workforces is likely due to a labor-substitution effect. Hiring and wages decline at firms whose workforces face more exposure to Generative AI. On the other hand, we show that data assets are complementary to Generative AI labor-force exposure. Relative returns are higher for firms with both more exposed workforces and more evidence of data assets.

Recent advances in Generative Artificial Intelligence are widely seen as a major technology shock with important implications for firm values. Relative to earlier artificial intelligence models, Generative AI models can digest more complex inputs, and can produce human-like output, making Generative AI more versatile and scalable than prior innovations in AI and machine learning.¹ We document the effects of a key event in the advancement and dissemination of Generative AI technology in almost real time. While firms may progressively adopt the technology, the unmatched media attention and user base that ChatGPT has garnered within just months indicates that firms and investors are actively assessing the potential fast diffusion of this technology. We show that Twitter mentions of Generative AI and earnings call mentions by exposed firms increased substantially following the release of ChatGPT.

The labor effect of ChatGPT and related technologies on firm-level free cash flows could potentially work through two channels. First, firms whose labor force can be substituted for with cheaper Generative AI-based capital can experience higher free cash flows by lowering input costs.² Second, firms could generate higher cash flows due to the technological improvement in an input that is complementary to their workforce.³ We provide evidence

¹Recent studies of Generative AI include Eloundou, Manning, Mishkin, and Rock (2023) who study the impact of Generative AI on industries’ labor forces, Noy and Zhang (2023) who study the displacement effects of Generative AI on professional writing tasks, and Brynjolfsson, Li, and Raymond (2023) who study the effects of Generative AI on customer support agents, and Felten, Raj, and Seamans (2023) who consider heterogeneity in occupational exposure.

²See Autor, Levy, and Murnane (2003) and Zhang (2019) for measures of firm exposure to automation and Webb (2019) and Lane and Saint-Martin (2021) for the impact of AI on firms.

³See Krusell, Ohanian, Ríos-Rull, and Violante (2000) and Eisefeldt, Falato, and Xiaolan (2022).

consistent with Generative AI being, on balance, more of a substitute for labor inputs using large-scale data on job postings as well as occupation-level wages following the release of ChatGPT.

We measure firm-level exposure to Generative AI in three steps, starting with tasks and then aggregating first to occupations and then to firms. For task and occupation exposure, we build on Eloundou et al. (2023). We use ChatGPT itself to assess whether each of the 19,265 tasks currently performed by various occupations can be done by the current ChatGPT or by future ChatGPT after investment in additional capabilities. We aggregate the task-level exposure measures to the occupations using the O*NET database. Finally, we map occupations to publicly-traded firms using data from Revelio Labs. This dataset is constructed from millions of public employee profiles such as LinkedIn. Our firm-level exposure measure thus captures the ability of the tasks currently performed by labor at those firms to be performed (or made more efficient) by Generative AI. To the best of our knowledge, our study is the first to create a firm-level measure of exposure to Generative AI, and thus it is the first to document the valuation effects of Generative AI and to assign such valuation effects to a labor-input channel.

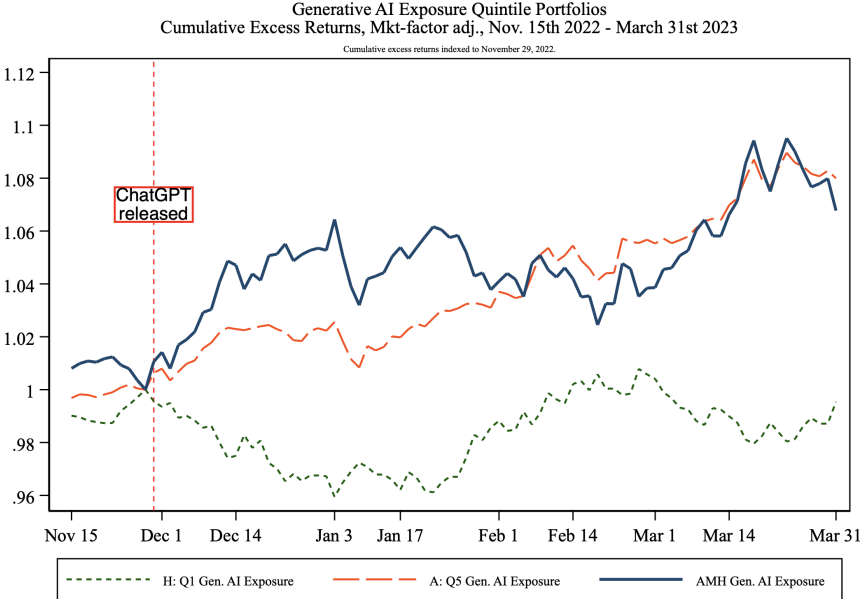
Our results for task-level exposures are intuitive. ChatGPT can code, but it can't rotate tires. To aid in biometric research ChatGPT can help, but it requires the add-on of custom software or data. For occupations, the most affected are those that involve non-routine cognitive tasks. This is in stark contrast with findings for prior waves of automation. Automation mainly displaces occupations involving routine tasks (Autor et al. (2003)). Indeed, the most Generative-AI affected occupations are those with a high share of non-routine cognitive analytical tasks or routine cognitive tasks, while manual physical tasks are relatively unaffected. Interpersonal tasks lie in between cognitive and manual tasks in terms of their exposure to Generative AI. Consistent with Generative AI being able to execute high-level cognitive tasks, we show that occupations with higher wages also have higher exposure to Generative AI.⁴ At the industry level, more exposed sectors have higher wages, consistent with those sectors employing more workers in higher-paid occupations that also tend to be more exposed to Generative AI.

The effect of the release of Chat GPT on firm values is very large. We document the returns of firms with higher and lower occupational exposure to Generative AI during and outside the two-week window following the release of ChatGPT on November 30, 2022. In this sense, our study utilizes an event study methodology combined with a difference in dif-

⁴Our result is consistent with recent findings by Kogan, Papanikolaou, Schmidt, and Seegmiller (2019), who find that technological advances impact workers at the higher end of the wage distribution. On the other hand, other studies (Krusell et al. (2000), Eisfeldt, Falato, and Xiaolan (2023)) document substitutability between low-skilled labor and capital but complementarity between high-skilled labor and capital.

ferences comparison across the returns of firms with higher and lower labor-force exposures. Figure 1 plots the returns of the highest-exposure quintile, the lowest-exposure quintile, and a long-short portfolio, which we denote AMH for “Artificial Minus Human”. The effects of Generative-AI exposure on firm value are substantial, with the highest-exposure quintile of firms earning positive excess returns of 28 basis points per day on average during the release period. The overall market declined during the release window, and the lowest exposure quintile experienced negative excess returns of 17 basis points per day. During the two week release period window, the AMH portfolio earned returns of 43 basis points per day, or a two-week return of nearly 5%.

Figure 1: Generative AI exposure quintile portfolio returns over time: market factor-adjusted. The graph shows the cumulative excess realized returns on portfolios based on value-weighted sorts. All portfolio returns shown are net of the risk-free rate. The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023. The figure shows returns adjusted for market factor exposure.



The AMH returns hold within industries. Comparing Generative AI-exposure terciles within industries we find that firms with workforces executing tasks that Generative AI can execute or assist with outperform firms with workforces less exposed to Generative AI by an average of 28 basis points daily. The fact that these strong effects exist *within* industries supports the idea that Generative AI can indeed have a broad impact on firm values through the effects on their labor inputs.

Firm-level exposure to Generative AI through workforce exposure has an intuitive relationship to firm characteristics. Exposure is positively related to Tobin’s Q, return on

assets, labor intensity, and the ratio of intangible or organization capital to total capital.⁵ By contrast, firms with higher asset tangibility have lower workforce exposure to Generative AI. Importantly, we also observe similar relationships between firms' exposure to Generative AI and firm characteristics within industry sectors.

We investigate whether firms with higher exposure to Generative AI operate technologies in which Generative AI is likely to be a complement to or substitute for existing labor inputs. Our results show that the initial effect is one of substitution. Firms with higher Generative AI exposure post fewer total jobs and, in particular, fewer jobs in occupations with high Generative AI exposure, after the release of ChatGPT. We use CPS data to analyze national trends in occupational wages and document the fact that wages have declined relatively more in the most exposed occupations since the release of ChatGPT. Indeed, wages in the top tercile of exposed occupations have experienced a decline of 1.2% since November 2022 relative to all other occupations.

One important concern with our labor-focused analysis is that the ChatGPT release period might happen to coincide with higher returns for firms that are benefiting from the AI boom through their *product market exposure* rather than through the impact of Generative AI on the productivity potential for the tasks done by their workers. While we agree that Generative AI is expected to impact expected corporate free cash flows through both product-market and a labor-market channels, we argue that the labor-market channel is likely to have a broader impact on firm values because it can impact firm productivity both across and within a broad array of industries. In addition, we show that our results are robust to several ways of controlling for product-market effects. First, we emphasize that industry effects explain only 18% of firm-level workforce exposures to Generative AI, and our return results hold for sorts conducted within industries. Our main results also hold excluding information-related industries, as well as controlling for product-market effects using measures based on textual analysis of firms 10-K filings and an ex-post classification by Goldman Sachs of firms that are direct beneficiaries of the AI boom to identify the firms with large product market impacts. We find that, as expected, greater product market AI exposure also predicts higher returns during our sample period. However, the returns to exposure in the form of task-based productivity potential are robust to allowing for these product market-related impacts. We also show that release period returns associated with standard firm characteristics used to explain the cross section of stock returns (and that may be correlated with both product and labor force exposure to Generative AI) do not explain our results. These findings suggest that there are two distinct channels for the impact of Generative-AI on firm values, namely labor productivity potential and product

⁵See Eisfeldt and Papanikolaou (2014) and Eisfeldt and Papanikolaou (2013).

market impact.

We further validate our labor-based measure of firms’ exposure to Generative AI by examining firms’ earnings call transcripts in recent years. We document a strong relationship between our measure of exposure to Generative AI and firms’ discussions of Generative AI and related technologies in firms’ earnings calls that only emerged following the release of ChatGPT and is not evident for discussions of other technical topics. Moreover, these findings remain even after we exclude all firms from the most IT-related sectors,⁶ suggesting that firms’ recent discussions about Generative AI go beyond its impact on related products, and extend to the impact on operations including labor inputs.

One interesting question is why we might expect incumbent firms to benefit from Generative AI, vs. Generative AI enabling displacement by new entrants. Firms that have a higher share of occupations exposed to Generative AI experience gains in value across a wide array of industries. At the same time, the effect of the release of ChatGPT on firm values varies widely across industries, as well as within industries across firms. Indeed, we find a significantly negative impact from the release of ChatGPT for some industries. Value losses for incumbents are consistent with the idea that for some industries Generative AI will lead to new entrants and displacement of existing firms. The fact that the overall impact of the arrival of ChatGPT on firms with more exposure to Generative AI is significantly positive is consistent with recent studies showing that it is increasingly more difficult for new entrants to displace incumbent firms.⁷ In addition, we argue that one advantage certain incumbents may have is internal corpuses of data that may be complementary to Generative AI.⁸ We show empirically that firms that have both higher labor-market exposures and are more likely to have large data assets experience larger increases in valuation following the release of ChatGPT. We argue that the complementarity between data and Generative AI that our results suggest may be one way to understand winners and losers from Generative AI within industries.

Our study contributes to the literature studying the impact of disruptive technologies on firm valuations.⁹ Papanikolaou (2011) and Kogan and Papanikolaou (2014) study the effects of investment-specific technological changes on asset prices. Zhang (2019) studies firms’ exposure to routine-biased automation. In a series of papers, Babina, Fedyk, He, and Hodson (2020), Babina, Fedyk, He, and Hodson (2021), and Babina, Fedyk, He, and Hodson

⁶To be precise, we exclude the NAICS 51 “Information” and NAICS 54 “Professional, Scientific, and Technical Services” sectors.

⁷See, for example, Gutiérrez and Philippon (2019) and Akcigit and Ates (2020).

⁸See Veldkamp and Chung (2019) and the cites therein for a summary of research on the value of data. Farboodi, Mihet, Philippon, and Veldkamp (2019) studies when data might benefit old vs. new firms.

⁹See Greenwood, Hercowitz, and Krusell (1997) for an early contribution on the long-run impacts of investment-specific technological change.

(2022) study the effects of AI on firm growth, compensation, and workforce composition. See also Webb (2019) for the impact of AI on firms, and Acemoglu, Autor, Hazell, and Restrepo (2022) for evidence on the impact of firm exposure to AI on hiring and skill demand over the 2010-2018 period. Kelly, Papanikolaou, Seru, and Taddy (2021) study firms' exposure to disruptive technological shocks using patent textual data, and Kogan et al. (2019) assesses worker displacement from technological change over a very long sample. These studies offer important insights into investors' and firms' responses to technological shocks in historical samples.

Our study departs from these works by focusing on measuring firms' exposure to Generative AI and assessing investors' reaction to the technology shock *upon its arrival*. We argue that the release of ChatGPT in November of 2022 is an observable, large technology shock. We also highlight our contribution of measuring investors' reactions to this shock in real-time. Indeed, the information in market prices can potentially inform employees' and firms' ultimate responses to technological disruption. Timely assessment of the market's expectations of Generative AI's impact on firms can also help policy makers to effectively evaluate regulatory policies in response to the arrival of the new technology.

A large literature has provided insights into the impact of technology on firms' employees. A central message is that technology has heterogeneous effects on workers of different skills or tasks (see, for some examples, Krusell et al. (2000), Autor et al. (2003), Acemoglu and Autor (2011), Acemoglu and Restrepo (2018), Eisefeldt and Papanikolaou (2013), Tuzel and Zhang (2021), and Kogan, Papanikolaou, Schmidt, and Seegmiller (2023)). Given the mature state of theoretical work on the labor and production effects of technological progress (including investment-specific technological change), we rely on the theoretical insights from this literature to put our empirical findings into context. In particular, our results are consistent with Generative AI advancement having a positive impact on firm values by offering a higher productivity or lower cost substitute for more exposed labor inputs. These results can be understood using the theoretical framework of Krusell et al. (2000), for example, which uses a CES production function with three inputs to estimate the production function parameters that govern input substitutability. To further illustrate ideas, we outline a simple theoretical framework based on the model in Acemoglu et al. (2022), which specifies a parsimonious version of the "task-based" production function in Acemoglu and Autor (2011) in the Internet Appendix.

While other contemporaneous or recent studies such as Eloundou et al. (2023) also address the exposure of occupations to Generative AI advances, our paper is novel in its contributions to the effect on *firms*. Our use of the Revelio Labs data to link occupations to firms yields

a unique opportunity to study corporate outcomes.¹⁰.

The paper proceeds as follows: Section I describes our data and measure of firms’ exposure to Generative AI, including descriptive facts about Generative AI exposures across occupations, industries, and firms. Section II presents our results documenting the substantial changes in firm valuations following the introduction of ChatGPT. Section III documents corporate communications to investors regarding Generative AI, and the relationship between those communications and our measure of Generative AI exposures. Section IV provides evidence that Generative AI is, overall, a substitute for labor. Finally, Section V concludes.

I. Data and Measurement

We measure a firm’s labor exposure to Generative AI in three steps, starting with task-level exposures, and then aggregating to occupation and firm-level exposures respectively. Our starting point is the Occupational Information Network (ONET) data, which offers detailed descriptions of 19,265 tasks and 923 occupations that utilize subsets of these tasks. To aggregate to firm-level exposures we use Revelio data to measure firm-level occupation shares. A summary of our three-step procedure, which we detail further below, is:

1. Ask Open AI’s GPT 3.5 Turbo model to assign each of the 19,265 tasks in ONET an exposure to Generative AI.
2. Aggregate from tasks to occupation exposures by summing over task-level exposures within the 923 occupations based on the occupation’s task statements from the O*NET database.
3. Aggregate to firm-level exposures by summing over occupation-level exposures according to each firm’s occupational mix using the firm-occupational employment data from the Revelio Labs Workforce Dynamics database.

A. *Measuring task exposure to Generative AI*

Occupational task data We consider an occupation to be a portfolio of tasks-to-be-done and evaluate whether each task can be done more productively using ChatGPT and

¹⁰Indeed, right before we released the first draft of this study, IBM, the company ranked #1 in our exposure to Generative AI measure among the largest U.S. firms announced it would halt hiring of 7,800 jobs that could be replaced by AI. See <https://www.businessinsider.com/ibm-halts-hiring-for-7800-jobs-that-could-be-replaced-by-ai-report-2023-5>

similar large language models (LLMs) or future applications that will be built based on their capabilities.

We obtain information on the tasks involved in each occupation from the O*NET V27.2 database, which provides a list of task statements created by practitioners or occupational experts.¹¹ A task statement is usually one sentence, and an occupation consists of 22 tasks on average. We code each of the 19,265 tasks as being exposed to Generative AI technologies or not using the task statement.

Task scoring We score tasks’ exposure to Generative AI following the approach that was suggested and validated by Eloundou et al. (2023) in their study of occupational exposures to Generative AI vs. Large Language Models without a human-like interface. In particular, we use ChatGPT itself to score exposures of tasks based on whether the task can already be done directly using the ChatGPT interface, or can be done with additional tools built on top of it. If a task can already be done directly using the ChatGPT interface it receives a score of one. Tasks that can be done by overlaying software, data or other tools on top of ChatGPT receive a score of one half. The remainder of tasks receive a score of zero. Two advantages of using an LLM instead of humans to assess task statements are that it allows for better replicability of the research in terms of cost and speed of execution, and rapid scaling of the method to the full set of 19,265 task statements.¹²

Specifically, we use Open AI’s GPT 3.5 Turbo model to classify the full set of task statements and validate its reliability on a smaller subsample of tasks.¹³ The model is given an overall rubric for scoring Generative AI exposure, as well as two example interactions between a user and an assistant that showcase the kind of output it is expected to produce. Then, a task statement is submitted together with its O*NET occupation title, and the model returns a score. The scores indicate whether ChatGPT-like tools can shorten the time to complete the task by at least half for the same quality as human completion.

Each task $T \in [1 : 19,265]$ is assigned an exposure X^T according to the following scoring system: $X^T = 0$ indicates no exposure. The current ChatGPT tool is either not yet useful for this task, or cannot execute the task due to intrinsic nature of the task, e.g. the task

¹¹This data can be accessed via the O*NET website at <https://www.onetonline.org>

¹²While similar large-volume classification tasks in the past often relied on crowd-workers on online platforms such as Amazon Mechanical Turk (MTurk), ChatGPT has recently been shown to outperform human crowd-workers in accuracy in text classification tasks, while also exhibiting lower variability in scores across multiple runs of the program (Gilardi, Alizadeh, and Kubli, 2023). Economists have also recently used other large language models to classify unstructured text from job postings and found that they outperformed other machine learning methods Hansen, Lambert, Bloom, Davis, Sadun, and Taska (2023).

¹³The structure of the prompt submitted to the Open AI GPT API is shown in Appendix A. See also Lopez-Lira and Tang (2023) for using ChatGPT to extract sentiment from news headlines and predict stock returns.

involves physical activities. $X^T = 1$ is applied if a 50% reduction in completion time is already feasible with the existing large language model interfaces. Finally, $X^T = 0.5$ is assigned if such a productivity gain is feasible, but only when the current capabilities of the model require further inputs (e.g., access to Internet or proprietary databases), or if the model is trained on domain-specific issues or data. Importantly, the model is asked not only to respond with the score but also to explain its reasoning. The explanations allow us to audit whether GPT is in fact understanding the prompt as intended and interpreting the task correctly. This auditing capability is another advantage of LLMs as it is oftentimes unavailable in instances of human text classifications.¹⁴ A random sample of scored tasks together with the model’s explanations can be found in Appendix Table D1.¹⁵

B. Measuring occupation exposure to Generative AI

Scoring occupations’ exposure to Generative AI We next aggregate tasks’ exposures to Generative AI to the occupation level. For each 8-digit Standard Occupational Classification (SOC) occupation from the O*NET, we calculate the share of the total number of tasks for each occupation that have either a direct or “plus-overlay” exposure to Generative AI. Our measure of occupation-level exposure X^O is the sum of task-level exposures X^T for $T \in [0 : 19,265]$ within each occupation $O \in [1 : 923]$ divided by the total number of tasks in occupation O . That is,

$$X^O = \frac{\sum_{T \in O} X^T \mathbb{1}_{T \in O}}{\sum_{T \in O} \mathbb{1}_{T \in O}} \quad (1)$$

We then aggregate across the 923 8-digit O*Net occupation codes to the 778 6-digit SOC codes to match the occupation-level exposure measure to firms’ occupational data. Note that the Generative AI exposure measure is bounded by 100% on the upper end. An occupational exposure of 100% implies that all tasks in that occupation can be done at least 50% faster with the already-existing functionality of ChatGPT and similar tools. On the lower end, 0% indicates that none of the tasks involved in the occupation are likely to be more productively executed by GPT now, or even after additional applications have been built on top of current Generative AI technology.

