

NBER WORKING PAPER SERIES

GENERATIVE AI AND FIRM VALUES

Andrea L. Eisfeldt
Gregor Schubert
Miao Ben Zhang

Working Paper 31222
<http://www.nber.org/papers/w31222>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2023, Revised January 2026

We thank David Autor, Tania Babina, Thomas Chaney, Charlie Clarke, Olivier Dessaint, Anastassia Fedyk, Alex He, Chad Jones, Sophia Kazinnik, Mete Kilic, Elena Simintzi, Selale Tuzel, Stijn Van Nieuwerburgh, Junbo Wang, Paul Whelan, Emmanuel Yimfor, Paolo Zaffaroni, and participants in the UCLA Anderson Finance Brown Bag, the UCSD Finance Seminar, NYU Stern Finance Seminar, KDD Finance Day, Q Group 2023 Meeting, Duke/UNC Asset Pricing Conference, IMF-WIFPR Conference, the 2023 “Artificial Intelligence and the Economy” conference, Carey Finance Conference, TCU Finance Conference, the Chicago Booth Empirical Finance Conference, the OECD, the ITAM Finance Workshop, Macro-Finance Society Workshop, University of Kentucky Finance Conference, UW Summer Finance Conference, the IESE AI in Finance Workshop, the CEPR International Conference on Generative AI, the Purdue FinTech Conference, the Wharton AI and the Future of Work Conference, the NBER Big Data and Securities Markets 2023 Meeting, the Federal Reserve Board, the MIT Initiative on the Digital Economy, the Bank for International Settlements Seminar, and the Stanford Digital Economy Lab for helpful feedback. We also thank Colton Etheridge and Yi Li for their research assistance. All remaining errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Andrea L. Eisfeldt, Gregor Schubert, and Miao Ben Zhang. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Generative AI and Firm Values

Andrea L. Eisfeldt, Gregor Schubert, and Miao Ben Zhang

NBER Working Paper No. 31222

May 2023, Revised January 2026

JEL No. E0, G0

ABSTRACT

How do recent advances in Generative AI affect firm value? We construct the first measure of firms' workforce exposures to Generative AI and show that an "Artificial-Minus-Human" (AMH) portfolio earned 5% in the two weeks following the release of ChatGPT. The labor-exposure effect is more pronounced for firms with greater data assets and is distinct from the effect of firms' product exposures to AI. We assess whether exposed workforces are substituted or complemented by Generative AI based on whether their exposed tasks are core or supplemental. Examining firms' labor demand and profitability following the release of ChatGPT supports a labor-technology substitution channel.

Andrea L. Eisfeldt

University of California, Los Angeles

Anderson School of Management

and NBER

andrea.eisfeldt@anderson.ucla.edu

Miao Ben Zhang

University of Southern California

miao.zhang@marshall.usc.edu

Gregor Schubert

University of California, Los Angeles

schubert.gregor@gmail.com

Recent advances in Generative Artificial Intelligence are widely seen as a major technological breakthrough. We construct the first *firm-level* measure of the exposure of U.S. publicly traded firms' workforces to Generative AI. Using this measure, we show that the release of ChatGPT resulted in a significant and substantial change in the relative valuation of US firms. Firms whose workforces are highly exposed to Generative AI increased in market value by almost 5% relative to firms with a low exposure in the two weeks following the release of ChatGPT on November 30th, 2022. The divergence in value reflects the varied potential for Generative AI to execute the tasks currently accomplished by firms' workforces, and the resulting effects on firms' expected future free cash flows. It is not driven by firms' product-market exposure to Generative AI.

The release of ChatGPT can improve the expected future cash flows of firms with a high share of exposed workforces due to productivity improvements either from substituting labor with technology to save costs, or from complementing labor with Generative AI tools. We show that the labor substitution channel drives our results, and develop a new method for assessing an occupation's potential for Generative AI substitution versus complementarity. In particular, we show that the substitution effects measured by lower wages and job postings after the release of ChatGPT are concentrated in occupations that are most exposed to Generative AI through their *core* tasks. Examples include proofreaders and web developers. Occupations most exposed through *supplemental* tasks, such as financial managers and pharmacy aids, experience much smaller substitution effects. Our use of core versus supplemental task exposure to study substitution versus complementarity effects of technology on occupations can help to understand which occupations may ultimately benefit from Generative AI enhancements. However, our results for Generative AI's impact on firm value and future profitability illustrate that the main driver of the increase in value is the substitution channel. The observed increase in firm value for firms whose workforces are more exposed to Generative AI is driven by occupations whose core tasks are more easily executed by Generative AI technologies.

Generative AI is a general-purpose technology and is changing (and expected to change) the way work is conducted across a broad array of industries. Relative to earlier artificial intelligence models, Generative AI models can digest more complex inputs, and can produce human-like output, making Generative AI models (of which ChatGPT is an important example) more versatile and scalable than prior innovations in AI and machine learning. In the past, technology shocks diffused over long periods of time, but the large change in the accessibility of AI tools and the massive amount of analysis and attention generated by the release of ChatGPT allow us to study its impact almost in real-time.¹ Our study,

¹A large body of literature examined the historical diffusion of technologies and showed that the diffusion

which includes effects on the full cross-section of firm values, wages, job postings, and firm profits, thus sheds light on the quantitative effects of an important technology shock on key corporate and economic outcomes.

Recent work has documented variation in occupational exposure to Generative AI (see [Eloundou et al. \(2023\)](#)). However, without a connection to firm value, it is unclear whether higher exposure is value-increasing or value-decreasing for incumbent firms. If Generative AI does increase efficiency for existing firms, how much value do firms versus workers accrue? By documenting the increase in the value of firms with higher workforce exposure to Generative AI, while wages and job postings decline for exposed workers, we show that firm owners benefited relatively more than labor from the technology shock. This is an important finding regarding the economic impact, especially in the context of the recent literature on the decline in the share of value added accruing to labor. Our study is the first to document firms' workforce exposure to Generative AI, and to document the effect of these exposures on key labor market outcomes, including wages and job postings. Moreover, we link these labor outcomes to firm profits, providing additional support for a labor-substitution effect driving changes in firm value. By refining the measurement of task-level exposures using the distinction between core and supplementary task exposures, we are able to uncover both substitution and complementary effects of Generative AI on labor.

We build our measure of firms' exposure to Generative AI starting at the level of the tasks that workers in U.S. firms perform. From the O*Net database, we obtain information on the 19,265 tasks that constitute the activities performed in 923 SOC 8-digit occupations in the U.S. We use a large language model to classify each task's exposure into categories based on whether it can be done more effectively using ChatGPT based on descriptions of the tasks and known capabilities of Generative AI models. Aggregating to the SOC 6-digit level, our final sample includes 678 occupations with, on average, 23% of their tasks exposed to Generative AI, and 85% of the occupations having at least one task exposed to Generative AI. We aggregate this task-level exposure to construct average occupational exposures. We construct firm-level exposure estimates for publicly traded firms using data from Revelio Labs, which provides firms' occupational shares based on millions of individual public profiles. Our firm-level measure captures the potential for the tasks currently performed by labor at those firms to be done more efficiently using Generative AI.

process is remarkably slow, particularly in the initial periods after the technologies are available, constituting an S-curve for technology diffusion (see many examples in [Greenwood \(1999\)](#) and [Manuelli and Seshadri \(2014\)](#)). ChatGPT reached 1 million users in only 5 days after its release, according to OpenAI's CEO Sam Altman's [Tweet](#) on December 4, 2022. It reached 100 million users in only two months after its release, a milestone that took the World Wide Web 7 years, WhatsApp 3.5 years, Instagram 2.5 years, and Twitter 5 years to reach (see this [link](#)).

Our approach yields intuitive exposures that reflect the distinctive features of Generative AI compared to prior technologies. Unlike computerization that mainly disrupted routine-task jobs (e.g., [Autor et al. \(2003\)](#) and [Acemoglu and Autor \(2011\)](#)) and robots that mainly disrupted manual-task jobs ([Acemoglu and Restrepo \(2020\)](#)), occupations exposed to Generative AI are cognitive-task jobs, such as legal, financial, mathematical, and administrative professions.² Consistent with Generative AI being able to assist with high-level cognitive tasks, we show that occupations with higher wages also have higher exposure to Generative AI.³ Overall, the Generative AI exposure of the tasks in all white-collar occupations is 40% versus 9% for blue-collar and service occupations. Our firm-level measure also shows that Generative AI is indeed a general-purpose technology, which varies both across and within a broad array of industries, with 3-digit NAICS industry effects explaining less than a third of the variation in the firm-level exposure.

We derive our first main results from an event study that documents variation in the returns to firms with different labor-force exposures to Generative AI following the release of ChatGPT on November 30th, 2022. Since we study firms' relative returns following a major event, our study can be thought of as looking at the effect of the release of ChatGPT on difference-in-differences in the levels of firms' valuations. Sorting firms into five value-weighted portfolios based on their Generative AI exposures, we show that firms in the highest-exposure quintile, labeled the "Artificial" portfolio, earned 44 basis points higher daily returns than firms in the lowest-exposure quintile, labeled the "Human" portfolio, during the two weeks following the release of ChatGPT.⁴ This finding is robust to controlling for firms' exposure to the market factor and the Fama-French 5 factors, and also to controlling for firm characteristics that have been shown to predict returns. Figure 1 plots the cumulative abnormal returns during the time before and after the event for the zero-investment portfolio that goes long the artificial stocks and short the human stocks, which we denote as the "Artificial Minus Human" portfolio (AMH). The AMH portfolio returns do not show an obvious pre-trend before the event window, and also no reversal after the event window, supporting the validity of our treatment effect estimation.⁵

While the release of ChatGPT is also expected to affect firm values through effects on

²Our finding is consistent with [Webb \(2019\)](#), who studies broadly-defined AI on the labor market using patent texts and shows that AI tends to disrupt non-routine cognitive-analytical jobs.

³Our result is consistent with recent findings by [Kogan et al. \(2019\)](#), who find that technological advances impact workers at the higher end of the wage distribution. On the other hand, other studies ([Krusell et al. \(2000\)](#) and [Eisfeldt et al. \(2023\)](#)) document substitutability between low-skilled labor and capital but complementarity between high-skilled labor and capital.

⁴Our choice of the two-week event window is guided by the intensity of public attention to the release of ChatGPT on Twitter.

⁵See Appendix Figure IA.1 for a version of the figure with a longer pre-period.

