

Generative AI and Firm Values*

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

May 5, 2023

Abstract

What are the effects of recent advances in Generative AI on the value of firms? Our study offers a quantitative answer to this question for U.S. publicly traded companies based on the exposures of their workforce to Generative AI. Our novel firm-level measure of workforce exposure to Generative AI is validated by data from earnings calls, and has intuitive relationships with firm and industry-level characteristics. 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. Although this release was generally received by investors as good news for more exposed firms, there is wide variation across and within industries, consistent with the substantive disruptive potential of Generative AI technologies.

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

*We thank participants in the UCLA Anderson Finance Bag Lunch.

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

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

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

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. As a result, Generative AI has the potential for widespread corporate adoption, with implications for firm values both across and within a wide array of industries.

One of the biggest questions surrounding advances in Generative AI is what effect these technologies will have on corporate valuations as a result of the impact of Generative AI on firms' labor inputs. We construct a novel dataset containing firm-level workforce exposures to Generative AI. We provide a quantitative measure of the impact of Generative AI based on our firm-level exposure data combined with financial market data. Using this measure we compute the first estimates of the effect of Generative AI on firm values by studying the impact of the release of ChatGPT on firms with varying exposures to the technology shock.¹

We measure the impact of a major event in the advancement and dissemination of Generative AI technology, namely, the public release of ChatGPT, on equity returns at the firm level. This event had a substantial impact on firm returns, consistent with Generative AI advancement representing a major technological shock, one for which we can measure the arrival and impact in 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 and earnings call mentions of Generative AI increased substantially following the release of ChatGPT. Moreover, the massive information gathering and processing ability of ChatGPT itself allows us to assess each firm's exposure to ChatGPT's disruption in real-time.

Our key finding is that the arrival of ChatGPT had a sizable positive effect on the value of firms whose labor forces are more exposed to Generative AI and related Large Language Models (LLMs). Firms with higher exposure to the release of ChatGPT, as measured by the exposure of their labor force to being made more productive by tools like ChatGPT, outperform firms with lower exposures by over 40 basis points in daily excess returns during the two weeks following its release. Notably, these return differences are not only due to differences in labor force exposures across industries. Returns of firms with high labor force exposures also outperform firms with low exposures by about 40 basis points daily in

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.

industry-neutral portfolios.

Our methodology builds on the idea that ChatGPT and related technologies will increase firm-level free cash flows through a labor effect that can work through two potential channels. First, firms whose labor force can be substituted for with cheaper Generative AI-based capital will experience higher free cash flows by lowering input costs.² Second, firms whose labor inputs are more complementary to Generative AI will experience higher cash flows due to the technological improvement in an input that is complementary to their workforce.³ While we do not take a stand on whether (and for which workers) Generative AI is a substitute for, or a complementary to, labor, we are able to show that 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. While advances in Generative AI can have effects through the product market as well as through the labor market (for example, increasing demand for cloud computing services), our results support the idea that AI advances will have a broad impact on the economy through its effects on labor inputs. 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.⁴

We measure firm-level exposure to Generative AI in two steps. First, we build on Eloundou et al. (2023) and 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. Following Eloundou et al. (2023), we aggregate the task-level exposure measures to the occupations in the O*NET database. Second, and novel to our analysis, 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.

We next validate our labor-based measure of firms' exposure to Generative AI by examin-

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 Eisfeldt, Falato, and Xiaolan (2022).

⁴See, for example, Gutiérrez and Philippon (2019) and Akcigit and Ates (2020).

ing firms' earnings call transcripts in 2023. 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 following the release of ChatGPT. In contrast, firms with higher exposure to Generative AI do not increase discussions common technological topics such as Engineering following the release of ChatGPT. Moreover, these findings remain even after we exclude all firms from the most IT-related sectors,⁵ 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.

We start by showing the types of occupations that will be affected by advances in Generative AI. We find that the most affected occupations are those that involve non-routine cognitive tasks. This is in stark contrast with prior findings that automation mainly displaces occupations involving routine tasks (Autor et al. (2003)). Indeed, the most 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. Occupations with higher wages also have higher exposure to Generative AI. Our 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.

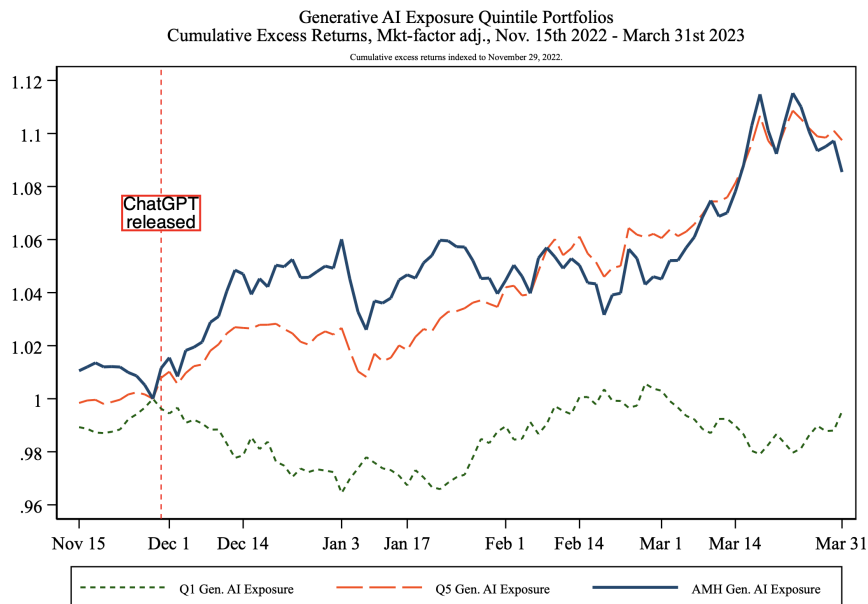
Exposure to Generative AI through firms' labor inputs has an intuitive relationship to average firm characteristics across and within industries. At the industry level, more exposed sectors have higher wages, consistent with those sectors employing more workers in higher-paid occupations that also tend to be more exposed to Generative AI. Regarding labor inputs, firms in more exposed industries tend to have higher labor intensity in terms of the number of employees per unit of capital, and lower asset tangibility. More exposed firms also have higher ratios of organizational to total capital.⁶ For the characteristics related to firm valuation, more exposed sectors have lower average firm size as measured by total assets and higher Tobin's Q. Importantly, we also observe similar relationships between firms' exposure to Generative AI and firm characteristics within industry sectors. The robust patterns of variation in industry and firm-level exposures with firm characteristics support our study of stock returns both across and within industries.

Firms with higher exposure to Generative AI experience higher volatility around the release of ChatGPT. However, it appears that it takes some time for the information in ChatGPT's release to be impounded into stock prices. The cumulative excess returns for

⁵To be precise, we exclude the NAICS 51 "Information" and NAICS 54 "Professional, Scientific, and Technical Services" sectors.

⁶See Eisfeldt and Papanikolaou (2014) and Eisfeldt and Papanikolaou (2013).

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 highest-exposure quintile of firms vs. the lowest-exposure quintile diverge for several weeks following the release of ChatGPT. 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”. Cumulative returns to holding the AMH portfolio that is long the highest-exposure quintile, and short the lowest-exposure quintile from the released date through March 31, 2023, are over 9%.

We study the effect of Generative AI on firm values by comparing 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. The effects are substantial, and monotonic, within industries across Generative AI-exposure quintiles. Adjusting for the market factor, the excess returns to quintile portfolios formed based on firm-level occupational exposure to Generative AI are monotonically increasing, with the highest-exposure quintile of firms within industries earning positive excess daily returns of over 40 basis points while the lowest exposure quintile experiences negative excess returns of around 25 basis points. The fact that these strong effects exist *within* industries for many industries provides evidence that Generative AI can have a broad impact on firm values through the effects on their labor inputs.

Across industries, the effects of Generative AI on firm value also vary widely. Publishing, information and computing-related industries have positive returns following the release of ChatGPT, while finance and transportation-related industries experience negative returns overall. Dispersion in industry returns is much higher during the two-week period following the release of ChatGPT than over the full sample from November 30, 2022 to March 31, 2023 overall.

Our within-industry results also display striking differences across industry sectors. Within finance, the return of more exposed firms relative to less exposed firms is substantially and significantly positive. Combined with the overall negative industry effect, this is consistent with some firms benefiting greatly from Generative AI advances while overall the impact of the release of ChatGPT was negative for value in the finance industry. Firms with higher exposures to Generative AI within manufacturing as well as the administrative support, waste management, and remediation services industry also significantly outperform firms with lower exposures. On the other hand, firms with higher exposures in the real estate and rental and leasing industry significantly underperform firms with lower exposures. This could mean that existing firms with large exposures to Generative AI may be displaced by new entrants in those industries. Finally, several industries do not display significant return spreads following the release of ChatGPT, including construction of buildings, mining, and heavy and civil engineering construction. The negligible impact in these industries is consistent with manual tasks' lower exposure to Generative AI.

Our study contributes to the literature studying the impact of disruptive technologies on firm valuations.⁷ Papanikolaou (2011) and Kogan and Papanikolaou (2014) study the effects of investment-specific technological changes on asset prices. Eisfeldt and Papanikolaou (2013) and Eisfeldt and Papanikolaou (2014) study firms' exposure to the organization capital technology frontier. 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 (2022) study the effects of AI on firm growth, compensation, and workforce composition. See also Webb (2019) for the impact of AI on firms. 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 Gen-

⁷See Greenwood, Hercowitz, and Krusell (1997) for an early contribution on the long-run impacts of investment-specific technological change.

erative 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.

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.⁸

The paper proceeds as follows: Section I describes our data and measure of firms’ exposure to Generative AI. Section II provides descriptive facts about Generative AI exposures across occupations, industries, and firms. 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 presents our results documenting the substantial changes in firm valuations following the introduction of ChatGPT. Finally, Section V concludes.

I. Data and Measurement

We measure a firm’s labor exposure to Generative AI in two steps. First, we measure each occupation’s exposure to Generative AI based on the occupation’s task statements from the O*NET database. Second, we aggregate the occupation-level Generative AI exposure measure to the firm level using the firm-occupational employment data from the Revelio Labs Workforce Dynamics database.

A. Measuring occupational exposure to Generative AI

Occupational task data To assess whether an occupation will likely experience a change in absolute or relative productivity as a result of Generative AI models becoming widely available and used, we use a task-based approach. That is, similar to Eloundou et al. (2023),

⁸Indeed, as we draft this study, IBM, the company ranked #1 in our exposure to Generative AI measure among the largest U.S. firms (see Table II) announced to 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?utm_source=superhuman.beehiiv.comutm_medium=newsletterutm_campaign=ibm-starts-replacing-jobs-with-ai = $USIR = T$

we consider an occupation to be a set of tasks-to-be-done and evaluate for each task whether it can be done more productively using ChatGPT and 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 database, which provides a list of task statements created by practitioners or experts.⁹ A task statement is usually one sentence, and an occupation has on average 22 tasks. The 19,265 pairs of task statements and the occupations that they belong to then need to be coded as being exposed to Generative AI technologies or not.

Task scoring We build on the approach for scoring tasks that was suggested and validated by Eloundou et al. (2023). In particular, we use GPT itself to score exposure 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. Two advantages of using an LLM 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.¹⁰

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.¹¹ The model is given an overall rubric for scoring LLM 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 title, and the model returns a score. The scores capture whether the time taken to complete task is reduced by at least half, at constant quality, if the worker can access ChatGPT-like tools. The scores fall into the following categories: *E0* indicates no exposure as the tool is either insufficiently useful for this task, or cannot be brought to bear as a result of the intrinsic nature of the task, e.g. if it involves physical activities; *E1* is applied if a 50% reduction in completion time is already feasible with the existing large language model interfaces; *E2* requires that such a productivity gain is feasible, but only once the current capabilities of the model can be deployed through applications with further inputs (e.g. access to internet or proprietary databases), or if it is trained on domain-specific issues or data; *E3* is applied when the productivity increase would

⁹This data can be accessed via the O*Net database at <https://www.onetonline.org/>

¹⁰While 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).

¹¹The 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.

require image processing capabilities in addition to current text processing. Importantly, the model is asked not only to respond with the score but also to explain its reasoning, which allows the researcher to audit whether GPT is in fact understanding the prompt as intended and interpreting the task correctly. Note that this exceeds the auditing capabilities that are available in many instances of human text classification - and that outputting this additional information is enabled by the feature of LLMs that text generation is relatively cheap in terms of time cost.¹² A random sample of scored tasks together with the model’s explanations can be found in Appendix Table C1.

Consistency of Generative AI scoring To validate the consistency and replicability of our procedure, 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. (3) “Broad exposure:” Any score other than 0 has been assigned (this includes exposure conditional on image capabilities becoming further developed). 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 C2. 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.

Scoring occupations’ exposure to Generative AI We next aggregate tasks’ exposure 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 are affected by Generative AI. We follow Eloundou et al. (2023) and focus on an aggregation that takes into account that scores of 1 represent the current direct feasibility of productivity improvements, while exposure scores of 2 rely on investment in additional capabilities, such as interaction through custom-built applications or the ability to search local or online databases, that complement the current LLM chat interface. Therefore, our main measure of the share of an occupation’s exposed tasks counts both the number of tasks with exposure rubric 1, (N_1) and those with exposure rubric 2

¹²The 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 the large majority of cases, the model expresses “high confidence” in its assessment.

(N_2) but applies half the weight to the latter. That is, our exposure score at the occupation level for each occupation o is:

$$E_o = \sum_{\text{tasks in } o} \frac{N_1 + 0.5 * N_2}{N_0 + N_1 + N_2 + N_3}.$$

Finally, we aggregate across 8-digit O*Net occupation codes to the 6-digit SOC level 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, which would represent 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 productive 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 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 ppt. The inter-quartile range extends from 6% to 38% exposure.

B. 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. 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,

$$E_f = \sum_{\text{occupations in } f} EmpShare_{f,o} * E_o,$$

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. Summary statistics for the distribution of Generative AI exposure across this set of firms are shown in Table I.

C. *Other data*

Firm earnings call transcripts data We manually collect firm earnings call transcripts from the Seeking Alpha website for the tickers of all S&P 500 firms over the 2018-2023 period. To account for variations in the word forms, we applied standard natural language pre-processing techniques to the text data, such as tokenizing and lowercasing. We then counted the frequency of each keyword in each transcript and created a panel data set that contained each ticker, quarter, year, and word counts for particular topics as variables.

Social media attention data To measure social media attention to GPT and related technologies, we obtain data on Twitter mentions of “GPT” and “ChatGPT” by day for 2022 and 2023 from the media search platform Media Cloud.¹³

II. Descriptive Facts about Generative AI Exposures

A. *Generative AI Exposure and occupation characteristics*

Appendix Table C3 shows an overview of the 20 occupations with the highest and lowest Generative AI exposure scores.¹⁴ Note that among the highest exposure occupations, many, such as “Telemarketers”, “Computer programmers”, and “Interpreters and translators”, map closely onto 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.¹⁵

To better understand which occupation characteristics are associated with higher exposure, Panel A of Figure 2 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 are the relatively low-wage “Office and Administrative Support” occupations, which are

¹³URL: <https://mediacloud.org/>

¹⁴Note that the lowest score category only shows a subset of a larger set of occupations with zero Generative AI exposure.

¹⁵While “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 afterwards).” (Source: <https://pandaily.com/mathematician-terence-tao-comments-on-chatgpt/>)

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 Figure 2, 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.

Occupational skills and Generative AI exposure Our measure of exposure to Generative AI technologies is based on the ability of tools like ChatGPT to make certain tasks more productive. Thus, we would like to understand how our exposure measure relates to other classifications of occupations which have been previously been defined based on the tasks that are involved in them. In particular, Acemoglu and Autor (2011) suggest that technology-based productivity changes of past decades can be understood by scoring occupations based on the skills they involve. They suggest a characterization based on the degree to which an occupation involves particular combinations of routine vs. non-routine, cognitive vs. manual, and analytical vs. interpersonal aspects. To understand how their categories map onto occupations that can be made more productive by ChatGPT-like technologies, we regress our 6-digit occupation Generative AI exposure measure jointly on the set of occupational skill scores defined by Acemoglu and Autor (2011). That is, we run the following regression:

$$E_o = \alpha + \sum_s \beta_s * Skill_o + \varepsilon_o$$

The results are shown in Figure 3. 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.