The full set of 6-digit occupations for which we compute Generative AI exposures consists of 778 occupations, of which 678 are also contained in the firm-level employment structure

¹⁴The model was also asked to return a confidence score (low/medium/high) for its prediction, but while this may have led the LLM to focus on refining its answer in this regard, we do not use this dimension of the response. In most cases, the model expresses “high confidence” in its assessment.

¹⁵We also perform consistency checks and show in Appendix A that GPT reliably provides classifications that are highly consistent across different runs.

data described below. The mean and median exposure in the latter set of occupations are 23% and 18%, respectively, with a standard deviation of 21%. The inter-quartile range of occupational exposure extends from 6% to 38%. Appendix Table D2 shows an overview of the 20 occupations with the highest and lowest Generative AI exposure scores.¹⁶ Note that among the highest exposure occupations, many, such as “Telemarketers”, “Computer programmers”, and “Interpreters and translators”, match closely with some of the key recent technological advances in Generative AI regarding its ability to hold natural text-based conversations, generate functioning code based on high-level descriptions of a programming task, and translate texts accurately between languages and styles.¹⁷

Occupations’ exposure to Generative AI differs substantially from their exposure to the types of automation studied in the prior literature. Acemoglu and Autor (2011) summarize the literature and characterize occupations’ skills by routine vs. non-routine, cognitive vs. manual, and analytical vs. interpersonal aspects. To understand how their categories map onto occupations’ exposure to ChatGPT-like technologies, we regress our 6-digit occupation Generative AI exposure measure on the set of occupational skill scores ($Skill_S^O$) defined by Acemoglu and Autor (2011). Specifically, we run the following regression:

$$X^O = \alpha + \sum_S \beta_S * Skill_S^O + \varepsilon^O.$$

Figure 2 shows the results. We find that occupations with higher Generative AI exposure are more likely to involve non-routine cognitive analytical skills or routine cognitive skills, and less likely to involve different kinds of manual skills, or interpersonal skills. These results demonstrate the unique features of Generative AI’s disruption to firms’ employees. The literature on previous waves of automation shows that the technology displaces routine-task jobs (i.e., repetitive and codifiable tasks) and complements cognitive-task jobs ((Autor et al., 2003)).

Different from computers, tools like ChatGPT can interpret and respond to relatively unstructured inputs, display a surprising amount of common sense in filling in gaps in instructions, and can respond with relatively complex outputs, such as texts in different styles, or – in the case of image generators like Stable Diffusion or Midjourney—even with new and

¹⁶Note that the lowest score category only shows a subset of a larger set of occupations with zero Generative AI exposure.

¹⁷While “Mathematicians” might seem out-of-place among occupations with high exposure to Generative AI, note that Fields Medal winner Terence Tao of UCLA remarked in April 2023, that “GPT-4 has saved [him] a significant amount of tedious work”, noting that “while these AI tools do not directly assist [him] in core tasks such as trying to attack an unsolved mathematical problem, they are quite useful for a wide variety of peripheral (but still work-related) tasks (though often with some manual tweaking afterward).” (Source: <https://pandaily.com/mathematician-terence-tao-comments-on-chatgpt/>)

original images (Bubeck, Chandrasekaran, Eldan, Gehrke, Horvitz, Kamar, Lee, Lee, Li, Lundberg, et al., 2023). As a result, this wave of technological change may differ from previous waves in that many tasks in non-routine cognitive analytical jobs that were safe from automation by previous technologies are now suddenly more likely to be substituted for by software and computers.¹⁸

Figure C1 plots the relationship between Generative AI exposure and wages by major occupation group and across industries. These figures show that the major occupation groups and industries with higher exposures to Generative AI have higher wages. The names of the occupational groups and industries in combination with their exposure scores and associated wages are intuitive. That Generative AI will impact higher-wage professions and industries is consistent with the finding in Kogan et al. (2023) that exposure to labor-augmenting technologies is associated with earnings declines for workers that are older and/or white collar.

C. Measuring firms’ exposure to Generative AI

To estimate a firm’s exposure to Generative AI, we use data on firms’ occupational structure to aggregate our occupational exposure measure to the firm level. We obtain data on firms’ occupational employment from Revelio Labs, which collects information on job titles and employers from LinkedIn and other resume profiles and constructs occupation-by-firm employment counts. Our customized data define a firm at the unique Compustat identifier *gvkey* and define an occupation using the 6-digit SOC. We use the employment counts for each *gvkey*-SOC6 as of March 2022, which is the latest month in our data.

We construct a firm’s Generative AI exposure as the weighted average of its occupations’ Generative AI exposure, using the firm’s occupational employment as weights. That is, each firm *f* in Compustat’s labor-force exposure to Generative AI is computed as:

$$X^f = \sum_{O \in f} EmpShare_{f,O} * X^O \quad (2)$$

¹⁸To better understand which occupation characteristics are associated with higher exposure, Panel A of Appendix Figure C1 shows the relationship between the average wage level of each 2-digit major occupation group in 2021 and our estimated Generative AI exposure. As the graph shows, higher-wage occupations are generally more likely to be exposed to ChatGPT-like technological advances making their constituent tasks more productive. One notable exception is the relatively low-wage “Office and Administrative Support” occupations, which are also predicted to be highly exposed. The positive relation between Generative AI exposure and wage is also observed at the industry level. In Panel B of Appendix Figure C1, we aggregate occupational Generative AI exposure and occupational wages to the NAICS 2-digit industry level using the 2021 BLS National Employment Matrix. We observe a similar pattern that high-wage industries also tend to have high exposure to Generative AI.

where $EmpShare_{f,O} = \frac{emp_{f,O}}{emp_f}$ is the employment share of occupation O in firm f . The result of this procedure is a cross-section of 2,518 publicly-traded firms with predicted exposure to Generative AI and basic company characteristics from Compustat.

D. Summary statistics of firms' Generative AI exposure

Table I shows the summary statistics of our sample which includes 2,518 publicly traded firms in the cross-section of 2022. On average, the firms in our sample have a mean and median task exposure score of 35%, with a standard deviation of 8%. The 10th and 90th percentiles of firm-level exposures to Generative AI are 27% and 44% respectively, implying that labor exposure to Generative AI is indeed a broad phenomenon.

Table II lists the 15 firms with the highest and lowest exposure to Generative AI, respectively, among the top 100 largest U.S. firms by market capitalization in 2022. While many IT firms, such as IBM and Intuit, not surprisingly have a large fraction of employees exposed to Generative AI, we also observe manufacturing firms, such as 3M, and administrative conglomerates, such as S&P Global, in this high-exposure category. The large U.S. firms ranked at the bottom of the exposure distribution include restaurants, such as Starbucks and McDonald's, retail firms, such as Target and Walmart, transportation firms, such as UPS, and manufacturing firms, such as Tesla, suggesting that they have a smaller fraction of employees exposed to Generative AI.

Perhaps surprisingly, Figure 3 shows that the variation in firm-level exposure to Generative AI has a significant within-industry component. While industry sectors such as "Information" and "Professional, Scientific, and Technical Services" have an average firm exposure to Generative AI of about 13% greater than industries such as "Accommodation and Food Services," there is substantial variation of firms' exposure to GPT within each industry. A variance decomposition shows that industry differences explain only 18% of the firm-level variation in exposure to Generative AI. The rich within-industry variation in firms' exposure to Generative AI motivates us to explore our firm-level empirical analyses within industries. Importantly, the within-industry analyses distinguish our labor-based mechanism from the explanations based on products' exposure to Generative AI when studying the effects of firms' exposure to Generative AI on firm values.

How do firms' exposures to Generative AI relate to other firm characteristics? Panel A of Table III shows that firms with higher exposure to Generative AI tend to be smaller, have greater Tobin's Q , and are less profitable. These findings are consistent with the notion that such small and high-growth firms tend to focus their workforce on cognitive tasks. Moreover, we also observe that firms with higher labor intensity, higher organizational capital ratio

(Eisfeldt and Papanikolaou, 2013) and less tangible capital are more exposed to Generative AI. We will thus control for these firm characteristics in our later analyses. Panel B of Table III shows that the above associations between firms’ Generative AI exposure and firm characteristics are mainly a cross-industry feature but not within-industry, except for the associations with labor intensity and asset tangibility. Hence, other firm characteristics may cast a less confounding concern for our within-industry analyses later.

II. Stock Market Impact of Generative AI Exposure

A. *Realized returns and Generative AI exposure*

We identify the impact of Generative AI on firm value by comparing the differential stock returns for firms with high and low Generative AI exposure during the ChatGPT release period vs. during the non-release period. Our event study is thus an event-study version of a difference in differences-style analysis. We use the non-release period, as well as standard asset-pricing controls and firm characteristics to control for confounding information that may also differentially affect the cross section of firm returns. We form an “Artificial Minus Human” or “AMH” portfolio that is long firms in the highest quintile of exposures, and short firms in the lowest quintile of exposures and show that this portfolio has substantial abnormal returns within the release period but not outside it. Our findings are robust to the several ways we control for confounding information and other known drivers of the cross section of stock returns. Additionally, the ChatGPT release period is one in which the S&P 500 had *negative* returns of 4.4%. Finally, note that our study focuses on a cash-flow effect from labor-market exposure to Generative AI. While there may also be a discount rate effect, such an effect presumably requires a longer post-release sample to uncover it.

Defining the ChatGPT Release period ChatGPT 3.5 was released on November 30, 2022. This release was followed by a considerable spike in attention from the general public as well as investors that lasted over two weeks. We define the “Yes ChatGPT release period” as the trading dates starting with November 30, 2022 and ending on December 14, 2022. We define the rest of the days from December 15, 2022 to March 31, 2023 as the “Not ChatGPT release period”. Figure 4 shows that Twitter mentions climbed subsequently, peaking about a week after the release. Section III provides further support of for the increase in investor attention during the “Yes ChatGPT release period”. We choose a two-week window as a balance between allowing enough time for information to be digested and incorporated while not allowing so much time that other shocks affecting firm values become too much of a

confounding factor.

Forming Generative AI exposure portfolios To estimate whether Generative AI exposure affects the realized returns of stocks during the event window, we first form value-weighted high and low exposure quintile portfolios, and also a high-minus-low portfolio—which we refer to as the “Artificial Minus Human” (*AMH*) portfolio—that represents the zero net investment portfolio long high exposure (*H*) stocks and short low exposure (*L*) stocks. In particular, at the beginning of our sample period, we sort stocks traded on the NYSE exchange into five quintiles based on their Generative AI exposure, and we use the NYSE breakpoints to assign non-NYSE stocks into the quintile bins. We compute the value-weighted daily returns of each portfolio as the average daily returns of stocks in the portfolio weighted by their previous day’s market capitalization.¹⁹

Portfolio sorting results Panel A of Table IV shows the mean excess realized returns of the quintile portfolios sorted by firms’ Generative AI exposure and also the long-short AMH portfolio during the “ChatGPT release period” and the “Not ChatGPT release period.” We observe that the *AMH* high-minus-low exposure portfolio has positive daily returns of 0.43% (t -statistic = 3.54) on average during the two weeks post-ChatGPT release. However, outside this ChatGPT release period, we observe a negligible average return of 0.008% for the AMH portfolio.²⁰ The AMH portfolio’s average return during the ChatGPT release period is also statistically different from its average return outside the period, which is close to zero. This key finding—that companies with higher employee exposure to Generative AI show significantly higher average returns only during periods when advances in large language model technology became known but not in other periods—helps quantify the impact of Generative AI breakthroughs on firm value.

In Panel B of Table IV, we additionally control for each portfolio’s exposure to the market factor. The market adjusted alpha of the AMH portfolio shows very similar results as the returns of the AMH portfolio, suggesting that our main finding is not driven by firms’ heterogeneous exposure to the market. In Panel C, we further consider whether the excess returns on the Generative AI-exposed stocks after the release of ChatGPT can be explained by an association between this exposure and other risk factors. We construct the alphas of each portfolio using the Fama-French 5-factor model (Fama and French (2015)). Again,

¹⁹See Appendix Section A for more details on the portfolio construction.

²⁰In Appendix Table D5 we consider using the two weeks that occurred 1 month before ChatGPT’s release date (Oct. 31, 2022 - Nov. 15, 2022) as a placebo release period and find that the AMH portfolio has an estimated average return of 0.19% per day during this period, and its returns are not statistically significantly different from zero in any of the specifications that are analogous to those in Table IV.

we observe similar results: the AMH portfolio’s 5-factor adjusted alpha is 0.31% during the ChatGPT release period but -0.03% outside this period, and the difference between the two periods is statistically significant. We also consider whether the precise way that we construct our exposure measure in Equation 1 matters for our results: in Appendix Table D6, we show that these results are qualitatively very similar when we define occupation-level exposure either only as the share of tasks with “direct” (rubric 1) exposure without considering the more speculative productivity gains available through future applications, or when we weight task exposure by the “importance” rating of the task within each occupation according to O*NET.

Time series of Generative AI exposure portfolio returns To further visualize that the Generative AI exposure returns are associated with the advances in the related technology, we plot the time series of daily excess returns of the top and bottom quintile portfolios and also the AMH long-short portfolio in Figure 1. We observe that while the firms with low exposure to Generative AI had *negative* returns of 4% during the ChatGPT release period, firms with high exposure to Generative AI experienced a positive return of about 2%. Put together, the AMH long-short portfolio generates a positive 6% cumulative return during the ChatGPT release period. We note that the overall market actually declined substantially during this two-week window, consistent with the performance of the “Human” leg of our AMH portfolio.²¹ We extend the cumulative returns to the end of our sample and observe that the positive cumulative return of the AMH portfolio is sustained following the ChatGPT release period and reaches about 7% on March 31, 2023.

B. Firm characteristics and returns to Generative AI exposure

We next investigate whether our finding of increased value for companies with greater task-based exposure to Generative AI is confounded by our measure’s correlation with other firm characteristics. In this section, we run a horse race test using the following regression specification, estimated in a firm-level panel:

$$r_{it} - r_t^f = (\beta_1^x \mathbb{1}[\text{Release}]_t + \beta_2^x \mathbb{1}[\text{Non-Release}]_t) \times \text{GenAI Exp.}_i + (\beta_3 \mathbb{1}[\text{Release}]_t + \beta_4 \mathbb{1}[\text{Non-Release}]_t) \times \text{Firm Char.}_i + FEs + \chi' \Gamma + \varepsilon_{it},$$

²¹The S&P 500 index dropped by 4.4% from December 1, 2022 to December 15, 2022. See <https://www.google.com/finance/quote/.INX:INDEXSP?sa=X&sqi=2&ved=2ahUKEwiruL6e0fuAAxXGh-4BHW5MCd8Q3ecFegQILRAf&window=1>. Note that the overall market includes the middle three AMH quintiles, so A and H do not aggregate to the market.

where GenAI Exp. is firm i 's Generative AI exposure, and we are interested in β_1^x , which measures the impact of ChatGPT release on firms with higher Generative AI exposure relative to other firms, while allowing for a differential impact on firms with particular other characteristics Firm Char $_i$, which represent potential alternative explanations for our findings.²²

The results are shown in Table VI: Column (1) shows the baseline effect of Generative AI exposure at a firm on daily returns during the post-ChatGPT two week release period. The release period effect of greater exposure estimated in the firm-level panel is significant and very similar in magnitude to that found in the portfolio level analysis suggesting that a 10ppt higher exposure is associated with 26 basis points higher returns per trading day. Columns (2)-(6) control for differential returns during the release period associated with other firm characteristics, which consist of log size, Tobin's Q, ROA, labor Intensity, and asset tangibility. Column (7) controls for differential trends in all of these characteristics during the release period at the same time.

While these horse races suggest that returns during the release period were larger for firms that have higher ROA and are smaller in size, in no case does including a particular characteristic substantially affect the positive returns to Generative AI exposure during that period. Even when all potential confounder characteristics are included simultaneously, we still estimate a statistically significant daily outperformance during the release period of ~ 0.2 ppt per day for firms with 10 ppt higher exposure. That is, our main findings based on firms' exposure to Generative AI are unlikely to be confounded by firms' other financial characteristics.

C. *Product exposure vs. labor exposure to Generative AI*

An important question around the impact of Generative AI on different firms is whether our stock return findings are driven by firms' employees' exposure to ChatGPT disruption or their products' exposure to ChatGPT disruption. For example, the stock value of chip maker NVIDIA more than doubled from the beginning of 2023 to Sep. 19, 2023, likely due to the company being a key supplier of the chips used in data centers and for training Generative AI models. While our task-based measure naturally focuses on firms' labor exposure to Generative AI, we perform several tests to inspect the potential for confounding due to firms' product exposure to Generative AI.

²²We also include trading day by NAICS 3-digit industry fixed effects, and controls Γ that include the interactions between the firm characteristic and the non-release period indicator. For comparison, we also show a version of this table that uses tercile portfolio assignments instead of a continuous exposure measure in Appendix Table D7, which shows qualitatively very similar results.

C.1. Importance of within-industry heterogeneity

We have shown earlier in Figure 3 that firms’ exposure to Generative AI varies substantially within-industry. Here, we examine to what extent the AMH returns during the ChatGPT release period are driven by within-industry variation. To the extent that industry classification captures the separation of products among firms, within-industry variation of firms’ exposure to Generative AI is less affected by product differences.

We consider two industry classifications. The first one is the NAICS 3-digit industry classification, which includes 46 categories with at least 10 firms in our data. Hoberg and Phillips (2016) show that the rich text of firms’ descriptions of their products in their 10K filings can more precisely separate product markets. We thus also adopt their 10K-based FIC50 industry classification as an alternative measure of industries. Given the limited number of stocks within an industry, we form three tercile portfolios (instead of quintile portfolios) within each industry, and then value-weight stock returns in the industry-neutral tercile portfolios.

Table V shows the results. Panel A1 shows the market adjusted alphas for the tercile portfolios sorted by firms’ Generative AI exposure as the benchmark. Panels A2 and A3 show the market adjusted alphas for the industry-neutral tercile portfolios using the NAICS 3-digit industry classification and the FIC50 industry classification, respectively. In all three panels, we observe that the within-industry AMH portfolio has large and statistically significant alphas during the ChatGPT release period and small and negative alphas outside the period. More importantly, the AMH alphas are very similar whether we use the full cross-sectional variation or only within-industry variation in firms’ Generative AI exposure. Panels B1-B3 show a very similar result when we examine the FF5-factor adjusted alphas instead of the market adjusted alphas. If industry classifications capture the effects of improvements to Generative AI on the demand for firms’ products, then the within-industry findings suggest that a substantial portion of the positive effects of the Generative AI breakthrough on firm value are due to the technology’s impact on firms’ labor rather than on firms’ products. This interpretation may provide further support for our labor-task-based measure of firms’ Generative AI exposure.

C.2. Results outside the tech sector

To further investigate how much tech firms, in particular, are driving our main findings, we follow Acemoglu et al. (2022) and single out firms in two sectors (information sector (NAICS 51) and professional, scientific, and technical services sector (NAICS 54)) as likely tech firms whose outputs are especially affected by the release of ChatGPT. Note that we do

not claim that the ChatGPT release does not impact firms’ products and services. Instead, our goal is to provide support for the idea that the differential returns we document for firms with different labor-force exposures to Generative AI are indeed driven by the labor channel. Labor and product market channels are not mutually exclusive. In Panel D of Table IV, we redo our main quintile sorting analyses but exclude firms from the two tech sectors. We observe very similar results. The FF 5-factor adjusted alphas for the AMH portfolio using non-tech firms has an average daily return of 0.33% during the ChatGPT release period as compared to 0.31% using all firms. In Appendix Figure C2 we show the AMH portfolio returns during the ChatGPT release period for each NAICS 2-digit sector. We observe that positive significant AMH portfolio returns are mainly driven by sectors such as manufacturing, finance, mining, health care, administrative, etc., and *not* by the information or professional sectors. While this may seem surprising, note that these figures report AMH returns within industry, so the overall return to the tech sector during the release period can still be large.

C.3. A horserace against product exposure to Generative AI

We next construct proxies for firms’ product exposure to Generative AI and run horserace tests with our labor-based exposure to Generative AI. Our baseline measure of firms’ product exposure to Generative AI develops a similar methodology to that used for measuring tasks’ exposure to Generative AI in Section I. In particular, we follow Hoberg and Phillips (2016) and many other prior studies and infer information about a firm’s business products from the “Business” description section of a firm’s most recent 10-K annual report. We then ask GPT 3.5 to assess whether firms’ business products involve enabling or scaling new AI technologies or benefit from a direct incorporation of the new AI capabilities. We provide details in Appendix A. This procedure generates a dummy variable that equals to 1 if the answer is yes and 0 if the answer is no. We label this measure *10K-based Product GenAI Exp.* This measure, by using a similar methodology to that for our measurement of firms’ labor exposure to Generative AI, provides a close comparison between firms’ product exposure versus labor exposure to Generative AI.