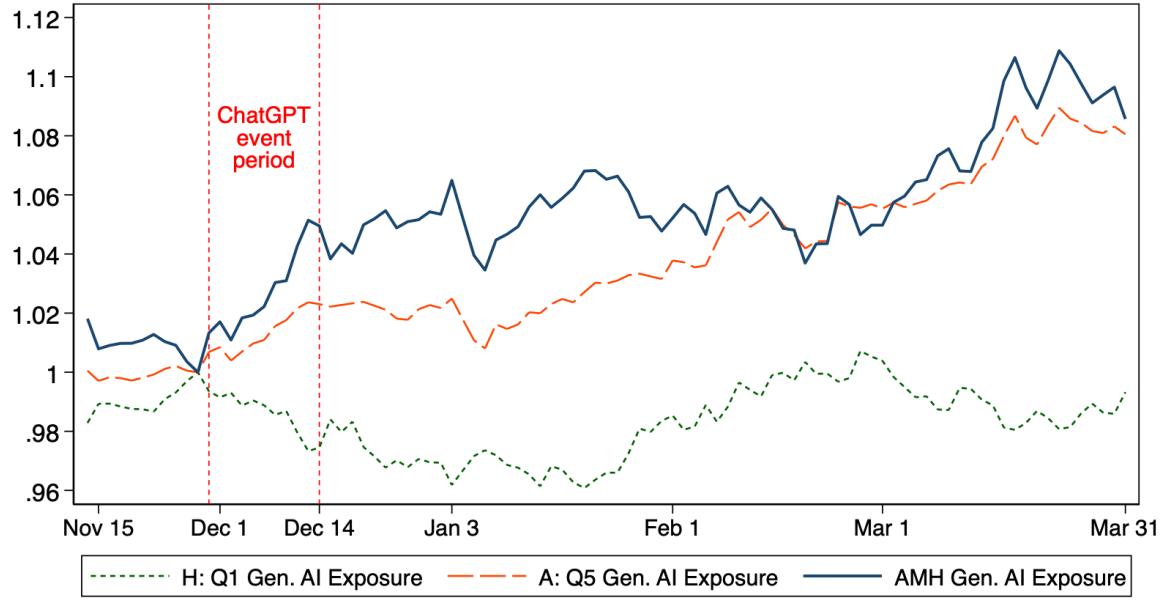


Figure 1. Cumulative Abnormal Returns by Generative AI Exposure. The figure plots the cumulative abnormal returns (CARs) of value-weighted quintile portfolios sorted by firms' labor-based Generative AI exposure. The graph shows the CARs of the lowest-exposure quintile portfolio, "Human" (H), the highest-exposure quintile portfolio, "Artificial" (A), and the zero investment portfolio that longs A and shorts H, "Artificial-minus-Human" (AMH). Market-adjusted daily abnormal returns are cumulated from November 29, 2022, the day before the release of ChatGPT, and are based on factor exposures computed over the 6-month period preceding the period shown in the graph. Daily stock returns are from Yahoo Finance. The dashed vertical lines indicate the "ChatGPT event period" from November 30, 2022, to December 14, 2022. See details of the definition of firms' Generative AI exposure in Section I and the construction of the portfolios, the calculation of portfolios' CARs, and the determination of the ChatGPT event period in Section II. GPT-4 was released on March 14, 2023.

products and services,⁶ we conduct three separate tests to show that our findings based on firms' labor exposure to Generative AI are distinct from any product-exposure channel. First, we show that the AMH returns during the event window hold when sorting portfolios within industries. We show this using both the NAICS 3-digit industry classification and the 10K text-based FIC 50 industry classification by [Hoberg and Phillips \(2016\)](#). Noting that only 31% of the firm-level variation in labor exposure to Generative AI is explained by 3-digit industry effects, our results are, not surprisingly, robust to these within-industry sorts. Second, the AMH returns remain the same if we exclude the tech sector, in which firms' products are most likely to be directly related to the Generative AI technology.⁷ Third, we

⁶For example, the stock value of chip maker NVIDIA more than doubled in early 2023 as the firm is a key chip supplier for training Generative AI models.

⁷Following [Acemoglu et al. \(2022\)](#), we identify the tech sector as the NAICS 51 "Information" and NAICS

use four proxies for firms' product exposure to Generative AI: a classification of AI-related business models based on firm annual reports; the count of AI-related keywords based on texts of firms' annual reports; a Goldman Sachs classification of AI beneficiaries; and a measure based on firms' share of AI-skilled workers adopted from [Babina et al. \(2024\)](#). We then run a standard event study test by regressing stock-level cumulative abnormal returns (CARs) on our labor-based Generative AI exposure measure while controlling for the proxies for product exposures. While firms' product exposures separately predict CARs, the effect of our labor-based Generative AI exposure remains significant, confirming the distinction between the labor and product-exposure channels.

We show that data is an important complement to Generative AI. This may help to understand the benefits of AI improvements for incumbent firms with existing data corpora. Examples highlighting the importance of data for using Generative AI can be seen in applications such as training customer service chatbots, automating workflows, improving predictions and analytics, and many others ([Caserta et al. \(2023\)](#)). Realizing the value associated with this technology at scale, therefore, may require a baseline level of data management capabilities and access to relevant input data. We thus expect the return effect of Generative AI exposure to be particularly pronounced among firms with readily available data. We find strong confirmation for this proposition. In particular, we construct two measures of firms' data assets following the prior literature ([Begenau et al. \(2018\)](#), [Farboodi et al. \(2019\)](#), [Eeckhout and Veldkamp \(2022\)](#), [Farboodi and Veldkamp \(2023\)](#), and [Abis and Veldkamp \(2023\)](#)), and we show that Generative AI exposure boosts firm value significantly more if the firm has greater data assets. This finding also explains why firms' potential benefits from Generative AI are not expected to be entirely competed away by new entrants, as data assets can be highly specific to, or proprietary property of, incumbent firms.

Our second set of main results documents the mechanism by which the labor channel generates our results for the impact of Generative AI on firm value. Importantly, we ask whether the effect of higher exposure to Generative AI reflects a greater substitution of, or complementarity to, firms' labor inputs. *Ex ante*, both channels could increase firms' future cash flows and boost their current market value. That is, firms whose labor force can be substituted with cheaper Generative AI-based capital could save costs and generate higher future cash flows. On the other hand, if the technology complements the firm's workers and increases their productivity, the firm may also experience an increase in future profitability.

To distinguish between these two channels, we refine our measure of occupational exposure. Specifically, we redefine exposure depending on whether the occupational tasks that are exposed to Generative AI are *core* or *supplemental* according to O*Net's classification.

⁵⁴ "Professional, Scientific, and Technical Services" and exclude firms from these two sectors.

Many occupations have exposure to Generative AI through both core and supplemental tasks, with 77% of the exposure from core tasks and 23% of the exposure deriving from supplemental tasks on average. Our key hypothesis is that an occupation is more likely to be substituted by Generative AI if its *core* tasks are more exposed to the productivity improvements enabled by the technology, while if an occupation’s *supplemental* tasks are more exposed, there is more opportunity for complementarity. Intuitively, core tasks represent the most fundamental duties an employee in that occupation is expected to perform. If Generative AI can more efficiently complete the core tasks at a much lower cost, then it is more likely that the technology can displace the occupation as a whole. In contrast, supplemental tasks are additional tasks or duties associated with the core tasks but which are not themselves considered critical or central to the occupation. If Generative AI can help workers more efficiently complete their supplemental tasks, this can free up time and effort for the worker to focus more on the core tasks, potentially increasing the worker’s productivity without making them replaceable. To the best of our knowledge, this is the first study that distinguishes between core and supplemental task exposures within occupations.

We conduct two sets of tests to investigate the mechanism for the labor-channel impact on firm value. Our first set of tests investigates the impact on occupation-level labor demand and wage rates as a result of higher overall Generative AI exposure. Next, we show that, holding total exposure constant, it is the occupations that derive their Generative AI exposure from core tasks that experience a decline in labor demand. We construct the number of job postings for each occupation-month from January 2022 to August 2023 from granular Lightcast job posting data. We also obtain individual-level hourly wage rates from January 2022 to October 2023 from the Census Current Population Survey data.

Our study contributes three key new findings regarding the labor market outcomes for different *occupations* before and after the release of ChatGPT. First, occupations with higher Generative AI exposure unconditionally experienced reduced labor demand and a lower relative wage rate after the release of ChatGPT: A one-standard-deviation increase in Generative AI exposure is associated with an 8% decline in job postings and a 0.6% decline in the hourly wage rate. Second, the magnitudes of the effects of Generative AI exposure on job posting and hourly wage rates are *66% and 97% greater*, respectively, if the occupations are more substitutable by Generative AI according to our approach, i.e. if occupational exposure is derived entirely from core tasks. Likewise, consistent with the greater potential for complementarities when supplemental tasks generate an occupation’s Generative AI exposure, the share of an occupation’s exposure stemming from supplemental tasks, as opposed to core tasks, significantly weakens the above associations. Third, the above findings remain when we examine *within-firm* occupational demand using granular job posting data, reinforcing

our interpretation that the occupation-level results are driven by firms reorganizing their operations after the technology shock instead of, for instance, variation in industry dynamics. In summary, these findings show that firms adjust their labor demand in response to the release of ChatGPT in a way that, on average, suggests a substitution effect in the initial months, but with significant heterogeneity depending on the kinds of tasks within an occupation that are affected.

Our second sets of tests examine the technology-labor substitution channel for explaining the cross-section of the labor demand of *firms* with high and low Generative AI exposure. This channel suggests that firms with a more exposed workforce may have a greater capacity to reduce reliance on exposed occupations, thereby decreasing costs, improving future profitability, and boosting firm value. To test this channel, we measure the shares of firms' Generative AI exposures that arise from their workers' supplemental tasks. We show that firms with higher Generative AI exposure unconditionally reduce job postings more for the occupations that are highly exposed to Generative AI. These firms see an increase in analyst forecasts of short-run and long-run earnings per share and also experience greater increases in actual profitability in quarterly earnings after the release of ChatGPT. Our estimates show that when firms' Generative AI exposure is derived entirely from the core tasks of their workforce, the effects of Generative AI exposure on labor demand and profitability are more than twice as large as the unconditional average effects. The impact of core-task Generative AI exposure on cumulative abnormal returns is also 75% greater than the impact of overall exposure. Our results for firms' labor demand, profitability, and market value strongly support that a labor-substitution channel is driving Generative AI's impact on firm value.

Our study contributes to the literature on disruptive technologies' impact on firm hiring and firm valuation.⁸ [Papanikolaou \(2011\)](#) and [Kogan and Papanikolaou \(2014\)](#) study the effects of investment-specific technological changes on asset prices. [Zhang \(2019\)](#) studies firms' exposure to routine-biased automation. [Babina et al. \(2024\)](#) and [Babina et al. \(2022\)](#) are key early contributions studying the effects of AI on firm growth, compensation, and workforce composition.⁹ See also [Webb \(2019\)](#) for the impact of AI on the workforce, and [Acemoglu et al. \(2022\)](#) for evidence of the effects of firm exposure to AI on hiring and skill demand over the 2010-2018 period. More recently, [Brynjolfsson et al. \(2025a\)](#) study the effects of generative AI on U.S. labor markets and find that employment in exposed occupations has declined in relative terms since the release of ChatGPT, especially among younger workers (see also in [Lichtinger and Hosseini Maasoum \(2025\)](#)). However, [Humlum](#)

⁸See [Greenwood et al. \(1997\)](#) for an early contribution on the long-run impacts of investment-specific technological change.

⁹These studies using job-posting data up to 2018 shed light on the contribution of AI to firm outcomes even before the advances offered by more recent Generative AI models.

and Vestergaard (2025) focus on effects within exposed occupations in Denmark and find no effects of generative AI adoption on hours worked or wages in that setting. Kelly et al. (2021) study firms' exposure to disruptive technological shocks using patent textual data, and Kogan et al. (2019) assess worker displacement from technological change over a very long sample. These two studies offer important insights into investors' and firms' responses to technological shocks using a long panel containing several innovation waves.

We measure firms' exposure to Generative AI and assess investors' reaction to the technology shock upon its arrival. We show that the release of ChatGPT in November 2022 is an observable and significant technology shock that created a substantive impact on firm valuation. As the stock market is forward-looking, the information contained in market prices can potentially inform firms and employees about where the technology is likely to be most disruptive. Indeed, in the same month we first released our study, IBM, the company ranked #1 in our exposure to Generative AI measure among the largest U.S. firms, announced it would halt hiring of 7,800 jobs.¹⁰ Timely assessment of the market's expectations of Generative AI's impact on firms can also help policymakers to effectively evaluate regulatory policies in response to the arrival of the new technology.