The literature on previous waves of computer-based automation argued that routine work was most likely to be substituted by computers and to complement non-routine communication and problem-solving tasks (Autor et al., 2003). This “routinization” hypothesis assumed that “computers and computer-controlled equipment are highly productive and reliable at performing the tasks that programmers can script - and relatively inept at everything else” (Acemoglu and Autor, 2011). While routine tasks and jobs were taken over by computers, workers skilled in “abstract” tasks were in high demand, leading to wage polarization. However, while in the past abstract jobs requiring “problem-solving, intuition, persuasion, and creativity” (Acemoglu and Autor, 2011) appeared safe from substitution by computers, Figure 3 suggests that the labor market impact due to recent Generative AI advances may be different. 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 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. Combined with the unprecedented speed with which this wave of computer-based automation tools is being adopted, this portends that the effects on wages and inequality across different demographic groups may look very different this time around than for the automation waves of the past.

B. Summary 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 ppt. The firm-level measure of exposure to Generative AI spans from 27% at the 10th percentile to 44% at the 90th percentile. Figure 4 shows that the variation of firm-level exposure to Generative AI has both across-industry and within-industry components. While industry sectors such as “Information” and “Professional, Scientific, and Technical Services” have an average firm exposure to Generative AI of about 13 ppt greater than industries such as “Accommodation and Food Services,” there are substantial variations of firms' exposure to GPT within each industry. A variance decomposition shows that industry differences explain about 18% of the firm-level variation in exposure to Generative AI.

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.

The rich cross-industry and within-industry variations in firms' exposure to Generative AI motivate us to explore our firm-level empirical analyses within and across industries. Importantly, our within-industry analyses also help to differentiate our labor-based mech-

anism from the product-based mechanism when studying the effects of firms' exposure to Generative AI on firm values.

How do firms' and industries' exposure to Generative AI relate to their other 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 such as R&D. Indeed, we also observe that firms with high R&D intensity are more exposed to Generative AI. 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. These cross-firm findings hold consistently within-industry and cross-industry as well. Panel B of Table III shows that the above findings are qualitatively similar after we include NAICS 2-digit industry fixed effects in the regressions. Figure 5 plots the relationship between industry sectors' mean exposure to Generative AI (averaged across Compustat firms) and industries' other characteristics. In line with the firm-level patterns, we observe a consistent picture that industries with higher exposure to Generative AI have firms that are smaller, have higher Tobin's Q and lower current profitability, and feature higher organizational capital and lower tangible capital.

III. Investor Communication and Exposure to Generative AI Technology

To understand whether firms communicate that they are affected by the technological change as a result of the evolution of large language models and other Generative AI tools, and to validate our bottom-up measure of language model exposure based on firms' employment structure, we first analyze the public communication between firms and investors. We focus on earnings conference calls, as firms use earnings conference calls to communicate their views regarding risks and opportunities, (see, e.g. Hassan, Hollander, Van Lent, and Tahoun (2019)) as well as past and expected future performance. Moreover, they respond to questions by analysts who may reflect investors' perspectives of which issues particular firms should focus on. If the technological change resulting from the recent rapid evolution of large language models is indeed affecting firms as we postulate, we should see both a time pattern of increasing communication regarding these issues that coincides with the launch of ChatGPT, as well as a cross-sectional pattern of larger increases in this kind of firm-investor communication among firms for which our bottom-up measure predicts higher Generative AI exposure.

A. Measuring earnings conference call mentions of technologies

We use the Seeking Alpha website to manually collect a panel of the earnings conference call transcripts for S&P 500 firms, for calls that were held 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).

Topic definitions We define four categories of words (incl. their plural form where applicable) that we compare to the list of unigrams and bigrams in each earnings call: (1) Machine learning-related words that are not specific to Generative AI technologies.¹⁶ (2) Generative AI -specific words: “LLM”, “ChatGPT”, “GPT”, “GPT3”, “GPT4”, “generative”, “language model”. (3) Generic engineering-related words: “engineer”, “engineering”. Our sample size and composition vary somewhat across years, and we may not fully capture all earnings calls for each firm, but we capture the vast majority of relevant earnings calls, with the number of unique S&P 500 companies with calls in our sample in each quarter varying from 377 to 474. Moreover, our sampling procedure of hand-collecting publicly available transcripts from Seeking Alpha is unlikely to be biased with regard to Generative AI exposure - which would be the selection bias of relevance.

B. Aggregate trends in earnings call mentions

To see how mentions of these topics evolve over time, we compute two aggregate variables for each month and topic: First, we capture the share of calls in our sample that mention a word from that category, which captures the extensive margin of whether there is *any* sufficiently large perceived firm exposure for investors or firm management to mention it. Second, we compute the mean number of category words per call in our sample, which additionally captures the intensive margin of longer discussions of the topics. The aggregate trend in these variables for the Generative AI topic, which is our proxy for firms discussing ChatGPT-like technologies, is shown in Figure 6. For comparison, each graph also shows the equivalent trends for other machine learning topics, and for generic mentions of engineering. To allay concerns that only investors in software companies are likely to be knowledgeable about, or have an interest in, the Generative AI topic, we also show the same trends excluding the “Information” (NAICS 51) and “Professional, Scientific, and Technical Services” (NAICS

¹⁶These are: “deep learning”, “ML”, “machine learning”, “deep learning”, “natural language”, “neural net”, “neural network”, “NLP”.

54) sectors in Panels B and D, as determined by Compustat industry codes.

While the sample only reflects companies that had earnings calls before March 2023, there is an unmistakable break in the trend shortly after the release of ChatGPT in November 2022: both measures show that an increasing number of firms seem to consider language models and generative AI as important enough that they are discussed in earnings conference calls. Moreover, these trends are separate from discussions of machine learning or engineering in general, which do not show a large increase in 2023 – so this is not simply a reflection of companies in general - or in our sample - becoming more technical in their language over time. Generative AI-related topics were discussed in 27% of all earnings conference calls by March 2023, and in 13% of calls outside the software-related sectors. The number of mentions per call in Panel C rises even more steeply than the share of firms discussing the topic, suggesting that discussions of the topic increase at the intensive margin in addition to the extensive margin. Overall, earnings call mentions show that the rapid rise in importance of language model technologies is reflected in the communication between firms and analysts - suggesting that they are also likely to influence valuations by investors.

C. Firm-level Generative AI exposure and earnings call topics

While Figure 6 shows that general interest in LLM-related topics is on the rise following the release of ChatGPT in November 2022, we also want to understand *which* firms are more likely to discuss related topics. To see how our firm-level predicted Generative AI exposure relates to these changes in earnings call mentions, we run repeated cross-sectional regressions of the form,

$$\mathbb{1}[\text{Topic X}]_{i,t} = \alpha_t + \beta_t^X E_i^f + \gamma \mathbb{1}[\text{Topic X}]_{i,2019} + \varepsilon_{i,t} \quad (1)$$

where the dependent variable is a binary indicator for whether company i 's earnings calls in quarter t have *any* mentions of topic X , and E_i^f is the firm's Generative AI exposure. We also control for whether the firm already mentioned the topic in any 2019 earnings calls. This means that the coefficient β_t^X here estimates the degree to which firms with higher Generative AI exposure are more likely to start (or stop) discussing topic X in earnings calls in comparison to 2019.

The quarterly β_t^X coefficients over 2020-2023 for each topic are shown in Figure 7. The findings validate our Generative AI exposure measure's ability to pick up on firms' exposure to a technology shock driven by Generative AI. Panels A and B confirm that exposed firms – according to our bottom-up measure – were more likely to discuss Generative AI topics than other firms, both in the full sample and when we exclude sectors with many software

and technology companies. Moreover, they only became more likely to discuss Generative AI topics after the recent wave of innovation in those fields in Q4 of 2022 - so our exposure measure does *not* identify companies that were involved in discussions about these technologies before ChatGPT was released. Put differently, the fact that our Generative AI exposure measure does not predict a higher likelihood of mentioning these technologies *before* the recent productivity shock suggests that it does a reasonable job of identifying companies that are newly exposed to this particular way of using AI to improve productivity. Moreover, the similar patterns in Panels B and D, where we exclude firms in the “Information” (NAICS 51) and “Professional, Scientific, and Technical Services” (NAICS 54) sectors, confirms that this pattern is not limited to investor communication for IT companies. The effect on mentions of generic engineering topics in Panels C and D shows that the break in trend with regard to the effect of our Generative AI exposure measure on whether a firm mentions Generative AI does not extend to higher mentions of engineering topics more generally. This validates our firm-level exposure measure capturing the potential impact of technologies like ChatGPT in particular.

The magnitude of these effects is large: The β_t^X coefficients indicate the percentage point change in the probability that the firm mentions the topic in an earnings call in that year, relative to 2019, for each percentage point change in the share of tasks in that firm that is exposed to Generative AI. That is, the March 2023 coefficient in Panel A suggests that a 1 ppt increase in firm exposure is associated with a more than 1 ppt higher likelihood in 2023 of the firm mentioning Generative AI relative to 2019.

Firm-level panel regressions of earnings call topics on Generative AI exposure

In order to be able to more flexibly control for different firm characteristics in determining the association between our Generative AI exposure measure and topics mentioned in firm earnings conference calls, we also estimate firm-level panel regressions. We aggregate our monthly data into a firm-quarter panel in order to ensure that most firms have a continuous time series of earnings calls over the period of Q3 2018–Q1 2023. Then, we estimate regressions of the form

$$\mathbb{1}[\text{Gen. AI Topic Mentioned}]_{i,t} = \alpha_t + \alpha_i + \beta_1 E_i^f + \beta_2 E_i^f \times \mathbb{1}[\text{Post-ChatGPT}] \quad (2)$$

$$+ \gamma \mathbb{1}[\text{Post-ChatGPT}] + \varepsilon_{i,t} \quad (3)$$

where we regress whether a firm mentions Generative AI related words in an earnings conference call in that quarter on the measure of firm exposure to the technology and its interaction with an indicator of whether the quarter is “post-ChatGPT”, which corresponds

to Q4 2022 and Q1 2023 in our sample. This specification allows us to quantify the precise degree to which higher exposure under our measure is associated with a higher likelihood of mentioning Generative AI technologies after the launch of ChatGPT compared to before – which was not possible in the repeated cross-sectional regressions in Figure 7.

The results of estimating Equation 3 in our sample are shown in Table IV. In the first four columns, we add increasingly stringent combinations of fixed effects to the baseline regression. In column 1 we control for calendar quarter fixed effects, which will capture the degree to which people are in general more likely to speak about the topic in particular periods. We also add first industry sector fixed effects, and then firm fixed effects, to control for the degree to which particular sectors and firms are more likely to mention particular technology topics at all times. The most stringent specification in column 4 suggests that a 10 ppt increase in Generative AI exposure is associated with a 7.6 ppt increase in the likelihood of mentioning the Generative AI topic – or that the interquartile range of exposure to Generative AI is associated with about a 7 ppt difference in the likelihood of talking about the topic. To ensure that this pattern is not limited to software companies, we repeat the most stringent estimation in column 5, and exclude NAICS sectors 51 and 54. The result shows that, even outside of these sectors, more exposed firms are significantly more likely to mention Generative AI after the ChatGPT launch, but the coefficient is smaller, suggesting a 3.1 ppt increase in the share of firms mentioning it for a 10 ppt increase in exposure. In general, our results on investor communication reveal that the technological advances in question are salient to investors, rapidly rising in importance based on their prevalence in communication, and more likely to be mentioned by firms that our bottom-up measure predicts to be more exposed.

IV. Stock Market Impact of Generative AI Exposure

A. *Stock return volatility, social media attention, and firms' Generative AI exposure*

Our measure of firm-level Generative AI exposure is intended to capture the relative degree to which tasks in a firm can be made more productive by advances in language models. However, it is ex-ante ambiguous whether greater exposure should be associated with higher returns: while higher productivity should lead to lower costs of producing the output of firms in a particular industry, the degree to which this cost improvement is captured by incumbent firms can vary. For instance, if new entrants are more likely to be flexible enough to take advantage of the benefits of large language models, incumbent profits could actually fall in

response to the technology shock.

Even if the sign of the impact of Generative AI exposure on stock returns is unclear, the release of new information about the technology can be expected to increase the volatility of returns for affected companies. Conversely, if our exposure measure is a valid proxy for firm-level characteristics that are specific to GPT-like technologies, higher return volatility associated with this exposure should be associated with new information about these technologies. First, we explore whether Generative AI exposure is in fact associated with higher stock return volatility on the days after the release of ChatGPT on November 30th, 2022, which was the beginning of the period when the technological advances around large language models started to become widely known. For each trading day from November 15th to March 31st, we run separate regressions of the form

$$|r|_{it} = \alpha_t + \alpha_{ind} + \beta_t \text{GenAIExposure}_i + \varepsilon_{it},$$

where β_t captures the degree to which our firm-level measure of Generative AI exposure predicts higher return volatility on that day, and we control for 2-digit industry sector fixed effects α_{ind} to capture the degree to which Generative AI exposure might simply correlate with industry news coming out around the same time. The time series of t-statistics testing whether this coefficient is zero on each day are shown in Panel A and B of Figure 8 for the specifications without and with industry fixed effects. There are only a small number of periods when the stock return volatility effect of Generative AI exposure clearly exceeds conventional significance thresholds with a t-statistic well above 2, but in both graphs they include November 30th, 2022, which is the day that ChatGPT was released to the public, and at least one other day in the two weeks after the release, which corresponds to the period when many major news outlets were first covering the surge in interest in large language models, that resulted from encounters with ChatGPT's capabilities. Note that, even though our exposure measure uses no information other than a firm's employment structure, its association with stock returns actually identifies the relevant period when the capabilities of ChatGPT were first becoming known, suggesting that it captures exposure to this particular technology.

Panels C and D of Figure 8 confirm that the days on which our firm employment-structure based measure of Generative AI exposure predicts higher stock return volatility around the release of ChatGPT closely coincide with surging social media attention to the technology: Twitter mentions of "ChatGPT" and "GPT" spike relative to their trend right after the release of ChatGPT, supporting the argument that the higher association between stock volatility and Generative AI exposure in this period indeed reflects investor incorporating

new information into stock prices. The fact that the volatility impact of our Generative AI exposure measure reliably identifies the key period when major related news was released validates that it captures a dimension of exposure to the associated technologies that is relevant to market participants.

Defining the ChatGPT “event” period In order to identify the effect of Generative AI exposure on firm returns in this relatively short time series, we want to focus on days when market participants were incorporating substantial news about related technologies into their valuation of firms. The method above suggests that Generative AI exposure was associated with high stock return volatility particularly during the initial ChatGPT launch and in the weeks thereafter as the market discerned the likely impact of this technological advance and aggregated relevant information. This period also coincides with a time period when the growth in Twitter mentions of the topic suggests a high level of interest and new information. We, therefore, focus on the “ChatGPT release period” consisting of November 30, 2022, and the two weeks following it in our main analysis of the stock return impact of the new technology.¹⁷

B. Realized returns and Generative AI exposure

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 (*H-L*) – which we will also 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.¹⁸ All portfolios are formed based on market capitalization weights as of October 31, 2022, and exposure measures based on March 2022 firm-level employment structures, and weights are adjusted daily to mimic passive buy-and-hold exposure. The returns data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31st, 2023, as well as Fama-French 5-Factor data and risk-free returns from Ken French’s website. We also form industry-neutral portfolios by first forming within-2-digit NAICS industry value-weighted quintile portfolios, and then averaging portfolio returns for the same quintiles across industries while applying industry market capitalization as weights.

Generative AI exposure returns after ChatGPT release We compute realized excess daily stock returns across Generative AI exposure portfolios during different time periods,

¹⁷To be precise, we use returns on all trading days from Nov. 30, 2022 to Dec. 14, 2022.

¹⁸See the Appendix for details on the portfolio construction.

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 represents either the mean return for the full period, or is allowed to vary with whether the day is part of the ChatGPT release period or not, as defined above. Table V shows the mean excess realized returns by portfolio for different specifications, comparing returns for the full time period of Nov. 15, 2022 - March 31, 2023, and returns on the days when we would expect the Generative AI exposure to matter for returns, i.e. the two weeks after the ChatGPT release.

Panel A of Table V shows raw excess returns: The *AMH* high-minus-low exposure portfolio has positive daily excess returns of 0.4% (t-statistic > 3) on average during the post-ChatGPT release dates, but not outside of that period. This key finding – that highly exposed companies have higher average returns only on days when advances in language model technology become known – also validates our bottom-up measure of technology exposure, as it predicts higher returns if, and only if, the dates are likely to be associated with updating about the potential of the recently released technology.

In Panel B of Table V, we additionally control for each portfolio’s exposure to the market return factor. The magnitude of the positive return effect is largely unaffected, and continues to be statistically highly significant and large. One additional concern might be that the Generative AI exposure quintiles load on particular industries, such as technology companies, and that the exposure to the relative performance of these industries drives the realized return variation across exposure quintiles. While the fact that the portfolio outperformance occurs only on GPT news release dates makes it unlikely that this is an important issue in our analysis, in Panel C we show the performance of sector-neutral factor portfolios, also adjusted for market factor exposure.¹⁹ In this specification, the excess returns to the *AMH* portfolio are almost identical to those in Panel B, suggesting that the Generative AI exposure factor is not simply driven by an association between exposure and particular industry sectors.