For robustness, we consider two popular alternative proxies. The first one counts the number of occurrences of AI related keywords in the firms’ business section in their 10-K report. This keywords count approach is widely used in prior studies (e.g., Webb (2019) and Babina et al. (2021)). The measure also generates a more continuous measure instead of our baseline binary measure. The second alternative proxy we consider is the list of firms that were recently classified as “near-term beneficiaries of AI” by the investment bank Goldman

Sachs.²³ This measure uses information beyond firms’ 10-K reports and thus may account for firms’ product exposure to AI beyond their 10-K discussion.

Table VII reports the results. Column (1) reports our baseline effect results in this firm-level panel regression without additional controls. Two salient findings stand out when we include controls for firms’ product exposure to Generative AI in Columns (2)-(5). First, unlike firms’ labor exposure to Generative AI which does not predict returns outside the ChatGPT release period, firms’ product exposure to Generative AI predicts returns in the “non-release” period across the board. Hence, investors may have expected some value-enhancement for the AI-related products before ChatGPT’s release, or ChatGPT’s impact on firms’ products are likely understood during the longer aftermath of ChatGPT release. In either cases, we do not observe a clean non-effect on firm value during the control period when inspecting firms’ product exposure to Generative AI, contrasting with the non-effect based on firms’ labor exposure to Generative AI.

Second, we observe that firms’ labor exposure to Generative AI robustly boosts firm values during the ChatGPT release period even after controlling for the three proxies for firms’ product exposure to Generative AI individually in Columns (2)-(4) or jointly in Column (5). This finding suggests that our stock return findings based on firms’ labor exposure are unlikely to be driven by their product exposure to Generative AI. This finding gives us confidence to further explore the mechanism behind our stock return findings via the labor-based channels, which we explore in Section IV.

III. Investor Attention and Generative AI Exposure

In this section, we highlight that the remarkably fast and pervasive salience of ChatGPT to investors, unlike many prior technologies, made our setting an unique opportunity to identify the impact of a technology shock on firm value.

Historically, understanding the impact of new technologies on businesses is a progressive and slow process. A strand of literature on the slow diffusion of technology shows that technological diffusion is subject to various frictions, including limited attention of firms to the new technology or firms’ unwillingness to employ the new technology (Greenwood (1999)).²⁴ Indeed, these papers show that it can take a long time for new technologies to have a widespread impact on production. As a result, many studies on the impact of

²³See Goldman Sachs US Equity Views, August 21, 2023, “The (AI) trade after the trade: Identifying potential long-term EPS beneficiaries of AI adoption”.

²⁴For instance, the average time length for a new technological product to diffuse from 10% to 90% (of the full adoption level) is over 10 years (see many examples in Greenwood (1999) and Manuelli and Seshadri (2014)).

technology are retrospective rather than occurring soon after the arrival of the technology.²⁵

In contrast, the release of ChatGPT rapidly garnered massive attention, potentially overcoming the information hurdles for past technologies.²⁶ As a result, this setting potentially offers an insight into the impact of a technology shock in “real time”. In addition, the fast spread of ChatGPT news and applications helps answer a crucial plausibility question about our main stock return findings: Can the stock market quickly pick up the impact of Generative AI on firms within weeks, as we argue? A key to answering this question is whether ChatGPT raises investors’ attention in our ChatGPT release period. In this section, we provide supporting evidence that ChatGPT, unlike prior technologies, quickly grabbed the attention of the general public. More importantly, the impact of the awareness by investors and firm executives appears disproportionately greater for firms with greater exposure to Generative AI.

A. Public attention from Twitter mentions of GPT

We begin by showing that the release of ChatGPT has garnered massive public awareness. Figure 4 shows that immediately after the release of ChatGPT on November 30, 2022, the Twitter mentions of “ChatGPT” or “GPT” reached over 100K daily. Mentions remain at 50K mentions daily even two weeks after the release.

B. Firm earnings conference call mentions of Generative AI

We next examine whether firms’ awareness of Generative AI’s impact is aligned with our measure of their exposure to Generative AI. We manually collect from the Seeking Alpha website a panel of the earnings conference call transcripts for S&P 500 firms from July 2018 to March 2023. For each of these earnings calls, we assign a calendar month, quarter, and year (as distinct from the fiscal year and quarter referenced in the call), based on the time stamp of the earnings call transcript. We process each transcript by converting it into a list of lower-case tokens, creating separate lists of unigrams (one-word tokens) and bigrams (two-word tokens).

We define a transcript as mentioning the Generative AI topic if it mentions the following keywords: “llm”, “chatgpt”, “gpt”, “gpt3”, “gpt4”, “generative”, or “language model”. For

²⁵For instance, the seminal work studying the impact of computer-based automation on the labor market since 1980 is Autor et al. (2003).

²⁶For instance, ChatGPT amassed 100 million users in just two months after its release, a milestone which took the World Wide Web 7 years, WhatsApp 3.5 years, Instagram 2.5 years, and Twitter 5 years to reach. See <https://economictimes.indiatimes.com/news/new-updates/chatgpt-witnesses-massive-rise-chatbot-gains-100-million-users-in-two-months/articleshow/98428443.cms?from=mdr>.

comparison, we define a transcript as mentioning a machine learning-related topic that is not specific to Generative AI technologies if it mentions “ml”, “machine learning”, “deep learning”, “natural language”, “neural net”, “neural network”, or “nlp”; and we define a transcript as mentioning a generic engineering topic if it mentions “engineer” or “engineering”.

Figure 5 shows that the share of firms mentioning Generative AI in their quarterly earnings calls rose substantially after ChatGPT’s release from less than 5% to 27%. In contrast, the percentage of firms mentioning machine learning or engineering does not show this increasing pattern after ChatGPT’s release.

Is the increase in firms’ mentions of Generative AI disproportionately more prominent among firms with higher Generative AI exposure? To answer this question, we run the following regression in each quarter:

$$\mathbb{1}[\text{Topic X}]_{it} = \alpha_t + \beta_t^X \text{GenAI Exposure}_i + \gamma \mathbb{1}[\text{Topic X}]_{i,2019} + \varepsilon_{it}, \quad (3)$$

where β_t^X represents the association between firms’ mentioning of topic X in quarter t and their exposure to Generative AI, while controlling for the firms’ past mentioning of the topic in any of the quarters of the year 2019. A positive association between firms’ awareness of Generative AI and our measure of firms’ exposure to Generative AI provides additional support for our labor-exposure measure.

Figure 6 shows the results. In Panel A, we observe that firms more exposed to Generative AI, as suggested by our measure, see a substantially stronger increase in mentions of Generative AI after ChatGPT’s release. A one standard deviation increase in Generative AI exposure among firms is associated with a 6.7 ppt larger likelihood of mentioning the related topics in an earnings call in Q1 2023. In contrast, we do not observe a strengthened association of firms’ mentions of non-Generative AI-related machine learning topics or engineering topics with firms’ exposure to Generative AI after ChatGPT’s release, as shown in Panels B and C of Figure 6. Moreover, to again ensure that this is not simply a function of only “tech sector” firms discussing the emerging Generative AI technology, we repeat the analysis after excluding the relevant industry sectors (NAICS 51 and 54) from the analysis. Even after this restriction, we find very similar results (shown in Appendix Figure C3), which confirm that our exposure measure’s ability to predict increased communication about Generative AI with a firm’s investors is not limited to the sectors where product market exposure is most likely to engender such attention.

In summary, the massive rise in public and corporate attention to ChatGPT after November 30, 2022, combined with the disproportionately greater reactions to firms with higher

Generative AI exposure, brings two important validation messages to our analyses. First, these findings reveal that the breakthrough of Generative AI, highlighted by the launch of ChatGPT, attracted attention from investors more rapidly than was the case for prior major technology breakthroughs. The swift and substantial attraction of investor attention makes detecting stock return responses to the launch of ChatGPT upon its release plausible. Second, the stronger reactions by investors and firms among firms with higher labor-based exposure to Generative AI supports the claim that firms' labor exposure to Generative AI is an important driver of the impact of Generative AI on firm values.

IV. Generative AI: Labor Complement to or Substitute?

In this section, we use data on wages and hiring of firms with higher and lower labor-force exposures to Generative AI to show that, on balance, Generative AI appears to be mainly a substitute for labor. We also show that firms with more data assets drive the positive returns of high labor-force exposure firms, suggesting that data is a complement to the productivity benefits to Generative AI. Thus, it can still be that some labor will benefit from driving the engines of Generative AI. For example, our results can still be consistent with Generative AI increasing the superstar effect of Rosen (1981) while on average being a substitute for most labor inputs. We provide a parsimonious theoretical framework consistent with our empirical results based on the model in Acemoglu et al. (2022) in the Internet Appendix, and note that our results can also be understood within the framework of Krusell et al. (2000) with the appropriate modifications to the specified production inputs.

While both substitution and complementary labor channels predict a greater increase in the market value for firms with higher Generative AI exposure in the ChatGPT release period, it is important to understand which effect is at play in the current period of Generative AI technology advances. Dissecting the two channels is also important for policy making, as it can tell whether the Generative AI boom leads to a divergence between exposed shareholders' welfare and exposed employees' welfare or if it aligns the two.

A. Evidence from the labor market

We explore the job posting and wage effects of the ChatGPT release to distinguish the labor displacement and productivity enhancement channels in the above framework.

We use the most recent Lightcast (formerly Burning Glass) data on firm job postings to shed light on firms' demand for various types of labor. From the Lightcast data, we

construct firms’ number of job postings for each occupation in each month from January 2010 to June 2023. We focus on changes in firms’ job posting patterns for occupations highly exposed to Generative AI according to our classification in Section I.B. In particular, we define occupations in the top tercile ranking by Generative AI exposure as the *High-GenAI Exposure* occupations. We next define all months after November 2022 as the *Post-ChatGPT* period and the earlier periods as the pre-ChatGPT period.

The conceptual framework predicts that if our stock return findings are mainly driven by Generative AI advances creating a productivity enhancement effect, we should expect an *increased* demand for Generative AI exposed occupations and an *increased* demand for overall employees by firms with higher Generative AI exposure. If our stock return findings are driven to a large degree by Generative AI creating a labor displacement effect, we should expect a *decreased* demand for Generative AI exposed occupations by firms with higher Generative AI exposure. We test the two predictions by estimating a difference-in-differences regression as follows:

$$\text{Job Posting}_{it} = \alpha_i + \alpha_t + \beta \mathbb{1}[\text{Post-ChatGPT}]_t \times \text{GenAI Exposure}_i + \text{FEs} + \varepsilon_{it},$$

where Job Posting_{it} represents firm i ’s job posting of *High GenAI Exposure* occupations or its total job posting numbers, and FEs represents year-month fixed effects and firm fixed effects.

Panel A of Table IX reports the estimates of β . Columns (1) and (2) show that firms with higher Generative AI exposure post fewer total jobs and also fewer jobs with high Generative AI exposure after the release of ChatGPT. This finding supports Generative AI creating a stronger labor displacement effect, at least in this early period of its deployment. Panel B shows that the *share* of job postings for occupations highly exposed to Generative AI in all job postings also declines more for firms with higher Generative AI exposure after ChatGPT’s release, validating that Generative AI displaces exposed tasks. The magnitude of the effect suggests that a 1 SD higher Generative AI exposure score (~ 8 ppt) is associated with a 0.5 ppt decline in the hiring share of highly exposed occupations.

We next analyze the wage rates for occupations highly exposed to Generative AI after the release of ChatGPT. We extract monthly individual-level data from the Census Current Population Survey (CPS).²⁷ We require the individuals to be between 18 and 65 years old, employed in the month of survey, and to have non-missing hourly wage rate. We extract each individual’s hourly wage rate, occupation code, gender, age, race, education level, and sampling weight in the survey. We next crosswalk our SOC 6-digit occupation codes to

²⁷The data can be downloaded at <https://cps.ipums.org/cps/>.

the more aggregate Census occupation codes, and we compute the Census occupations’ Generative AI exposure by averaging the SOC 6-digit occupations’ Generative AI exposure. Finally, we define a similar measure of *High-GenAI Exposure* occupations if the occupations are in the top tercile ranking by Generative AI exposure. Our sample period is from Jan 2019 to Jul 2023. We define similarly that all months after November 2022 are the *Post-ChatGPT* period and the earlier periods are the pre-ChatGPT period. We then run the following regression:

$$\ln(wage_{j,t}) = \chi \mathbb{1}[\text{Post-ChatGPT}]_t \times \text{GenAI Exp}_{occ} + \gamma X_{j,t} + \text{FEs} + \varepsilon_{j,t},$$

where we control for worker characteristics including gender, age, age squared, experience, experience squared, race (white vs. not), and years of education and fixed effects for occupation, year-month, and MSA. Following the literature, we weight the regression using each individual’s sampling weight.

The conceptual framework predicts that if the productivity enhancement effect dominates, we should expect an *increased* wage rate for Generative AI exposed occupations after ChatGPT’s release; and if the labor displacement effect dominates, we should expect a *decreased* wage rate for Generative AI exposed occupations after ChatGPT’s release.

Table X shows the results. Column (1) shows that occupations with higher Generative AI exposure see a significant decline in their hourly wage rate after the ChatGPT release. Consistent with the job posting results in Table IX, Column (2) shows that occupations with high-Generative AI exposure consistently see a lower hourly wage rate after the ChatGPT release, with wages in the top tercile of exposed occupations having decreased by 1.2 ppt since November 2022 relative to all other occupations. These results reinforce our prior findings that the labor displacement effect dominates the productivity enhancement effect in this round of Generative AI advances and in driving our stock return findings.

B. *The value of data for Generative AI*

One of the most important inputs to make effective use of Generative AI as a technology is data. Companies require access to a large amount of data that is specific to their business or customers in order to fine-tune generic LLMs for tasks specific to their own business. Moreover, as Agrawal, Gans, and Goldfarb (2022) argue, artificial intelligence has a particularly large impact by improving prediction as an input into decision-making, which also “requires the collection and organization of data” (Agrawal, Gans, and Goldfarb, 2019). Data has become an increasingly important asset in the economy, which has benefited firms that are more likely to generate valuable data as a result of their operations Farboodi et al.

(2019) and (Farboodi and Veldkamp, 2023). This firm-specific data can be used in conjunction with Generative AI models to create customer service chatbots, automate workflows, improve predictions and analytics, and for many other applications that replace the kinds of tasks that have high Generative AI exposure. As Generative AI models perform better the more data they have available, firms with greater access to data may be more likely to realize the potential productivity improvements as a result of these new technologies.

In this section, we test whether having access to valuable data is an important part of the mechanism through which greater Generative AI productivity potential—as identified by our exposure measure—leads to increases in firm value. We develop a new approach to measure proxies for a firm’s “data value” as a potential input into LLMs and quantify whether the returns to Generative AI exposure are higher in firms with access to better data.

Measuring firm “data value” in annual reports. Similar to our approach in Section II.C.3, we again use the text of the business description in a firm’s annual report to form an assessment of its access to relevant data. We assess whether a firm’s business description suggests access to relevant data based on its discussion of 6 topics: the general nature of the company’s business, the scale and reach of the firm, data collection mechanisms, data utilization, data infrastructure & management, and data regulation and privacy. We use an LLM to assign a score of 0 to 3 (no, little, moderate, or high) for the relevance of the firm’s business description in each of these categories, and also to form an overall score of 0 to 3 (no, little, moderate, or high) for the value of each firm’s data for LLM analytics. For comparison, we also deploy a simpler approach of creating an indicator of a non-zero count of “data” keyword mentions in the business description as a proxy for a firm having valuable data.²⁸

Data value and returns to Generative AI exposure. To estimate whether firms with access to more relevant data are more likely to see an impact on their value as a result of Generative AI productivity exposure, we estimate firm-level regressions with triple interactions, the coefficients of which reflect whether a firm with better data *and* higher Generative AI exposure experienced higher ChatGPT release period returns than a firm with the latter exposure but no access to relevant data.

The results are shown in Table VIII. We find that firms with better data—measured either by our 0-3 overall data quality assessment, or simply by an indicator of mentioning “data” in their annual report— have significantly higher returns to Generative AI exposure during the ChatGPT release period. In fact, the returns to Generative AI exposure during

²⁸See Appendix Section A for details. We also develop an alternative approach of measuring data value based on the predicted share of “data management” roles in a firm’s employment structure, which is inspired by the analysis in Abis and Veldkamp (2023).

the release period for companies *without* relevant data (the first two rows of columns 2 and 3) are not statistically significant. This suggests that having data to feed into Generative AI models is an important part of the mechanism through which financial markets expect task-level Generative AI exposure to result in greater firm value.²⁹

V. Conclusion

Market prices indicate that the arrival and diffusion of large language models and Generative AI represent a major technology shock with important effects on the overall value of firms, leading to winners and losers. This paper uses occupational exposures to Generative AI, along with firm-level measures of occupational composition, to assess the exposure to Generative AI innovations at the firm level for publicly traded U.S. corporations. We find that the effect of the release of ChatGPT on firm values was large, driving a difference in firm returns of approximately .4 ppt daily. Moreover, we provide evidence that firms with higher exposure to Generative AI-driven productivity increases are more likely to communicate with their investors about these technologies in their earnings calls, and differ in their hiring behavior in recent months, reducing their overall job postings, as well as the relative share of the highly exposed occupations in their job postings. We also document a relative decline in wages for highly exposed occupations at the national level in recent months. These findings raise important issues for policymakers: to the degree that shareholders benefit from increases in firm values while affected workers lose out, the new technology redistributes income in the economy, which may or may not be desirable from a welfare perspective. At the same time, our quantification of the productivity potential of Generative AI for different firms permits policymakers and entrepreneurs to better identify areas of opportunity and targets for disruption as this new technology reshapes the economic landscape. The degree to which this new technology will bring pain or plenty will depend on how firms and regulators can align in realizing the value which is promised by the financial market reaction to its release.

²⁹In Appendix Table D8 we show that these results are qualitatively similar if we measure access to valuable data instead as getting a score of “high” in at least one of the subcategories of the data value assessment of the business description, or by predicting the share of data management-intensive jobs in a firm’s 2022 employment structure.

REFERENCES

- Abis, Simona, and Laura Veldkamp, 2023, The changing economics of knowledge production, *Review of Financial Studies* .
- Acemoglu, Daron, and David Autor, 2011, Skills, tasks and technologies: Implications for employment and earnings, in *Handbook of labor economics*, volume 4, 1043–1171 (Elsevier).
- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, 2022, Artificial intelligence and jobs: evidence from online vacancies, *Journal of Labor Economics* 40, S293–S340.
- Acemoglu, Daron, and Pascual Restrepo, 2018, The race between man and machine: Implications of technology for growth, factor shares, and employment, *American economic review* 108, 1488–1542.
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, 2022, *Power and prediction: The disruptive economics of artificial intelligence* (Harvard Business Press).
- Agrawal, Ajay, Joshua S Gans, and Avi Goldfarb, 2019, Artificial intelligence: the ambiguous labor market impact of automating prediction, *Journal of Economic Perspectives* 33, 31–50.
- Akcigit, Ufuk, and Sina T Ates, 2020, Slowing business dynamism and productivity growth in the united states, in *Kansas City Federal Reserve. Paper prepared for the 2020 Jackson Hole conference*.
- Autor, David H, Frank Levy, and Richard J Murnane, 2003, The skill content of recent technological change: An empirical exploration, *The Quarterly journal of economics* 118, 1279–1333.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2021, Artificial intelligence,

- firm growth, and product innovation, *Firm Growth, and Product Innovation (November 9, 2021)* .
- Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson, 2020, Artificial intelligence, firm growth, and industry concentration, *Firm Growth, and Industry Concentration (November 22, 2020)*.
- Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson, 2022, Firm investments in artificial intelligence technologies and changes in workforce composition, *Available at SSRN* .
- Brynjolfsson, Erik, Danielle Li, and Lindsey R Raymond, 2023, Generative ai at work, Technical report, National Bureau of Economic Research.
- Bubeck, Sébastien, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al., 2023, Sparks of artificial general intelligence: Early experiments with gpt-4, *arXiv preprint arXiv:2303.12712* .
- Donangelo, Andres, 2014, Labor mobility: Implications for asset pricing, *The Journal of Finance* 69, 1321–1346.
- Eisfeldt, Andrea L., Antonio Falato, and Mindy Xiaolan, 2022, Human capitalists, in Martin Eichenbaum, Erik Hurst, and Valery Ramey, eds., *NBER Macroeconomics Annual*, chapter 1, 2 (University of Chicago Press, Chicago).
- Eisfeldt, Andrea L, Antonio Falato, and Mindy Z Xiaolan, 2023, Human capitalists, *NBER Macroeconomics Annual* 37, 1–61.
- Eisfeldt, Andrea L., and Dimitris Papanikolaou, 2013, Organization capital and the cross-section of expected returns, *The Journal of Finance* 68, 1365 – 1406.