This paper contributes to a large body of literature examining the heterogeneous effects of technological shocks on workers with different skills or tasks (see, for examples, Krusell et al. (2000), Autor et al. (2003), Acemoglu and Autor (2011), Acemoglu and Restrepo (2018), Eisfeldt and Papanikolaou (2013), Tuzel and Zhang (2021), Eisfeldt et al. (2022), Kogan et al. (2023), among others). A common empirical approach to measuring how substitutable capital is for labor after a technology shock is to assess whether the technology can more efficiently complete the tasks in that occupation at a lower cost. We adopt this approach in our study and offer a novel extension.¹¹

We decompose our measure of task-level exposures to Generative AI into exposures of core versus supplemental tasks. We argue that occupations with higher exposure stemming from core tasks offer more opportunities for capital to replace labor. On the other hand, workers in occupations with exposure derived from supplemental tasks may exhibit more complementarity with Generative AI capital. Our empirical results analyzing job postings and wage rates show that recognizing the core versus supplemental task distinction offers an improvement to the traditional approach to measuring task-level exposures and substitution

¹⁰We released our first draft to the SSRN on May 9, 2023, while Bloomberg reported the IBM announcement on May 1, 2023. See <https://www.bloomberg.com/news/articles/2023-05-01/ibm-to-pause-hiring-for-back-office-jobs-that-ai-could-kill>

¹¹Since the initial release of our study, other researchers have found that financial markets rapidly respond to generative AI news: for example, Andrews and Farboodi (2025) find that long-run bond yields fall around model release dates, and Bertomeu et al. (2025) show that a ChatGPT ban in Italy led to a rapid drop in stock prices for AI-exposed firms.

effects. Our core-task-based Generative AI exposure shows significantly stronger results than the overall task-based Generative AI exposure. By decomposing occupations into core and supplemental tasks, we refine the distinction between task and occupation substitutability. By aggregating up to the firm level, we distinguish between occupational and workforce substitutability. We see both of these distinctions as key contributions of our work.

The paper proceeds as follows: Section I describes our data and measures of firms' exposures to Generative AI. Section II presents the impact of the release of ChatGPT on the market value of firms with different Generative AI exposures. Section IV supports a labor-technology substitution channel by proposing and testing a novel methodology to identify workers' substitutability by Generative AI. Section V concludes.

I. Data, Measurement, and Stylized Facts

We measure a firm's labor exposure to Generative AI in three steps, starting with task-level exposures, and then aggregating to the occupation- and firm-level exposures, respectively. The Occupational Information Network (O*Net) database describes 19,265 tasks that make up the 923 occupations of U.S. workers. Each occupation executes a subset of these tasks. The task and occupational exposure measurement follow [Eloundou et al. \(2023\)](#). To aggregate the exposures from the occupation level to the firm level, we use the Revelio Labs database to measure firm-level occupational employment shares. A summary of our three-step procedure, which we detail further below, is:

1. *Task-level exposure*: We use Open AI's GPT 3.5 Turbo model to assign each of the 19,265 tasks in the O*Net an exposure to Generative AI by evaluating whether each task can be more efficiently completed by having access to the capabilities of LLM-based tools like ChatGPT.
2. *Occupation-level exposure*: Aggregate from tasks to occupation exposures by averaging the task-level exposures within each of the 923 occupations.¹²
3. *Firm-level exposure*: Compute the firm-level exposures by weighted averaging the occupation-level exposures using each firm's occupational employment share from the Revelio Labs database as the weight.

¹²We begin by equally weighting tasks within an occupation. In Section IV.D, we exploit the weights and distinguish between core and supplemental tasks to shed light on the substitution and complementarity of occupation to the Generative AI technology.

A. Measuring task exposure to Generative AI

Occupational task data We consider an occupation to be a portfolio of tasks to be done. From the O*Net V27.2 database, we obtain the task statement for each task in an occupation, documented by practitioners or occupational experts.¹³ A task statement is usually one sentence, and an occupation consists of 22 tasks on average. We code each of the 19,265 tasks as being exposed to Generative AI technologies or not using the task statement.

Task scoring We use OpenAI’s GPT 3.5 Turbo model and the following algorithm to score each task T ’s exposure X^T to Generative AI, following the approach suggested and validated by [Eloundou et al. \(2023\)](#). This approach categorizes each task into one of the following three categories $X^T \in \{1, 0.5, 0\}$ in terms of its Generative AI exposure:

- *Direct Exposure* ($X^T = 1$) if using ChatGPT reduces the time required to complete the task by at least half.
- *Indirect Exposure* ($X^T = 0.5$) if using ChatGPT would not reduce the time required to complete the task by at least half, but additional software could be developed on top of the existing capabilities of ChatGPT or related LLMs that could reduce the time it takes to complete the task with equal quality by at least half.
- *No Exposure* ($X^T = 0$) if using ChatGPT does not reduce the time required to complete the task by half while maintaining equivalent quality, or using ChatGPT reduces the quality of the task’s output.

Our classification uses the “few-shot prompting” technique and takes three steps. First, we create a system prompt that explains the classification exercise and describes each category. Second, we create two examples of user-assist prompts that showcase the expected output, where the user prompt asks ChatGPT to classify an example task and to explain its reasoning, while the assist prompt provides the example answers. Third, for each of the 19,265 tasks, T , we feed the GPT model with the system prompt, the two examples of user-assist prompts, and a new user prompt that includes the text of task T ’s statement and occupation title. The model produces answers similar to the assist prompt in the examples, including the classification and a short explanation, allowing us to audit whether the LLM actually understands the prompt as intended and interprets the task correctly. The Appendix [IA.3.C.1](#) details our prompts and the classification procedure.

Note that this classification method should *not* be interpreted as requiring ChatGPT to have any kind of correct “knowledge of its own capabilities.” Instead, the categories for what

¹³This data can be accessed via the O*Net website at <https://www.onetonline.org>.

kinds of capabilities state-of-the-art LLMs have are given by us to the model, as they were pre-defined by researchers in collaboration with OpenAI in [Eloundou et al. \(2023\)](#). That is, the right way to think about the role of ChatGPT here is as a research assistant mapping task statements into existing categories, which relies on its ability to interpret language, understand occupational contexts, and reason about which known LLM capabilities would be relevant for the task at hand. There is thus no need for ChatGPT to have oracular powers of prediction for “which tasks *it* could automate.” Instead, the LLM is scaling an approach that could, in theory, be implemented using only human research assistants with sufficient occupational context for each task.

However, the approach using GPT instead of human labeling has several important advantages. First, the GPT model is less subject to individual idiosyncrasies compared to human labeling because it leverages a vast amount of prior information that may be difficult for individual humans to master ([Gilardi et al. \(2023\)](#)).¹⁴ Second, the GPT approach permits rapid scaling of the method to categorize the complete set of 19,265 task statements, which would likely be very time-consuming for human labeling ([Frey and Osborne \(2017\)](#)). Third, the approach also provides explanations for labeling, allowing for an auditing capability often unavailable in instances of human labeling.

Of the 19,265 tasks, 14% were categorized as directly exposed to ChatGPT, $X^T = 1$. Table [I](#) provides examples of tasks in this category and explanations. For instance, “adjust sales scripts” for telemarketers and “write supporting codes for web applications” for web developers. We give these tasks an exposure score $X^T = 1$ as they are directly exposed to Generative AI. Another 22% were categorized as indirectly exposed to ChatGPT if appropriate software or applications based on the technology were developed, with $X^T = 0.5$. For instance, “review financial transactions” for food service managers and “identify (trading) opportunities” for financial services sales agents. Following [Eloundou et al. \(2023\)](#), we give a one-half exposure score for these tasks since their exposure to Generative AI is not as direct and requires further development of software and applications.¹⁵ The rest of the tasks are categorized as not exposed, $X^T = 0$. For instance, “connect heating or AC equipment” for installers and “mentor new faculty” for postsecondary business teachers.

¹⁴We further check the model consistency and show in Appendix Section [IA.3.C.2](#) that GPT reliably provides classifications that are highly consistent across different runs. We also check the confidence of the classification by requiring GPT to return a confidence score for its prediction, which shows “high confidence” in most cases. The GPT model has been increasingly used to classify content in recent academic studies, e.g., [Hansen et al. \(2023\)](#) and [Lopez-Lira and Tang \(2023\)](#). Appendix Section [IA.3.C.3](#) also validates the LLM-based scoring by comparing it to scores assigned by research assistants for a sub-sample, and showing that there is high alignment between human labels and LLM labels, with higher consistency between LLM runs than between human labelers.

¹⁵The results are similar if we use alternative discounts for the exposure score of tasks in the E2 category.

— Insert Table I about here —

B. Measuring occupation exposure to Generative AI

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

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

We then aggregate the 8-digit O*Net occupation codes to 6-digit SOC codes using equal-weighted averages of each 6-digit code’s 8-digit sub-codes to match the occupation-level exposure measure to firms’ occupational employment data. Generally speaking, this measure captures the percentage of an occupation’s tasks that can be more efficiently completed using ChatGPT and similar tools.¹⁶

C. Occupation-level stylized facts

In this section, we present summary statistics describing occupational exposures. We note that 15% of occupations in our data have zero exposure. The average exposure for white-collar occupations is 40%, while it is 9% for blue-collar and service occupations.¹⁷ We also show that Generative AI exposures are higher for occupational skill sets that are more cognitive, and that more exposed occupations tend to be at the higher end of the occupational wage distribution. The latter facts show how the Generative AI substitution effect is likely to operate on a very different part of the labor force relative to prior automation waves.

Panel A of Table II shows that the mean and median Generative AI exposure of occupations, X^O , are 23% and 18%, respectively, with a standard deviation of 21%. The 10th and

¹⁶After aggregating from O*Net occupations, we have exposure scores for 778 SOC 2010 occupations, of which 709 can also be found in the LinkedIn data and are used in our analysis. The summary statistics below refer to this set.

¹⁷This calculation is based on the commonly used classification of major SOC occupation groups 11-29 (Management, Business, Science, and Arts Occupations) and 41-43 (Sales and Office Occupations) as “white collar”, and all other major occupation groups as blue-collar or service occupations. The reported percentages are 2022 employment-weighted averages of 6-digit SOC exposures within each group. For details on this aggregation, see https://www.bls.gov/soc/soc_2010_class_and_coding_structure.pdf.

90th percentiles of occupational exposure are 0% and 53%. The Appendix Table IA.1 lists the 20 occupations with the highest and lowest Generative AI exposure scores.¹⁸ Occupations such as “telemarketers”, “proofreaders and copy markers”, and “computer programmers” have the highest Generative AI exposure scores. This is intuitive as recent technological advances in Generative AI feature natural language-based conversations, translating texts between languages and styles, and generating functioning code based on high-level descriptions of a programming task. On the other hand, occupations requiring more physical manual tasks, such as “shampooers”, “installers and repairs”, and “stonemasons” seem not to be exposed to Generative AI.

— Insert Table II about here —

A large body of literature has explored how prior technologies such as computers, automation, and robots have affected different occupations over the past decades (e.g., Autor et al. (2003), Acemoglu and Autor (2011), Acemoglu and Restrepo (2020), among many others). Computerization has been shown to replace primarily routine-task jobs, which are jobs that perform repetitive and codifiable tasks, such as production line assemblers and record keepers (Autor et al. (2003) and Acemoglu and Autor (2011)). Robots have been shown to replace primarily programmable manual-task jobs such as machinists, material handlers, and welders (Acemoglu and Restrepo (2020)) and Webb (2019)).

To examine whether occupations exposed to Generative AI are distinct from routine and manual jobs, we follow Acemoglu and Autor (2011) and construct scores for the level of non-routine-analytical, non-routine-interpersonal, routine-cognitive, routine-manual, non-routine-interpersonal, and non-routine-manual skills required for each occupation. For any given occupation, the skill-measure scores essentially measure how important each of the six skills is for executing its tasks. Unlike our measure of exposure to Generative AI, which is truly task-based, Acemoglu and Autor (2011) use O*Net’s importance metrics for several pre-defined skills within occupations to construct the scores for these six skill measures for each occupation. See the Appendix IA.3.C.4 for the detailed construction of the scores for these six skill measures.¹⁹

¹⁸Note that the lowest score category only shows an alphabetically sorted subset of a larger set of occupations with zero Generative AI exposure. Overall, 15% of occupations have zero exposure.