Note that, across the different models in Table V, mean daily excess returns across the period as a whole and returns outside of the ChatGPT release period are small, statistically indistinguishable from zero, and almost flat across Generative AI exposure quintiles for the

¹⁹These portfolios are formed by sorting firms into quintiles *within* each industry and then forming quintile portfolios by taking a weighted average of the corresponding quintile in each industry sector, thereby preventing the quintiles from loading excessively on particular industries. See the Appendix for further details on the portfolio construction.

period as a whole. This reduces concerns that Generative AI exposure spuriously correlates with other firm characteristics that are driving differential returns over the period in question. In contrast, on the days when exposure to large language model productivity changes should actually matter, returns vary significantly with whether a firm is more or less exposed based on our bottom-up task-based measure. In the two week period after the ChatGPT release, excess returns monotonically increase across Generative AI exposure quintiles, as shown in Figure 9 for the market-factor adjusted, industry-neutral exposure quintiles.²⁰

As a robustness check for our analysis, we also 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 replicate the previous portfolio returns analysis but additionally control for the returns on the factors in the Fama French 5-factor model (Fama and French, 2015). The results are shown in Appendix Table C7. We find that the 5-factor-adjusted excess returns (and the industry-neutral version) on the *AML* portfolio during the ChatGPT release period are smaller, albeit still sizable, in this specification, suggesting daily excess returns of 0.3%, and statistically highly significant. Moreover, Appendix Figure B2 shows that excess return variation across Generative AI exposure portfolios is highly monotonic even in this specification.

Time series of Generative AI exposure portfolio returns To further validate that the Generative AI exposure factor returns are associated with the advances in the related technology, we also consider the time series of daily excess returns on industry-neutral portfolios in Figures 1 and 10 – which show the results for global sorts and industry-neutral sorts – comparing the cumulative excess returns of the high, the low, and the high-minus-low *AMH* portfolios. The graphs of the high and low exposure portfolios show that, if anything, that the high exposure portfolio had lower returns in the two weeks preceding the release of ChatGPT. On the other hand, the returns of the “A” and “H” portfolios diverge rapidly thereafter, with most of the gains to the high exposure portfolio concentrated in the two-week period after the release. The cumulative returns to the high-minus-low zero net investment portfolio over the full time period through March 31st are large. Cumulative returns to holding a portfolio that is long the highest-exposure quintile, and short the lowest-exposure quintile from the release date (Nov. 30, 2022) through March 31, 2023 are 9.4% for the global sort across all stocks, and 5.1% for the sector-neutral sorts (see Appendix Figure B4 for the time series of the cumulative excess returns without adjusting for market factors). The corresponding values when adjusting for the market factor are 8.5% and 5.0%.

²⁰See Figure B1 for the variation across portfolios sorted jointly across all stocks (not industry-neutral) in market-factor adjusted returns.

C. Heterogeneity in returns to Generative AI exposure

The results in the previous section showed our key result that, on average, a higher exposure to Generative AI technology advances due to a prevalence of affected tasks in a firm was associated with higher realized stock returns after the release of ChatGPT. However, there is reason to believe that these effects might vary in size, and perhaps also in sign, across different industries. On the one hand, our exposure measure focuses on the labor cost and productivity dimensions of the impact of these technology shocks. However, there will likely also be simultaneous effects on product market competitiveness for different firms, depending on whether their product complements, or substitutes for, the capabilities of products like ChatGPT for end-consumers. Similarly, some companies may have valuable intangible capital that is made obsolete – e.g. because ChatGPT obviates the productivity advantage of proprietary internal software and processes – or increases in value, for instance if the company controls access to data sources that can be more effectively analyzed in conjunction with large language model capabilities.

In addition, industry structure and contestability might affect the impact of higher Generative AI exposure on publicly traded incumbents. If entry into the product market becomes easier as Generative AI technology lowers the cost of putting together minimum viable products, or if start-ups in an industry are more likely to have the organizational flexibility to quickly incorporate these new technologies into their workflows, then incumbents with high exposure to Generative AI technologies may lose out in relative terms, even if the overall impact on the industry’s productivity is positive. While it is too early to make definitive statements about which firms will be winners or losers as a result of recent technology advances, this section will provide early suggestive evidence of what industry characteristics have been associated with higher return effects.

Industry portfolio heterogeneity in ChatGPT release period returns One way of looking at industry heterogeneity is to consider which industry portfolios exhibit returns after the release of ChatGPT that exceed or fall short of that predicted by their Generative AI exposure. Figure 11 shows the variation in market factor-adjusted excess realized returns of 3-digit industry portfolios either during the two weeks after the ChatGPT release (upper panel), or on all other days of the Nov 15, 2022 - March 31, 2023, period (lower panel). There is a positive association across industries between average industry level Generative AI exposure and returns during the weeks after the ChatGPT release. Moreover, as the lower panel confirms, this is *not* due to the fact that high Generative AI exposure industries incidentally experienced systematically higher or lower returns over the nearly 5 months in our study: the average excess daily returns of almost all 3-digit industry portfolios are small

and do not vary with Generative AI exposure over this period. That is, the Generative AI -related returns effect is evident at the industry level only on the event days when ChatGPT news is most likely being incorporated into valuations.

However, Figure 11 also shows that some industries with high Generative AI exposure perform even better on news days than we would expect based on a simple linear relationship between exposure and returns, i.e. their news day returns lie above the red line. For example, the highest returns on news days among large subsectors accrue to firms in the industry subsector “Publishing industry (except internet)”, in which the largest firms by market cap are Microsoft, Salesforce, and Intuit. While Microsoft is directly associated with the success of ChatGPT and related products through its investment in Open AI, the company releasing this tool, Salesforce and Intuit operate software businesses focused on customer relationships, tax, personal finance, and accounting data. It is intuitive that better analytical tools based on Generative AI can, for instance, make the proprietary data resulting from these business lines more valuable.

Within-industry Generative AI exposure effects While there is substantial variation across industries in their overall exposure and performance after the release of ChatGPT, there may also be differences in the degree to which high-exposure firms outperform low-exposure firms relative to other firms *within* the same industry. That is, in some industries the ability to convert *relatively* higher exposure to the technology shock into higher firm-level returns may be better than in others, independent of whether the industry as a whole is benefiting from higher Generative AI exposure or not. Moreover, exploring how the within-industry outperformance varies also provides evidence for how widespread the Generative AI technology effects are outside of the “Information” sector that is most likely to see impacts distorted by simultaneous product market effects.

Figures 12 and 13 show the average daily return alpha during the ChatGPT release period for H-L portfolios within each sector or subsector, respectively, for quintiles at the NAICS 2-digit sector and terciles at the 3-digit subsector level.²¹

Although the standard errors for these within-industry portfolio regressions are necessarily larger, as some industries have a limited number of firms in the sample, we do indeed find that the Generative AI exposure effect is significantly different from zero in several sectors and subsectors. Focusing on the more detailed categories in Figure 13, we find significant

²¹Each graph only includes sector or subsectors that have (1) a large enough sample of firms in the data, consisting of at least 10 firms across the high and low exposure quantile, and (2) a large enough Generative AI exposure spread between the high and low exposure portfolio, with cutoffs set at a 10 ppt spread between the high and low quintiles for 2-digit sectors, and at 5 ppt between the high and low tercile for 3-digit subsectors.

positive returns within ‘credit intermediation & related activities’, which contains several large banks and brokerages, “publishing industries (except internet)”, which contains Microsoft, Salesforce, and Intuit, as noted above, and also the “administrative and support services” subsectors. In contrast, the “real estate”, “other information services”, and “food manufacturing” subsectors” all show negative within-industry returns to higher Generative AI exposure. Contrasting this with non-release period within-industry returns shown in Appendix Figures B5 and B6, we can see that the latter are more precisely estimated and small, as well as not significant with few exceptions. This suggests that the variation in within-industry returns to Generative AI exposure are likely to be driven by information incorporated during the ChatGPT release period, rather than generally during our 4.5 month sample period.

Industry characteristics and within-industry Generative AI exposure effects The previous analysis raises the question of what industry characteristics can explain the observed heterogeneity in whether industries are more likely to exhibit a high return on the within-industry *AMH* exposure portfolio. We explore this question by running regressions of the form

$$\alpha_i^{\text{news},AMH} = \eta + \gamma \text{Characteristic}_i + \varepsilon_i,$$

which relate the subsector *AMH* Generative AI exposure portfolio alpha to various industry characteristics that are computed as either the average of the characteristics for the industry (left panel of Figure 14), or the difference between the high and low Generative AI exposure portfolios in the characteristic within that industry (right panel of Figure 14). In either case, both the dependent and independent variables are transformed into standardized Z-scores, such that coefficients capture the relationship between a change in the characteristic by one standard deviation and a standardized change in within-industry returns to Generative AI exposure.

As shown in Figure 14, industry subsectors with larger firms are more likely to have high within-industry returns to differences in Generative AI exposure. Moreover, high Generative AI exposure firms are more likely to outperform the low exposure firms in the same industry, if the former have higher ROE, ROA, market capitalization, gross profitability, or organizational capital. However, these associations are only suggestive: the precise characteristics that allow some firms to take better advantage of language model-driven technological advances will hopefully be considered in future studies.

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, as well 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% daily, translating to over 100% on an annualized basis. These differences were realized both within and across industries, and display wide variation which is correlated with firm characteristics such as organizational capital or gross profitability. According to investors, ChatGPT represents an important shock to corporate valuations.

REFERENCES

- 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).
- 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., 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.
- 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 crowd-workers for text-annotation tasks, *arXiv preprint arXiv:2303.15056* .
- 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.

- Hassan, Tarek A, Stephan Hollander, Laurence Van Lent, and Ahmed Tahoun, 2019, Firm-level political risk: Measurement and effects, *The Quarterly Journal of Economics* 134, 2135–2202.
- 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* .
- 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.
- 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.

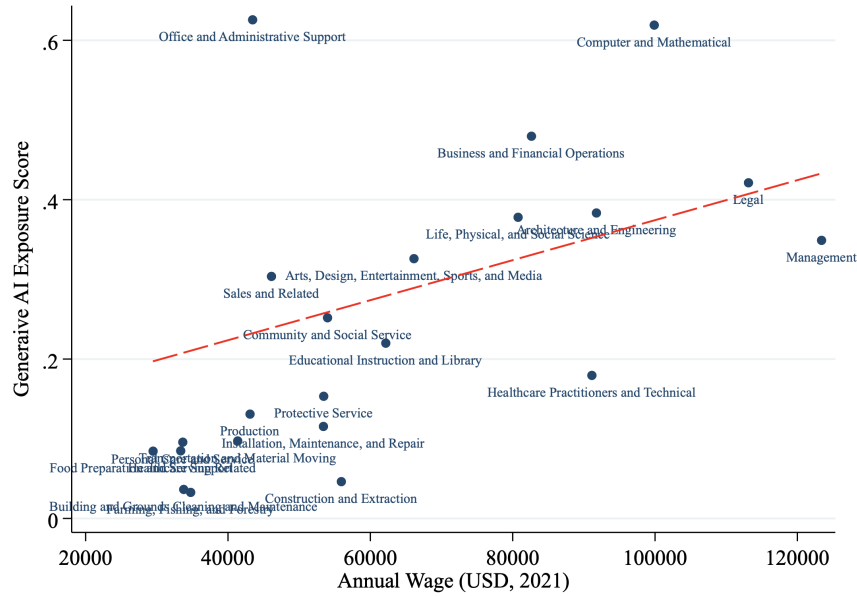
Underwood, Ted, 2023, Using gpt-4 to measure the passage of time in fiction, <https://tedunderwood.com/2023/03/19/using-gpt-4-to-measure-the-passage-of-time-in-fiction/>.

Webb, Michael, 2019, The impact of artificial intelligence on the labor market, *Available at SSRN 3482150* .

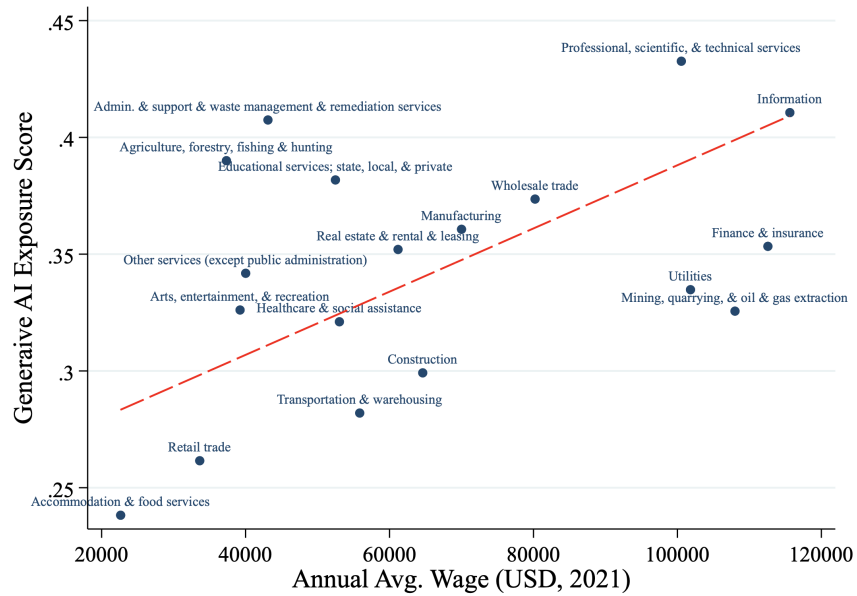
Willison, Simon, 2023, A simple python wrapper for the chatgpt api, <https://til.simonwillison.net/gpt3/chatgpt-api>.

Zhang, Miao Ben, 2019, Labor-technology substitution: Implications for asset pricing, *The Journal of Finance* 74, 1793–1839.

Figure 2: Generative AI exposure and wages by major occupation group and industry sectors.



(A) Major occupation groups



(B) NAICS 2-digit industry sectors

Figure 3: Occupational skills and Generative AI exposure. The graph below shows the results of regressing our 6-digit occupation's Generative AI exposure measure jointly on a set of occupational skill scores defined in Acemoglu and Autor (2011), converted into standard z-scores. That is, we run the following regression:

$$\text{Exposure}_o^{GPT} = \alpha + \sum_S \beta_S \text{Skill}_o + \varepsilon_o$$

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

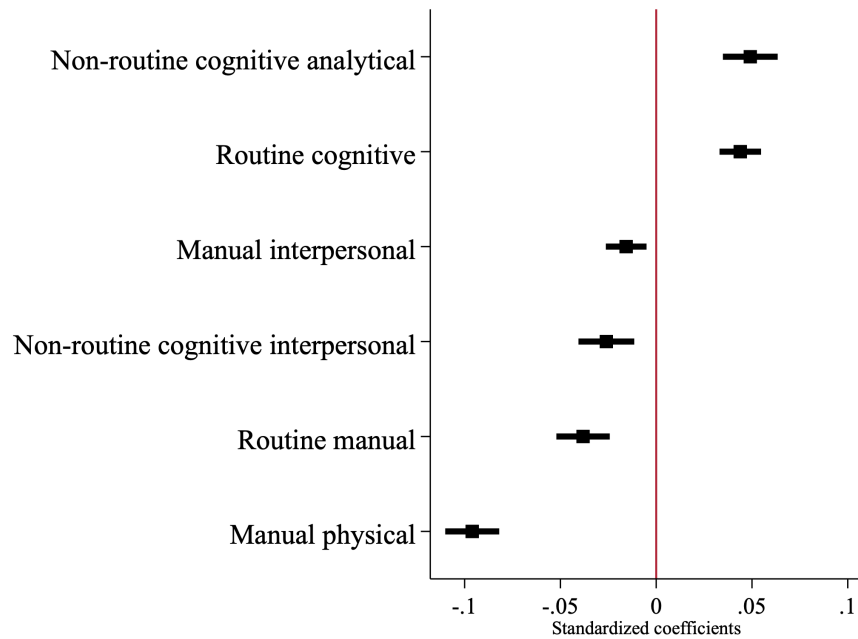


Figure 4: Generative AI exposure Across and Within Industries This figure plots the average and the standard deviation of Compustat firms' Generative AI exposure within each NAICS 2-digit sector.