- Eisfeldt, Andrea L., and Dimitris Papanikolaou, 2014, The value and ownership of intangible capital, *American Economic Review* 104, 189–94.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, 2023, Gpts are gpts: An early look at the labor market impact potential of large language models, *arXiv preprint arXiv:2303.10130* .
- Fama, Eugene F, and Kenneth R French, 2015, A five-factor asset pricing model, *Journal of financial economics* 116, 1–22.
- Farboodi, Maryam, Roxana Mihet, Thomas Philippon, and Laura Veldkamp, 2019, Big data and firm dynamics, in *AEA papers and proceedings*, volume 109, 38–42, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.
- Farboodi, Maryam, and Laura Veldkamp, 2023, Data and markets, *Annual Review of Economics* 15, 23–40.
- Felten, Edward W, Manav Raj, and Robert Seamans, 2023, Occupational heterogeneity in exposure to generative ai, *Available at SSRN 4414065* .
- Gilardi, Fabrizio, Meysam Alizadeh, and Maël Kubli, 2023, Chatgpt outperforms crowdworkers for text-annotation tasks, *arXiv preprint arXiv:2303.15056* .
- Greenwood, Jeremy, 1999, The Third Industrial Revolution: Technology, Productivity, and Income Inequality, *Economic Review* 35, 2–2.
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, 1997, Long-run implications of investment-specific technological change, *The American economic review* 342–362.
- Gutiérrez, Germán, and Thomas Philippon, 2019, The failure of free entry, Technical report.
- Hansen, Stephen, Peter John Lambert, Nicholas Bloom, Steven J Davis, Raffaella Sadun, and Bledi Taska, 2023, Remote work across jobs, companies, and space, Technical report, National Bureau of Economic Research.

- Hoberg, Gerard, and Gordon Phillips, 2016, Text-based network industries and endogenous product differentiation, *Journal of Political Economy* 124, 1423–1465.
- Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, 2021, Measuring technological innovation over the long run, *American Economic Review: Insights* 3, 303–320.
- Kogan, Leonid, and Dimitris Papanikolaou, 2014, Growth opportunities, technology shocks, and asset prices, *The journal of finance* 69, 675–718.
- Kogan, Leonid, Dimitris Papanikolaou, Lawrence Schmidt, and Bryan Seegmiller, 2019, Technology-skill complementarity and labor displacement: Evidence from linking two centuries of patents with occupations, *Available at SSRN 3585676* .
- Kogan, Leonid, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller, 2023, Technology and labor displacement: Evidence from linking patents with worker-level data .
- Krusell, Per, Lee E Ohanian, José-Víctor Ríos-Rull, and Giovanni L Violante, 2000, Capital-skill complementarity and inequality: A macroeconomic analysis, *Econometrica* 68, 1029–1053.
- Lane, Marguerita, and Anne Saint-Martin, 2021, The impact of artificial intelligence on the labour market: What do we know so far? .
- Lopez-Lira, Alejandro, and Yuehua Tang, 2023, Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models, SSRN Scholarly Paper 4412788, Available at SSRN 4412788, Rochester, NY.
- Loughran, Tim, and Bill McDonald, 2016, Textual analysis in accounting and finance: A survey, *Journal of Accounting Research* 54, 1187–1230.
- Manuelli, Rodolfo E., and Ananth Seshadri, 2014, Frictionless Technology Diffusion: The Case of Tractors, *American Economic Review* 104, 1368–1391.

- Noy, Shakked, and Whitney Zhang, 2023, Experimental evidence on the productivity effects of generative artificial intelligence, *Available at SSRN 4375283* .
- Papanikolaou, Dimitris, 2011, Investment shocks and asset prices, *Journal of Political Economy* 119, 639–685.
- Rosen, Sherwin, 1981, The economics of superstars, *The American economic review* 71, 845–858.
- Tuzel, Selale, and Miao Ben Zhang, 2021, Economic Stimulus at the Expense of Routine-Task Jobs, *Journal of Finance* 76, 3347–3399.
- Underwood, Ted, 2023, Using gpt-4 to measure the passage of time in fiction, <https://tedunderwood.com/2023/03/19/using-gpt-4-to-measure-the-passage-of-time-in-fiction/>.
- Veldkamp, Laura, and Cindy Chung, 2019, Data and the aggregate economy, *Journal of Economic Literature* .
- Webb, Michael, 2019, The impact of artificial intelligence on the labor market, *Available at SSRN 3482150* .
- Willison, Simon, 2023, A simple python wrapper for the chatgpt api, <https://til.simonwillison.net/gpt3/chatgpt-api>.
- Zhang, Miao Ben, 2019, Labor-technology substitution: Implications for asset pricing, *The Journal of Finance* 74, 1793–1839.

Figure 2: Occupational skills and Generative AI exposure. The graph below shows the results of regressing our SOC 6-digit occupation’s Generative AI exposure measure on a set of standardized occupational skill scores defined in Acemoglu and Autor (2011). That is, we run the following regression:

$$\text{GenAI Exp}_o = \alpha + \sum_S \beta_S \text{Skill}_{S,o} + \varepsilon_o$$

The regression sample contains 690 occupations. The bars around each coefficient show 95% confidence intervals based on heteroskedasticity-robust standard errors.

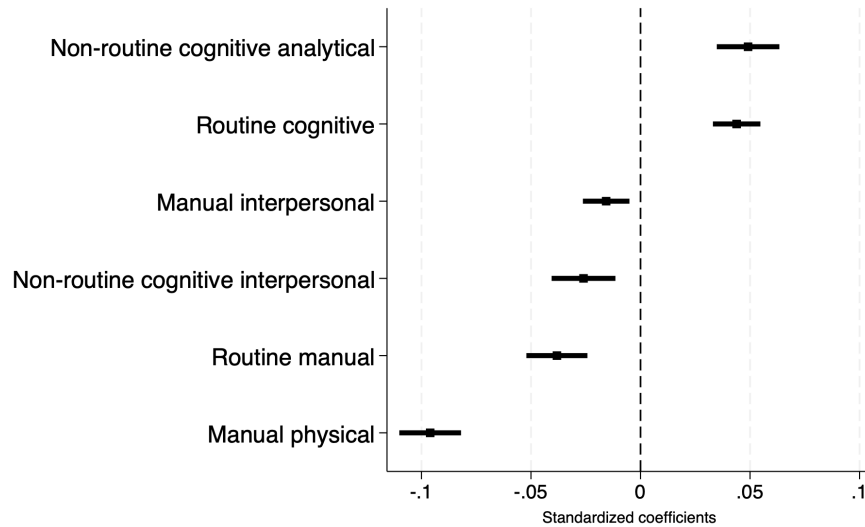


Figure 3: Generative AI exposure across and within sectors. This figure plots the average and the standard deviation of Compustat firms' Generative AI exposure within each NAICS 2-digit sector.

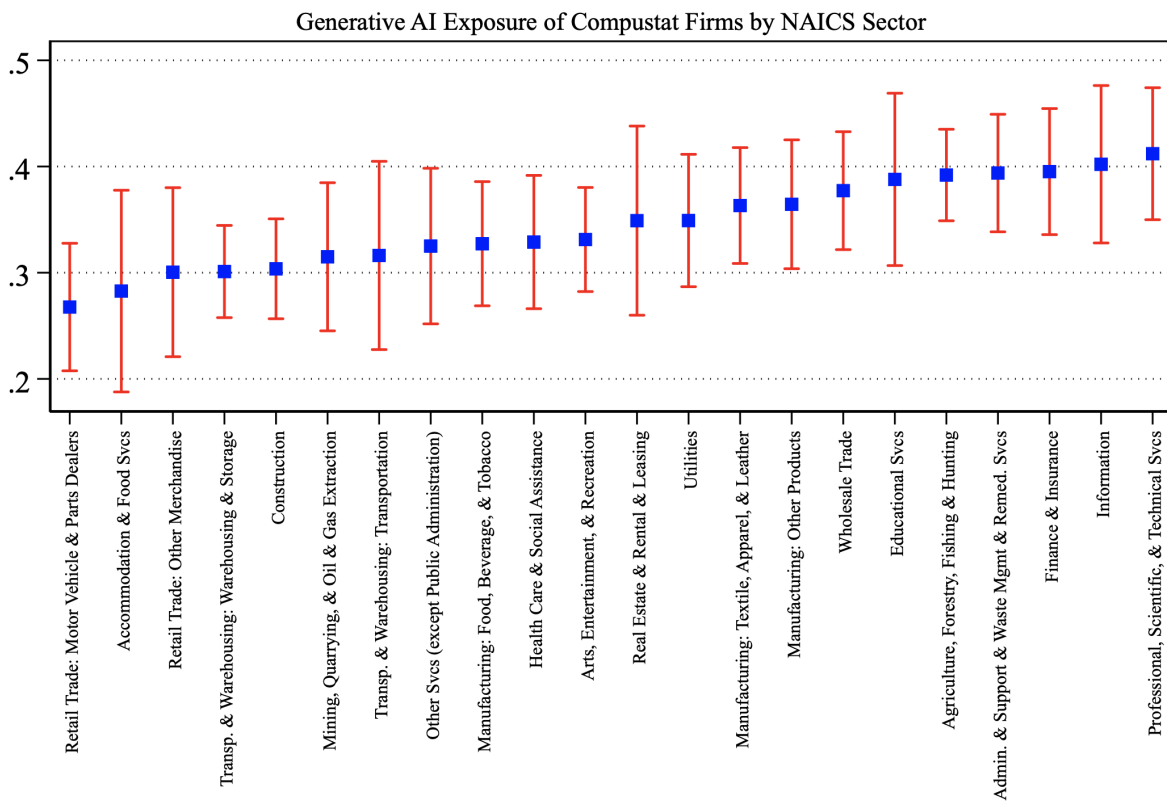
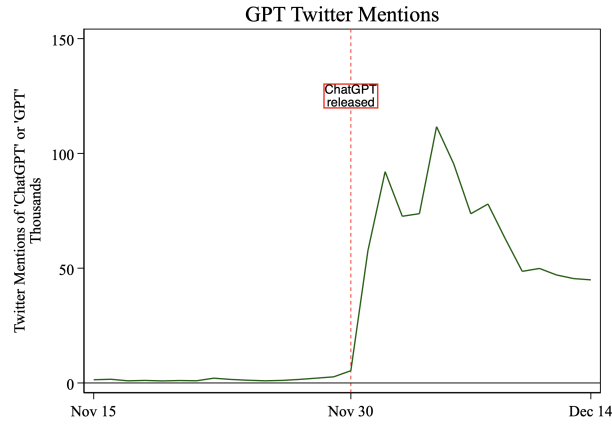
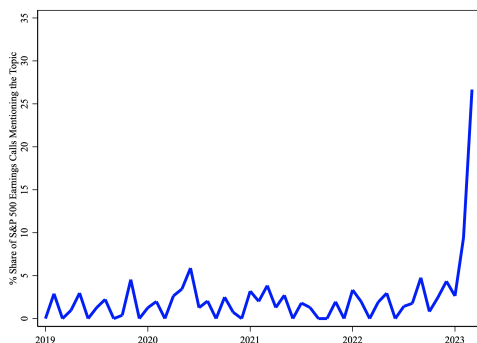


Figure 4: Investor attention to ChatGPT release. Panel A reports the t -statistic for β_t in the following cross-sectional regression $|r|_{it} = \alpha_t + \beta_t \text{GenAI Exp}_i + \varepsilon_{it}$, estimated at the firm level, for each trading day from Nov 15, 2022, to Dec 14, 2023 (two weeks before and after the release of ChatGPT). The dependent variable in the regression is the absolute value of daily stock returns, and the independent variable is the firm’s Generative AI exposure (see Section I). The dashed horizontal line indicates the t -statistics for p -value = 0.01. Orange dashed lines indicate the average t -statistic in the two weeks pre- and post-ChatGPT release on Nov 30, 2022. Panel B shows the total count of Twitter mentions of “ChatGPT” or “GPT” reported by Media Cloud.

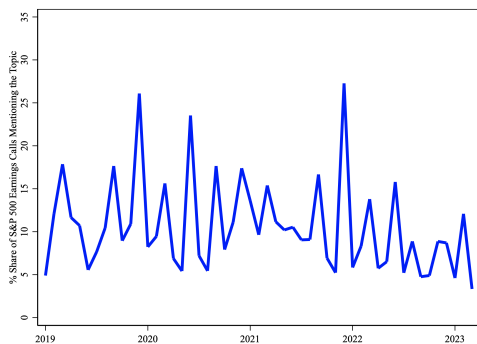


(A) Twitter mentions of GPT

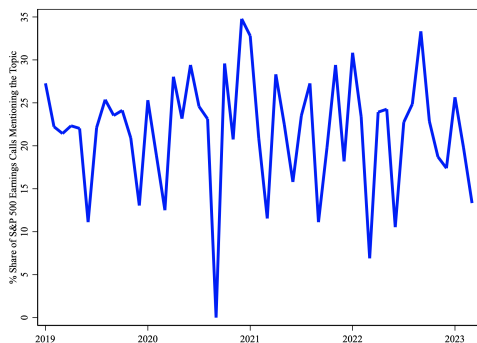
Figure 5: Topic mentions in company earnings conference calls. This figure shows the share of S&P 500 firms’ earnings calls that mention each topic. The quarterly earnings conference call transcripts for S&P 500 firms’ are manually collected from the *Seeking Alpha* website. Each call transcript is converted into a list of lower-case unigrams and bigrams. Panel A shows the share of earnings calls mentioning Generative AI keywords: “llm”, “chat-gpt”, “gpt”, “gpt3”, “gpt4”, “generative” and “language model”. In comparison, Panel B shows the share of earnings calls mentioning keywords of machine learning, deep learning or natural language processing: “ml”, “machine learning”, “deep learning”, “natural language”, “neural net”, “neural network” and “nlp”. Panel C shows the share of earnings calls mentioning keywords about engineering: “engineer” and “engineering”. See Section I for more details about our data collection and variable construction.



(A) Share of earnings calls mentioning Generative AI



(B) Share of earnings calls mentioning ML/DL/NLP

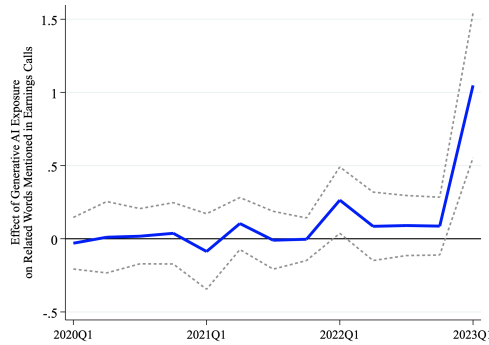


(C) Share of earnings calls mentioning Engineering

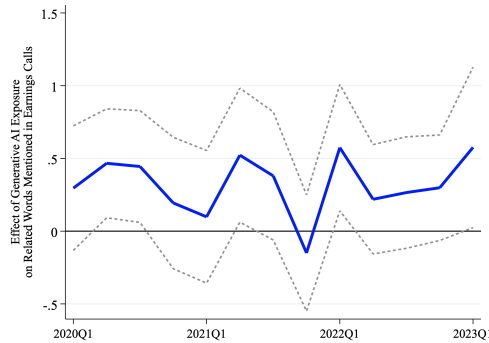
Figure 6: Firm-level Generative AI exposure and topic mentions in company earnings conference calls. Each graph show the result of estimating regression specifications of the form

$$\mathbb{1}[\text{Topic X}]_{i,t} = \alpha_t + \beta_t^X \text{GenAI Exp}_i + \gamma \mathbb{1}[\text{Topic X}]_{i,2019} + \varepsilon_{i,t}$$

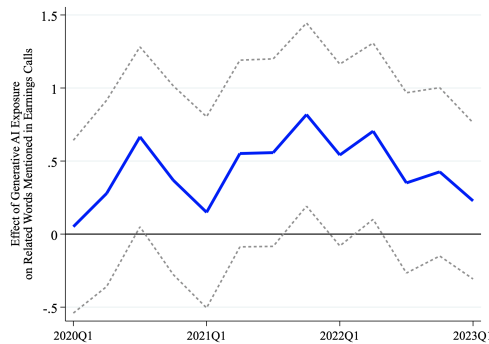
for each fiscal quarter from 2019 Q1 to 2023 Q1 for each topic using a manually collected panel of quarterly earnings conference call transcripts for S&P 500 firms from the *Seeking Alpha* website. See Table 5 for keywords corresponding to the Generative AI topic, ML/DL/NLP topic, and Engineering topic. Dotted lines show 95% confidence intervals based on heteroskedasticity-robust standard errors.



(A) Firm mentioning Generative AI and firm Generative AI exposure



(B) Firm mentioning ML/DL/NLP and firm Generative AI exposure



(C) Firm mentioning Engineering and firm Generative AI exposure

Table I: Summary statistics for selected firm characteristics. This table reports the summary statistics. *Generative AI exposure* is our bottom-up task-based measure of occupational exposure, aggregated to the firm-level based on the firm’s occupational employment structure. *Log Size* is the natural logarithm of total assets. *Labor Intensity* is the logarithm of the ratio of Property, Plant & Equipment (PP&E) to Total Assets, based on Donangelo (2014). *Tangibility* is the ratio of PP& E to Total Assets. *Org. Capital Ratio* is the organizational capital stock from Eisfeldt and Papanikolaou (2013) divided by Total Assets.

	Mean	Std. Dev.	p10	p50	p90	Obs.
Generative AI Exposure	0.373	0.078	0.284	0.388	0.446	5,907
Log Size	1.876	2.381	-1.091	2.063	4.636	2,517
Tobin’s Q	3.667	10.081	1.354	2.176	5.832	2,380
ROA	-0.011	0.891	-0.181	0.093	0.216	2,513
Labor Intensity	0.761	1.761	-2.157	1.126	2.512	2,387
Org. Capital Ratio	1.190	4.416	0.137	0.712	2.179	1,571
Tangibility	0.301	0.266	0.036	0.198	0.758	2,515

Table II: Major U.S. firms with the highest and lowest exposure to ChatGPT.

This table shows the 15 large U.S. publicly traded firms with the highest employee exposure to ChatGPT in Panel A and the 15 firms with the lowest exposure in Panel B. We select the large U.S. publicly traded firms as the top 100 firms with the largest market capitalization as of November 1, 2022, which also have headquarters in the U.S. *Generative AI exposure* is the firm’s labor exposure to ChatGPT-like technologies defined in Section I. *MktCap* is the firm’s market capitalization as of November 1, 2022, in \$B. *Sector* is defined at the NAICS 2-digit level.

Panel A: Top 15 Large U.S. Companies with Highest Exposure to ChatGPT			
Company Name	Generative AI exposure	MktCap	Sector
Int. Business Machines Corp	0.488	125	Information
Intuit Inc.	0.480	111	Information
QUALCOMM Inc.	0.479	132	Manufacturing
Fiserv Inc.	0.475	66	Information
NVIDIA Corporation	0.468	337	Manufacturing
S&P Global Inc	0.452	103	Admin. & Support Services
Broadcom Inc	0.449	195	Manufacturing
Verizon Communications Inc	0.444	157	Information
Microsoft Corp	0.442	1,700	Information
3M Co	0.442	69	Manufacturing
Advanced Micro Devices Inc	0.441	96	Manufacturing
ServiceNow Inc	0.434	85	Information
Adobe Inc	0.427	147	Information
PayPal Holdings Inc	0.418	96	Information
Thermo Fisher Scientific Inc	0.411	203	Manufacturing
Panel B: Bottom 15 Large U.S. Companies with Lowest Exposure to ChatGPT			
Company Name	Generative AI exposure	MktCap	Sector
Starbucks Corp	0.119	100	Accommodation & Food Svcs
McDonald’s Corp	0.194	201	Accommodation & Food Svcs
Dollar General Corporation	0.212	57	Retail Trade
Target Corp	0.235	76	Retail Trade
Walmart Inc	0.235	385	Retail Trade
Lowe’s Cos Inc	0.238	120	Retail Trade
TJX Companies Inc	0.243	83	Retail Trade
Costco Wholesale Corp	0.252	221	Retail Trade
Union Pacific Corp	0.253	121	Transportation & Warehousing
CSX Corp	0.256	61	Transportation & Warehousing
United Parcel Service Inc	0.256	123	Transportation & Warehousing
Home Depot Inc	0.261	303	Retail Trade
Tesla Inc	0.283	719	Manufacturing
Northrop Grumman Corp	0.291	83	Manufacturing
Mondelez International Inc	0.292	85	Manufacturing

Table III: Firm Generative AI exposure and firm characteristics. This table regresses our firms' Generative AI exposure measure on firm characteristics using the cross-section of U.S. publicly traded firms in 2022. See Table I for variable definitions. Panel B controls for fixed effects at the NAICS 3-digit level. All variables are winsorized at the top and bottom 1%. Standard errors are clustered at the industry level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Panel A: Across All Firms						
	(1)	(2)	(3)	(4)	(5)	(6)
Log Size	-2.653** (1.199)					
Tobin's Q		3.076*** (0.798)				
ROA			-21.516* (11.100)			
Labor Intensity				7.892*** (2.325)		
Org. Capital Ratio					9.139*** (2.506)	
Tangibility						-89.931*** (20.961)
Observations	2517	2380	2513	2387	1571	2515
Adjusted R^2	0.006	0.013	0.006	0.038	0.022	0.107
Panel B: Within-Industry						
	(1)	(2)	(3)	(4)	(5)	(6)
Log Size	-0.730 (1.021)					
Tobin's Q		0.524 (0.661)				
ROA			-7.303 (10.112)			
Labor Intensity				5.593*** (1.559)		
Org. Capital Ratio					1.111 (1.870)	
Tangibility						-61.794** (25.082)
Observations	2517	2380	2513	2387	1571	2515
Adjusted R^2	0.280	0.292	0.279	0.313	0.312	0.298

Table IV: Portfolio sorts on Generative AI exposure. This table reports daily excess stock returns of five value-weighted portfolios of firms sorted on Generative AI exposure. *AMH* is the “Artificial Minus Human” zero net investment portfolio long high exposure (*H*) stocks and short low exposure (*L*) stocks. On Nov 14, 2022, we sort NYSE stocks into five quintiles by their Generative AI exposure, and we assign all stocks into the portfolios using the NYSE breakpoints. We aggregate stock returns within each portfolio using the stocks’ market capitalization on the prior trading day as the weight. *ChatGPT release period* is from Nov 30, 2022 - Dec. 14, 2022, and *Not ChatGPT release period* include other days from Nov 15, 2022 to March 31, 2023. Panel A reports the portfolios’ excess returns net of the risk free rate. Panel B reports the portfolios’ alphas after adjusting for the market factor. Panel C reports the Fama-French 5-factor adjusted alphas. Panel D reports the Fama-French 5-factor adjusted alphas for a sample that excludes firms in the tech sectors NAICS 51 (information) and NAICS 54 (business services). *t*-statistics in parentheses are computed using Newey-West standard errors with five day lags.