¹⁹For each occupation, the O*Net database not only provides the textual statements of each task the occupation performs, which is the data that we used to construct our Generative AI exposure measure, the database also provides a numerical “importance” score for a large number of pre-defined skills for each occupation. Acemoglu and Autor (2011) select certain pre-defined skills and aggregate them to measure an occupation’s skill requirement in each of the six skill dimensions mentioned above. For instance, the measure of non-routine-analytical skill for an occupation is constructed as the average of the standardized importance score of three detailed O*Net skills, *analyzing data/information*, *thinking creatively*, and *interpreting information from others*.

Next, we study whether occupations' exposures to Generative AI are distinct from previously documented occupational exposures to computerization and robots. Specifically, we regress an occupation's Generative AI exposure measure on its six occupational skill scores (Skill_S^O) in the following cross-sectional regression:

$$X^O = \alpha + \sum_S \beta_S * \text{Skill}_S^O + \varepsilon^O. \quad (2)$$

Figure 2 shows the results. We show that occupations with higher Generative AI exposure are more likely to involve non-routine cognitive analytical skills and routine cognitive skills, and less likely to involve other non-cognitive skills. Hence, distinct from computerization which disrupts routine jobs, and robots which disrupt manual jobs, Generative AI tends to primarily affect cognitive jobs, in particular those requiring non-routine analytical skills.²⁰ Hence, one would expect that firms with more cognitive jobs in their occupational portfolio will be more disrupted by Generative AI.

We show in Figure 3 that major occupation groups with higher exposures to Generative AI also tend to have higher wages. This is in contrast to the recent automation wave that impacted low-wage occupations more. Consistent with our finding, important recent work by [Kogan et al. \(2023\)](#) studies the longer-term impact of breakthrough patents on workers' wages, and high-wage workers are impacted more negatively.

— Insert Figure 2 about here —

— Insert Figure 3 about here —

D. Measuring firms' exposures to Generative AI

To measure a firm's exposure to Generative AI, 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.²¹ For each Compustat firm, we use its employment counts at the 6-digit SOC occupation level as of March 2022, the latest month in our Revelio data. A firm's Generative AI exposure is the weighted average of its occupations' Generative AI exposure, X^O . That is, for each firm

²⁰Using a similar regression setting, [Webb \(2019\)](#) shows that broadly-defined AI technologies before the recent advances in Generative AI also tended to disrupt cognitive non-routine analytical jobs, but not routine cognitive jobs.

²¹See, for example, [Li et al. \(2022\)](#) and [Chen et al. \(2023\)](#) for more descriptions of the Revelio Labs data.

f , labor force exposure to Generative AI is computed as:

$$X^f = \sum_{O \in f} EmpShare_{f,O} \times X^O, \quad (3)$$

where $EmpShare_{f,O} = \frac{emp_{f,O}}{emp_f}$ is the employment share of occupation O in firm f . This procedure generates the Generative AI exposures for 2,518 publicly traded firms in 2022. Intuitively, our exposure measure captures the fraction of labor tasks in the firm that can be more efficiently completed using Generative AI.

E. Firm-level stylized facts

Panel B of Table II reports the summary statistics. The mean and median of firms' Generative AI exposures are both 35%, with a standard deviation of 8%. The 10th and 90th percentiles of firm-level exposures to Generative AI are 27% and 44%, respectively, implying that labor exposure to Generative AI is a broad phenomenon.

Table III lists the 15 firms among the top 100 largest U.S. firms by market capitalization with the highest and lowest exposure to Generative AI, respectively.²² Although many IT firms, such as IBM and Intuit, not surprisingly have a large fraction of employees exposed to Generative AI, we also find manufacturing firms, such as 3M, and administrative conglomerates, such as S&P Global, in the high-exposure list. 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 these firms' activities require more manual or interpersonal tasks and have a smaller fraction of employees exposed to Generative AI.

— Insert Table III about here —

We next formally investigate the industry component of firms' Generative AI exposures. Figure 4 shows that the variation in firm-level Generative AI exposure has a substantial within-industry component where the industry is classified at the NAICS 3-digit level. While firms in industries such as "data processing, hosting, and related services" and "professional, scientific, and technical services" have high average exposures to Generative AI and firms in industries such as "clothing and clothing accessories stores" and "food service and drinking places" have low average exposures, there is considerable variation of firms' exposures within each industry. A variance decomposition shows that NAICS 3-digit industry differences explain only 31% of the firm-level variation in exposure to Generative AI. The rich

²²See Appendix Table IA.2 for the full list of the exposures for the 100 largest firms in 2022, and Table IA.3 for the exposures aggregated to the NAICS 3-digit industry level.

within-industry variation in firms’ exposures to Generative AI suggests that our labor-based measure captures firms’ exposures to Generative AI that can be distinct from their products’ exposures to Generative AI. We explore labor versus product exposure effects in-depth in the next section.

In the Appendix Table IA.4, we also show that our measure of firms’ exposures to Generative AI is not explained by other firm characteristics that have been shown to predict stock returns in the cross-section, such as size, Tobin’s Q, ROA, Labor Intensity, and Asset Tangibility. Cross-sectional regressions of firms’ Generative AI exposures on each of these variables yield adjusted R^2 ’s between 0.6% and 10.7%.

— Insert Figure 4 about here —

II. Generative AI Exposures and Firm Value

A. *Empirical design*

We examine the impact of Generative AI on firm value by analyzing the relative returns, or changes in market value, of firms with high versus low exposure to Generative AI immediately following the release of ChatGPT. We sort portfolios into quintiles based on stocks’ labor exposures to Generative AI and compare the returns of the top “Artificial” quintile to those of the bottom “Human” quintile. We show that an AMH portfolio that is long the highest exposure firms and short the lowest exposure firms earns 0.45% daily in the two weeks following the release of ChatGPT (0.44% if we control for market exposures). We argue, and provide supporting evidence, that these changes in value are consistent with changes in the expected future cash flows of firms based on labor cost savings.

By exploring changes in firms’ market value, i.e., stock returns, our empirical design uses a first-difference method to control for time-invariant firm-level heterogeneity. By studying the alphas of the portfolios from asset pricing factor models, our setting controls for firms’ exposures to time-varying priced risk factors. Finally, because we study relative returns, our study highlights the differences in differences of firm value following the release of ChatGPT across firms whose labor forces are more versus less exposed to Generative AI. To rule out alternative channels, we corroborate our portfolio sorting results with standard event study regressions of stocks’ cumulative abnormal returns (CAR) on firms’ Generative AI exposures while controlling for other firm characteristics.

We start by defining our event window. We choose the two weeks from November 30, 2022, to December 14, 2022 (the release date and subsequent ten trading days) as our event window, and we label these two weeks as the “**ChatGPT event period**.” In choosing the

event window, we make a tradeoff. On the one hand, we want to have a window that is long enough for investors to digest the full cross-section of labor implications on firm value. On the other hand, we want the window to be short enough to prevent our estimated stock returns from being contaminated by other related events. The two-week window ensures enough power to ascertain the market’s absorption of the complex effects of ChatGPT on firm values through their labor force exposures, however we provide results for a one-week window in the Appendix Table IA.5.²³

Our test examines a joint hypothesis that the large language model can assess various tasks’ exposure to ChatGPT (see Section I.A) and that the financial market can respond to the assessment. Whether the joint hypothesis holds is an empirical question, however several facts indicate that the hypothesis is plausible for testing. First, the machine-generated task labels, validated in Eloundou et al. (2023) and by ourselves, show high agreement with human labels (see Appendix Section IA.3.C.3). Second, the exposed tasks and occupations as listed in Table I and Appendix Table IA.1 appear intuitive; for example, proofreaders are expected to be highly affected by the release of ChatGPT. Third, the release of ChatGPT has immediately garnered massive public attention. Figure 5 plots the daily mentions of “ChatGPT” or “GPT” on Twitter around the ChatGPT public release date on November 30, 2023. The figure shows that the daily mentions reached over 100K in just a few days after the release of ChatGPT and remained at 50K two weeks after the release. This pattern is consistent with both immediate public attention and continued discussion and information processing over time.²⁴ Finally, in the Appendix, we show that our results are robust to using a shorter (one-week) event window, or a longer (three-week) event window.²⁵

— Insert Figure 5 about here —

B. Realized returns and Generative AI exposures

Portfolio sorts To measure the effect of firms’ workforce exposures to Generative AI on changes in value during the ChatGPT event period, we form five portfolios based on our firm-level Generative AI exposure measure. Specifically, at the beginning of our sample period on November 29, 2022, we sort stocks traded on the NYSE exchange into five quintiles

²³In the Appendix Figure IA.4 we also show that the impact of the release of ChatGPT on IT producers occurred in one day, consistent with the industry effect being simpler than, and distinct to, the labor effects.

²⁴ChatGPT amassed 100 million users in just two months after its release, a milestone which took the World Wide Web 7 years, WhatsApp 3.5 years, Instagram 2.5 years, and Twitter 5 years to reach. See <https://economictimes.indiatimes.com/news/new-updates/chatgpt-witnesses-massive-rise-chatbot-gains-100-million-users-in-two-months/articleshow/98428443.cms?from=mdr>.

²⁵See Appendix Tables IA.5 and IA.6.

based on their Generative AI exposures. We use these NYSE breakpoints to assign non-NYSE stocks into the quintile bins. We compute the value-weighted daily returns of each portfolio as the average daily returns of stocks in the portfolio weighted by their previous day's market capitalization.²⁶ We refer to the high-minus-low portfolio as the “Artificial Minus Human” (*AMH*) portfolio, which represents the zero-net-investment portfolio that goes long high-exposure (*A*) stocks and shorts the low-exposure (*H*) stocks.

Panel A of Table IV shows the realized excess returns, i.e., the raw daily returns minus the daily risk-free rate, of the quintile portfolios sorted by firms' Generative AI exposures and also the long-short *AMH* portfolio during the ChatGPT event period. The *AMH* portfolio yields positive daily returns of 0.45% (*t*-statistic = 3.53) on average during the two weeks after ChatGPT's release.

One concern might be that our results are driven by firms' differential exposures to risk factors. We show that this is not the case. In Panel B of Table IV, we report the alpha of each portfolio after controlling for the market factor (i.e., the CAPM model). The market-adjusted alpha of the *AMH* portfolio shows very similar results with a point estimate of 0.44% per day (*t*-statistic = 4.70), suggesting that our main finding is not driven by firms' heterogeneous exposures to the market. In Panel C, we further control for the Fama-French 5-factor model (Fama and French (2015)) and again find similar results. The 5-factor adjusted alpha of the *AMH* portfolio is 0.35% per day (*t*-statistic = 3.85).

A related issue arises if the differential stock returns between the high and low Generative AI exposure portfolios are driven by different trends in the stock returns of the two portfolios ex-ante. We show that this is also not the case. In the Appendix Tables IA.7 and IA.8, we conduct a placebo test and examine the excess returns and alphas during periods immediately before and after our ChatGPT event period. We do not find significant differences between the high and low Generative AI exposure portfolios. This absence of statistically significant differences in returns outside the ChatGPT event period supports our main findings during the ChatGPT event period being due to the impact of the release of ChatGPT.