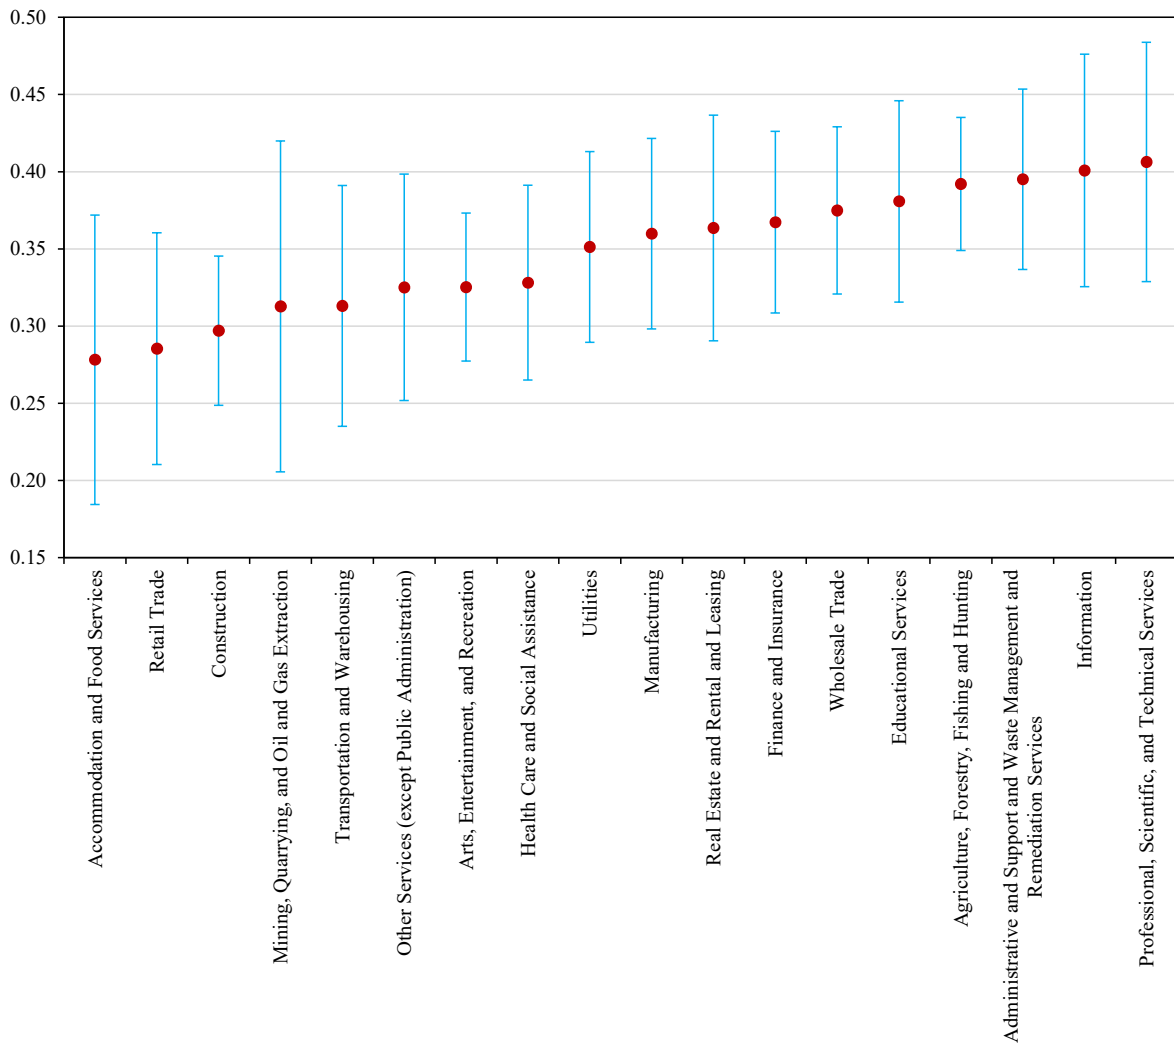


Figure 5: Generative AI exposure and average firm characteristics by industry sector. The graphs show the market capitalization weighted average of the characteristic among firms in our Generative AI exposure sample in each NAICS sector, and plots against it the weighted average of our Generative AI exposure measure in the same sector. The firm characteristics are based on Compustat fiscal year 2021 data., and the wages are from the Bureau of Labor Statistics. The red line is the market-capitalization-weighted linear fit across the sector level aggregates.

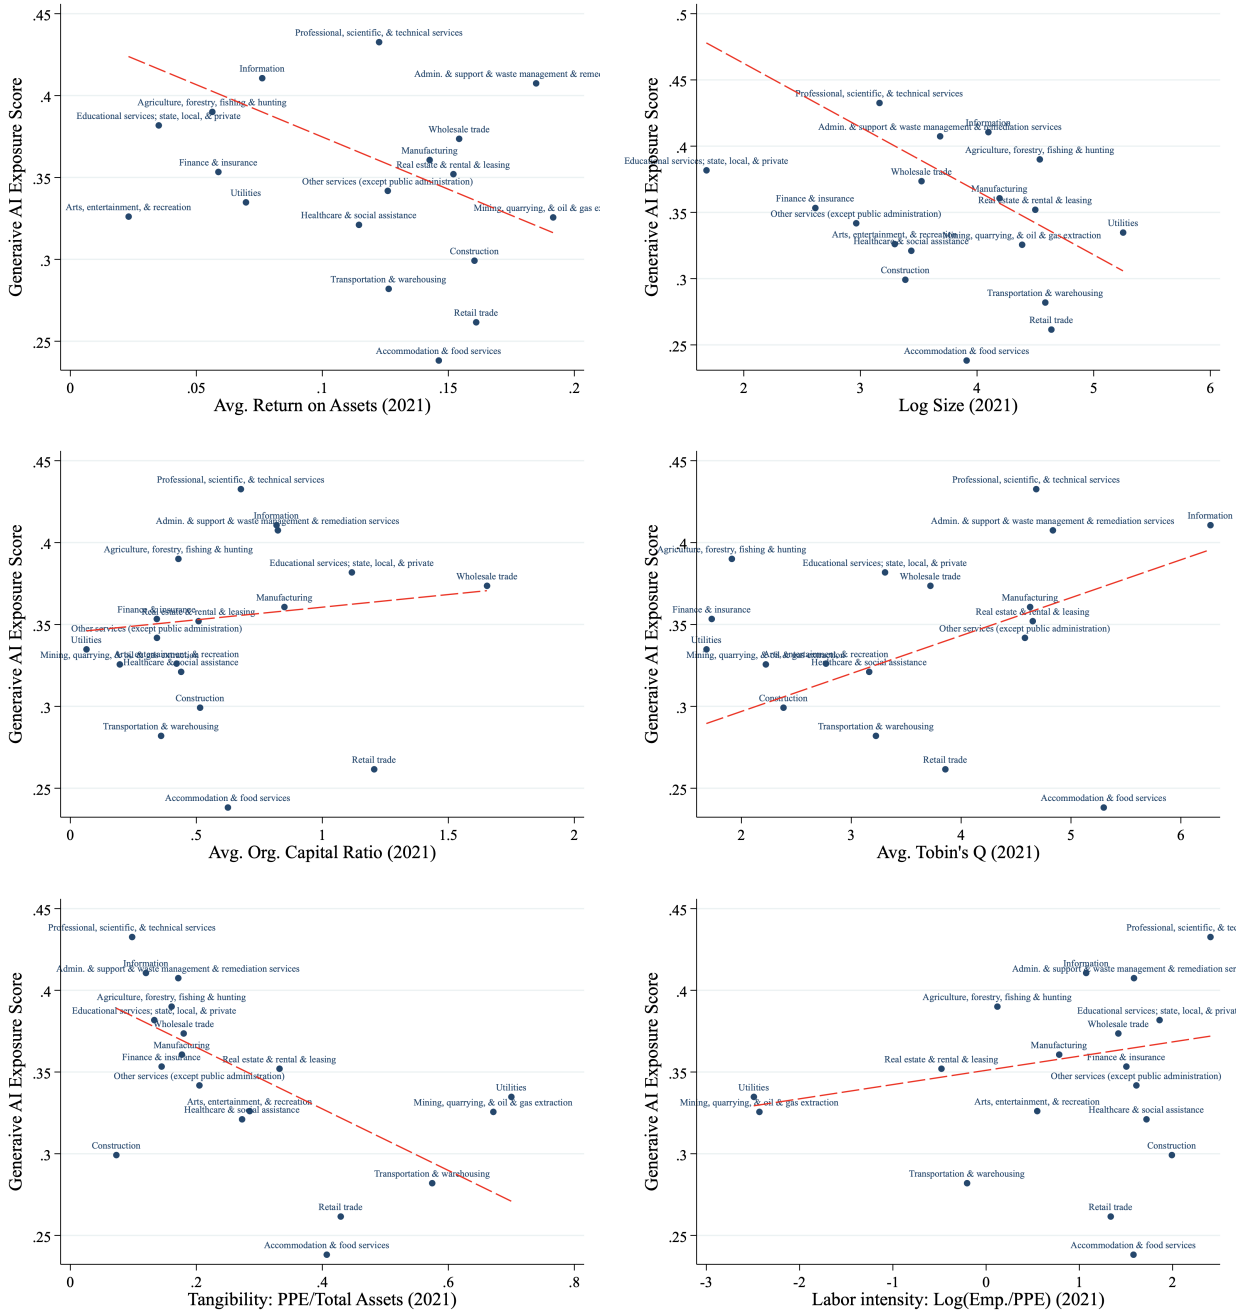
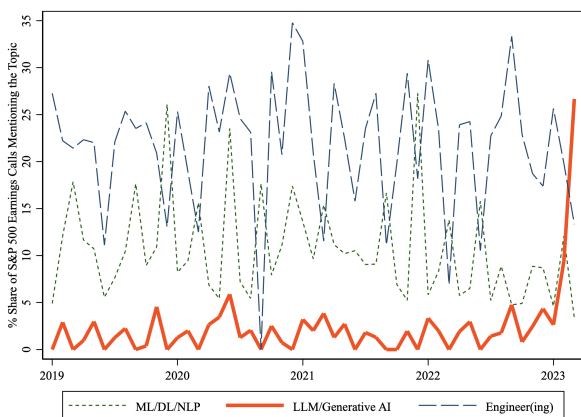
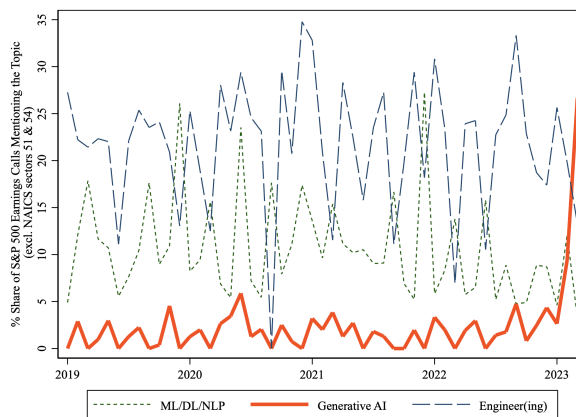


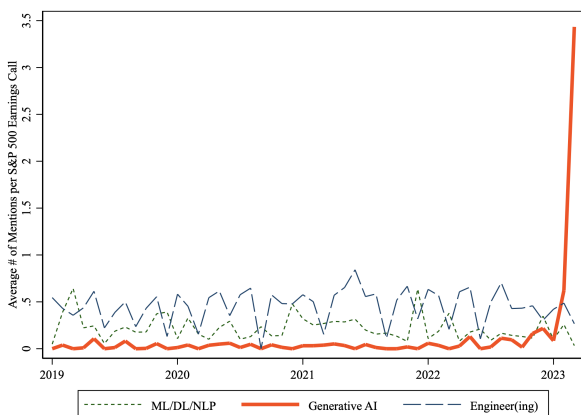
Figure 6: Topic mentions in company earnings conference calls. The graphs below show data on the share of calls mentioning Generative AI related words (Panels A and B) or the average number of such mentions per call across the sample (Panels C and D). The data set is a manually collected panel of earnings conference call transcripts for S&P 500 firms’ tickers from the *Seeking Alpha* website. The graphs show monthly statistics for calls that were held from Jan 2019 to March 2023. Calls are assigned a calendar month, quarter and year based on the time stamp at the beginning of the transcript. Each transcript is converted into a list of lower-case unigrams and bigrams. Word tokens in these lists are counted as Generative AI-specific words if they are in the following set: “llm”, “chatgpt”, “gpt”, “gpt3”, “gpt4”, “generative”, “language model”. Trends in other machine learning-related words and generic engineering terms are shown for comparison - see text for details. in Panels B and D, the sample of firms excludes the “Information” (NAICS 51) and Professional, Scientific, and Technical Services (NAICS 54) sectors, as determined by Compustat industry codes.



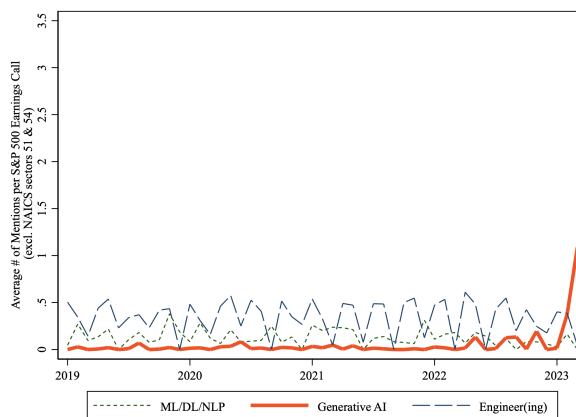
(A) Share of calls: all firms



(B) Share of calls: excl. Information & Prof., Sci., and Tech. Svcs. sectors



(C) Mentions per call: all firms



(D) Mentions per call: excl. Information & Prof., Sci., and Tech. Svcs. sectors

Figure 7: Firm-level Generative AI exposure and topics in company earnings conference calls: Each graph show the result of estimating regression specifications of the form

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

for each topic for the fiscal quarters 2019 Q1-2023 Q1 for a manually collected panel of earnings conference call transcripts for S&P 500 firms' tickers from the *Seeking Alpha* website. Calls are assigned a calendar month, quarter and year based on the time stamp at the beginning of the transcript. Each transcript is converted into a list of lower-case unigrams and bigrams. Word tokens in these lists are counted as LLM/Generative AI-specific words if they are in the following set: “llm”, “chatgpt”, “gpt”, “gpt3”, “gpt4”, “generative”, “language model”. The association of Generative AI exposure with generic engineering terms is shown for comparison - see text for details of other topic definitions. in Panels B and D, the sample of firms excludes the “Information” (NAICS 51) and Professional, Scientific, and Technical Services (NAICS 54) sectors, as determined by Compustat industry codes. Quarterly sample sizes for the regressions are 341-416 for the full sample, and 299-368 firms for the sample excluding NAICS 51 and 54. Dotted lines show 95% confidence intervals based on heteroskedasticity-robust standard errors.