Sample	Portfolios					AMH
	Q1	Q2	Q3	Q4	Q5	
<i>A: Excess returns (%)</i>						
Not ChatGPT release period	0.071 (0.59)	0.030 (0.26)	0.038 (0.34)	0.045 (0.38)	0.096 (0.77)	0.008 (0.14)
ChatGPT release period	-0.169 (-0.65)	0.020 (0.06)	0.141 (0.43)	0.035 (0.12)	0.281 (0.76)	0.434 (3.54)
<i>B: Market-factor-adjusted alpha (%)</i>						
Not ChatGPT release period	0.047 (1.09)	0.003 (0.09)	0.011 (0.30)	0.018 (0.44)	0.065 (1.90)	0.002 (0.04)
ChatGPT release period	-0.230 (-5.49)	-0.045 (-0.44)	0.075 (1.01)	-0.031 (-0.88)	0.207 (4.86)	0.420 (5.65)
<i>C: FF 5-factor-adjusted alpha (%)</i>						
Not ChatGPT release period	0.067 (1.57)	0.019 (0.56)	-0.023 (-0.81)	0.033 (1.13)	0.052 (1.94)	-0.031 (-0.55)
ChatGPT release period	-0.167 (-2.38)	0.014 (0.16)	-0.020 (-0.57)	0.026 (0.66)	0.157 (3.53)	0.309 (3.56)
<i>D: FF 5-factor-adjusted alpha ex. tech sector (%)</i>						
Not ChatGPT release period	0.068 (1.56)	0.013 (0.35)	-0.023 (-0.68)	-0.019 (-0.74)	0.073 (2.35)	-0.011 (-0.20)
ChatGPT release period	-0.171 (-2.40)	0.022 (0.22)	0.011 (0.32)	0.048 (1.01)	0.171 (1.94)	0.327 (3.51)

Table V: Within-industry portfolio sorts on Generative AI exposure. This table reports daily excess stock returns of three value-weighted portfolios of firms sorted on Generative AI exposure. *AMH* is the “Artificial Minus Human” zero net investment portfolio long high exposure (*H*) stocks and short low exposure (*L*) stocks. On Nov 14, 2022, we sort NYSE stocks into three terciles by their Generative AI exposure across all stocks in Panels A1 and B1, within NAICS 3-digit industry in Panels A2 and B2, and within the Hoberg and Phillips (2016) FIC 50-industry in Panels A3 and B3, and we assign all stocks into the portfolios using the NYSE breakpoints. We aggregate stock returns within each portfolio using the stocks’ market capitalization on the prior trading day as the weight. *ChatGPT release period* is from Nov 30, 2022 - Dec. 14, 2022, and *Not ChatGPT release period* include other days from Nov 15, 2022 to March 31, 2023. Panels A1-A3 report the portfolios’ alphas after adjusting for the market factor. Panels B1-B3 report the Fama-French 5-factor adjusted alphas. *t*-statistics in parentheses are computed using Newey-West standard errors with five day lags.

Sample	Portfolios			
	Q1	Q2	Q3	AMH
<i>A1: Market-factor-adjusted alpha (%)</i>				
Not ChatGPT release period	0.032 (1.13)	0.028 (1.15)	0.036 (1.72)	-0.013 (-0.32)
ChatGPT release period	-0.170 (-3.19)	0.069 (1.41)	0.123 (2.83)	0.277 (3.71)
<i>A2: NAICS 3-digit ind.-neutral market-factor-adjusted alpha (%)</i>				
Not ChatGPT release period	0.035 (1.36)	0.020 (0.59)	0.035 (2.07)	-0.017 (-0.58)
ChatGPT release period	-0.158 (-3.97)	0.035 (0.62)	0.144 (9.96)	0.286 (5.77)
<i>A3: FIC 50 ind.-neutral market-factor-adjusted alpha (%)</i>				
Not ChatGPT release period	0.044 (0.90)	-0.009 (-0.56)	0.020 (0.94)	-0.040 (-0.79)
ChatGPT release period	-0.181 (-3.13)	0.046 (2.09)	0.124 (2.21)	0.290 (2.82)
<i>B1: FF 5-factor-adjusted alpha (%)</i>				
Not ChatGPT release period	0.051 (2.07)	0.008 (0.48)	0.034 (1.69)	-0.033 (-0.86)
ChatGPT release period	-0.102 (-2.58)	0.005 (0.21)	0.123 (2.47)	0.210 (2.51)
<i>B2: NAICS 3-digit ind.-neutral FF 5-factor-adjusted alpha (%)</i>				
Not ChatGPT release period	0.052 (2.50)	0.041 (1.82)	0.032 (1.88)	-0.036 (-1.32)
ChatGPT release period	-0.113 (-2.95)	0.077 (1.34)	0.128 (6.22)	0.226 (4.61)
<i>B3: FIC 50 ind.-neutral FF 5-factor-adjusted alpha (%)</i>				
Not ChatGPT release period	0.075 (2.00)	-0.004 (-0.37)	0.021 (1.02)	-0.070 (-1.58)
ChatGPT release period	-0.113 (-2.23)	0.079 (2.88)	0.117 (2.36)	0.214 (2.50)

Table VI: Panel regressions of firm returns on Generative AI exposure. This table reports the results of regressing firms’ daily excess returns on their Generative AI exposure and other firm characteristics interacted with ChatGPT and Not ChatGPT release period dummies, using the following specification:

$$r_{it} - r_t^f = (\beta_1 \mathbb{1}[\text{ChatGPT release period}]_t + \beta_2 \mathbb{1}[\text{Not ChatGPT release period}]_t) \times \text{GenAI Exp}_i + (\beta_3 \mathbb{1}[\text{ChatGPT release period}]_t + \beta_4 \mathbb{1}[\text{Not ChatGPT release period}]_t) \times \text{Firm Char}_i + \text{Ind-Day FEs} + \varepsilon_{it}.$$

The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023. Risk free returns are obtained from Ken French’s website. *Ind-Day FEs* represents fixed effects of NAICS 3-digit industry interacted with trading day. The regressions are estimated using the stock’s prior trading day’s market capitalization as weights. See Table I for variable definitions. *t*-statistics in parentheses are computed using standard errors clustered at the firm level.

Dep. var.:	Daily Firm-level Excess Returns (ppt)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{GenAI Exp}_i$	2.540*** (2.685)	2.451*** (2.690)	2.528*** (2.693)	2.545*** (2.691)	2.480*** (2.598)	2.276** (2.339)	2.115** (2.391)
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{GenAI Exp}_i$	0.448 (0.834)	0.406 (0.776)	0.333 (0.759)	0.443 (0.815)	0.566 (1.070)	0.712 (1.347)	0.534 (1.279)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{Log Size}_i$		-0.017 (-1.021)					-0.021 (-1.147)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{Tobin's } Q_i$			0.002 (0.138)				0.003 (0.182)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{ROA}_i$				-0.040 (-0.119)			0.047 (0.110)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{Labor Intensity}_i$					0.023 (0.648)		-0.032 (-0.893)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{Tangibility}_i$						-0.405 (-1.117)	-0.588 (-1.567)
R-squared	0.52	0.52	0.52	0.52	0.52	0.52	0.52
Observations	183272	183272	183272	182802	181110	182426	180922
	<i>Controls & fixed effects</i>						
Three-digit industry \times trading day FEs	X	X	X	X	X	X	X
Non-release period \times Characteristic	X	X	X	X	X	X	X

Table VII: Product exposure vs. labor exposure to Generative AI. This table reports the results of regressing firms’ daily excess returns on their Generative AI exposure and other firm characteristics interacted with ChatGPT and Not ChatGPT release period dummies, using the following specification:

$$r_{it} - r_t^f = (\beta_1 \mathbb{1}[\text{ChatGPT release period}]_t + \beta_2 \mathbb{1}[\text{Not ChatGPT release period}]_t) \times \text{GenAI Exp}_i \\ + (\beta_3 \mathbb{1}[\text{ChatGPT release period}]_t + \beta_4 \mathbb{1}[\text{Not ChatGPT release period}]_t) \times \text{Product GenAI Exp}_i + \text{Ind-Day FEs} + \varepsilon_{it}.$$

The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023. Risk free returns are obtained from Ken French’s website. *Ind-Day FEs* represents fixed effects of NAICS 3-digit industry interacted with trading day. The regressions are estimated using the stock’s prior trading day’s market capitalization as weights. *10K-based Product GenAI Exp* is a dummy variable of whether the firm’s products would benefit from a Generative AI boom, as assessed using a similar methodology to our main measure by reviewing the business description section in the firm’s recent annual report using a large language model (see more details in Appendix A). *10K-based AI Mentions* is the count of mentions of “AI” or “artificial intelligence” in the business description in the firm’s recent annual report. *Goldman Sachs GenAI Exp* is a dummy variable identifying the “near-term AI beneficiaries” from a Goldman Sachs report on Aug. 21, 2023. *t*-statistics in parentheses are computed using standard errors clustered at the firm level.

Dep. var.:	Daily Firm-level Excess Returns (ppt)				
	(1)	(2)	(3)	(4)	(5)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{Labor GenAI Exp}_i$	2.540*** (2.685)	2.511*** (2.719)	1.601** (2.132)	2.304*** (3.056)	1.616** (2.136)
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{Labor GenAI Exp}_i$	0.448 (0.834)	0.231 (0.595)	-0.175 (-0.669)	0.198 (0.782)	0.033 (0.149)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{10K-based Product GenAI Exp}_i$		0.132 (0.893)			-0.027 (-0.176)
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{10K-based Product GenAI Exp}_i$		0.184*** (3.086)			0.078* (1.927)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{10K-based AI Mentions}_i$			0.017*** (6.206)		0.018** (2.423)
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{10K-based AI Mentions}_i$			0.010*** (4.881)		0.002 (0.641)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{Goldman Sachs GenAI Exp}_i$				0.368** (2.268)	0.016 (0.063)
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{Goldman Sachs GenAI Exp}_i$				0.335*** (4.481)	0.251*** (2.773)
R-squared	0.52	0.53	0.53	0.52	0.53
Observations	183272	168230	168230	183210	168230
	<i>Controls & fixed effects</i>				
Three-digit industry \times trading day FEs	X	X	X	X	X
Non-release period \times Characteristic	X	X	X	X	X

Table VIII: Generative AI exposure and firm data value. This table reports the results of regressing firms’ daily excess returns on their Generative AI exposure interacted with ChatGPT and Not ChatGPT release period dummies, and proxies for the value of a firm’s data using the following specification:

$$\begin{aligned}
 r_{it} - r_t^f &= (\beta_1 \mathbb{1}[\text{ChatGPT release period}]_t + \beta_2 \mathbb{1}[\text{Not ChatGPT release period}]_t) \times \text{GenAI Exp}_i + \text{Ind-Day FEs} + \varepsilon_{it} \\
 &+ (\beta_3 \mathbb{1}[\text{ChatGPT release period}]_t + \beta_4 \mathbb{1}[\text{Not ChatGPT release period}]_t) \times \text{Data Value Proxy}_i \\
 &+ (\beta_5 \mathbb{1}[\text{ChatGPT release period}]_t + \beta_6 \mathbb{1}[\text{Not ChatGPT release period}]_t) \times \text{Data Value Proxy}_i \times \text{GenAI Exp}_i
 \end{aligned}$$

The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023. Risk free returns are obtained from Ken French’s website. *Ind-Day FEs* represents fixed effects of NAICS 3-digit industry interacted with trading day. The regressions are estimated using the stock’s prior trading day’s market capitalization as weights. The first column shows the baseline effect of higher firm Generative AI exposure on returns during the release period and outside of it. All other columns show the triple interactions between this Generative AI exposure effect in different periods and different proxies for a firm’s access to data with value for LLM analytics: *10K Data Value: Overall* is a score of 0, 1, 2, or 3 (no, little, moderate, or high value) for the degree to which a firm is likely to have data that can be used as an input into LLM analytics, as assessed using a similar methodology to our main measure by reviewing the business description section in the firm’s recent annual report using a large language model (see more details in Appendix A). In particular, the LLM is asked to first assess whether a firm’s business description suggests access to relevant data based on its coverage of 6 categories: the general nature of the company’s business, the scale and reach of the firm, data collection mechanisms, data utilization, data infrastructure & management, and data regulation and privacy. *10K Data Value: Overall* then reflects the LLM’s assessment of the overall relevance of the company’s data for LLM analytics. *10K Mentions Data* indicates a non-zero count of mentions of “Data” in the business description in the firm’s annual report filed in 2022. *t*-statistics in parentheses are computed using standard errors clustered at the firm level.

Dep. var.:	Daily Firm-level Excess Returns (ppt)		
	(1)	(2)	(3)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{Labor GenAI Exp}_i$	2.540*** (2.685)	-0.484 (-0.613)	-0.994 (-0.701)
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{Labor GenAI Exp}_i$	0.448 (0.834)	0.389 (1.192)	0.937* (1.702)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{Labor GenAI Exp}_i \times 10\text{K Data Value: Overall}_i$		2.506*** (3.646)	
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{Labor GenAI Exp}_i \times 10\text{K Data Value: Overall}_i$		0.044 (0.111)	
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{Labor GenAI Exp}_i \times 10\text{K Mentions Data}_i$			3.688** (2.484)
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{Labor GenAI Exp}_i \times 10\text{K Mentions Data}_i$			-0.456 (-0.699)
R-squared	0.52	0.53	0.53
Observations	183272	168230	168230
	<i>Controls & fixed effects</i>		
Three-digit industry \times trading day FEs	X	X	X
Release / non-release period \times Characteristic	X	X	X

Table IX: Firm hiring and Generative AI exposure: The table reports the results of regressing firms’ number of job postings in each month on the interaction of their Generative AI exposure and a post-ChatGPT dummy in the following specification:

$$\text{Job posting}_{i,t} = \alpha_i + \alpha_t + \chi \mathbb{1}[\text{Post-ChatGPT}]_t \times \text{Gen. AI Exp}_i + \text{FEs} + \varepsilon_{it}.$$

Data on firm-level monthly job postings from Jan. 2010 to July 2023 are from Lightcast (formerly Burning Glass). *High-GenAI Exp Occ* refers to occupations that rank in the top tercile distribution of their Generative AI exposure (see more details of occupations’ Generative AI exposure in Section I). The *Post-chatGPT* period is defined as all months after November 2022. Panel A shows the results using raw number of job postings as the dependent variable, while Panel B uses the share of job postings out of the total number of job postings as the dependent variable. All regressions control for year-month fixed effects and firm fixed effects. *t*-statistics in parentheses are computed with standard errors clustered at both the year-month level and the firm level.

Job category:	All (1)	High-GenAI Exp Occ (2)
<i>Panel A: Number of firm’s job postings in the category</i>		
$\mathbb{1}[\text{Post-ChatGPT}]_t \times \text{GenAI Exp}_i$	-599.548** (-2.304)	-297.326*** (-3.635)
<i>Panel B: Percentage of firm’s job postings in the category</i>		
$\mathbb{1}[\text{Post-ChatGPT}]_t \times \text{GenAI Exp}_i$		-6.644** (-2.064)
Observations	132,964	132,964
	<i>Addl. controls & fixed effects</i>	
Month-Year FEs	X	X
Firm FEs	X	X

Table X: Worker wages and occupational Generative AI exposure. The table reports the results of regressing individual wages on their occupations’ Generative AI exposure in the following specification:

$$\ln(wage_{j,t}) = \chi \mathbb{1}[\text{Post-ChatGPT}]_t \times \text{GenAI Exp}_{occ} + \gamma X_{j,t} + \text{FEs} + \varepsilon_{j,t}.$$

Individual wage and characteristics are from the Census Current Population Survey (CPS) monthly panels from Jan. 2019 to July 2023. The dependent variable is the natural logarithm of the individual’s hourly wage rate. *GenAI Exp_{occ}* is the Generative AI exposure of the individual’s occupation, where we crosswalk our SOC 6-digit Generative AI exposure measure to the Census occupation codes. *High-GenAI Exp_{occ}* is a dummy variable that equals to 1 if the occupation’s Generative AI exposure ranks in the top tercile across all occupations (see also in Table IX). The *Post-chatGPT* period is defined as all months after November 2022. All regressions control for worker characteristics including gender, age, age squared, experience, experience squared, race (white vs. not), and years of education. All regressions also control for occupation fixed effects, year-month fixed effects, and MSA fixed effects. *t*-statistics in parentheses are computed with standard errors clustered at both the year-month level and the occupation level.

Dep. var.:	Log wage	
	(1)	(2)
$\mathbb{1}[\text{Post-ChatGPT}]_t \times \text{GenAI Exp}_{occ}$	-0.027** (-2.046)	
$\mathbb{1}[\text{Post-ChatGPT}]_t \times \text{High-GenAI Exp}_{occ}$		-0.012** (-2.427)
Observations	258,851	258,851
Occupation FE	X	X
Month FE	X	X
MSA FE	X	X
Individual Characteristics	X	X

Appendix A. Appendix: Methodology Notes

Generative AI exposure portfolio construction.

Portfolios for the main realized return analysis are formed from quintiles of stocks that have Yahoo Finance data for Nov. 15, 2022 - March 31, 2023. Quintile thresholds that define value-weighted portfolios within industries or for all stocks are solely based on the sample of stocks listed on NYSE as of the sorting date. All portfolios are formed based on equal weighted sorts on November 14, 2022, and weights for computing portfolio returns are adjusted based on daily returns to mimic passive buy-and-hold exposure.

Industry-neutral portfolios are computed by first forming within-industry equal-weighted tercile portfolios, and then averaging portfolio returns for the same terciles across industries. Returns for within-industry terciles and for all global (not industry-neutral) portfolio quintile sorts are value-weighted, while across-industry averages are industry market-cap. weighted. *AMH* is the "Artificial Minus Human" (also referred to as H-L) zero net investment portfolio long highest exposure quantile (*H*) stocks and short lowest exposure quantile (*L*) stocks.

The data set for estimating portfolio returns consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023. The Fama-French factors, including the market factor and risk free returns are obtained from Ken French's website. The portfolio return tables show alphas for each portfolio i estimated from regressions of the form

$$r_{it}^{pf} - r_t^f = \alpha_i^{\text{release}} \mathbb{1}[\text{ChatGPT release period}]_t + \alpha_i^{\text{not release}} \mathbb{1}[\text{Not ChatGPT release period}]_t + \sum_x \beta_i F_t^x + \varepsilon_{it},$$

where the intercept is allowed to vary with whether the day is in the ChatGPT release period consisting of Nov 30, 2022 - Dec. 14, 2022, or is one of the other days in the sample - the "Not ChatGPT release period". As noted in the tables, regressions may include controls for the returns to different factors F_t^x .