— Insert Table IV about here —

Time series of cumulative abnormal returns We further visualize the differential impact of the release of ChatGPT on high and low-Generative AI exposure portfolios by plotting the difference in cumulative abnormal returns (CAR) between portfolios in Figure 1. To do so, we first compute the cumulative abnormal returns of the top and bottom

²⁶See Appendix Section IA.3.C.5 for more details on the portfolio construction.

quintile portfolios in Table IV using November 29, 2022 as the reference date.²⁷ This figure represents the holding period returns of the long-short AMH portfolio at any given date in the two weeks before and three months after the ChatGPT event period. It is reassuring that we do not find a strong trend in the AMH portfolio either immediately before or after the ChatGPT event period, suggesting that the market actively reassessed the impact of Generative AI on firm value via our labor channel during the two weeks after the release of ChatGPT, when the Twitter mentions of ChatGPT were exceptionally high.²⁸ In Appendix Figure IA.1, we extend the pre-period out to 3 months before the release and also observe no systematic pre-trend in the AMH returns.

Prior to the announcement date, the correlation between the A and H portfolio abnormal returns is -0.54, while after the announcement, it is -0.31. The negative correlation prior to the announcement may be due to different ex-ante exposures to technology shocks, as in Kogan and Papanikolaou (2014) and Kogan et al. (2023). The change in the correlation between the two portfolios may also be due to a compositional effect of overall systematic risk as generative AI becomes a more important determinant of market-level cash flows (see Babina et al. (2023)). Our cash flow analysis in Section III using profits and analyst forecasts to show the magnitudes of the changes in realized and expected future cash flows suggests that the main effect on valuations is due to cash flows rather than discount rates, but we cannot rule out that both cash flow and discount rate effects may be present. Indeed, the discount rate channel is an interesting avenue for future research.

The cumulative abnormal returns of the AMH portfolio remain high and do not reverse after the event period, suggesting that the revaluation is not short-lived. In fact, the AMH portfolio seems to have experienced another period of outsized returns in early March leading up to the release of GPT 4, in line with the anticipation of the release and functionality of this more capable model representing an additional increase in the expected productivity impact of Generative AI on exposed firms.²⁹

²⁷Following the literature, we compute each portfolio’s daily abnormal return as the portfolio’s daily return minus the product of each portfolio’s market beta (computed over the 6 months before the time window shown in the figure) and the daily market return, and we compute the CAR for each portfolio by accumulating its daily abnormal returns from November 29, 2022.

²⁸The effects on valuations from the release of ChatGPT among tech firms seem to be mostly realized on the announcement day (Appendix Figure IA.4), while the labor-exposure-based “Artificial-minus-Human” portfolio returns continue to increase throughout the 10 trading day release period (Figure 1), in line with the more complex information about labor effects taking longer to be incorporated into prices than product effects.

²⁹While the time series of returns suggests that the revaluation impact for the AMH portfolio may have been similar in the run-up to the GPT 4 release, we cannot repeat our event study analysis in that case, as the existence of the model was widely known in advance, and only the uncertainty about the exact capabilities and precise release date for the more advanced model was resolved during the run-up.

Event study regressions controlling for firm characteristics We show that our finding of increased value for firms with greater labor exposures to Generative AI during the ChatGPT event period is robust to controlling for firm characteristics commonly used to explain the cross-section of returns. Specifically, we run an event study regression of firms' cumulative abnormal returns in excess of market returns during the ChatGPT event period on their labor-based Generative AI exposure, and also firm characteristics, in the following specification:

$$CAR_i = \beta \times \text{GenAI Exp}_i + \gamma \times \text{Char}_i + \varepsilon_i, \quad (4)$$

where CAR_i is firm i 's cumulative abnormal return during the ChatGPT event period (from November 30, 2022 to December 14, 2022), GenAI Exp_i is the firm's labor-based Generative AI exposure, and Char_i is firm characteristics, including size, Tobin's Q, ROA, Labor Intensity, and Asset Tangibility. We are interested in β , which measures the impact of the release of ChatGPT on the value of firms with a higher labor-based Generative AI exposure relative to other firms while allowing a differential impact on firms with particular characteristics. We utilize weighted least squares, weighting each stock by its market capitalization as of November 29, 2022, to be consistent with the value-weighted portfolio sorts.

Table V shows the results. Column (1) provides the benchmark result without controlling for firm characteristics, which shows that firms with higher Generative AI exposures experienced a greater boost in firm value during the ChatGPT event period, consistent with our portfolio-sorting results. The magnitude of the estimate suggests that a one-standard-deviation increase in the firm-level Generative AI exposure is associated with 1.7 pp higher cumulative abnormal returns over the event period, which is comparable to the effect size estimated using the portfolio sorts.³⁰ Columns (2)-(8) show that the effect of our Generative AI exposure remains highly robust with coefficients similar to the benchmark estimate after controlling for the effects of each firm characteristic separately or the effects of all firm characteristics jointly.³¹

— Insert Table V about here —

³⁰The returns for a one-standard-deviation change are computed as $0.078 \times 21.6 = 1.7$. Note that the market cap weighted exposure scores of the firms in the highest and lowest quintile portfolios are 44.5% and 26.1%, respectively, so the firm-level $CAR[-1,10]$ event study estimate in Table V would predict a daily return difference of $(.445 - .261) \times 21.6/11 = 0.36$ pp for the portfolios, which is not far from the 0.45 pp long-short returns estimated in Panel A of Table IV.

³¹The Appendix Table IA.10 shows that the results are also robust to controlling for NAICS 3-digit industry fixed effects and using alternative models for computing abnormal returns.

C. Product exposure versus labor exposure to Generative AI

A key concern is that labor market effects may be correlated with firms' product market exposures to Generative AI. The stock-market value of chip-maker NVIDIA more than doubled from the beginning of 2023 to September 19, 2023, which has been widely attributed to the company being a key supplier of chips used in data centers and for training Generative AI models. We perform several tests to show that our labor-exposure results are robust to controlling for firms' product-market exposures to Generative AI. To be clear, we do not claim that the release of ChatGPT does not impact firms' products and services. Instead, we provide strong support for the idea that the differential returns we document for firms with different labor force exposures to Generative AI are not driven by the product exposure channel.

C.1. Importance of within-industry heterogeneity

Our first robustness check, which controls for the product exposure channel, examines the within-industry variation in firms' Generative AI exposure in our portfolio sorting test. This test takes the industry classification as a categorization of the product market. Hence, to the degree that product markets align with detailed industry definitions, a significant industry-neutral AMH portfolio return would suggest that our main findings in Section II.B are not driven by firms' product exposures to Generative AI.

We have shown earlier in Figure 4 that firms' exposures to Generative AI vary substantially within industry. Indeed, only 31% of firm-level variation in labor force exposure to Generative AI can be explained by 3-digit industry fixed effects. Here, we sort firms into portfolios based on their Generative AI exposures within industries. We consider two industry classifications. The first one is the NAICS 3-digit industry classification, which includes 46 categories with at least ten firms in our data. [Hoberg and Phillips \(2016\)](#) show that the rich texts of firms' product descriptions in their 10K filings can more precisely separate product markets. We thus also adopt their 10K-based FIC50 industry classification as an alternative measure of industries. Given the more limited number of stocks within industries, we form tercile portfolios (instead of quintile portfolios) within each industry and then value-weight stock returns in the industry-neutral tercile portfolios.

Table VI shows the results. Panels A1-A3 show the market-adjusted alphas for the tercile portfolios sorted by Generative AI exposures across all firms unconditionally, within NAICS 3-digit industries, and within FIC50 industries, respectively. Using Panel A1 as the unconditional benchmark for the tercile portfolios, Panels A2 and A3 show that the within-industry AMH portfolios using both industry definitions also have large and statistically

significant market-adjusted alphas during the ChatGPT event period. The magnitudes of the within-industry alphas are also very similar to the alpha of the unconditionally sorted AMH portfolio in Panel A1. Panels B1-B3 show similar results when we examine the FF5-factor-adjusted alphas instead of the market-adjusted alphas.

In summary, if industry classifications capture the segmentation of firms' products and services, then the within-industry findings suggest that the impact of Generative AI on firm value, as we described in our main findings in Section II.B is not driven by firms' products' differential exposures to Generative AI.

— Insert Table VI about here —

C.2. Results excluding the tech sector

Our results are also robust to excluding the tech sector altogether. We follow [Acemoglu et al. \(2022\)](#) and identify firms as potential tech firms if they are in the following two sectors: information sector (NAICS 51) and professional, scientific, and technical services sector (NAICS 54).³² In Panel D of Table IV, we report our baseline quintile sorting results excluding firms from these two sectors. We document a strikingly similar result to our main results when we exclude the tech sectors. In particular, the FF 5-factor-adjusted alpha for the AMH portfolio using non-tech firms is 0.38% per day during the ChatGPT event period as compared to 0.35% per day using all firms. These results, along with our within-industry results, confirm that the labor impact of Generative AI is broad, and that Generative AI is a general-purpose technology.

C.3. Horseraces between labor versus product exposures

Next, we run horse race tests between our labor exposure to Generative AI and direct measures of product market exposures. By constructing these measures, we aim to directly identify firms that can benefit from an AI boom either because they are selling related products (e.g., cloud computing hardware and graphics processing units) or because they are using existing versions of AI-related technologies as direct inputs into their products and services. Our product exposure measures may be of independent interest since we show they are related to firm returns; however, importantly, our labor exposure effects remain even when we control for these product market effects.

Our first measure of firms' product exposure to Generative AI follows [Hoberg and Phillips \(2016\)](#) and many other prior studies and infers information about a firm's products from the

³²The 2023 U.S. Census Business Trends and Outlook Survey shows that these two sectors have the highest use of AI in their products. See <https://www.census.gov/library/stories/2023/11/businesses-use-ai.html>.

business description section of a firm’s most recent 10-K annual report. We then use GPT 3.5 Turbo to assess whether firms’ business products involve enabling or scaling new Generative AI technologies or benefit from a direct incorporation of the new Generative AI capabilities as a functionality of their products.³³ We provide details in the Internet Appendices [IA.3.C.6](#) and [IA.3.C.7](#). This procedure generates a dummy variable that equals 1 if the answer is yes—the firm’s annual report suggests product market exposure to Generative AI, and 0 otherwise. We label this measure *GPT10K Product Exposure*.

Our second measure counts the number of occurrences of AI-related keywords in the firm’s business section in its 10-K report. This keyword-count approach is consistent with the approach used in prior studies (e.g., [Webb \(2019\)](#) and [Babina et al. \(2024\)](#)), and it generates a continuous measure instead of the binary value obtained in our first product exposure measure. We label this AI keyword-based measure *Count10K Product Exposure*.

Our third measure adopts the list of firms recently classified as “near-term beneficiaries of AI” by the investment bank Goldman Sachs.³⁴ This measure identifies stocks with business models that are directly exposed to the development of AI and which, as a result, are expected to experience an immediate increase in earnings. The list includes makers of semiconductors and related equipment, and large technology companies with extensive cloud computing infrastructure or business models that are likely to benefit from incorporating AI capabilities into their products. This measure is explicitly focused on *ex post* identifying stocks with high recent returns due to their earnings potential increasing as a result of the AI boom, and it may thus account for firms’ product exposure to AI beyond firms’ discussion in past annual reports. We label this measure from Goldman Sachs *GS Product Exposure*.

Finally, our fourth measure uses the share of workers at a firm who have AI skills on their resumes constructed by [Babina et al. \(2024\)](#), who show that this measure effectively predicts AI-related product innovations and R&D spending during the pre-ChatGPT wave of AI advancements. We use a firm’s most recently available AI skill share in the data from [Babina et al. \(2024\)](#) to proxy for its prior investment in using AI-related tools pre-ChatGPT and its susceptibility to adopt new AI tools in their products post-ChatGPT.³⁵ We label this

³³Conceptually, this product-exposure channel differs from our labor-exposure channel as the former channel does not directly imply an impact of Generative AI on the firm’s operation and production process.