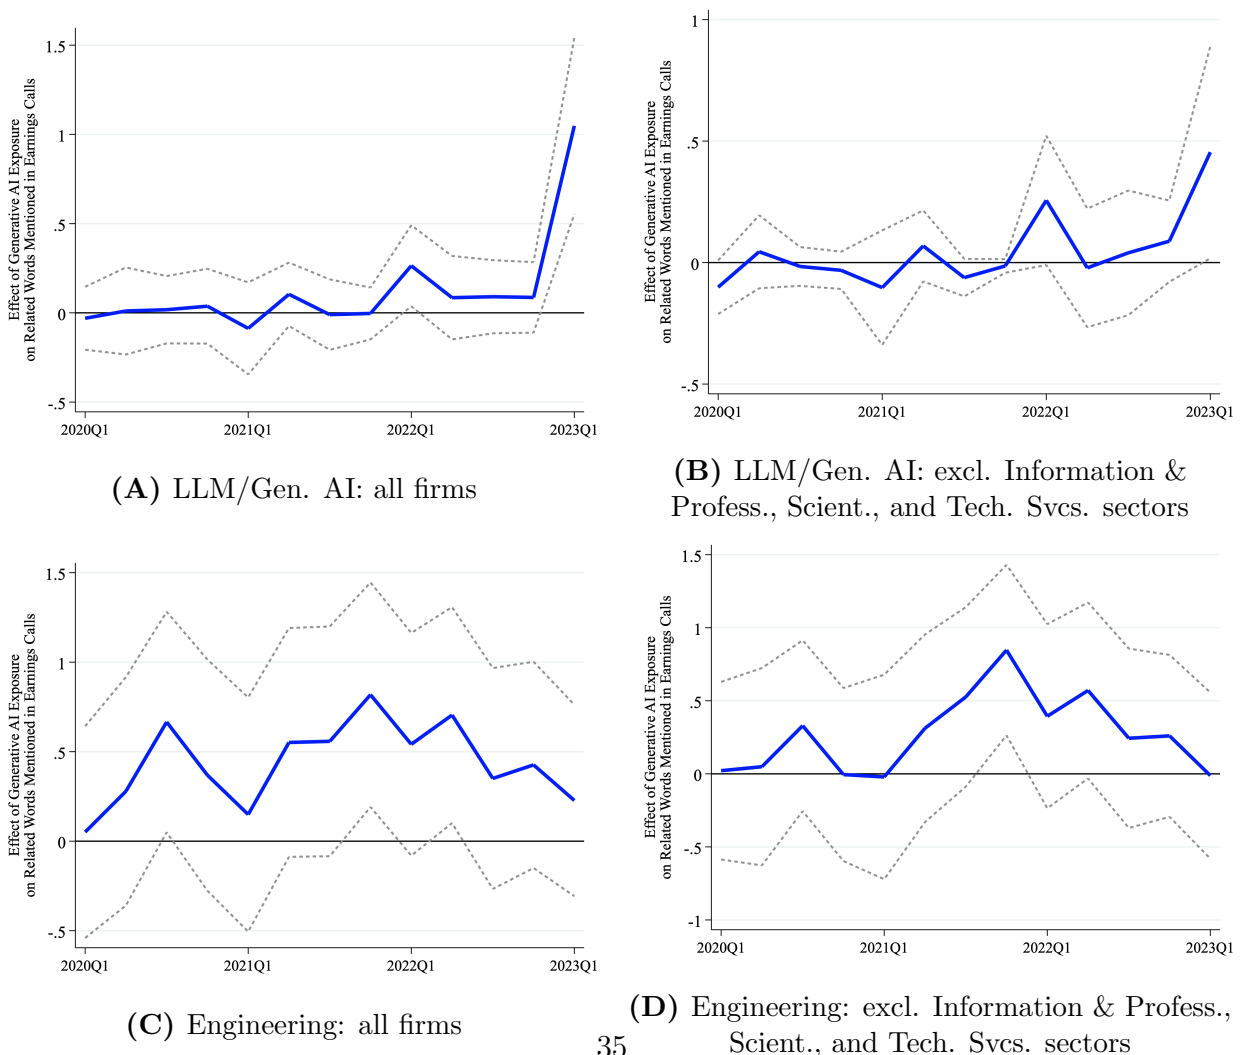
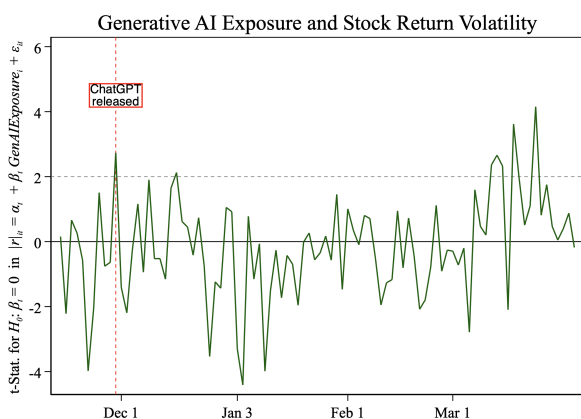
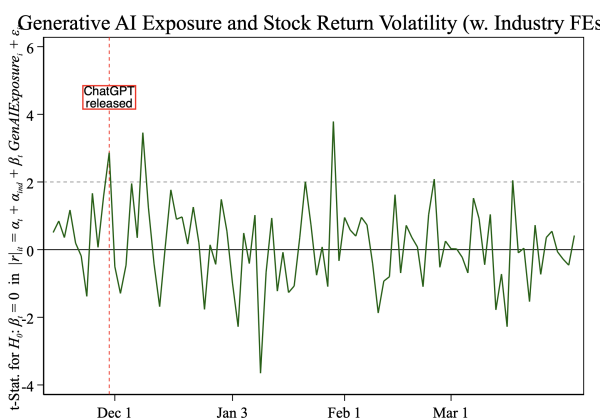


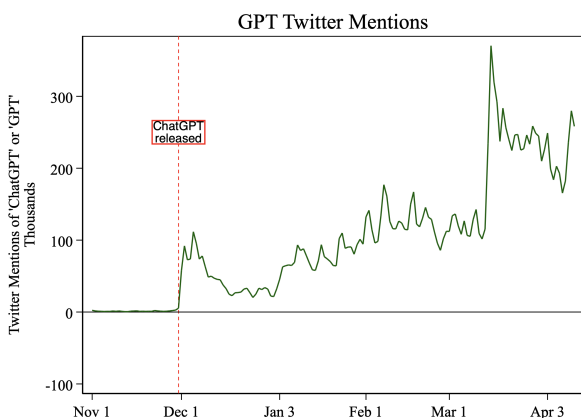
Figure 8: Stock return volatility, Generative AI exposure, and social media attention: Panel A shows the t-statistic for the hypothesis $\beta_t = 0$ in the regression $|r|_{it} = \alpha_t + \alpha_{ind} + \beta_t \text{GenAIExposure}_i + \varepsilon_{it}$, estimated at the firm level, for each trading day from Nov. 15, 2022, to March 31, 2023. Panel shows the estimates without, and Panel B including, 2-digit NAICS sector fixed effects. The dependent variable is the absolute value of daily stock returns, and the independent variable is the firm-level measure of task exposure to Generative AI technology productivity changes. Standard errors are heteroskedasticity-robust. The dashed horizontal line indicates where t-statistics exceed 2. Panel C shows the total count of Twitter mentions of “ChatGPT” or “GPT” reported by Media Cloud, and Panel D shows the growth rate (using the average base method) of these Twitter mentions.



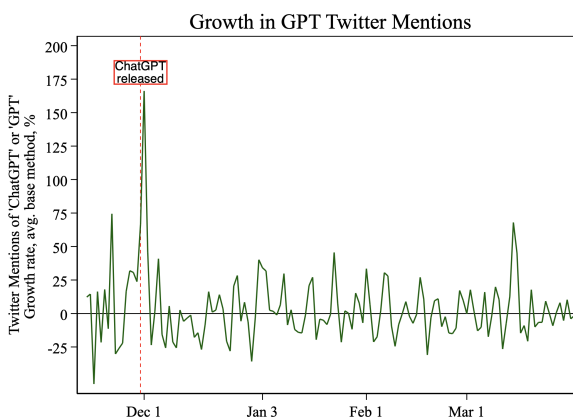
(A) Exposure & stock return volatility: no industry FEs



(B) Exposure & stock return volatility: incl. industry FEs



(C) Twitter mentions: raw count



(D) Twitter mentions: growth rate

Figure 9: Industry-neutral Generative AI exposure quintile portfolio returns: market-factor adjusted. Each bar shows the average daily return alpha during the ChatGPT release period and also for all days, for portfolios based on industry-neutral Generative AI exposure sorts, and also for the net zero investment high-minus-low exposure portfolio that represents the “Artificial Minus Human” (*AMH*) factor. Industry-neutral portfolios are computed by first forming within-industry value-weighted quintile portfolios, and then averaging portfolio returns for the same quintiles across industries. Returns within-industry quintiles are value-weighted, while across-industry averages are industry market-cap. weighted. The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023, and factors from Ken French’s website. The figure shows alphas estimated from portfolio-level 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 bars show either the intercept for the “ChatGPT release period”, or the intercept from a regression where all days have the same intercept. Error bars indicate 95% confidence intervals computed using Newey West standard errors with five lags.

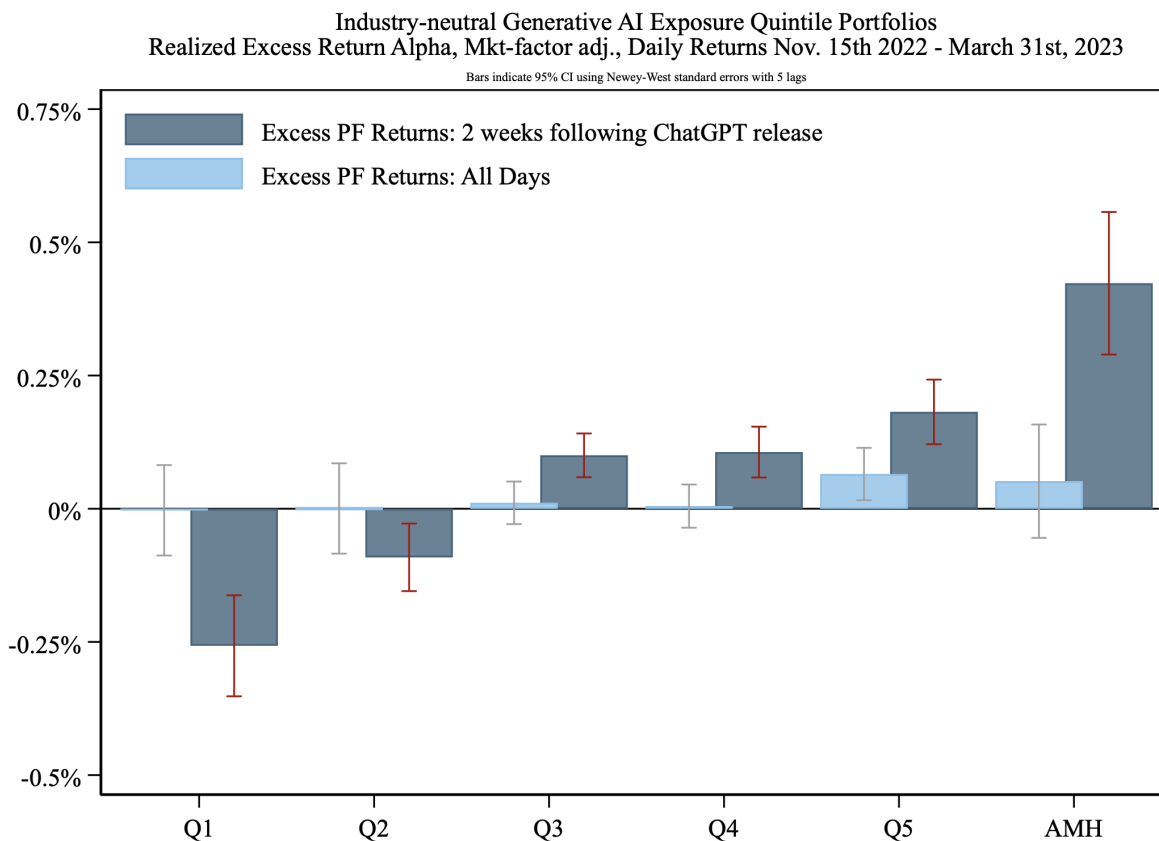


Figure 10: Generative AI exposure industry-neutral quintile portfolio returns over time: market factor-adjusted. The graph shows the cumulative excess realized returns on portfolios based on industry-neutral Generative AI exposure sorts. Industry-neutral portfolios are computed by first forming within-industry value-weighted quintile portfolios, and then averaging portfolio returns for the same quintiles across industries. Returns within-industry quintiles are value-weighted, while across-industry averages are industry market-cap. weighted. 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.

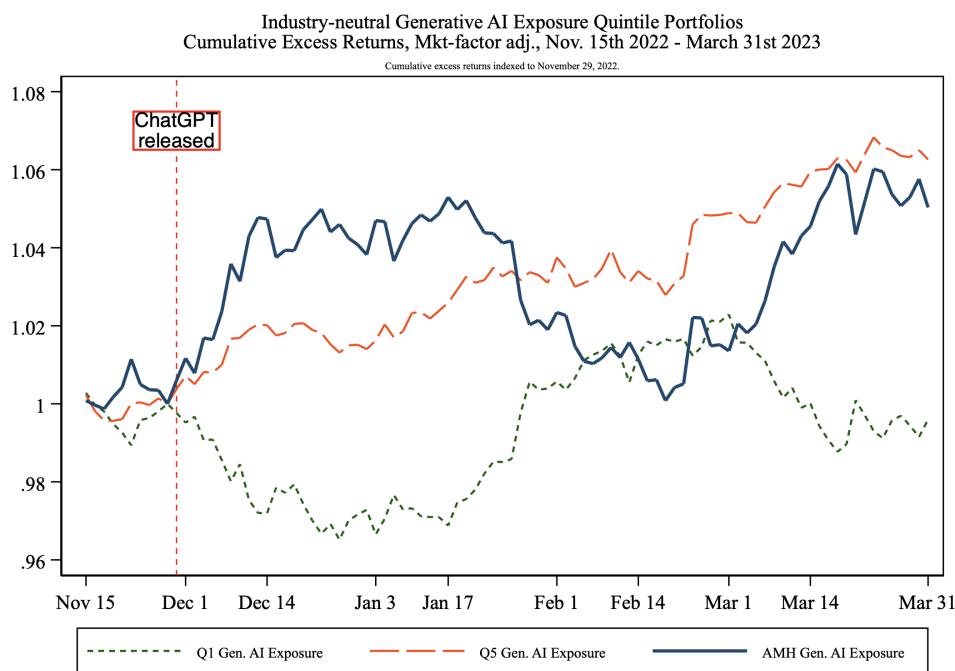


Figure 11: Industry Portfolio Stock Returns in ChatGPT Release Period. Each graph shows the average daily return alpha for the 3-digit NAICS industry portfolios corresponding to the 50 largest subsectors. The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 -March 31, 2023. The table shows 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 part of the “ChatGPT release period” (Nov. 30, 2022, and the following two weeks), as defined in the text, or not. The upper panel shows the excess daily returns estimates for the GPT release period, while the lower panel shows excess returns on all other days for the Nov. 15, 2022 to March 31 2023, period. The ten largest 3-digit industry subsectors by market capitalization are indicated in red and labeled. Grey lines indicate 95% confidence intervals computed using Newey West standard errors with five lags.

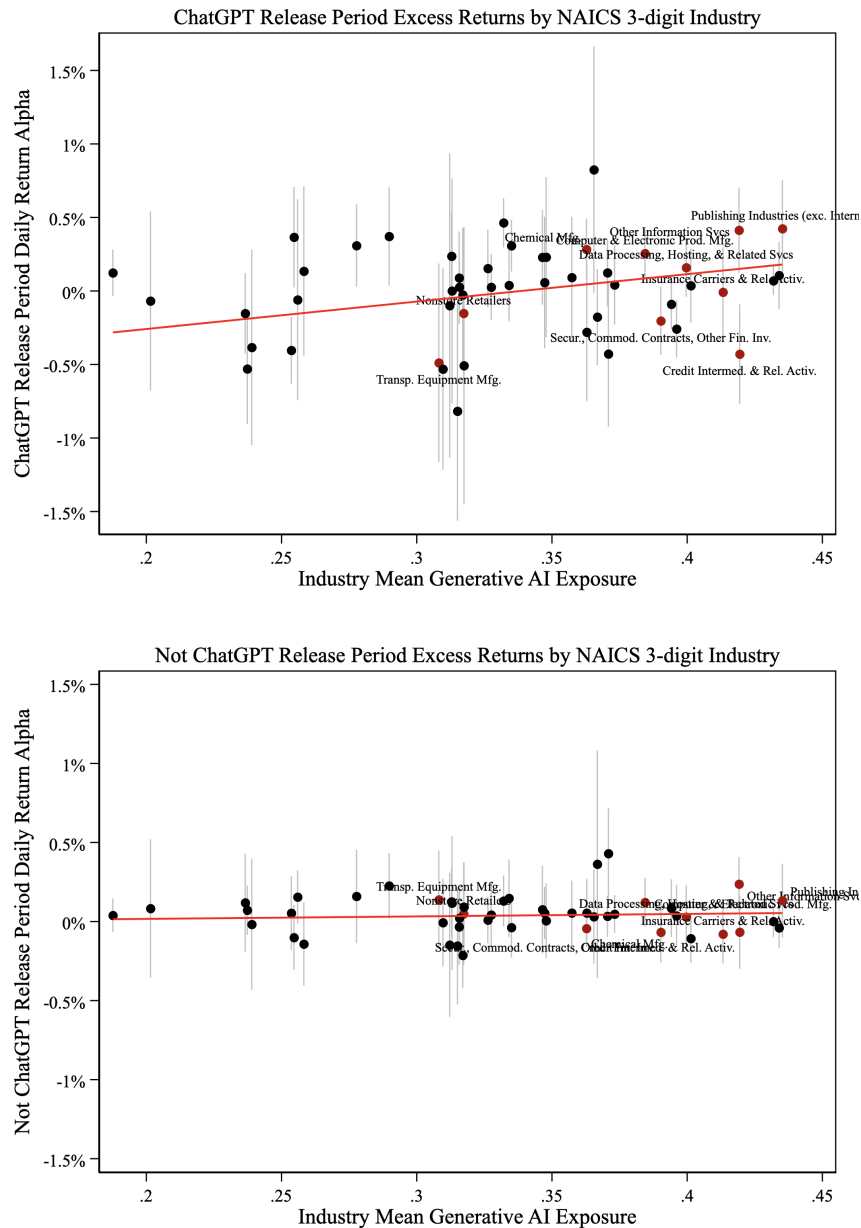


Figure 12: Within-Sector AMH Generative AI exposure Portfolio Realized Returns: ChatGPT Release Period. The graph shows the average daily return alpha for the ChatGPT release period for AMH Generative AI exposure portfolios within each industry, at the 2-digit NAICS sector level. Each AMH portfolio is 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, and the following two weeks, 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.