GPT prompt for task-based exposure scoring

The following prompt structure was based on the rubric language by Eloundou et al. (2023), as well as insights by Willison (2023) and Underwood (2023) about how to best structure API calls for GPT classification. Here are the instruction prompts submitted before asking GPT 3.5 Turbo to classify each task statement (using the version as of March 28th, 2023). Note that the order in which the two user-assistant interactions are provided to the API is randomized for each task, and the GPT temperature is set to 0:

systemprompt = "Consider the most powerful OpenAI large language model (LLM). This model can complete many tasks that can be formulated as having text input and text output where the context for the input can be captured in 2000 words. The model also cannot draw up-to-date facts (those from <1 year ago) unless they are captured in the input. Assume you are a worker with an average level of expertise in your role trying to complete the given task. You have access to the LLM as well as any other existing software or computer hardware tools mentioned in the task. You also have access to any commonly available technical tools accessible via a laptop (e.g. a microphone, speakers, etc.). You do not have access to any other physical tools or materials. You are a helpful research assistant who wants to label the given tasks according to the rubric below. Equivalent quality means someone reviewing the work would not be able to tell whether a human completed it on their own or with assistance from the LLM. If you aren't sure how to judge the amount of time a task takes, consider whether the tools described exposed the majority of subtasks associated with the task.

Exposure rubric:

E1 - Direct exposure: Label tasks E1 if direct access to the LLM through an interface like ChatGPT or the OpenAI playground alone can reduce the time it takes to complete the task with equivalent quality by at least half. This includes tasks that can be reduced to: - Writing and transforming text and code according to complex instructions, - Providing edits to existing text or code following specifications, - Writing code that can help perform a task that used to be done by hand, - Translating text between languages, - Summarizing medium-length documents, - Providing feedback on documents, - Answering questions about a document, - Generating questions a user might want to ask about a document, - Writing questions for an interview or assessment, - Writing and responding to emails, including ones that involve refuting information or engaging in a negotiation (but only if the negotiation is via written correspondence), - Maintain records of written data, - Prepare training materials based on general knowledge, or - Inform anyone of any information via any written or spoken medium.

E2 - Exposure by LLM-powered applications: Label tasks E2 if having access to the LLM alone may not reduce the time it takes to complete the task by at least half, but it is easy to imagine additional software that could be developed on top of the LLM that would reduce the time it takes to complete the task by half. This software may include capabilities such as: - Summarizing documents longer than 2000 words and answering questions about those documents, - Retrieving up-to-date facts from the Internet and using those facts in combination with the LLM capabilities, - Searching over an organization's

existing knowledge, data, or documents and retrieving information, - Retrieving highly specialized domain knowledge, - Make recommendations given data or written input, - Analyze written information to inform decisions, - Prepare training materials based on highly specialized knowledge, - Provide counsel on issues, and - Maintain complex databases. ## E3 - Exposure given image capabilities: Suppose you had access to both the LLM and a system that could view, caption, and create images as well as any systems powered by the LLM (those in E2 above). This system cannot take video as an input and it cannot produce video as an output. This system cannot accurately retrieve very detailed information from image inputs, such as measurements of dimensions within an image. Label tasks as E3 if there is a significant reduction in the time it takes to complete the task given access to a LLM and these image capabilities: - Reading text from PDFs, - Scanning images, or - Creating or editing digital images according to instructions. The images can be realistic but they should not be detailed. The model can identify objects in the image but not relationships between those options

E0 - No exposure: Label tasks E0 if none of the above clearly decrease the time it takes for an experienced worker to complete the task with high quality by at least half. Some examples: - If a task requires a high degree of human interaction (for example, in-person demonstrations) then it should be classified as E0. - If a task requires precise measurements then it should be classified as E0. - If a task requires reviewing visuals in detail then it should be classified as E0. - If a task requires any use of a hand or walking then it should be classified as E0. - Tools built on top of the LLM cannot make any decisions that might impact human livelihood (e.g. hiring, grading, etc.). If any part of the task involves collecting inputs to make a final decision (as opposed to analyzing data to inform a decision or make a recommendation) then it should be classified as E0. The LLM can make recommendations. - Even if tools built on top of the LLM can do a task, if using those tools would not save an experienced worker significant time completing the task, then it should be classified as E0. - The LLM and systems built on top of it cannot do anything that legally requires a human to perform the task. - If there is existing technology not powered by an LLM that is commonly used and can complete the task then you should mark the task E0 if using an LLM or LLM-powered tool will not further reduce the time to complete the task. When in doubt, you should default to E0.”

user_prompts = [””, ””, ””, ””] user_prompts[0] = ”Read the following occupation title and one of the tasks that belong to that occupation. Then do three things. 1: Reason step by step to decide which of the labels (E0/E1/E2/E3) from the exposure rubric you were given applies to the task’s exposure to LLM. Report. Give an explanation for which label applies and report the label that you think fits best. Do not say zero or N/A.

2: Report only the label that you determined for the task, which should match the label in step 1. Do not reply N/A.

3: Given the amount of speculation required in step 1, describe your certainty about the estimate—either high, moderate, or low.

The task to label follows:

Occupation: Inspectors, Testers, Sorters, Samplers, and Weighers. Task: Adjust, clean, or repair products or processing equipment to correct defects found during inspections.”

assistant_prompts = [””, ””, ””, ””] assistant_prompts[0] = ”1: Label (E0/E1/E2/E3): E0. Explanation: The model does not have access to any kind of physicality, and more than half of the task (adjusting, cleaning and repairing equipment) described requires hands or other embodiment.

2: E0.

3: High confidence.”

user_prompts[1] = ”Read the following occupation title and one of the tasks that belong to that occupation. Then do three things. 1: Reason step by step to decide which of the labels (E0/E1/E2/E3) from the exposure rubric you were given applies to the task’s exposure to LLM. Report. Give an explanation for which label applies and report the label that you think fits best. Do not say zero or N/A.

2: Report only the label that you determined for the task, which should match the label in step 1. Do not reply N/A.

3: Given the amount of speculation required in step 1, describe your certainty about the estimate—either high, moderate, or low.

The task to label follows:

Occupation: Computer and Information Research Scientists. Task: Apply theoretical expertise and innovation to create or apply new technology, such as adapting principles for applying computers to new uses.”

assistant_prompts[1] = ”1: Label (E0/E1/E2/E3): E1. Explanation: The model can learn theoretical expertise during training as part of its general knowledge base, and the principles to adapt can be captured in the text input to the model.

2: E1.

3: Medium confidence.”

Consistency of Generative AI scoring

To validate the consistency and replicability of our procedure that employs the GPT API for classification, we compare the scores assigned across 3 different GPT runs (which may

vary in results due to the randomized order of example cases provided, or non-deterministic features of the underlying LLM) for a randomly selected subsample of 100 task statements. We compare the different sets of scores as follows: First, we construct 3 different classifications for each task based on the assigned score: (1) “Current exposure”: score 1 has been assigned. (2) “Expected exposure:” Either score 1 or 2 has been assigned. Then, we compute the agreement between different scoring runs with regard to which tasks belong in these categories. The comparison between different runs is shown in Appendix Table A1. We find that the agreement between different GPT runs is very high - they arrive at the same score for at least 88% of all cases independent of the exposure classification considered. This validates that GPT reliably provides classifications that are highly consistent across different runs.

Table A1: Exposure score variation across GPT scoring runs

Score comparison	Agreement %	
	Current Exposure	Expected exposure
GPT #1 vs. GPT #2	95	90
GPT #1 vs. GPT #3	93	88
GPT #2 vs. GPT #3	96	88

Product Market AI Exposure Measures

We create a number of new measures to try and capture the degree to which a company has product market exposure to AI technology innovations. The data sources and construction of these measures is detailed below.

Company annual report data. The 10-K annual reports filed by companies at the SEC’s EDGAR system are obtained in pre-cleaned text files from Bill McDonald’s “Software Repository for Accounting and Finance” website,³⁰ based on the work in Loughran and McDonald (2016). We then use regular expressions to break up the text into the different “items” contained within the report.³¹, focusing only on reports filed in 2022, and on the “Business” section of each report.

AI-related business description keywords. For our first product market AI exposure measures we follow Hoberg and Phillips (2016) and use the “Business” description section of

³⁰See <https://sraf.nd.edu/>

³¹We built on code provided by Yu Zhu at <https://yuzhu.run/how-to-parse-10x/> and coding support from ChatGPT in this analysis.

a firm’s 10-K annual report to infer information about its product markets. We construct a simple measure of AI relevance for the firm’s business by tokenizing the text and counting the number of mentions of “AI” or “artificial intelligence”. We use the total count of mentions as a proxy for a firm’s products either using AI or depending on the use of it by other actors in the value chain. We do this by first lower-casing and tokenizing the text in the “Business” section, constructing lists of all possible unigrams and bigrams contained in the text, and then counting the occurrences of “ai” and “artificial intelligence” in the text.

GPT assessment of business product exposure to AI. The method of counting keywords potentially discards relevant information contained in the full text and the context and interdependence of the AI-related concepts discussed in it. Therefore, we also use a method that uses a large language model to evaluate the text of the business description in the firm’s annual report, asking it to determine whether there would be direct positive product market impact of a Generative AI boom on the described company. This allows us to keep the text in its original form and take into account the context within which AI or related topics are discussed. We develop a rubric that provides guidance to the model in the form of asking it to consider whether the firm’s products might be involved in enabling or scaling AI technologies, or might benefit from a direct incorporation of the new AI capabilities (the full text of the prompt is shown below). The model is then given two examples of scores applied to company business descriptions and given a new business description and asked to apply a binary label of whether the firm is “directly product market exposed to AI” or not, and also to provide an explanation for the score, which allows for an audit of the model’s reasoning. To economize on computing resources and API costs, we only do this analysis for annual reports in our sample that belong to firms for which we have previously computed a task-based AI exposure score, which have stock price data in our sample, and for which the “Business” section contains the word “and” at least three times (which removes annual reports that omitted the “Business” section or where our text extraction parsed a snippet that is too small). We submit the prompt for scoring together with a “Business” section to OpenAI’s GPT 3.5 Turbo API for evaluation. We are limited by the context window to evaluating only the first 3000 words of the business description for each firm, which is rarely binding and allows for ample business description for almost all companies - as the beginning of the section tends to provide the general overview of the company. The result of this procedure are binary product market AI exposure scores for ~ 2.2 K companies. Moreover, in early trial runs we audited the explanations provided by the model to ensure that the prompt leads to scoring that closely corresponds to a human scorer’s interpretation of the business impact.

Goldman Sachs “near-term AI beneficiaries”. This classification is based on the

list of firms in the report by Goldman Sachs US Equity Views, August 21, 2023, “The (AI) trade after the trade: Identifying potential long-term EPS beneficiaries of AI adoption”, which identifies 11 large firms “directly exposed to the development of AI technology.” While this report is produced ex post, this should, if anything, bias it towards better capturing the (at that point) observed product market exposure that leads to a stock market reaction. We use the report to code a binary variable on whether or not a firm is on the following list: NVIDIA (NVDA), Meta (META), Amazon (AMZN), Salesforce (CRM), Marvell Technology (MRVL), Adobe (ADBE), Alphabet (GOOGL), ServiceNow (NOW), Microsoft Corporation (MSFT), Intuit (INTU) and Credo Technology (CRDO).

Product Market AI Exposure: GPT Prompt for Scoring

General context prompt submitted with the completion task to GPT 3.5 Turbo to score company annual report text from the “Business” section for whether it suggests the company has direct product market AI exposure:

```

systemprompt = "You are a financial analyst who is an expert in evaluating company filings for their business
implications. \
  You are evaluating whether a given company's business description in their annual report suggests that the
company's products will be in higher demand as a result of a recent boom in generative artificial intelligence
applications.\
  You are only interested in companies whose product sales might be directly impacted by the Generative AI
boom, not companies that might benefit from labor productivity gains due to incorporating Generative AI into
their processes.\
Among the companies that are likely to see their product market positively affected by the Generative AI boom,
ALL of the following might be reasons why they have positive direct exposure:\
  1) Empowered users. Technology companies leveraging Generative AI technology to amplify their product
functionality will be directly positively exposed.\n \
  2) Enablers: Select makers of semiconductors and related equipment or other infrastructure needed to build
or deploy Generative AI technology will be directly positively exposed.\n \
  3) Hyperscalers: Large companies using their extensive cloud computing infrastructures to commercialize
Generative AI on a large scale will be directly positively exposed. \n \
  You are reasoning through how a company's product as described in the business description section of the
annual report might be affected by a boom in demand for Generative artificial intelligence applications, taking
into account the possible impact channels just described.\
  Then, you use that reasoning to decide if a company is: \n \
    P1: directly positively affected in its product market, which requires that it either supplies products
(e.g. equipment, chips, infrastructure) that will be needed to support the Generative AI boom, or has immediate
benefits in terms of adding AI capabilities into its product, e.g. because of access to proprietary data, or
providing solutions to user problems for which Generative AI is highly complementary. The company should fall
into one of the categories of Empowered User, Enabler, or Hyperscaler to get a P1 exposure score. \n \
    P0: no direct positive impact in its product market, as it does not sell inputs into the Generative AI
activities of other firms and does not have immediate ability to incorporate new AI capabilities into its
products. This includes more remote indirect impacts from generative AI leading to broader increases in
innovation or demand for the industry but not necessarily for this company in particular. \n \
    Note that you should assign P1 only when there is immediate product market impact (e.g. the company
product involves to a large degree chatbots or text generation that Generative AI can easily improve, or
supplies key inputs), but should assign P0 (not exposed) when the AI impact on the firm's product capabilities
is more remote, or only somewhat probable, such as requiring future AI-based tools to be built by other
companies that might then complement the firm's product. Also note that companies that provide critical inputs
into Generative AI usage, such as infrastructure or hardware can also be directly affected and should get a
score of P1, as those components are essential for using AI and will see higher demand in a boom."

```

(Example interactions submitted with the prompt are omitted - available from authors upon request)

Prompt preceding the “Business” text submission”:

```

prefix = "Read the following business description from a company's annual report. Then do two things. \
1: Reason step by step to apply a score of P0 (not directly positively exposed) or P1 (directly positively \
exposed) to rate its product market exposure. Use the rubric and generative AI product market exposure \
categories provided to you before (considering whether the company might be an Empowered User, Enabler, or \
Hyperscaler, all of which positively exposed) to decide whether the company's product market is going to be \
positively affected by the boom in Generative artificial intelligence. \
Report the label that you think fits best and provide up to two sentences that summarize your reasoning for the \
product market exposure score that you assigned. Do not say N/A.\n \
2: Report only the product market exposure score that you determined for the company, which should match \
the score in step 1. Do not reply N/A. \n\
The business description follows:\n "

```

Overall prompt structure (here, *fulltext* is the “Business” description to be scored):

```

prompt = [{"role": "system", "content": systemprompt},
{"role": "user", "content": user_prompts[0]},
{"role": "assistant", "content": assistant_prompts[0]},
{"role": "user", "content": user_prompts[1]},
{"role": "assistant", "content": assistant_prompts[1]}]

```

Measuring Firm “Data Value” for Generative AI

In order to measure the “data value” of a company that might contribute to its ability to productively deploy Generative AI tools and their analytics capabilities, we develop a number of new measures to quantify the amount of data that a company has effective access to. Similar to the product market Generative AI exposure measurement approach described above, we again review the business description section in the firm’s recent annual report using both traditional NLP approaches and a large language model. Again, we follow Hoberg and Phillips (2016) and use the “Business” description section of a firm’s 10-K annual report to infer information about its product markets. We also develop an alternative approach of measuring data value based on the predicted share of “data management” roles in a firm’s employment structure, which is inspired by the analysis in Abis and Veldkamp (2023).

Data-related business description keywords. We construct a simple measure of data value for the firm’s business by tokenizing the text of the business description in the firm’s annual report filed in 2022 and counting the number of mentions of “data”. We use the total count of mentions as a proxy for the importance of data in firm’s existing business. We do this by first lower-casing and tokenizing the text in the “Business” section, constructing lists of all possible unigrams contained in the text, and then counting the occurrences of “data” in the text. In the regression analysis we use an indicator of a non-zero count of “data” mentions as our proxy for a firm having valuable data.

GPT assessment of business “data value” for LLMs. Again, we also use a method that uses a large language model to evaluate the text of the business description in the firm’s annual report. We develop a rubric that provides guidance to the model in the form of asking it to first consider subcategories of data relevance. In particular, the LLM is asked to first assess whether a firm’s business description suggests access to relevant data based on its coverage of 6 categories: the general nature of the company’s business, the scale and reach

of the firm, data collection mechanisms, data utilization, data infrastructure & management, and data regulation and privacy (full prompt shown below). The model is asked to assign a score of 0, 1, 2, or 3 (no, little, moderate, or high relevance) in each category. Only then is the model asked to also provide an overall score for the degree to which a firm is likely to have data that can be used as an input into LLM analytics (low, moderate, or high data value).

We submit the prompt for scoring together with a “Business” section to OpenAI’s GPT 3.5 Turbo 16k API for evaluation. We submit only the first 3000 words of the business description for each firm, which is rarely binding and allows for ample business description for almost all companies - as the beginning of the section tends to provide the general overview of the company. In our analysis based on this scoring, we use both the LLM’s 0-3 assessment of the “overall” relevance of the company’s data for LLM analytics, as well as a binary indicator for whether any of the subcategory scores was assessed as a 3 (high relevance).

Data Management Skill. As an alternative measure of a company’s effective ability to leverage data in combination with LLM analytical capabilities, we build on the insight in Abis and Veldkamp (2023) that the prevalence of “data management” skills in a firm’s employment indicates the accumulation of valuable data. First, we predict the likelihood that a U.S. job posting for a 6-digit SOC occupation mentions at least 3 “data management skills”. These skills are classified based on a list of words indicating relevant skills in job postings from Abis and Veldkamp (2023), which we fuzzy match to skill tags in Lightcast job posting data using Stata’s *matchit* command. We retain all matches with a similarity score above 0.7, and then manually inspect whether the matched Lightcast skill tag actually corresponds to a data management skill or was a spurious match.

Then, we count the number of such data management skills present in each job posting for 2017-2021—the 5 years preceding the year when ChatGPT was released. We aggregate the posting-level data into occupations to compute the probability that a U.S. job posting for a 6-digit SOC occupation mentions at least 3 data management skills in this time period—which we define as the likelihood that a job in this occupation has “high data management skill.”

Last, we again use the LinkedIn occupational employment distribution at each firm in 2022, together with these occupation-level expected shares of high data management skill jobs, to predict the probability that a job at a firm is a “high data management skill” position. That is, this measure represents the predicted share of a firm’s positions on LinkedIn that are “data management”-intensive. In our analysis, we use both this predicted share and a binary variable for firms in the top tercile of this predicted data management intensity

among its workers—as proxies for the degree to which company is likely to have valuable data based on its employment structure.

Data Value Assessment: GPT Prompt for Scoring

General context prompt submitted with the completion task to GPT 3.5 Turbo 16k to score company annual report text from the “Business” section for whether it suggests the company has data that would be valuable as an input for LLM analytics:

```
systemprompt = "You are an equity analyst who is an expert in assessing companies' intangible assets. Analyze the 'Business Description' section of the provided 10-K annual report. You are to evaluate specifically the company's potential data value as it would pertain to complementing Large Language Model-driven analytics. \n \n To ensure a thorough analysis, consider the following dimensions in which the company's business description might reveal the value of a company's data : \n \n 1) Nature of Business (NOB): Examine and specify if the company is in a data-centric industry, such as e-commerce, or a more traditional industry. Detail the nature of its data interactions. \n \n 2) Scale and Reach (SAR): Assess the geographical scope and customer breadth of the company. Does it operate globally with a diverse user base, or is its impact more regional or niche? Provide specific examples. \n \n 3) Data Collection Mechanisms (DCM): Elaborate on how the company collects its data. Highlight if it actively collects through proprietary platforms or devices. Identify any passive collection methods or partnerships focused on data. \n \n 4) Data Utilization (DU): Dive deep into the company's application of its data. Describe if there are methods of direct monetization or if the data supports business processes or provides insights. Identify and discuss any future plans mentioned about leveraging data. \n \n 5) Infrastructure & Management (IM): Extract and comment on any mentions of investments in data-centric infrastructure or tools, collaborations related to data, or any dedicated internal teams focused on data. \n \n 6) Regulation & Privacy (RP): Detail the company's approach to data regulation and privacy. Outline mentions of data security, potential legal hurdles related to data, or initiatives ensuring ethical data usage. \n \n For each of the six points above, provide a clear reasoning for your conclusions. \n \n Score the company on a scale of 0-3 for each point, where 0 is 'No mention or not applicable', 1 is 'Limited evidence', 2 is 'Moderate evidence', and 3 is 'Strong evidence'. Conclude only after you've addressed all points, and base your summary on the evidence provided, specifically focusing on the company's data value potential for LLM analytics. \n \n Put your overall summary on a scale of 0-3, where 0 is 'No valuable data', 1 is 'Limited data value', 2 is 'Moderate data value', 3 is 'High data value'. "
```

(Example interactions submitted with the prompt are omitted - available from authors upon request)

Prompt preceding the “Business” text submission”:

```
prefix = "Read the following business description from a company's annual report. Then do two things. \n 1: Reason step by step to apply a score of 0,1,2, or 3 in each of the six categories provided, which highlight different ways that data value relevant for LLM analytics might be revealed in the business description. Also apply a summary data value score of 0,1, 2, or 3. \n Report your scores for each category, and your summary score, and provide one sentence that summarizes your reasoning for each category. \n \n 2: Report only the scores for each of the six categories and the overall score, preceding them by the abbreviated category names (NOB, SAR, DCM, DU, IM, RP, Overall), which should match the scores reported in step 1. Do not reply N/A. \n \n The business description follows:\n \n "
```

Appendix B. Appendix: Example Theoretical Framework

We offer a parsimonious model consistent with our empirical results documenting the positive effect of Generative AI labor exposure on firm values. Our simple framework builds on Acemoglu et al. (2022), which specifies a parsimonious version of the “task-based” production function in Acemoglu and Autor (2011). In this framework, a firm’s output is produced by combining three sets of tasks: tasks that are currently performed by humans and unexposed to AI displacement, labeled u ; tasks that are currently performed by humans and exposed to AI displacement, labeled e ; and tasks that are already performed by rented capital powered by AI, labeled a . Acemoglu et al. (2022) argue that even tasks performed by AI-powered capital still need labor, albeit a small number of workers, to operate, program, or maintain the AI capital. Put together, the firm’s production function can be specified in the following Cobb-Douglas form:

$$y = A (\gamma_u l_u)^{\alpha_u} (\gamma_e l_e)^{\alpha_e} [\gamma_a \min\{k_a, l_a\}]^{\alpha_a}, \quad (\text{B1})$$

where the γ s are the productivities of inputs for each set of tasks, A is the firm’s total factor productivity, and $\alpha_u + \alpha_e + \alpha_a = 1$.³²

In this framework, an advance in AI, such as the launch of ChatGPT, is modeled as a shock that allows the exposed human task e to be performed by AI. As the productivity-to-cost ratio for AI tasks are greater than that for exposed tasks, the firm then finds it economical to adopt AI.³³ Important to our study, Acemoglu et al. (2022) show that an AI advance may either reduce or increase the employment of AI-exposed workers, depending on whether the displacement effect or the productivity effect dominates (see Section II.C of Acemoglu et al. (2022)). The intuition is analogous to considering the balance between the substitution effect and the income effect on the purchases of other goods when the price of a consumption goods falls.