³⁴See Goldman Sachs US Equity Views, August 21, 2023, “*The (AI) trade after the trade: Identifying potential long-term EPS beneficiaries of AI adoption*” available from <https://research.gs.com/>. For a publicly accessible write-up of the results of the study, see, for instance, <https://markets.businessinsider.com/news/stocks/goldman-sachs-reveals-long-term-ai-portfolio-here-are-the-50-stocks-to-monitor-1032573922>.

³⁵We thank [Babina et al. \(2024\)](#) for making the data publicly available. The last year of available skill share data is 2021 for most firms in the sample. While their measure also derives from firms’ labor heterogeneity, their approach uses keywords to account for firms’ employees’ AI skills, and their measure aims to capture firms’ AI investment. In contrast, our labor-based Generative AI exposure accounts for firms’ employees

measure from Babina et al. (2024) *BFHH Product Exposure*.

How well do these four proxies for firms' product exposure to Generative AI capture the potential for firms to adopt recent AI in their products and services after the release of ChatGPT? To answer this question, we use data from the U.S. Census Business Trends and Outlook Survey, which reports data on the share of firms using "Artificial Intelligence (AI) in producing goods and services" for a number of industries after August 2023.³⁶ We thus aggregate our four proxies of product exposure to Generative AI to the industry level and compare our proxies with the intensity of industries' actual use of AI in their products after the release of ChatGPT. Appendix Table IA.9 shows that all four proxies correlate strongly with the Census survey results at the industry level. Hence, we argue that these four proxies capture firms' product exposure to Generative AI reasonably well. Moreover, they capture different dimensions of product exposure, as evidenced by the fact that they are only modestly correlated with one another - see Appendix Table IA.11.

We investigate whether our finding of increased value for companies with greater labor-based Generative AI exposure is confounded by our measure's relation with firms' product exposure to Generative AI. Specifically, we run an event study regression of firms' cumulative abnormal returns during the ChatGPT event period on their labor-based Generative AI exposure and also various proxies for firms' product exposure to Generative AI:

$$CAR_i = \beta \times \text{GenAI Exp}_i + \gamma \times \text{Product Exp}_i + FEs + \varepsilon_i, \quad (5)$$

where CAR_i is firm i 's cumulative abnormal return during the ChatGPT event period, GenAI Exp_i is the firm's labor-based Generative AI exposure, and Product Exp_i is the proxy for the firm's product exposure to Generative AI. We are interested in β , which measures the impact of the release of ChatGPT on the value of firms with a higher labor-based Generative AI exposure relative to other firms, while allowing for a differential impact on firms with a particular product exposure to Generative AI. Similar to our estimation in equation (4), we weigh each stock by its market capitalization as of November 29, 2022. Finally, we include NAICS 3-digit industry fixed effects so that the horseraces are run within industries, further controlling for any remaining product market effects. Moreover, all product exposure measures are standardized, except *GS Product Exposure* and *GPT10K Product Exposure*, which are binary variables.

Table VII reports the results of the cross-sectional regression in equation (5). Column (1) restates the within-industry importance of our effects documented in Section II.C.1, showing

whose tasks can be more efficiently completed by using Generative AI, and thus the firms' *labor exposure* to Generative AI.

³⁶See details about the survey and data at <https://www.census.gov/hfp/btos/about>.

that the coefficients of our labor-based Generative AI exposure are similar when controlling for NAICS 3-digit industry fixed effects. Columns (2)-(6) show the horse race results by controlling for firms' product exposure to Generative AI and industry fixed effects. Two observations stand out: First, three out of four proxies for firms' product exposure positively boost firm value during the ChatGPT event period. This result supports findings in the prior literature that AI technologies improve firms' product functionality and innovations ([Babina et al. \(2024\)](#)).

Second, Table [VII](#) shows that firms' labor exposure to Generative AI is robustly related to increases in firm value during the ChatGPT event period, even after controlling for the proxies for firms' product exposure to Generative AI individually in Columns (3)-(5) or jointly in Column (6). The economic magnitude of the CAR associated with the labor exposure to Generative AI is only modestly attenuated, suggesting a 1.4 pp higher event period CAR for firms with a 1 SD higher Generative AI exposure, even when controlling for all product market exposure proxies jointly.

— Insert Table [VII](#) about here —

D. The role of firms' data assets

A salient feature of Generative AI is its superior capability to learn from and work with firms' data to upgrade their operations and save labor costs, such as by training customer service chatbots, automating workflows, improving predictions and analytics, and performing many other applications ([Caserta et al. \(2023\)](#)). For instance, to adopt Generative AI tools and improve the chatbots for internal customer service knowledge management, firms need internal data from past interactions and the technology to upgrade existing labor-intensive customer service systems.

In this section, we test whether our stock return results are more pronounced among firms relying more heavily on data assets, which can better incorporate Generative AI technologies. Measuring firms' data assets is challenging ([Veldkamp \(2023\)](#)).³⁷ We construct two proxies for firms' reliance on data assets that capture different aspects of firms' ability to use Generative AI capabilities. Our first measure uses the text of the business description in a firm's 10-K annual report. We again apply the GPT model to assign a score of 0 to 3 to indicate whether the firms' business descriptions provide no, little, moderate, or high evidence of data that would be valuable for Generative AI-based analysis. The model is instructed to base its assessment on mentions of six data-related topics such as data collection,

³⁷See seminal works on the measurement of firms' data assets from [Begenau et al. \(2018\)](#), [Farboodi et al. \(2019\)](#), [Eeckhout and Veldkamp \(2022\)](#), [Farboodi and Veldkamp \(2023\)](#), [Abis and Veldkamp \(2023\)](#), among others.

data utilization, data infrastructure & management, and data regulation and privacy (see more details in the Internet Appendices [IA.3.C.8](#) and [IA.3.C.9](#)). The overall score assigned by the model is our first measure of the firm’s reliance on data assets. We label this proxy *10K Data Assets*. We think of this measure as capturing the degree to which the nature of a firm’s business practices and operational choices expose it to large volumes of customer data and require data management capabilities.

Our second measure follows [Abis and Veldkamp \(2023\)](#) and calculates firms’ share of labor skills in data management. In particular, we first measure the likelihood that an occupation is a “data management” position by analyzing whether the requirements listed in each occupation’s job postings in recent years correspond to data management skills classified by [Abis and Veldkamp \(2023\)](#). Job postings are classified as suggesting a data management-intensive role if at least three of these skills are mentioned. For each occupation, we compute the share of data management-intensive job postings as a proxy for the propensity of that occupation to be a data management position.³⁸ Then, we aggregate the data management intensity from the occupation level to the firm level using the firm’s occupational employment shares from the Revelio Labs database, similar to our calculation of firms’ Generative AI exposure (see more details in the Appendix [IA.3.C.8](#)). The intuition for this measure is that whether a firm has many workers responsible for managing data is a good proxy for whether it has a lot of valuable data to be managed—and for the firm’s ability to handle data-based technologies in general. We label this proxy *AV Data Assets*.³⁹ We think of this measure as capturing the degree to which a firm’s labor force is likely to have the skills to handle the data management needs of Generative AI models deployed in the firm at scale. That is, where the *10K Data Assets* capture the firm’s access to data, *AV Data Assets* capture its workforce’s data handling capabilities.

Equipped with the two proxies for firms’ reliance on data assets, we test whether our stock return results are more pronounced among firms with a greater reliance on their data assets. In particular, we regress firms’ cumulative abnormal returns on the interaction between their Generative AI exposure and the proxies, one at a time. Table [VIII](#) shows the results. The positive coefficients of the interaction terms suggest that firms with greater access to data assets have significantly higher returns to Generative AI exposure during the ChatGPT event period. This finding supports our conjecture that firm value disproportionately accrues to firms with labor exposure to Generative AI *and* significant data assets, allowing them to better adopt Generative AI technologies and upgrade their labor operations. An alternative

³⁸Occupations’ data management intensity is 9% correlated with their Generative AI exposure.

³⁹See Appendix Table [IA.11](#) for the raw correlation of the data asset measures with the firm-level Generative AI exposure, and the product market exposure measures.

explanation for these data asset effects could be that larger firms are better able to pay the fixed costs of adopting Generative AI technologies and also tend to have greater data assets. In Appendix Table [IA.12](#), we show that this does not seem to be the case: first, column (1) shows that the exposure effects do not vary significantly across firm sizes; second, in a horserace that allows the exposure effects to vary with firm employment size and data assets (see columns (2) and (3)), there is no significant variation of the effects with firm size, and the data asset interaction coefficients are basically unchanged relative to Table [VIII](#). On the other hand, if firms with greater data assets face lower fixed costs when adopting new technologies, then fixed costs may still partially account for the observed heterogeneity.

We also note that data assets can be firm-specific, making existing incumbent firms' potential gains from adopting Generative AI larger, and less likely to be competed away by new entrants. Our finding that the stock returns to Generative AI exposure during the ChatGPT event period are greater among firms with greater reliance on data assets is consistent with this view.

— Insert Table [VIII](#) about here —

III. Inspecting the Mechanism: Future Cash Flows

The divergence in values between firms with high and low workforce exposure to Generative AI upon the release of ChatGPT reflects investors' expectations of a differential effect on firms' discounted future cash flows. Prior studies suggest that new technologies like AI can potentially help firms organize more efficiently and generate more future cash flows (e.g., [Babina et al. \(2022\)](#), [Babina et al. \(2024\)](#), [Hampole et al. \(2025\)](#), and [Mishra et al. \(2022\)](#)). In this section, we inspect this future cash flow channel for our firm value results. In particular, we provide evidence that both analyst forecasts of firms' earnings and actual firm profitability increase more after the release of ChatGPT for firms with greater Generative AI exposures.

The release of ChatGPT may also differentially affect firms' risk premia and thus their discount rates, which could potentially contribute to our firm value results.⁴⁰ Assessing potential changes in firms' expected returns during the shock period directly is challenging due to the short time series. However, our evidence from cash flow forecasts suggests that

⁴⁰A large body of literature shows that new technologies can cause significant employment and income risk for employees in exposed jobs through job displacement (e.g., [Autor et al. \(2003\)](#), [Acemoglu et al. \(2022\)](#), [Kogan et al. \(2021\)](#), [Kogan et al. \(2023\)](#), [Gârleanu et al. \(2012\)](#), and many others) or empowering the less talented workers to challenge more talented peers (e.g., [Dell'Acqua et al. \(2023\)](#) and [Brynjolfsson et al. \(2025b\)](#)). Such risk can incentivize employees to hedge their income risks by purchasing the stocks of firms that benefit from new technologies, which reduces those firms' discount rates.

a discount rate channel is unlikely to be the main driver of our firm value results.⁴¹ The magnitude of ChatGPT’s cash flow impact we estimate is nearly identical to the impact on firm value across firms with high and low Generative AI exposure, consistent with the cash flow channel having dominant explanatory power. Additionally, in theory, discount rates for high-exposure firms could increase (due to an increase in systematic risk) or decrease (if high-exposure firms are a hedge for investors exposed through their labor income). In recent work, [Babina et al. \(2023\)](#) find that firms with higher AI investing opportunities see a greater increase in their systematic risk. An increase in systematic risk for exposed firms would be expected to lead to a corresponding decrease in their value due to increased discount rates, along with higher average returns in the longer run. Thus, these recent findings also cut against a discount rate channel for our findings regarding the increase in value of highly exposed firms upon the release of ChatGPT.