**Within-Ind. 2-digit Sector AMH Quintile Generative AI Exposure Portfolio
ChatGPT Release Period Alpha, Mkt-factor adj., Daily Returns Nov. 15, 2022 - March 31, 2023**

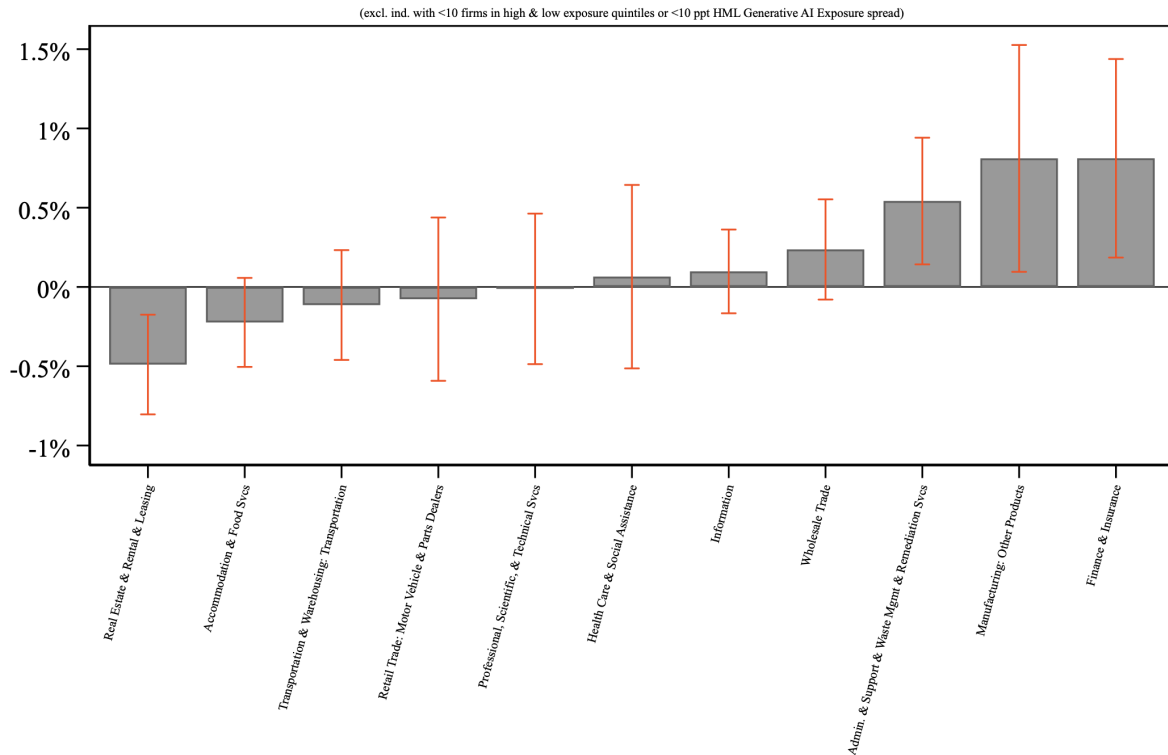


Figure 13: Within-Subsector AMH Generative AI exposure Portfolio Realized Returns: ChatGPT Release Period. Each graph shows the average daily return alpha for the ChatGPT release period for AMH Generative AI exposure portfolios within each industry, at the 3-digit subsector level. Each AMH portfolio is formed by taking the value-weighted highest and lowest terciles of Generative AI exposure within each industry (based on NYSE stock cutoffs) and forming zero net investment H-L 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 subsectors with fewer than 10 firms combined in the highest and lowest quantiles in the sample, as well as subsectors with less than a 5 ppt Generative AI exposure spread between the high and low tercile. 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, and the following two weeks, 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.

**Within-Ind. 3-digit Subsector AMH Tercile Generative AI Exposure Portfolio
ChatGPT Release Period Alpha, Mkt-factor adj., Daily Returns Nov. 15, 2022 - March 31, 2023**

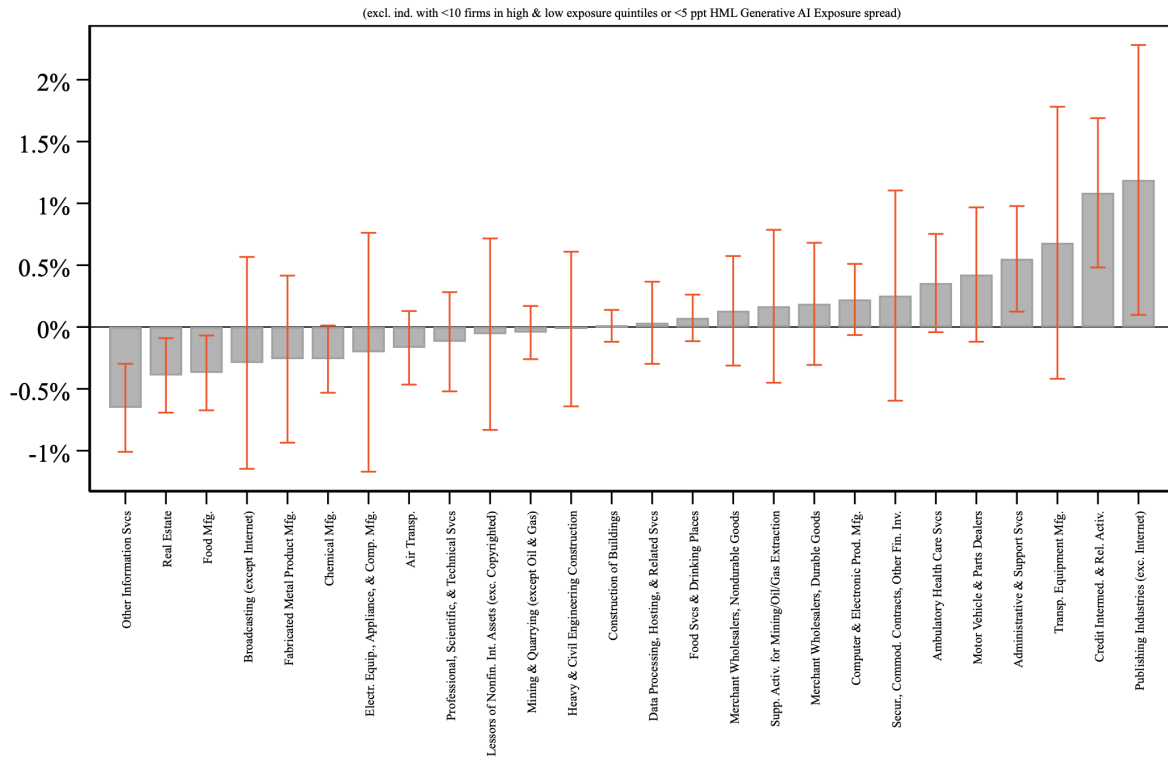


Figure 14: Industry Characteristics and Heterogeneity in Within-Industry Returns to Generative AI exposure. Each graph shows the result of a regression where the dependent variable is the average daily return alpha during the ChatGPT release period on the *AMH* tercile portfolio within each industry subsector, corresponding to the estimates shown in Figure 13. The independent variable is either the average of the characteristics for the industry (left panel), or the difference between the high and low Generative AI exposure portfolios in the characteristic within that industry, in either case transformed into standardized Z-scores. Each coefficient in the graph represents an estimate of $\hat{\gamma}$ from a regression of the form

$$z_i^{\alpha^{\text{release}}} = \eta + \gamma \text{Characteristic}_i + \varepsilon_i,$$

where $z_i^{\alpha^{\text{release}}}$ is the within-industry H-L Generative AI exposure portfolio return for each 3-digit industry subsector i as shown in Figure 13, transformed into a standardized Z-score. Each characteristic Characteristic_i has also been transformed into a standard Z-score. Only subsectors with a total of at least 10 firms in the high and low tercile are included, as well as with a Generative AI exposure spread of at least 5 ppt between the highest and lowest tercile. Depending on the characteristic, the sample size varies from 19 and 22 subsectors for regressions involving R& D or organizational capital to 26 subsectors for other characteristics. Error bars indicate 95% confidence intervals computed using heteroskedasticity-robust standard errors.

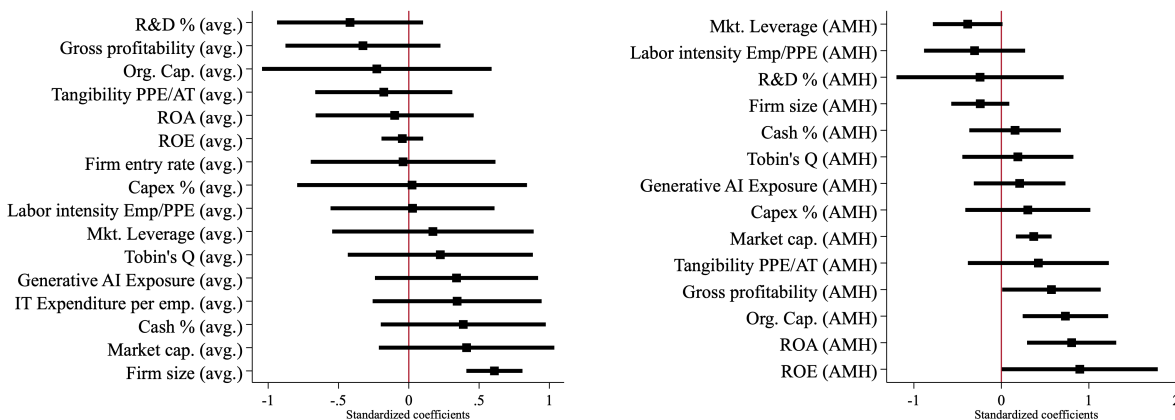


Table I: Summary statistics for selected firm characteristics. *Generative AI exposure* is our bottom-up task-based measure of occupational exposure, aggregated to the firm-level based on the firm's occupational employment structure. *Log Size* is the natural logarithm of total assets. *Labor Intensity* 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 ratio of the organizational capital stock from Eisfeldt and Papanikolaou (2013) divided by Total Assets.

Measure	Mean	Std. Dev.	p10	p50	p90	Obs.
Generative AI Exposure	0.354	0.078	0.268	0.353	0.442	2,518
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 GPT. 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 \$. *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 2-digit level. 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** (1.230)					
Tobin's Q		3.076** (1.217)				
ROA			-21.516* (11.882)			
Labor Intensity				7.892** (2.785)		
Org. Capital Ratio					9.139*** (3.081)	
Tangibility						-89.931*** (22.426)
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	-2.101*** (0.665)					
Tobin's Q		1.499*** (0.490)				
ROA			-13.850 (8.733)			
Labor Intensity				5.476** (2.586)		
Org. Capital Ratio					3.742** (1.668)	
Tangibility						-72.084** (25.866)
Observations	2517	2380	2513	2387	1571	2515
Adjusted R^2	0.200	0.210	0.197	0.227	0.211	0.227

Table IV: Firm-level effects of Generative AI exposure on earnings call topic mentions. The table shows firm-quarter panel regressions over the period of Q3 2018 - Q1 2023 of the form

$$\mathbb{1}[\text{Gen. AI Topic}]_{i,t} = \alpha_t + \alpha_i + \beta_1 E_i^f + \beta_2 E_i^f \times \mathbb{1}[\text{Post-ChatGPT}] + \gamma \mathbb{1}[\text{Post-ChatGPT}] + \varepsilon_{i,t}$$

where we regress whether a firm mentions Generative AI related words in an earnings conference call in that quarter on the measure of firm exposure to Generative AI technologies and its interactions with an indicator of whether the quarter is “post-ChatGPT”, which corresponds to Q4 2022 and Q1 2023 in our sample. The firm sample consists of all S&P 500 firms for which we were able to collect earnings call transcripts in columns 1-4. In column 5, we exclude the “Information” (NAICS 51) and “Professional, Scientific, and Technical Services” (NAICS 54) sectors, as determined by Compustat industry codes. The panel includes 437 firms for the most restrictive specification with all firms, and 385 firms when excluding NAICS 51 and 54. T-statistics in parentheses based on standard errors clustered at the NAICS sector level.

Dep. var.:	$\mathbb{1}[\text{Generative AI Topic Mentioned}]_{i,t}$				
	(1)	(2)	(3)	(4)	(5)
$\mathbb{1}[\text{Post-ChatGPT}]_t \times \text{Gen. AI Exposure}_i$	0.660** (2.307)	0.659** (2.303)	0.644** (2.319)	0.761** (2.407)	0.310** (2.474)
R-squared	0.01	0.01	0.04	0.27	0.19
Observations	6235	6235	6235	6228	5502
	<i>Fixed Effects & Controls</i>				
Quarter FEs		X	X	X	X
Industry Sector FEs			X		
Firm FEs				X	X
Excl. NAICS 51 and 54					X

Table V: Realized returns of portfolios sorted on Generative AI exposure after ChatGPT release. This table reports daily excess stock returns of 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. Quintile thresholds that define value-weighted portfolios are solely based on the sample of stocks listed on NYSE as of the sorting date. All quintile portfolios are formed based on value weights on October 31, 2022, and weights are adjusted based on daily returns to mimic passive buy-and-hold exposure. Industry-neutral portfolios are computed by first forming within-industry value-weighted quintile portfolios, and then averaging portfolio returns for the same quintiles across industries. Returns for within-industry quintiles are value-weighted, while across-industry averages are industry market-cap. weighted. 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 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 either for the full sample, or is allowed to vary with whether the day is in the ChatGPT release period consisting of Nov 30, 2022 - Dec. 14, 2022. Panel A does not include the market factor. T-statistics in parentheses are computed using Newey-West standard errors with five lags.

Sample	Portfolios					
	Q1	Q2	Q3	Q4	Q5	AMH
<i>A: Excess returns (%)</i>						
All days	0.042 (0.38)	0.027 (0.25)	0.050 (0.54)	0.037 (0.31)	0.134 (1.17)	0.076 (1.33)
Not ChatGPT release period	0.066 (0.57)	0.030 (0.25)	0.038 (0.36)	0.037 (0.28)	0.110 (0.85)	0.028 (0.44)
ChatGPT release period	-0.137 (-0.53)	0.000 (0.00)	0.137 (0.43)	0.034 (0.11)	0.316 (0.82)	0.437 (3.08)
<i>B: Market factor-adjusted alpha (%)</i>						
All days	0.014 (0.38)	-0.005 (-0.15)	0.020 (0.56)	0.003 (0.07)	0.098 (2.50)	0.068 (1.11)
Not ChatGPT release period	0.042 (1.07)	0.003 (0.08)	0.013 (0.32)	0.008 (0.17)	0.079 (1.83)	0.021 (0.32)
ChatGPT release period	-0.196 (-4.88)	-0.066 (-0.55)	0.076 (0.82)	-0.036 (-0.63)	0.240 (5.29)	0.420 (5.42)
<i>C: Ind.-neutral mkt. factor-adjusted alpha (%)</i>						
All days	-0.003 (-0.07)	0.000 (0.01)	0.011 (0.54)	0.005 (0.24)	0.065 (2.60)	0.052 (0.95)
Not ChatGPT release period	0.031 (0.69)	0.013 (0.26)	-0.001 (-0.03)	-0.008 (-0.41)	0.050 (1.86)	0.002 (0.05)
ChatGPT release period	-0.257 (-5.32)	-0.091 (-2.82)	0.100 (4.78)	0.106 (4.37)	0.182 (5.88)	0.423 (6.20)

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 value weights on October 31, 2022, and weights are adjusted based on daily returns to mimic passive buy-and-hold exposure. Industry-neutral portfolios are computed by first forming within-industry value-weighted quintile portfolios, and then averaging portfolio returns for the same quintiles across industries. Returns for within-industry quintiles and for all global (not industry-neutral) portfolio 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) is the zero net investment portfolio long highest exposure quintile (*H*) stocks and short lowest exposure quintile (*L*) stocks.

GPT prompt for 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.”

Appendix B. Appendix Figures

Figure B1: Generative AI exposure quintile portfolio returns: market-factor adjusted. Each bar shows the average daily return alpha during the ChatGPT release period and also for all days, for portfolios based on Generative AI exposure sorts across all stocks in the sample, and also for the net zero investment high-minus-low exposure portfolio that represents the “Artificial Minus Human” (*AMH*) factor. The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023, and factors from Ken French’s website. The figure shows alphas estimated from portfolio-level 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 bars show either the intercept for the “ChatGPT release period”, or the intercept from a regression where all days have the same intercept. Error bars indicate 95% confidence intervals computed using Newey West standard errors with five lags.