We now discuss how AI advances affect firm value in this framework. A firm is an entity that hires labor and rents AI capital to generate profits. In a two-period setting, the firm’s value at the initial period equals its discounted future profits. Following Acemoglu et al. (2022) and also our empirical measure, we specify a firm with a high Generative AI exposure in this framework as one that possesses a higher share of exposed-task workers, l_e . The prediction from Acemoglu et al. (2022), albeit ambiguous on the employment effects,

³²See more details on page 298 of Acemoglu et al. (2022).

³³In other words, while each firm is a wage and price taker in this framework, firms as a group extract rents from employees and capital leasers due to barriers to entry.

unambiguously predicts a more positive impact of AI advance on the value of firms with greater Generative AI exposure.

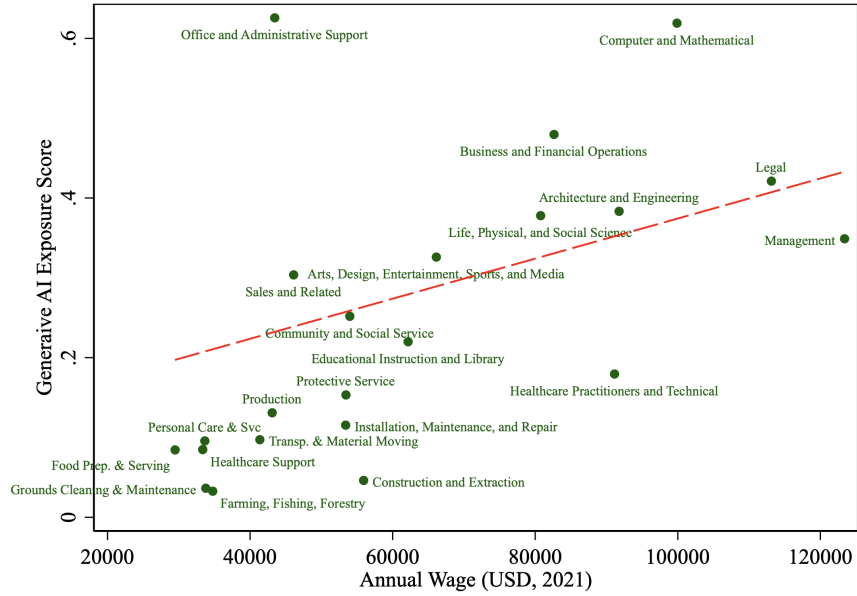
The reason that the employment results are ambiguous is similar to the general fact that investment-specific technological change can have both income and substitution effects. Firms can replace more exposed-task workers with Generative AI to reduce costs. If an AI advance creates a stronger displacement effect than productivity effect, we expect firms to hire fewer workers in Generative AI exposed occupations, and for the wages to decline in occupations with high Generative AI exposure.³⁴ In this case, the reduced costs resulting from the Generative AI advance drive an increase in firm value while wages and hiring for exposed jobs declines. This is what we find in the data so far following the release of ChatGPT 3.5 and 4.

Even if Generative AI is a labor substitute it is theoretically possible in this framework that hiring and wages could increase following an advance in AI. Such increases can occur if the advance in AI results in a productivity effect strong enough to overwhelm the displacement effect. In this case, the firm finds it so beneficial to use AI that it expands its production enough to also hire more workers in occupations with high Generative AI exposure (but even more in other occupations), and the wages for these occupations may increase. In either case, the productivity enhancement from the Generative AI advances is boosting firm value.

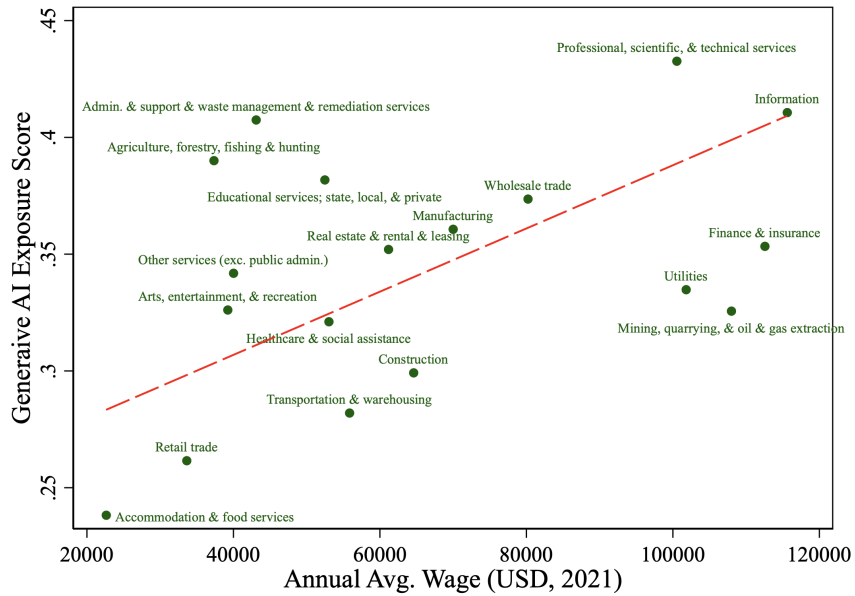
³⁴See further discussion of the AI impact on wages in Acemoglu and Restrepo (2018).

Appendix C. Appendix Figures

Figure C1: Generative AI exposure and wages by major occupation group and industry sectors.



(A) Major occupation groups



(B) NAICS 2-digit industry sectors

Figure C2: Within-Sector AMH Generative AI exposure Portfolio Realized Returns: ChatGPT Release Period. The graph shows the average daily return alpha for the ChatGPT release period for AMH Generative AI exposure portfolios within each industry, at the 2-digit NAICS sector level. Each AMH portfolio is formed by taking the value-weighted highest and lowest quintiles of Generative AI exposure within each industry (based on NYSE stock cutoffs) and forming zero net investment AMH portfolio returns as the equal-weighted difference in the daily realized returns between these portfolios, and then subtracting the daily risk-free return. The industries shown omit any sectors with fewer than 10 firms combined in the highest and lowest quintiles in the sample, as well as sectors with less than a 10 ppt Generative AI exposure spread at the sector level between the high and low quintile. The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023. The graphs show market-factor adjusted alphas estimated from regressions of the form

$$r_{it}^{pf} - r_t^f = \alpha_i^{\text{release}} \mathbb{1}[\text{ChatGPT release period}]_t + \alpha_i^{\text{not release}} \mathbb{1}[\text{Not ChatGPT release period}]_t + \beta_i(Mkt_t - r_t^f) + \varepsilon_{it},$$

where the intercept is allowed to vary with whether the day is in the release period defined by Nov. 30, 2022 - Dec. 14, 2022, or is one of the other trading days in the sample. The returns are shown in units of average daily excess realized returns (controlling for the market factor). Red error bars indicate 95% confidence intervals computed using Newey-West standard errors with five lags.

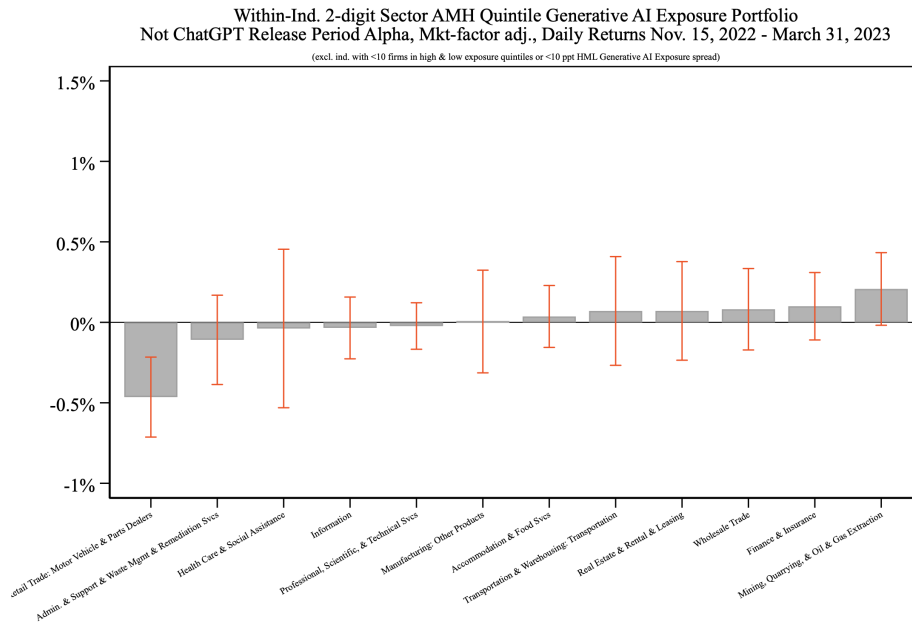
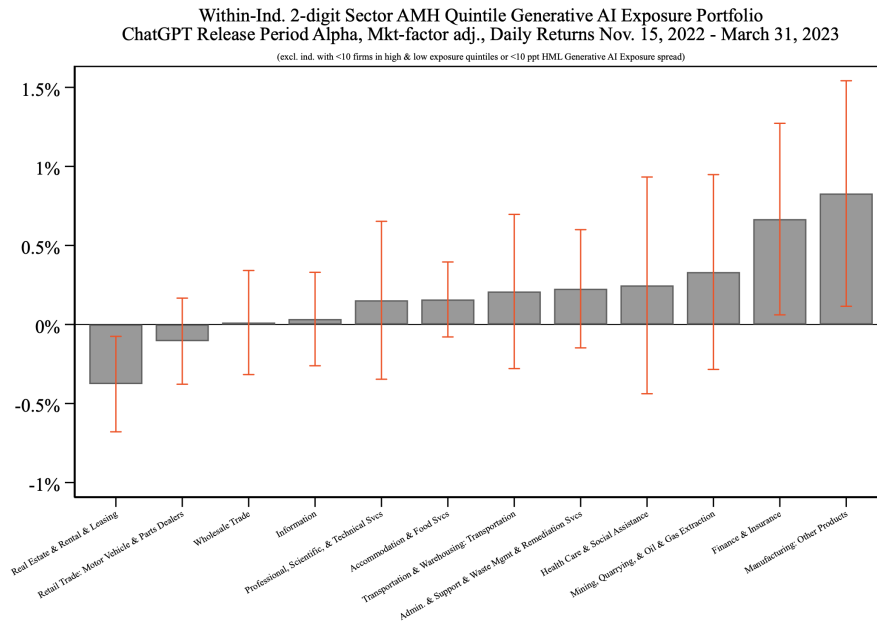
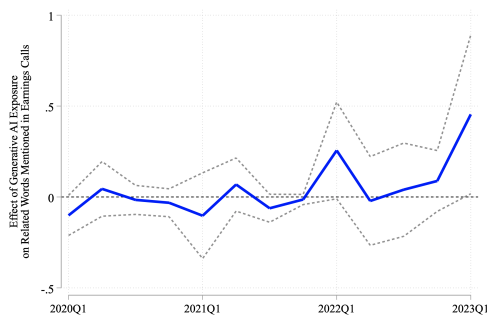


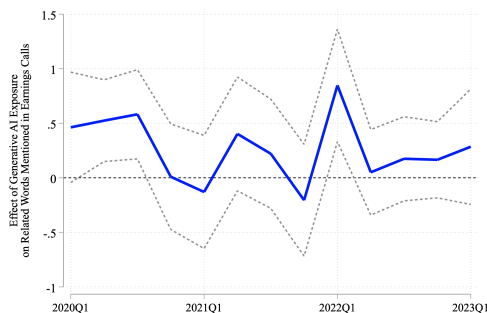
Figure C3: Firm-level Generative AI exposure and topic mentions in company earnings conference calls: excluding “tech sector”. Each graph show the result of estimating regression specifications of the form

$$\mathbb{1}[\text{Topic X}]_{i,t} = \alpha_t + \beta_t^X \text{GenAI Exp}_i + \gamma \mathbb{1}[\text{Topic X}]_{i,2019} + \varepsilon_{i,t}$$

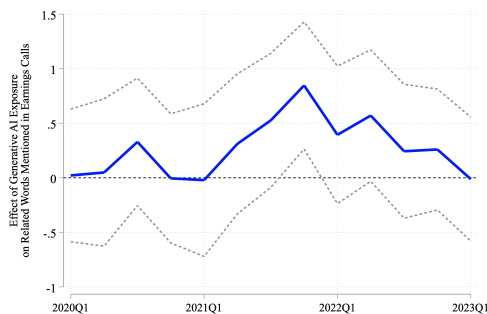
for each fiscal quarters from 2019 Q1 to 2023 Q1 for each topic using a manually collected panel of quarterly earnings conference call transcripts for S&P 500 firms from the *Seeking Alpha* website. See Table 5 for keywords about Generative AI topic, ML/DL/NLP topic, and Engineering topic. *The sample excludes all firms in the industry “tech sectors” most likely involved in producing AI-related products, which we define following Acemoglu et al. (2022) as NAICS 51 (information) and 54 (professional and business services).* Dotted lines show 95% confidence intervals based on heteroskedasticity-robust standard errors.



(A) Firm mentioning Generative AI and firm Generative AI exposure



(B) Firm mentioning ML/DL/NLP and firm Generative AI exposure



(C) Firm mentioning Engineering and firm Generative AI exposure

Appendix D. Appendix Tables

Occupation	Task ID	Task	GPT Score	GPT Explanation
Secretaries and Administrative Assistants, Except Legal, Medical, and Executive Nuclear Technicians	20283	Operate office equipment, such as fax machines, copiers, or phone systems and arrange for repairs when equipment malfunctions.	0	The task requires physical interaction with office equipment, which cannot be done by the LLM.
Nuclear Technicians	22315	Monitor instruments, gauges, or recording devices under direction of nuclear experimenters.	0	The task requires monitoring physical instruments and gauges, which cannot be done by the LLM.
Veterinary Technologists and Technicians	4180	Administer emergency first aid, such as performing emergency resuscitation or other life saving procedures.	0	The task requires hands-on physical intervention and cannot be completed through text input/output alone.
Penetration Testers	21754	Document penetration test findings.	1	The model can be trained on penetration testing and can generate text output based on the input.
Gambling Dealers	4453	Open and close cash floats and game tables.	0	The task requires physical interaction with cash and game tables, and the LLM cannot perform any physical actions.
Agricultural Engineers	5337	Meet with clients, such as district or regional councils, farmers, and developers, to discuss their needs.	0	The task requires human interaction and communication skills, which cannot be replaced by the LLM.
Farm Equipment Mechanics and Service Technicians	13750	Maintain, repair, and overhaul farm machinery and vehicles, such as tractors, harvesters, and irrigation systems.	0	The task requires physical maintenance and repair of machinery and vehicles, which cannot be done by the LLM.
Adult Basic Education, Adult Secondary Education, and English as a Second Language Instructors	6846	Meet with other professionals to discuss individual students' needs and progress.	1	The model can assist in writing and responding to emails, including those that involve discussing student progress with other professionals.
Payroll and Timekeeping Clerks	2526	Distribute and collect timecards each pay period.	0	The task does not involve writing or transforming text, nor does it require any complex decision-making or analysis that the LLM could assist with.
Environmental Engineering Technologists and Technicians	3647	Obtain product information, identify vendors or suppliers, or order materials or equipment to maintain inventory.	2	The model can help identify vendors or suppliers by searching the internet and retrieving information. It can also help order materials or equipment by generating text that can be sent to suppliers.

Table D1: Examples of GPT scores assigned to task statements and GPT-provided explanations.

SOC Code	Occupation Title	Exposure Score
41-9041	Telemarketers	.96
43-9081	Proofreaders and copy markers	.95
43-3031	Bookkeeping, accounting, and auditing clerks	.87
15-2021	Mathematicians	.86
15-1251	Computer programmers	.85
43-9022	Word processors and typists	.85
43-3011	Bill and account collectors	.83
27-3091	Interpreters and translators	.82
43-9111	Statistical assistants	.82
15-1254	Web developers	.81
43-6011	Executive secretaries and executive administrative assistants	.77
43-3051	Payroll and timekeeping clerks	.77
43-6014	Secretaries and administrative assistants, except legal, medical, and executive	.77
43-5061	Production, planning, and expediting clerks	.76
15-1212	Information security analysts	.75
43-6013	Medical secretaries and administrative assistants	.75
27-3043	Writers and authors	.75
43-4021	Correspondence clerks	.74
43-9061	Office clerks, general	.74
41-3091	Sales representatives of services, except advertising, insurance, financial services, and travel	.73
:	:	:
39-5093	Shampooers	0
51-6041	Shoe and leather workers and repairers	0
51-6042	Shoe machine operators and tenders	0
51-3023	Slaughterers and meat packers	0
47-2022	Stonemasons	0
47-2221	Structural iron and steel workers	0
51-2041	Structural metal fabricators and fitters	0
29-9093	Surgical assistants	0
51-6052	Tailors, dressmakers, and custom sewers	0
47-2082	Tapers	0
49-9052	Telecommunications line installers and repairers	0
47-2053	Terrazzo workers and finishers	0
51-6064	Textile winding, twisting, and drawing out machine setters, operators, and tenders	0
47-2044	Tile and stone setters	0
51-9197	Tire builders	0
49-3093	Tire repairers and changers	0
51-4194	Tool grinders, filers, and sharpeners	0
39-3031	Ushers, lobby attendants, and ticket takers	0
49-9064	Watch and clock repairers	0
53-7073	Wellhead pumpers	0

Table D2: Highest and lowest Generative AI exposure score occupations

NAICS Code	Industry Title	Exposure Score
54	Professional, Scientific, and Technical Svcs	.41
51	Information	.4
52	Finance and Insurance	.4
56	Admin. and Support and Waste Mgmt and Remed. Svcs	.39
11	Agriculture, Forestry, Fishing and Hunting	.39
61	Educational Svcs	.39
42	Wholesale Trade	.38
33	Manufacturing: Other Products	.36
32	Manufacturing: Textile, Apparel, and Leather	.36
22	Utilities	.35
53	Real Estate and Rental and Leasing	.35
71	Arts, Entertainment, and Recreation	.33
62	Health Care and Social Assistance	.33
31	Manufacturing: Food, Beverage, and Tobacco	.33
81	Other Svcs (except Public Administration)	.33
48	Transp. and Warehousing: Transportation	.32
21	Mining, Quarrying, and Oil and Gas Extraction	.32
23	Construction	.3
49	Transp. and Warehousing: Warehousing and Storage	.3
45	Retail Trade: Other Merchandise	.3
72	Accommodation and Food Svcs	.28
44	Retail Trade: Motor Vehicle and Parts Dealers	.27

Table D3: Generative AI exposure scores by industry

Table D4: Generative AI exposure for the Largest 100 U.S. Firms This table lists

the Generative AI exposure scores for the largest 100 publicly-traded firms with headquarters in the U.S., where size is measured as the market capitalization as of November 1, 2022. *Generative AI exposure* is the firm’s labor exposure defined in Section I. *MktCap* is the firm’s market capitalization as of November 1, 2022, in \$B. *Sector* is defined at the NAICS 2-digit level.