A. Firms’ cash flow forecasts and Generative AI exposure

We obtain monthly forecasts of firms’ future earnings from the Summary Statistics of the I/B/E/S database. Following the standard approach in the literature (e.g., [Bordalo et al. \(2024\)](#)), we measure the consensus analyst forecasts for each firm’s future earnings based on the median forecasts in the month. We examine forecasts for firms’ next year earnings per share (EPS) (i.e., the EPS for firms’ fiscal year ending in December 2023) as well as the long-term forecast of the growth rate in EPS (LTG).⁴²

With this firm-month level data, we examine changes in forecasts for firms’ future earnings after the release of ChatGPT on November 30th, 2022. We regard the consensus forecasts reported in October 2022 as the pre-period forecasts. We use the consensus forecasts reported in January 2023 as the post-period forecasts to allow all updates to be fully reflected in the data.⁴³ Next, we estimate a difference-in-differences model that regresses the forecasts for either short-term EPS or long-term earnings growth on an interaction between a post-

⁴¹The literature typically proxies for firms’ expected returns using average realized returns. Such an approach requires a long time series of data to ensure that the realized return, on average, reflects investors’ expected returns. An important study is [Babina et al. \(2023\)](#), who examines changes in the systematic risk of firms with high and low investment potential in AI from 2010-2018. In contrast, our study focuses on firms’ realized returns shortly after the release of ChatGPT. [Pástor et al. \(2022\)](#) warn that realized returns do not reflect expected returns during the shock periods in the context of explaining green stocks’ higher returns over brown stocks.

⁴²We choose only firms where the fiscal year ends in December to ensure the forecast horizons are unaffected by different fiscal year ending months across firms. Over 80% of firms in our sample have fiscal years that end in December.

⁴³We inspect the forecasts only one month after ChatGPT’s release (i.e., January 2023) to ensure the changes in earnings forecasts are not affected by news about ChatGPT’s future updates. The results are robust if we choose a wider pre- and post-period window though.

ChatGPT dummy and a continuous measure of Generative AI exposure at the firm level. We control for firm fixed effects to ensure our estimate reflects the within-firm *update* on forecasts and month fixed effects to control for aggregate economic conditions. We cluster all standard errors by firm.

The results are shown in Table IX: Column (1) shows that the consensus forecasts for firms' short-term future EPS update more positively after the release of ChatGPT for firms with higher Generative AI exposure. Firms with a one-standard-deviation higher Generative AI exposure experience a \$0.076 ($= 0.066 \times 1.14$) higher update in the forecast of their 2023 year-end cash flows after the release of ChatGPT. Relative to the sample mean of the respective forecasts prior to the release of ChatGPT, this effect corresponds to a 1.8% higher increase in the first year-end EPS forecasts. This magnitude is nearly identical to our firm value results that firms with a one-standard-deviation higher Generative AI exposure experienced a 1.7% greater stock returns, as shown in Section II.B.

We can interpret this finding in a simple discounted future cash flow framework: if we assume that the release of ChatGPT does not affect the low-GenAI exposure firms, and that it permanently raises each future year's cash flows for the high exposure firms by the same percentage, then the above estimate suggests that the cash flow channel fully explains our firm value results. On the other hand, if the large impact on the first year's cash flow is transitory, there can still be room for the discount rate channel. While 65% of our final sample does not have the long-term earnings growth forecast data, we inspect the impact on LTG across firms with high and low Generative AI exposure using the remaining 35% of the sample. Column (2) of Table IX shows that the forecast update on long-term earnings growth is also more positive for firms with higher Generative AI exposure. Hence, it is unlikely that analysts expect the impact of ChatGPT on firms' future cash flows to be transitory.

In summary, these findings suggest that the heterogeneous impact of the release of ChatGPT for firms with high and low Generative AI exposures is consistently observed in both firms' values and their forecast future cash flow.

B. Firms' realized cash flows and Generative AI exposure

We next provide additional support by examining the impact of the release of ChatGPT on realized cash flows for firms with high and low Generative AI exposure. We compute firms' quarterly profitability using Compustat quarterly data, where profitability is the firm's revenue minus the cost of goods sold normalized by its total assets (Novy-Marx (2013)).⁴⁴

⁴⁴Novy-Marx (2013) highlights that this gross profitability is "the cleanest accounting measure of true economic profitability. The farther down the income statement one goes, the more polluted profitability measures become."

Our sample includes 22,403 firm-quarter observations of firms in our stock return analyses from 2022Q1 to 2024Q3.

Consistent with our findings on cash flow forecasts and firm values, Column (3) in Table [IX](#) shows that the firms with higher Generative AI exposure show higher gross profitability after the release of ChatGPT. A one-standard-deviation increase in Generative AI exposure is associated with a 0.13% ($= 0.078 \times 1.71\%$) increase in the gross profitability rate. While many factors can lead expected and realized profitability to differ, our consistent findings on realized profitability, forecast earnings, and stock market valuation support the cash flow channel for our stock return results.

Finally, the Appendix Table [IA.13](#) further shows that these profitability results are unlikely to be driven by a revenue channel, as we do not observe significant responses of firms' forecast and actual revenues to the release of ChatGPT associated with their Generative AI exposure. This finding supports a cost-saving mechanism, which we examine in the next section.

— Insert Table [IX](#) about here —

IV. Generative AI: Labor Complement or Substitute?

A firm's workforce exposure to Generative AI could mean that their workers can be substituted by the technology or complemented by it.⁴⁵ Both substitution and complementarity channels can potentially improve future cash flows and market values for firms with higher Generative AI exposure. We find that, at the firm-level, the labor substitution channel is dominant and tends to be most consistent with our findings on cash flows and stock returns. Dissecting the labor complementarity versus substitution channels shortly after the arrival of the new technology is challenging. We develop a conceptual framework offering a novel approach for interpreting Generative AI exposure as labor-complementing or substituting. We combine recent job posting data and individual wage data to validate our approach, and show, using our refined measures of core versus supplementary exposures, that core task exposure leads to lower wages and fewer jobs posted while supplementary task exposures is associated with higher wages and better employment outcomes.

⁴⁵A large body of literature has studied the complementarity and substitution effects of technologies on workers in the context of automation ([Autor et al. \(2003\)](#), [Autor et al. \(2006\)](#), [Frey and Osborne \(2017\)](#), and [Acemoglu and Autor \(2011\)](#)), robots ([Acemoglu and Restrepo \(2020\)](#)), artificial intelligence ([Webb \(2019\)](#), [Agrawal et al. \(2019\)](#), and [Babina et al. \(2022\)](#)), and disruptive technologies in general ([Krusell et al. \(2000\)](#), [Kogan et al. \(2019\)](#) and [Bloom et al. \(2021\)](#)).

A. Challenges for timely assessing new technologies' impact on workers

Detecting whether a new technology will substitute for or complement labor shortly after the arrival of new technologies is challenging for at least two reasons. First, a large body of literature shows that the diffusion of major technologies in history is usually remarkably slow due to adoption frictions, such as limited attention of businesses to the new technologies (Greenwood (1999)).⁴⁶ This lack of attention problem appears less relevant for the diffusion of ChatGPT. Previously, we showed that the release of ChatGPT garnered massive immediate public attention on social media. Here, we further document that the release of ChatGPT also attracted the attention of *firms*. In Figure 6, we measure firms' mentions of Generative AI in their quarterly earnings calls. Panel A shows that the share of firms mentioning Generative AI rose substantially from less than 5% before ChatGPT's release to 27% in the first quarter after the release.⁴⁷ More importantly, Panel B shows that firms with a *higher* Generative AI exposure see a *stronger* increase in the likelihood of mentioning Generative AI after ChatGPT's release. These findings suggest that the release of ChatGPT likely garnered the attention of firms to evaluate how Generative AI can be applied to affect their labor, operations, and profits.⁴⁸

Second, the fact that a worker's *tasks* are substituted by Generative AI does not necessarily mean the worker's *occupation* can or will be substituted by Generative AI. Even if many tasks in a given occupation are highly exposed to being executed by Generative AI, an occupation overall may be less exposed to these substitution effects because the exposed tasks are not core tasks, but are rather supplementary tasks through which Generative AI can increase a worker's productivity. For instance, while proofreaders and product managers both have some tasks that can be completed more efficiently by ChatGPT, ChatGPT may substantially reduce the demand for proofreaders but may make product managers who use it to complete their supplementary tasks more productive. The distinction between task exposure and occupation exposure represents an under-appreciated challenge in the study of the effects of automation. Our utilization of the core versus supplementary task distinction offers a novel way to uncover differences between occupational exposures and the raw share of exposed tasks within occupations that has been studied previously (eg. Eloundou et al. (2023)). Prior studies on automation also compute an occupation's susceptibility

⁴⁶The diffusion of new technologies is particularly slow in the initial periods after the technologies are available, constituting an S-curve for technology diffusion (see many examples in Greenwood (1999) and Manuelli and Seshadri (2014)). As a result, many seminal studies on the impact of major technology shocks were highly retrospective rather than conducted soon after the arrival of the technologies.

⁴⁷In contrast, the share of firms mentioning other related topics, such as engineering, does not increase. See the Appendix Figure IA.3 for this result.

⁴⁸In fact, Schubert (2025) shows that firm exposure strongly predicts firm-level hiring for generative AI skills after the release of ChatGPT.

to computer substitution based on an occupation’s routine task share (e.g., [Autor et al. \(2003\)](#), [Acemoglu and Autor \(2011\)](#), and [Zhang \(2019\)](#)). In what follows, we propose that an occupation’s substitutability by Generative AI depends not only on the share of its tasks that are substitutable by Generative AI, but also on whether the exposed tasks are *core* or *supplemental* to the occupation.

— Insert Figure 6 about here —

B. Conceptual framework: Task-exposure vs. job-exposure to Generative AI

Conceptualizing core versus supplemental tasks An occupation is a bundle of tasks. As such, some tasks being automated are not necessarily equivalent to the full job being automated. Our key point is that whether automation substitutes for or complements an occupation depends on how essential the automatable tasks are to the occupation’s expected functionality. In particular, *core tasks* represent the most fundamental duties an employee in the occupation is expected to perform. For instance, “correct or record omissions, errors, or inconsistencies found” is a core task for proofreaders. If firms can adopt ChatGPT to complete this task at a lower cost, proofreaders may face a potential decline in demand and future wages as they lose their competitive advantage in performing their most expected fundamental duties to ChatGPT.⁴⁹

In contrast, *supplemental tasks* are additional tasks or duties associated with the core tasks but are not considered critical or central to the occupation. For instance, “prepare operating and maintenance manuals, studies, or reports” is a supplemental task for architects. If firms can adopt ChatGPT to complete this type of task more efficiently, architects may have a greater bandwidth to work on their core tasks such as “represent clients in obtaining bids or awarding construction contracts.” As long as the demand for architects’ services is not fully satiated, architects may produce more output per unit of effort with ChatGPT helping with the supplemental tasks.⁵⁰ In this case, architects may see increased demand in the labor market and higher future wages.

This distinction between core and supplemental tasks’ exposure to Generative AI suggests that the impact of Generative AI on jobs has two layers—it is “*not just how many tasks but which tasks*” as it was put by David Autor in his 2024 [Joseph Schumpeter lecture](#).

⁴⁹From the firm’s perspective, this labor-technology substitution could happen either through a reduction in the number of workers in that occupation or by “unbundling.” [Agrawal et al. \(2023\)](#) note that if AI replaces some tasks that were previously bundled into jobs requiring scarce skilled workers, firms may unbundle the rest of the tasks into new jobs for workers without those specialized skills, enabling the firm to employ fewer skilled workers.

⁵⁰If the overall demand for architects’ services is fully satiated, then ChatGPT will result in an oversupply of architects’ services and may still reduce the demand for the occupation.

Specifically, while an occupation’s Generative AI exposure in Equation (1) characterizes the intensity of the technology’s impact on the occupation, the share of the exposure deriving from the occupation’s core versus supplemental tasks can help to capture the substitution versus complementarity of the technology’s impact on the job.