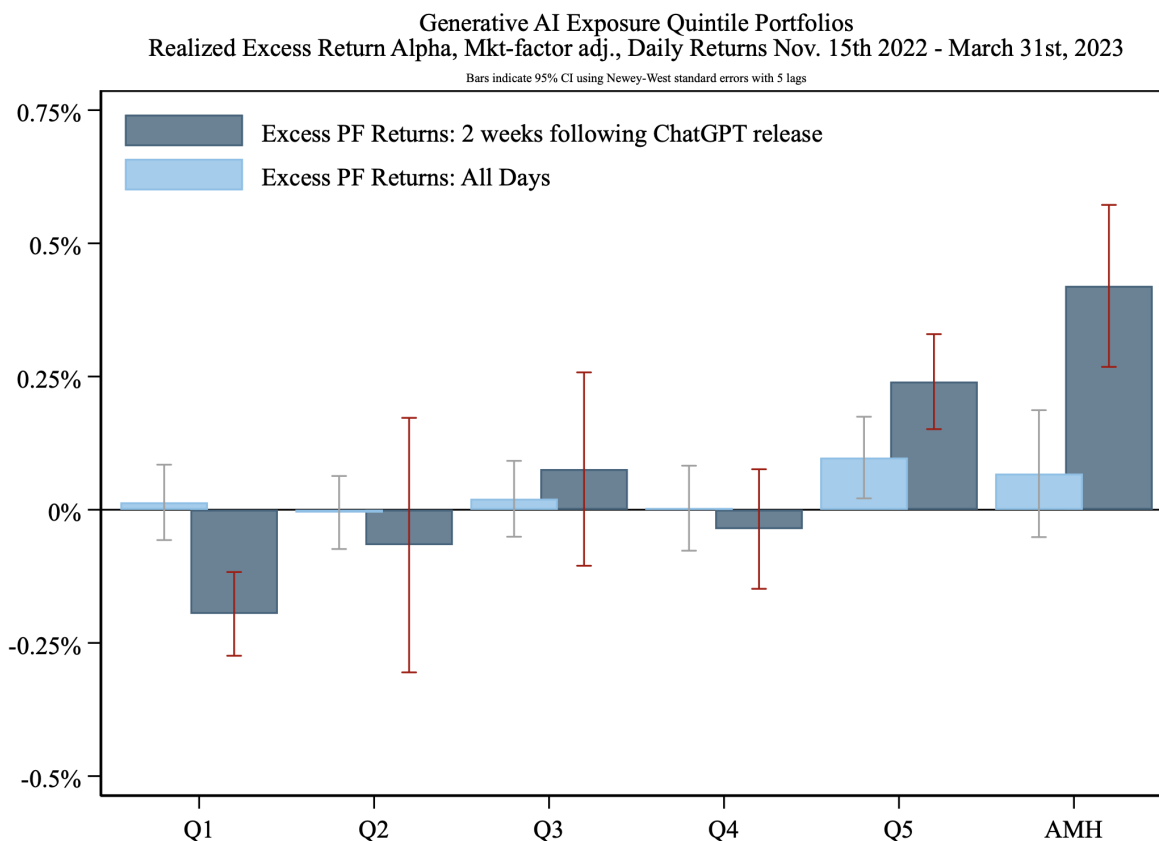


Figure B2: Industry-neutral Generative AI exposure quintile portfolio returns: Fama French 5-factor adjusted. Each bar shows the average daily return alpha during the ChatGPT release period and also for all days, for portfolios based on industry-neutral Generative AI exposure sorts, and also for the net zero investment high-minus-low exposure portfolio that represents the “Artificial Minus Human” (*AMH*) factor. Industry-neutral portfolios are computed by first forming within-industry value-weighted quintile portfolios, and then averaging portfolio returns for the same quintiles across industries. Returns within-industry quintiles are value-weighted, while across-industry averages are industry market-cap. weighted. The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023, and Fama French factors from Ken French’s website. The figure shows alphas estimated from portfolio-level regressions of the form

$$r_{it}^{pf} - r_t^f = \alpha_i^{\text{release}} \mathbb{1}[\text{ChatGPT release period}]_t + \alpha_i^{\text{not release}} \mathbb{1}[\text{Not ChatGPT release period}]_t + \sum_{fac \in FF5} \beta_i^{fac} r_t^{fac} + \varepsilon_{it},$$

where the bars show either the intercept for the “ChatGPT release period”, or the intercept from a regression where all days have the same intercept. Error bars indicate 95% confidence intervals computed using Newey West standard errors with five lags.

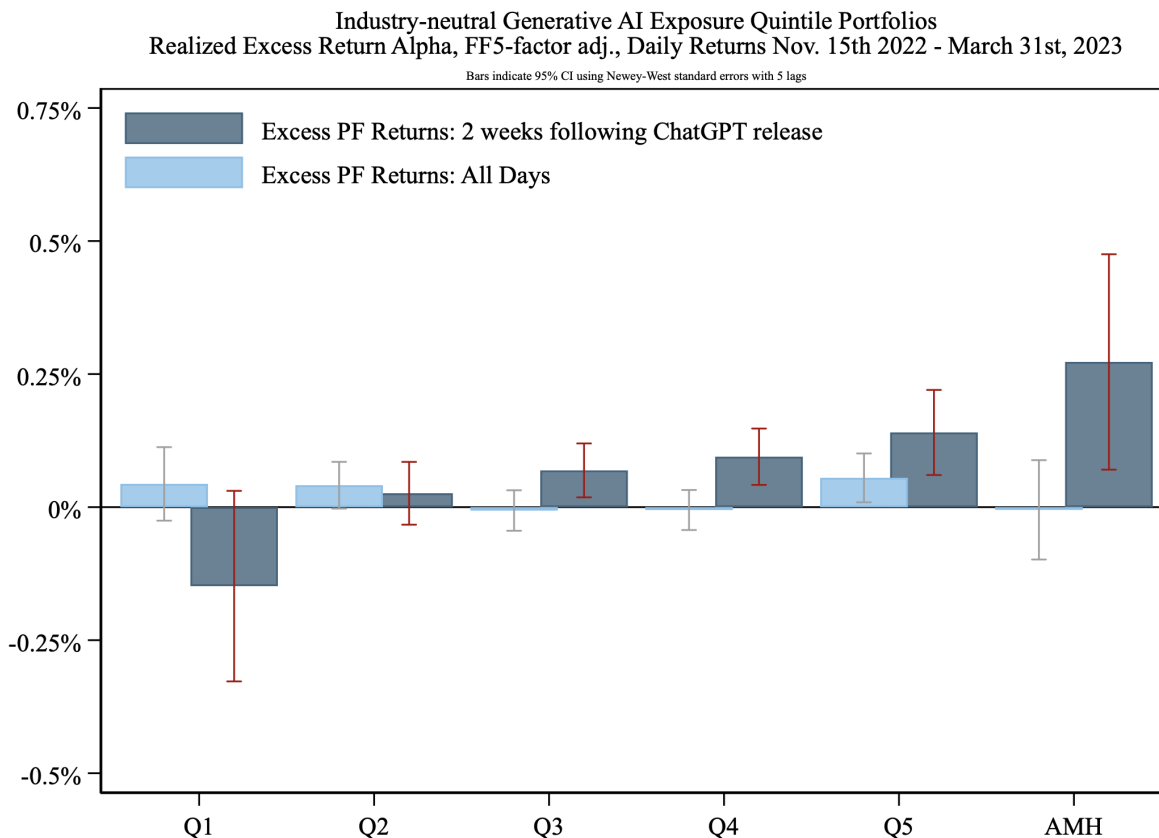


Figure B3: Generative AI exposure quintile portfolio returns over time: Fama French 5-factor adjusted. The top graph shows Generative AI exposure quintile sorts across all stocks in the sample. The bottom graph shows the cumulative excess realized returns on portfolios based on industry-neutral Generative AI exposure sorts. Industry-neutral portfolios are computed by first forming within-industry value-weighted quintile portfolios, and then averaging portfolio returns for the same quintiles across industries. Returns within-industry quintiles are value-weighted, while across-industry averages are industry market-cap. weighted. 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, and Fama French factors from Ken French's website. The Fama French 5-factor model robust regressions include the market factor, *HML*, *SMB*, *RMW*, and *CMA* factor returns in the regression.

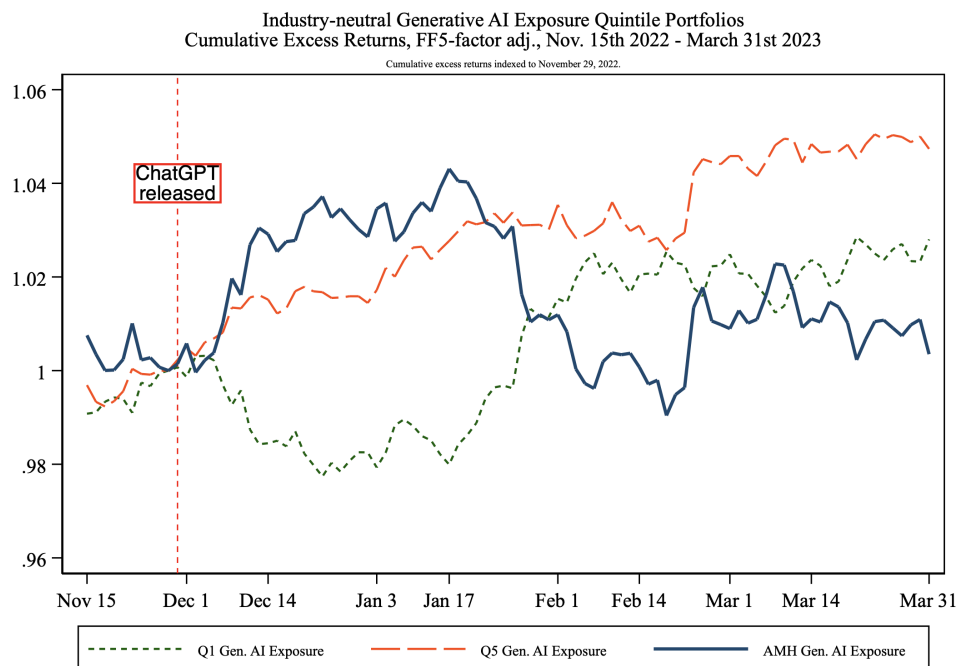
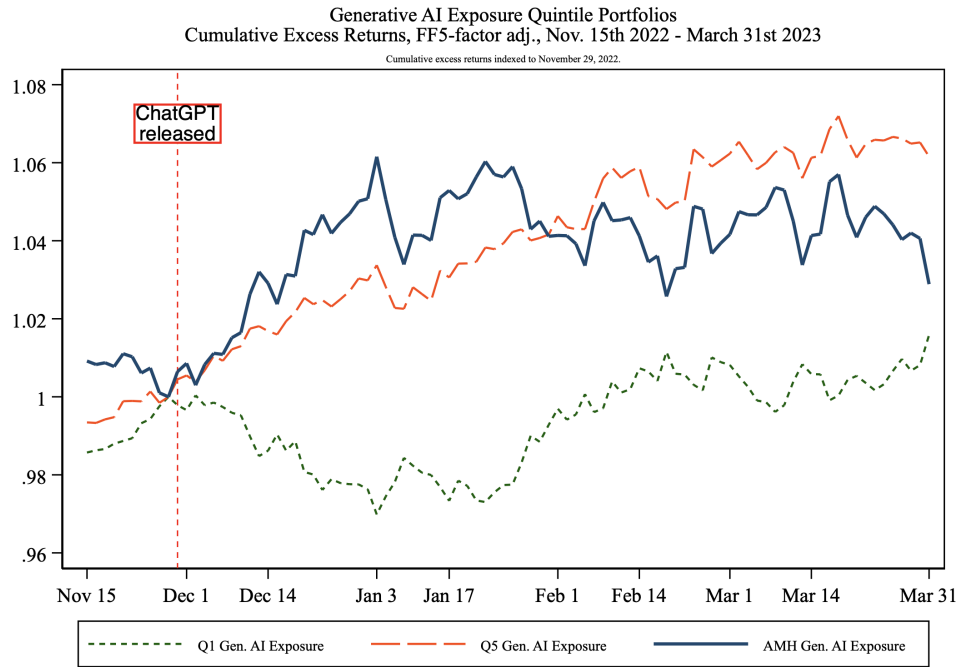


Figure B4: Generative AI exposure quintile portfolio returns over time. The top graph shows Generative AI exposure quintile sorts across all stocks in the sample. The bottom graph shows the cumulative excess realized returns on portfolios based on industry-neutral Generative AI exposure sorts. Industry-neutral portfolios are computed by first forming within-industry value-weighted quintile portfolios, and then averaging portfolio returns for the same quintiles across industries. Returns within-industry quintiles are value-weighted, while across-industry averages are industry market-cap. weighted. 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.

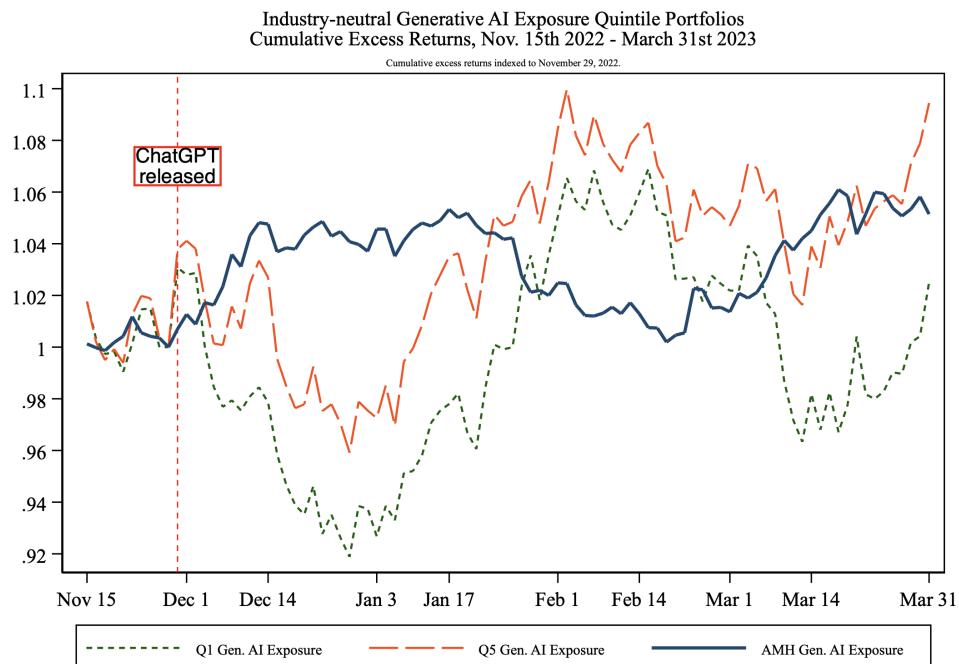
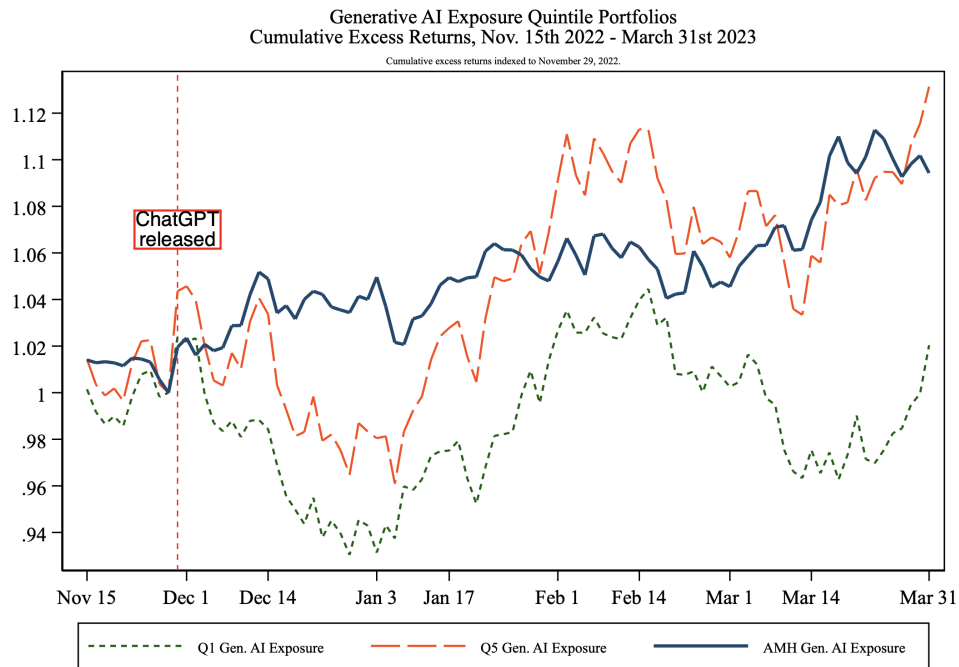


Figure B5: Within-Sector H-L Generative AI exposure Portfolio Realized Returns: Not ChatGPT Release Period. The graph shows the average daily return alpha for days that are not in the ChatGPT release period for H-L Generative AI exposure portfolios within each industry, at the 2-digit NAICS sector level. Each H-L 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 H-L 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, and the following two weeks, 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.