Company Name	Gen. AI exposure	MktCap	Sector
International Business Machines Corp	0.488	125	Information
Intuit Inc.	0.480	111	Information
QUALCOMM Inc.	0.479	132	Manufacturing
Fiserv Inc	0.475	66	Information
NVIDIA Corporation	0.468	337	Manufacturing
S&P Global Inc	0.452	103	Administrative and Support and Waste Management and Remediation Services
Broadcom Inc	0.449	195	Manufacturing
Verizon Communications Inc	0.444	157	Information
Microsoft Corp	0.442	1,701	Information
3M Co	0.442	69	Manufacturing
Advanced Micro Devices Inc	0.441	96	Manufacturing
ServiceNow Inc	0.434	85	Information
Adobe Inc	0.427	147	Information
PayPal Holdings Inc	0.418	96	Information
Thermo Fisher Scientific Inc	0.411	203	Manufacturing
Intuitive Surgical Inc	0.404	87	Manufacturing
Automatic Data Processing Inc	0.398	101	Information
Comcast Corp	0.396	136	Information
Vertex Pharmaceuticals Inc	0.395	81	Manufacturing
Analog Devices Inc	0.392	74	Manufacturing
AbbVie Inc	0.391	260	Manufacturing
Regeneron Pharmaceuticals Inc	0.390	81	Manufacturing
Gilead Sciences Inc	0.388	99	Manufacturing
Micron Technology Inc.	0.388	60	Manufacturing
Intel Corp	0.386	117	Manufacturing
Bristol-Myers Squibb Co	0.385	165	Manufacturing
Illinois Tool Works Inc.	0.382	66	Manufacturing
Netflix Inc	0.381	128	Real Estate and Rental and Leasing
Meta Platforms Inc	0.381	217	Information
Lam Research Corp	0.380	56	Manufacturing
SALESFORCE INC	0.379	160	Information
General Dynamics Corp	0.378	69	Manufacturing
Abbott Laboratories	0.376	174	Manufacturing
AT&T Inc	0.375	131	Information
Applied Materials Inc	0.374	77	Manufacturing
Booking Holdings Inc	0.373	75	Information
General Electric Co	0.373	85	Wholesale Trade
Merck & Co Inc	0.372	253	Manufacturing
T-Mobile US Inc	0.371	189	Information
Johnson & Johnson	0.371	453	Manufacturing
Honeywell International Inc	0.368	137	Manufacturing
Alphabet Inc	0.366	546	Information
Amgen Inc	0.365	146	Manufacturing
Eli Lilly and Co	0.364	335	Manufacturing
Apple Inc	0.364	2,397	Manufacturing
Philip Morris International Inc	0.364	142	Manufacturing
DEERE & COMPANY	0.364	117	Manufacturing
Texas Instruments Inc	0.363	148	Manufacturing
Caterpillar Inc	0.358	115	Manufacturing
CVS Health Corp	0.356	124	Health Care and Social Assistance
Cisco Systems Inc	0.355	187	Manufacturing
Zetis Inc	0.353	71	Manufacturing
Pfizer Inc	0.352	269	Manufacturing
Southern Co (The)	0.351	71	Utilities
DanaHER Corp	0.350	186	Manufacturing
Procter & Gamble Co (The)	0.342	329	Manufacturing
Raytheon Technologies Corp	0.339	140	Manufacturing
Colgate-Palmolive Co	0.337	62	Manufacturing
Becton Dickinson and Co	0.331	67	Manufacturing
Dominion Energy Inc	0.330	58	Utilities
NextEra Energy Inc	0.329	154	Utilities
Walt Disney Co (The)	0.328	193	Information
Altria Group Inc	0.327	83	Manufacturing
Air Products and Chemicals Inc.	0.326	56	Manufacturing
Waste Management Inc.	0.325	64	Administrative and Support and Waste Management and Remediation Services
Duke Energy Corp	0.322	72	Utilities
EOG Resources Inc.	0.322	80	Mining, Quarrying, and Oil and Gas Extraction
Exxon Mobil Corp	0.320	466	Manufacturing
Amazon.com Inc	0.317	987	Retail Trade
Stryker Corp	0.317	83	Manufacturing
Schlumberger Ltd	0.316	73	Mining, Quarrying, and Oil and Gas Extraction
Conocophillips	0.316	163	Mining, Quarrying, and Oil and Gas Extraction
HCA Healthcare Inc	0.312	63	Health Care and Social Assistance
Marathon Petroleum Corp	0.308	59	Manufacturing
Occidental Petroleum Corp	0.307	69	Mining, Quarrying, and Oil and Gas Extraction
Coca-Cola Co (The)	0.306	258	Manufacturing
Boston Scientific Corp	0.305	61	Manufacturing
PepsiCo Inc	0.303	249	Manufacturing
Chevron Corp	0.301	357	Manufacturing
Berkshire Hathaway Inc	0.300	324	Finance and Insurance
Lockheed Martin Corp	0.299	127	Manufacturing
Boeing Co	0.298	85	Manufacturing
Sherwin-Williams Co (The)	0.296	58	Manufacturing
Activision Blizzard Inc	0.295	57	Information
Pioneer Natural Resources Co	0.294	60	Mining, Quarrying, and Oil and Gas Extraction
Mondelez International Inc	0.292	85	Manufacturing
Northrop Grumman Corp	0.291	82	Manufacturing
Tesla Inc	0.283	719	Manufacturing
Home Depot Inc. (The)	0.261	303	Retail Trade
United Parcel Service Inc	0.256	123	Transportation and Warehousing
CSX Corp	0.256	61	Transportation and Warehousing
Union Pacific Corp	0.253	121	Transportation and Warehousing
Costco Wholesale Corp	0.252	221	Retail Trade
TJX Companies Inc (The)	0.243	83	Retail Trade
Lowe's Cos Inc	0.238	120	Retail Trade
Walmart Inc	0.235	385	Retail Trade
Target Corp	0.235	76	Retail Trade
Dollar General Corporation	0.212	57	Retail Trade
McDonald's Corp	0.194	201	Accommodation and Food Services
Starbucks Corp	0.119	100	Accommodation and Food Services

Table D5: Portfolio sorts on Generative AI exposure—Placebo test. This table reports daily excess stock returns of five value-weighted portfolios of firms sorted on Generative AI exposure during the placebo periods. *AMH* is the "Artificial Minus Human" zero net investment portfolio long high exposure (*H*) stocks and short low exposure (*L*) stocks. On Nov 14, 2022, we sort NYSE stocks into five quintiles by their Generative AI exposure, and we assign all stocks into the portfolios using the NYSE breakpoints. We aggregate stock returns within each portfolio using the stocks' market capitalization on the prior trading day as the weight. *ChatGPT release period* in this placebo test is from Oct 31, 2022 - Nov. 15, 2022 (prior to the actual release of ChatGPT), and *Not ChatGPT release period* include other days from Oct. 31, 2022 to March 31, 2023. Panel A reports the portfolios' excess returns net of the risk free rate. Panel B reports the portfolios' alpha after controlling for the market factor. Panel C reports the Fama-French 5-factor adjusted alphas. Panel D reports the Fama-French 5-factor adjusted alphas for a sample that excludes firms in the tech sectors NAICS 51 (information) and NAICS 54 (business services). *t*-statistics in parentheses are computed using Newey-West standard errors with five day lags.

Sample	Portfolios					
	Q1	Q2	Q3	Q4	Q5	AMH
<i>A: Excess returns (%)</i>						
Not ChatGPT release period	0.026 (0.22)	0.020 (0.18)	0.041 (0.41)	0.033 (0.30)	0.110 (0.97)	0.068 (1.31)
ChatGPT release period	0.183 (0.45)	0.155 (0.30)	0.135 (0.21)	0.371 (0.81)	0.389 (0.60)	0.192 (0.60)
<i>B: Market-factor-adjusted alpha (%)</i>						
Not ChatGPT release period	0.007 (0.17)	-0.002 (-0.05)	0.020 (0.57)	0.013 (0.36)	0.086 (2.66)	0.063 (1.07)
ChatGPT release period	0.005 (0.03)	-0.045 (-0.61)	-0.066 (-0.67)	0.182 (2.36)	0.160 (1.58)	0.142 (0.61)
<i>C: FF 5-factor-adjusted alpha (%)</i>						
Not ChatGPT release period	0.038 (0.95)	0.019 (0.60)	-0.021 (-0.83)	0.034 (1.29)	0.063 (2.32)	0.008 (0.15)
ChatGPT release period	0.051 (0.55)	-0.035 (-0.51)	-0.111 (-1.85)	0.155 (2.51)	0.172 (1.79)	0.107 (0.68)
<i>D: FF 5-factor-adjusted alpha ex. tech sector (%)</i>						
Not ChatGPT release period	0.039 (0.95)	0.015 (0.42)	-0.016 (-0.53)	-0.009 (-0.40)	0.086 (2.74)	0.031 (0.54)
ChatGPT release period	0.037 (0.38)	-0.013 (-0.21)	-0.119 (-2.19)	0.113 (1.56)	0.260 (2.22)	0.210 (1.56)

Table D6: Portfolio sorts on Generative AI exposure—Alternative exposure measures. This table reports daily excess stock returns of five value-weighted portfolios of firms sorted on different measures of Generative AI exposure. *AMH* is the “Artificial Minus Human” zero net investment portfolio long high exposure (*H*) stocks and short low exposure (*L*) stocks. On Nov 14, 2022, we sort NYSE stocks into five quintiles by their Generative AI exposure, and we assign all stocks into the portfolios using the NYSE breakpoints. We aggregate stock returns within each portfolio using the stocks’ market capitalization on the prior trading day as the weight. *ChatGPT release period* is from Nov 30, 2022 - Dec. 14, 2022, and *Not ChatGPT release period* include other days from Nov 15, 2022 to March 31, 2023. Panels A and B differ from the baseline approach to constructing the Generative AI exposure measure by using only the share of tasks within each occupation that got a “direct exposure” categorization (score of 1) to measure the exposure. Panels C and D differ from the baseline approach in that we weight each tasks within an occupation by the level of its “importance” rating for that occupation assigned by O*Net. The aggregation to the firm level from occupation-level exposure scores is unchanged from the baseline in both alternative exposure score approaches. Panels A and C reports the portfolios’ alphas after adjusting for the market factor. Panel B and D report the Fama-French 5-factor adjusted alphas. *t*-statistics in parentheses are computed using Newey-West standard errors with five day lags.

Sample	Portfolios					
	Q1	Q2	Q3	Q4	Q5	AMH
<i>A: Direct exposure only Mkt.-factor adjusted returns (%)</i>						
Not ChatGPT release period	0.037 (0.88)	-0.009 (-0.23)	0.006 (0.18)	0.000 (0.01)	0.086 (2.25)	0.033 (0.49)
ChatGPT release period	-0.293 (-9.11)	0.118 (2.98)	0.007 (0.21)	-0.043 (-0.80)	0.196 (4.49)	0.473 (6.53)
<i>B: Direct exposure only FF 5-factor-adjusted alpha (%)</i>						
Not ChatGPT release period	0.069 (1.95)	-0.002 (-0.07)	-0.003 (-0.13)	-0.004 (-0.14)	0.068 (2.46)	-0.017 (-0.32)
ChatGPT release period	-0.177 (-2.77)	0.167 (3.13)	-0.049 (-0.81)	-0.039 (-0.81)	0.129 (2.84)	0.291 (3.47)
<i>C: Imp.-weighted exposure Mkt.-factor adjusted returns (%)</i>						
Not ChatGPT release period	0.047 (1.10)	0.003 (0.08)	0.011 (0.30)	0.018 (0.44)	0.066 (1.91)	0.003 (0.04)
ChatGPT release period	-0.228 (-5.23)	-0.048 (-0.45)	0.073 (0.98)	-0.025 (-0.71)	0.205 (4.77)	0.418 (5.42)
<i>D: Imp.-weighted exposure FF 5-factor-adjusted alpha (%)</i>						
Not ChatGPT release period	0.068 (1.62)	0.018 (0.56)	-0.022 (-0.78)	0.032 (1.09)	0.053 (1.95)	-0.032 (-0.56)
ChatGPT release period	-0.163 (-2.20)	0.009 (0.09)	-0.018 (-0.51)	0.028 (0.72)	0.156 (3.51)	0.303 (3.34)

Table D7: Robustness in firm-level effects of Generative AI exposure on returns: exposure terciles. The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023. The market factor and risk free returns are obtained from Ken French's website. The table shows selected coefficients estimated from regressions of the form

$$r_{it} - r_{ft} = \sum_{x \in \{2,3\}} (\beta_1^x \mathbb{I}[\text{Release Period}]_t \times \mathbb{I}[\text{GenAI Exposure Tercile} = x]_i) + \beta_2 \mathbb{I}[\text{Release Period}]_t \times \text{Firm Char}_i + FE_{s,t} + \chi' \Gamma + \varepsilon_{it},$$

where $\mathbb{I}[\text{GenAI Exposure Tercile} = x]$ is a dummy variable indicating that firm i 's Generative AI exposure is in a particular tercile among equal-weighted NYSE firms sorted within each NAICS 3-digit industry, β_1^x measures the impact of ChatGPT release on firms in the x th tercile of Generative AI exposure relative to other firms, β_2 measures the differential impact on firms with particular other characteristics, $FE_{s,t}$ represents trading day fixed effects and NAICS 3-digit industry fixed effects, and the additional controls Γ include the Fama French 5-factor returns on each trading day, as well as interactions between the firm characteristic and the firm exposure tercile with the non-release period indicator. The intercept for the panel is allowed to vary with whether the day is in the ChatGPT release period consisting of Nov 30, 2022 - Dec. 14, 2022. All regressions also control for the interaction of the firm characteristic of interest and the exposure tercile indicators with the non-release period indicators. Column (1) shows the baseline effect of high exposure at a firm on daily returns during the post-ChatGPT two week release period. Column (2) allows for differential returns during the release period in NAICS sectors 51 (information) and 54 (business services) - the "tech sectors" - relative to the rest of the sample. Columns (3)-(7) control for differential returns during the release period associated with characteristics computed from Compustat data: Log Size ($\ln(AT)$), Tobin's Q ($\frac{AT - CEQ + PRCC + CSHO}{AT}$), Return on Assets (ROA), Labor Intensity ($\ln \frac{EMP}{PPEVT}$), and Tangibility ($\frac{PPEVT}{AT}$). Column (8) controls for differential trends in all of these characteristics during the release period at the same time. Note that the reference category for the release period tercile coefficients is the first tercile return during the release period. The regressions are estimated with firm value weights for each firm-by-trading-day cell. T-statistics in parentheses are computed using standard errors clustered at the firm level.

Dep. var.:	Daily Firm-level Excess Returns (ppt)								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
$\mathbb{I}[\text{Release Period}]_t \times [\text{Q3 GAI Exp.}]_i$	0.274** (2.542)	0.260** (2.551)	0.253** (2.227)	0.276*** (2.615)	0.278*** (2.684)	0.279** (2.523)	0.246** (2.344)	0.195** (2.212)	
$\mathbb{I}[\text{Release Period}]_t \times [\text{Q2 GAI Exp.}]_i$	0.136 (1.442)	0.133 (1.405)	0.120 (1.246)	0.165** (2.072)	0.165* (1.705)	0.142 (1.481)	0.127 (1.410)	0.124 (1.613)	
$\mathbb{I}[\text{Release Period}]_t \times [\text{Tech Sector}]_i$		0.126 (1.264)						0.124 (1.044)	
$\mathbb{I}[\text{Release Period}]_t \times [\text{Log Size}]_i$			-0.040** (-2.025)					-0.050** (-2.546)	
$\mathbb{I}[\text{Release Period}]_t \times [\text{Tobin's Q}]_i$				0.024 (0.975)				0.010 (0.338)	
$\mathbb{I}[\text{Release Period}]_t \times [\text{ROA}]_i$					0.660** (2.503)			0.468 (0.797)	
$\mathbb{I}[\text{Release Period}]_t \times [\text{Labor Intensity}]_i$						-0.005 (-0.200)		-0.059* (-1.804)	
$\mathbb{I}[\text{Release Period}]_t \times [\text{Tangibility}]_i$							-0.106 (-0.543)	-0.477* (-1.683)	
R-squared	0.30	0.30	0.30	0.30	0.30	0.30	0.30	0.30	
Observations	183272	183272	183272	183272	182802	181110	182426	180922	
				<i>Controls & fixed effects</i>					
Three-digit industry FEs	X	X	X	X	X	X	X	X	
Trading day FEs	X	X	X	X	X	X	X	X	
Non-release period × Characteristic	X	X	X	X	X	X	X	X	

Table D8: Generative AI exposure and firm data value—Additional data value proxies. This table reports the results of regressing firms’ daily excess returns on their Generative AI exposure interacted with ChatGPT and Not ChatGPT release period dummies, and proxies for the value of a firm’s data using the following specification:

$$r_{it} - r_t^f = (\beta_1 \mathbb{1}[\text{ChatGPT release period}]_t + \beta_2 \mathbb{1}[\text{Not ChatGPT release period}]_t) \times \text{GenAI Exp}_i + \text{Ind-Day FEs} + \varepsilon_{it} \\ + (\beta_3 \mathbb{1}[\text{ChatGPT release period}]_t + \beta_4 \mathbb{1}[\text{Not ChatGPT release period}]_t) \times \text{Data Value Proxy}_i \\ + (\beta_5 \mathbb{1}[\text{ChatGPT release period}]_t + \beta_6 \mathbb{1}[\text{Not ChatGPT release period}]_t) \times \text{Data Value Proxy}_i \times \text{GenAI Exp}_i$$

The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023. Risk free returns are obtained from Ken French’s website. *Ind-Day FEs* represents fixed effects of NAICS 3-digit industry interacted with trading day. The regressions are estimated using the stock’s prior trading day’s market capitalization as weights. The first column shows the baseline effect of higher firm Generative AI exposure on returns during the release period and outside of it. All other columns show the triple interactions between this Generative AI exposure effect in different periods and different proxies for a firm’s access to data with value for LLM analytics: We assess the degree to which a firm is likely to have data that can be used as an input into LLM analytics, using a similar methodology to our main measure by reviewing the business description section in the firm’s recent annual report using a large language model (see more details in Appendix A). In particular, the LLM is asked to assess whether a firm’s business description suggests access to relevant data based on its coverage of 6 categories: the general nature of the company’s business, the scale and reach of the firm, data collection mechanisms, data utilization, data infrastructure & management, and data regulation and privacy. *10K Data Value: High Subscore* in column 1 is a binary indicator for whether any of the subcategory scores was scored as a 3 (high relevance). *Data Mgmt Skill Share* is computed from the LinkedIn employment distribution at the firm in 2022, by first predicting the likelihood that a U.S. job posting for a 6-digit SOC occupation mentions at least 3 data management skills (classified based on Abis and Veldkamp (2023)) in 2017-2021 to get the probability that a job is a “data management” position, and then aggregating those probabilities for each firm using its 2022 employment distribution. That is, the variable represents the predicted share of a firm’s positions on LinkedIn that are “data management”-intensive. *textitHigh Data Mgmt* correspondingly defines a binary variable for firms in the top tercile of this predicted data management intensity among its workers. *t*-statistics in parentheses are computed using standard errors clustered at the firm level.

Dep. var.:	Daily Firm-level Excess Returns (ppt)		
	(1)	(2)	(3)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{GenAI Exp}_i$	0.649 (0.885)	-1.186 (-1.100)	-1.267* (-1.697)
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{GenAI Exp}_i$	0.349 (0.947)	-0.775 (-1.317)	0.079 (0.267)
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{GenAI Exp}_i \times 10\text{K Data Value: High Subscore}_i$	2.904*** (3.301)		
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{GenAI Exp}_i \times 10\text{K Data Value: High Subscore}_i$	0.163 (0.366)		
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{GenAI Exp}_i \times \text{Data Mgmt Skill Share}_i$		115.169** (2.007)	
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{GenAI Exp}_i \times \text{Data Mgmt Skill Share}_i$		8.414 (0.225)	
$\mathbb{1}[\text{ChatGPT release period}]_t \times \text{GenAI Exp}_i \times \text{High Data Mgmt Skill}_i$			5.738*** (4.113)
$\mathbb{1}[\text{Not ChatGPT release period}]_t \times \text{GenAI Exp}_i \times \text{High Data Mgmt Skill}_i$			0.242 (0.343)
R-squared	0.53	0.52	0.52
Observations	168230	179512	179512
	<i>Controls & fixed effects</i>		
Three-digit industry \times trading day FEs	X	X	X
Release / non-release period \times Characteristic	X	X	X