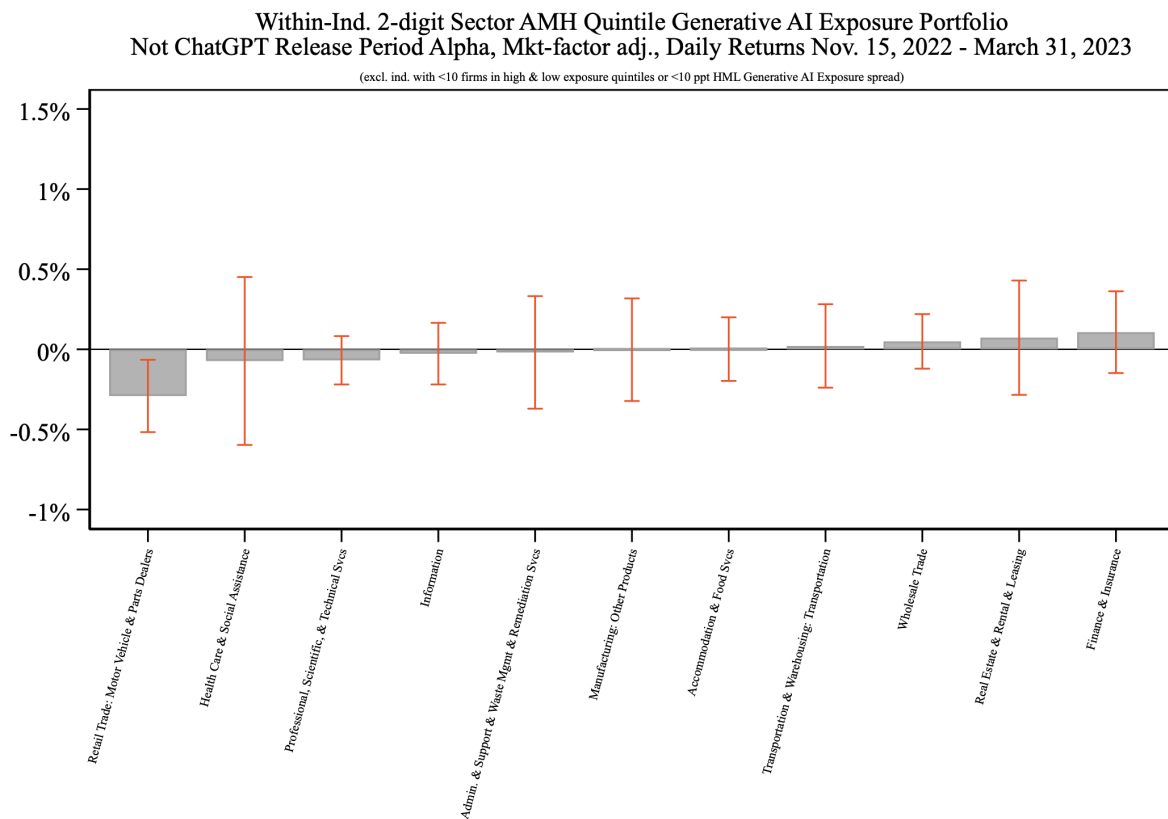
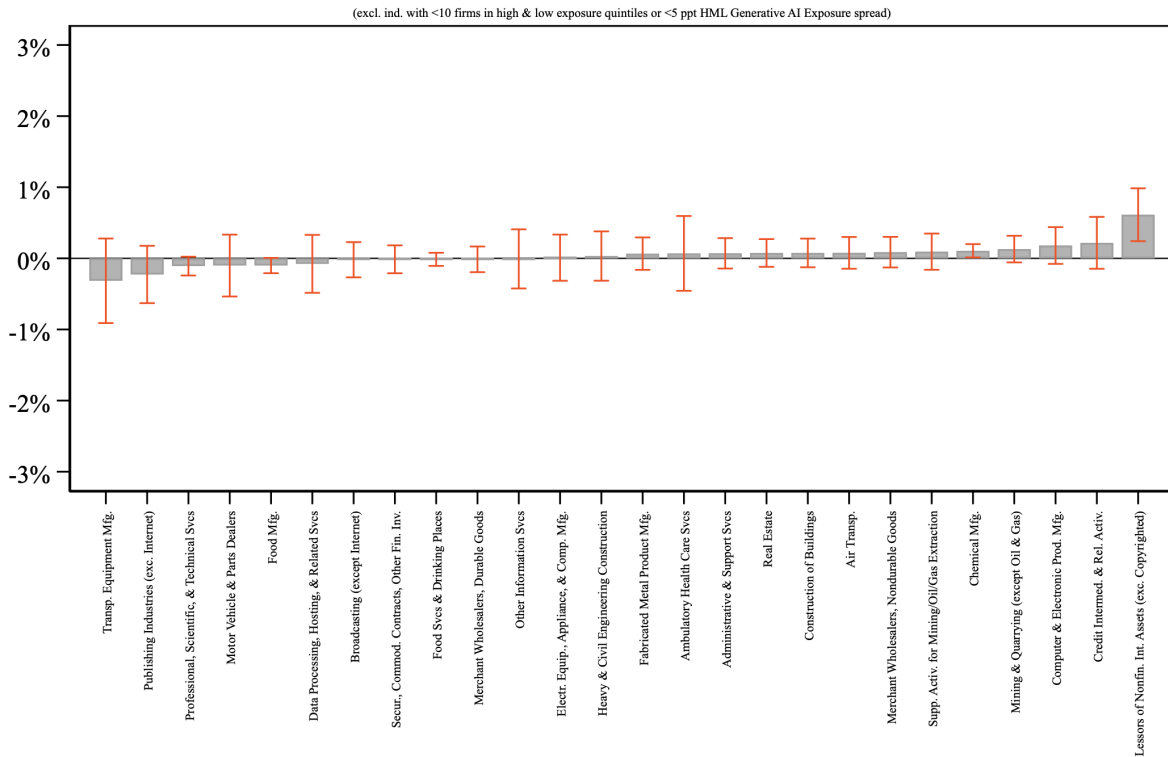


Figure B6: Within-Subsector H-L Generative AI exposure Portfolio Realized Returns: Not ChatGPT Release Period. Each graph shows the average daily return alpha for days that are not in the ChatGPT release period for H-L Generative AI exposure portfolios within each industry, at the 3-digit subsector level. Each H-L portfolio is formed by taking the value-weighted highest and lowest terciles of Generative AI exposure within each industry (based on NYSE stock cutoffs) and forming zero net investment H-L 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 subsectors with fewer than 10 firms combined in the highest and lowest quantiles in the sample, as well as subsectors with less than a 5 ppt Generative AI exposure spread between the high and low tercile. 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, and the following two weeks, 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.

**Within-Ind. 3-digit Subsector AMH Tercile Generative AI Exposure Portfolio
No ChatGPT Release Period Alpha, Mkt-factor adj., Daily Returns Nov. 15, 2022 - March 31, 2023**



Appendix C. 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 C1: Examples of GPT scores assigned to task statements and GPT-provided explanations.

Table C2: Exposure score variation across GPT scoring runs

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

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 C3: Highest and lowest Generative AI exposure score occupations

NAICS Code	Industry Title	Exposure Score
52	Finance and insurance	.49
54	Professional, scientific, and technical services	.49
55	Management of companies and enterprises	.48
51	Information	.47
42	Wholesale trade	.35
91	Federal government	.34
53	Real estate and rental and leasing	.33
90	Government	.3
22	Utilities	.29
61	Educational services; state, local, and private	.29
56	Administrative and support and waste management and remediation services	.27
81	Other services (except public administration)	.24
31-33	Manufacturing	.24
44-45	Retail trade	.22
62	Healthcare and social assistance	.22
71	Arts, entertainment, and recreation	.22
21	Mining, quarrying, and oil and gas extraction	.21
48-49	Transportation and warehousing	.2
23	Construction	.17
72	Accommodation and food services	.11
11	Agriculture, forestry, fishing and hunting	.086

Table C4: Generative AI exposure scores by industry

Table C5: Within-industry H-L Generative AI exposure returns: 2-digit sectors.

The table shows the realized excess return on GPT news days and no-GPT-news days on the within-industry H-L Generative AI exposure quintile portfolio for the Nov. 15, 2022 - March 31, 2023, period, in daily returns data from Yahoo Finance, controlling for average market returns. Each H-L portfolio is formed by taking the value-weighted highest and lowest quintiles of Generative AI exposure within each industry sector and forming zero net investment H-L 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 are 2-digit NAICS sectors, omitting any sectors with fewer than 10 firms combined in the high and low quintiles in the sample, as well as industries with less than a 10 ppt Generative AI exposure spread between the high and low quintile. Standard errors shown are Newey-West standard errors with a five period lag bandwidth.

Ind. Code	Industry Title	News Day α	News Day t-Stat.	No News Day α	No News Day t-Stat.
52	Finance and Insurance	0.81	2.54	0.11	0.82
33	Manufacturing: Other Products	0.81	2.22	-0.00	-0.02
56	Admin. and Support and Waste Mgmt and Remediation Svcs	0.54	2.66	-0.02	-0.11
42	Wholesale Trade	0.24	1.47	0.05	0.57
51	Information	0.10	0.73	-0.03	-0.28
62	Health Care and Social Assistance	0.06	0.22	-0.07	-0.27
54	Professional, Scientific, and Technical Svcs	-0.01	-0.05	-0.07	-0.89
44	Retail Trade: Motor Vehicle and Parts Dealers	-0.08	-0.29	-0.29	-2.53
48	Transportation and Warehousing: Transportation	-0.11	-0.64	0.02	0.16
72	Accommodation and Food Svcs	-0.22	-1.56	0.00	0.01
53	Real Estate and Rental and Leasing	-0.49	-3.05	0.07	0.40

Table C6: Within-industry H-L Generative AI exposure returns: 3-digit sub-sectors. The table shows the realized excess return on GPT news days and no-GPT-news days on the within-industry H-L Generative AI exposure tercile portfolio for the Nov. 15, 2022 - March 31 2023, period, in daily returns data from Yahoo Finance, controlling for average market returns. Each H-L portfolio is formed by taking the value-weighted highest and lowest terciles of Generative AI exposure within each industry subsector and forming zero net investment H-L 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 are 3-digit NAICS subsectors, omitting any sectors with fewer than 10 firms combined in the high and low quintiles in the sample, as well as industries with less than a 5 ppt Generative AI exposure spread between the high and low tercile. Standard errors shown are Newey-West standard errors with a five period lag bandwidth.

Ind. Code	Industry Title	News Day α	News Day t-Stat.	No News Day α	No News Day t-Stat.
511	Publishing Industries (exc. Internet)	1.19	2.14	-0.23	-1.10
522	Credit Intermed. and Rel. Activ.	1.09	3.52	0.22	1.17
336	Transp. Equipment Mfg.	0.68	1.21	-0.32	-1.04
561	Administrative and Support Svcs	0.55	2.53	0.07	0.65
441	Motor Vehicle and Parts Dealers	0.42	1.53	-0.10	-0.46
621	Ambulatory Health Care Svcs	0.36	1.75	0.07	0.26
523	Secur., Commod. Contracts, Other Fin. Inv.	0.25	0.59	-0.01	-0.14
334	Computer and Electronic Prod. Mfg.	0.22	1.52	0.18	1.37
423	Merchant Wholesalers, Durable Goods	0.19	0.74	-0.01	-0.15
213	Supp. Activ. for Mining/Oil/Gas Extraction	0.17	0.53	0.09	0.72
424	Merchant Wholesalers, Nondurable Goods	0.13	0.58	0.09	0.79
722	Food Svcs and Drinking Places	0.07	0.77	-0.01	-0.30
518	Data Processing, Hosting, and Related Svcs	0.03	0.20	-0.08	-0.37
236	Construction of Buildings	0.01	0.15	0.08	0.74
237	Heavy and Civil Engineering Construction	-0.02	-0.05	0.03	0.18
212	Mining and Quarrying (except Oil and Gas)	-0.04	-0.41	0.13	1.36
533	Lessors of Nonfin. Int. Assets (exc. Copyrighted)	-0.06	-0.15	0.61	3.24
541	Professional, Scientific, and Technical Svcs	-0.12	-0.58	-0.11	-1.62
481	Air Transp.	-0.17	-1.11	0.08	0.68
335	Electr. Equip., Appliance, and Comp. Mfg.	-0.20	-0.41	0.01	0.06
325	Chemical Mfg.	-0.26	-1.87	0.11	2.24
332	Fabricated Metal Product Mfg.	-0.26	-0.75	0.07	0.57
515	Broadcasting (except Internet)	-0.29	-0.66	-0.02	-0.16
311	Food Mfg.	-0.37	-2.40	-0.10	-1.85
531	Real Estate	-0.39	-2.55	0.08	0.76
519	Other Information Svcs	-0.65	-3.60	-0.01	-0.04

Table C7: Realized returns of portfolios sorted on Generative AI exposure after ChatGPT release.: adjusted for Fama French 5 factors. This table reports daily excess stock returns of value-weighted portfolios of firms sorted on Generative AI exposure. *AMH* is the "Artificial Minus Human" is the zero net investment portfolio long high exposure (*H*) stocks and short low exposure (*L*) stocks. Quintile thresholds that define value-weighted portfolios are solely based on the sample of stocks listed on NYSE as of the sorting date. All quintile portfolios are formed based on value weights on October 31, 2022, and weights are adjusted based on daily returns to mimic passive buy-and-hold exposure. Industry-neutral portfolios are computed by first forming within-industry value-weighted quintile portfolios, and then averaging portfolio returns for the same quintiles across industries. Returns for within-industry quintiles are value-weighted, while across-industry averages are industry market-cap. weighted. The data set consists of daily stock returns from Yahoo Finance for Nov. 15, 2022 - March 31, 2023. Daily market returns, risk free rates, and additional factors are obtained from Ken French's website. The table shows alphas estimated from regressions of the form

$$r_{it}^{pf} - r_t^f = \alpha_i^{\text{release}} \mathbb{1}[\text{ChatGPT release period}]_t + \alpha_i^{\text{not release}} \mathbb{1}[\text{Not ChatGPT release period}]_t + \sum_{fac \in FF5} \beta_i^{fac} r_t^{fac} + \varepsilon_{it},$$

where the intercept is either for the full sample, or is allowed to vary with whether the day is in the ChatGPT release period consisting of Nov 30, 2022 - Dec. 14, 2022. The Fama French 5-factor model robust regressions include the market factor, *HML*, *SMB*, *RMW*, and *CMA* factor returns r_t^{fac} in the regression. T-statistics in parentheses are computed using Newey-West standard errors with five lags.

Sample	Portfolios					
	Q1	Q2	Q3	Q4	Q5	AMH
<i>A: Excess returns (%)</i>						
All days	0.042 (0.38)	0.027 (0.25)	0.050 (0.54)	0.037 (0.31)	0.134 (1.17)	0.076 (1.33)
Not ChatGPT release period	0.066 (0.57)	0.030 (0.25)	0.038 (0.36)	0.037 (0.28)	0.110 (0.85)	0.028 (0.44)
ChatGPT release period	-0.137 (-0.53)	0.000 (0.00)	0.137 (0.43)	0.034 (0.11)	0.316 (0.82)	0.437 (3.08)
<i>B: Fama French 5-factor-adjusted alpha (%)</i>						
All days	0.040 (1.11)	0.021 (0.71)	-0.024 (-1.05)	0.036 (1.27)	0.068 (2.40)	0.011 (0.22)
Not ChatGPT release period	0.061 (1.57)	0.024 (0.76)	-0.025 (-0.94)	0.034 (1.09)	0.057 (1.91)	-0.020 (-0.36)
ChatGPT release period	-0.132 (-1.96)	-0.000 (-0.00)	-0.021 (-0.43)	0.053 (0.86)	0.155 (3.37)	0.272 (3.09)
<i>C: Ind.-neutral Fama French 5-factor-adjusted alpha (%)</i>						
All days	0.044 (1.24)	0.041 (1.84)	-0.006 (-0.33)	-0.005 (-0.28)	0.055 (2.35)	-0.005 (-0.11)
Not ChatGPT release period	0.067 (2.00)	0.043 (1.69)	-0.016 (-0.75)	-0.018 (-0.99)	0.045 (1.80)	-0.039 (-0.84)
ChatGPT release period	-0.148 (-1.63)	0.026 (0.86)	0.069 (2.67)	0.095 (3.51)	0.140 (3.44)	0.273 (2.64)

Table C8: 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