PREDICTION: *Occupations with a higher Generative AI exposure and also a higher share of the exposure deriving from core (supplemental) tasks experience lower (higher) labor demand and wages after the release of ChatGPT.*

Implications for firm valuation In principle, both the labor-complementing and labor-substituting mechanisms could contribute to the greater increase in market value for firms with a higher Generative AI exposure after the release of ChatGPT. If the firm’s Generative AI exposure is primarily due to its high fraction of occupations with core tasks substitutable by Generative AI, the firm can potentially save labor costs by substituting the new technology. The cost-saving can potentially increase their future profits and thus positively impact their current firm value. If the firm’s Generative AI exposure is primarily due to its high fraction of occupations with supplemental tasks substitutable by Generative AI, the firm can expect a boost from the productivity of its labor, allowing the firm to expand and potentially increase its market value. Which channel drives the post-GPT labor market and our firm value results is an empirical question. We thus rely on data to shed light on the mechanism.

Before we head to further empirical analyses, we emphasize an important note that the above arguments rely on the assumption that increases in future cash flows cannot be fully appropriated by workers or competed away in the product market. Literature suggests that labor markets and product markets may not be perfectly competitive. For instance, many prior studies document that firms possess certain monopsony power in the labor market (e.g., Schubert et al. (2020), Berger et al. (2022), Yeh et al. (2022), and Seegmiller (2021)) and that measures of product market competition in the U.S. have been low and declining (e.g., Gutiérrez and Philippon (2017a) and Akcigit and Ates (2021)).

C. Measuring core and supplemental task exposure to Generative AI

Task definition We obtain the classification of each task as core or supplemental for each occupation directly from the O*Net database. According to the [O*Net description](#), core tasks are *critical to the occupation*, while supplemental tasks are *less relevant and/or important to the occupation*.⁵¹

⁵¹O*Net defines a task as *core* for an occupation if the task has a relevance score above 67% and an importance rating above 3.0 for the occupation and as *supplemental* if otherwise. See details at <https://www.onetonline.org/help/online/scales#score>. For an average occupation in the O*Net V27.2

Occupation-level measure To test the prediction from our conceptual framework above, we distinguish whether an occupation’s Generative AI exposure derives from its core or supplemental tasks. Specifically, we measure the share of an occupation’s Generative AI exposure deriving from its supplemental tasks as

$$ShareSupp_o = \frac{\sum_{T \in (o|Supplemental)} X^T}{\sum_{T \in o} X^T}, \quad (6)$$

where the numerator is the sum of task exposures to Generative AI for occupation O ’s supplemental tasks, and the denominator is the sum of task exposures to Generative AI from all of the occupation’s tasks. Correspondingly, the share of exposure from core tasks is $ShareCore_o = 1 - ShareSupp_o$. For occupations with a positive Generative AI exposure, on average, 23% of the exposure derives from supplemental tasks and 77% comes from core tasks (see Panel A of Table II). Many occupations have exposure to Generative AI through both core and supplemental tasks (see Appendix Figure IA.6). We thus focus on variations in occupations’ continuous $ShareSupp_o$ instead of creating binary labels for occupations. Intuitively, occupations in which core tasks involve textual or quantitative analysis (or coding) are more likely to see stronger substitution effects. Occupations in which core tasks involve manual or interpersonal activities, but in which supplemental tasks involve textual or quantitative analysis, are more amenable to complementarity.

Firm-level measure Finally, we measure the share of a *firm*’s Generative AI exposure deriving from its workers’ supplemental tasks as

$$ShareSupp_f = \frac{\sum_{o \in f} EmpShare_{f,o} \times X^o \times ShareSupp_o}{\sum_{o \in f} EmpShare_{f,o} \times X^o}, \quad (7)$$

where X^O is occupation O ’s Generative AI exposure defined in equation (1), and $ShareSupp_o$ is the occupation’s share of Generative AI exposure derived from its supplemental tasks’ exposure to Generative AI defined in equation (6). The firm’s share of Generative AI exposure from core tasks is $ShareCore_f = 1 - ShareSupp_f$.

database, 76% of tasks are core and 24% of tasks are supplemental.

D. Validation: Evidence on occupational substitution vs. complementarity

We test our prediction about the impact of Generative AI on occupations in our conceptual framework by running several variants of the following regression specification:

$$y_{j,o,t} = \theta \cdot \text{Post}_t \times \text{GenAI Exp}_o + \eta \cdot \text{Post}_t \times \text{ShareSupp}_o + \lambda \cdot \text{Post}_t \times \text{GenAI Exp}_o \times \text{ShareSupp}_o + \text{FEs} + \varepsilon_{j,o,t}. \quad (8)$$

Here, $y_{j,o,t}$ is the labor demand or wage variables to be detailed later, Post_t is an indicator of the time period after the release of ChatGPT in November 2022, GenAI Exp_o is the occupation's Generative AI exposure defined in equation (1), and ShareSupp_o is the share of exposure to Generative AI that comes from supplemental tasks in equation (6), and FEs include occupation and time fixed effects which capture other terms that vary only by occupation or time.

Evidence on occupational labor demand To test Generative AI's heterogeneous impact on occupational labor demand, we use data from Lightcast (formerly Burning Glass Technologies) and construct firms' number of job postings for each SOC 6-digit occupation in each month from January 2022 to August 2023. We aggregate the number of job postings to the occupation-month level, resulting in a final sample with a balanced panel of 43,512 observations for occupations that ever posted jobs in our time period, where occupation-months without any job postings are coded as zeros. We run three variants of the specification (8) above to test our prediction.

First, we inspect the overall effect of ChatGPT on labor demand for occupations with high and low Generative AI exposure by estimating the regression without the ShareSupp_o term, i.e.,

$$y_{j,o,t} = \chi \cdot \text{Post}_t \times \text{GenAI Exp}_o + \text{FEs} + \varepsilon_{j,o,t}. \quad (9)$$

This specification does not distinguish occupations' Generative AI exposure deriving from their core versus supplemental tasks. Instead, it provides a benchmark estimate for us to highlight the labor-substitution and labor-complementarity effects later. To present the estimates as a percentage change in job postings but without dropping observations when log-transforming the dependent variable, we follow [Cohn et al. \(2022\)](#) and estimate the above specification using a Poisson model with fixed effects.⁵²

⁵²A commonly adopted approach is to take the natural logarithm of the variable before estimating a linear model. However, the natural logarithm of zero is undefined. Prior studies therefore often transform the dependent variable into the logarithm of a constant plus the original variable. However, [Cohn et al. \(2022\)](#)

Column (1) of Table X reports the results. Occupations with a one-standard-deviation higher Generative AI exposure experience about an 8% ($= 0.209 \times 0.387$) decline in job postings after the release of ChatGPT, relative to less exposed occupations. This result suggests the substitution effect tends to dominate the impact on occupations with high Generative AI exposure.

One potential concern is that more exposed occupations may happen to be more prevalent in firms that are not performing well during this time period (e.g., some tech firms experienced downsizing in 2023). In that case, our occupation-level finding might be unrelated to firms choosing to shift hiring away from more exposed occupations, but rather might be due to firms that typically hired more exposed occupations reducing their overall job postings. To address this concern, we re-aggregate the detailed job posting data to the occupation-firm-month level for the publicly traded firms for which Lightcast provides a crosswalk of job postings to Compustat firm identifiers. We then estimate the heterogeneous labor demand across occupations within a firm by running the regression in equation (9) while controlling for firm-month fixed effects. Column (4) shows similar results, suggesting that firms, on average, shift their labor demand away from more exposed occupations after the release of ChatGPT.

Second, we run the specification in equation (8) which includes $ShareSupp_o$. We focus on the coefficient θ , which measures the labor-substitution mechanism in our prediction. In particular, θ estimates the impact of ChatGPT on the spectrum of occupations with Generative AI derived solely from core tasks, i.e., $ShareSupp_o = 0$. Supporting our prediction, we observe in Column (2) that the point estimate of θ is 74% ($= 0.672/0.387 - 1$) greater than the estimate of χ in the benchmark case in Column (1). Hence, if we restrict the exposure to be solely from core tasks, an increase in occupations' Generative AI exposure corresponds to a much stronger substitution effect from ChatGPT. Similarly, Column (5) shows that this inference holds within a firm-month. In addition, in both columns, the point estimate of λ for the triple interaction is positive and statistically significant, confirming that the labor-substitution effect is mitigated as occupations' Generative AI exposure derives less from core tasks and more from supplemental tasks.

Third, we run the specification in equation (8) but substitute $ShareSupp_o$ with $ShareCore_o$. In this case, the coefficient θ measures the labor-complementarity mechanism in our prediction. That is, θ estimates the impact of ChatGPT on the spectrum of occupations with Generative AI exposure derived solely from supplemental tasks, i.e., $ShareCore_o = 0$. Sup-

show that estimates using this approach do not generate the correct interpretation and can even produce the wrong signs, while the Poisson model with fixed effects generates the correct interpretation of the estimated coefficient as representing percentage changes in the dependent variable. In particular, we used the PPML Stata model developed by [Santos Silva and Tenreyro \(2015\)](#).

porting our prediction, we observe in Columns (3) and (6) that the point estimate of θ is positive and statistically significant. Hence, if we restrict the exposure to be solely from supplemental tasks, an increase in occupations' Generative AI exposure corresponds to a positive impact of ChatGPT on the occupations' demand, consistent with the complementarity mechanism.

Evidence on the occupational wage rate. We corroborate our labor demand findings by further analyzing the impact of Generative AI on workers' wage rates. We use the monthly individual-level data from the Census Current Population Survey (CPS) from January 2022 to October 2023.⁵³ We compute the more aggregated Census occupations' Generative AI exposure by applying a crosswalk and then averaging over the exposure of related SOC 6-digit occupations. Our final sample includes the natural logarithm of hourly wage rates of about 271,000 individual-by-month observations, along with their occupation and other demographic characteristics, including gender, age, age squared, work experience, work experience squared, race (white versus non-white), and years of education.

To test our prediction on occupational wages in the conceptual framework, we run three variants similar to specification (8), where the dependent variable is the hourly wage rate of the individual. We further control for the above individual demographics, and we weight the observations using each individual's CPS sampling weight following the literature. We still run Poisson regressions for the above specifications, and hence the coefficients should be interpreted as percentage changes in hourly wages.

Column (7) of Table X shows the benchmark results consistent with our labor demand findings: overall, occupations with higher Generative AI exposure experience a greater relative decline in wage rates after the release of ChatGPT, although the point estimate is not statistically significant. As we have seen in our labor demand results, examining occupations across their Generative AI exposure alone may conceal substantial heterogeneity in the effect.

Columns (8)-(9) show the heterogeneous wage effects when occupations' exposure derives from core or supplemental tasks. For occupations with exposure derived solely from core tasks, a one-standard-deviation higher Generative AI exposure corresponds to a significantly greater decline in relative wage rates by $1.4\% = 0.209 \times 0.065$ after the release of ChatGPT. This result supports the labor-substitution effect. In contrast, the wage effect is estimated to be positive and significant for occupations with exposure coming only from supplemental

⁵³From the CPS data, we extract each individual's hourly wage rate, occupation code, gender, age, race, education level, and sampling weight in the survey. We require the individuals to be between 18 and 65 years old, employed in the month of the survey, and to have a non-missing hourly wage rate and a non-missing occupation code. The data can be downloaded at <https://cps.ipums.org/cps/>.

tasks (first row of Column (9)), consistent with the labor-complementarity effect.

In summary, our findings on occupational job postings and wages validate that our approach based on core versus supplemental tasks can capture Generative AI's labor-substitution and labor-complementarity effects. The overall impact of ChatGPT on occupations with varying Generative AI exposure is dominated by the substitution effect, but the average effect estimates fail to capture substantial heterogeneity in these effects based on which of an occupation's tasks are exposed.

— Insert Table X about here —

E. Firms' labor reaction and Generative AI exposure

We next demonstrate the implications of the labor market findings above for firms with high and low labor exposure to Generative AI. In particular, if occupations with higher exposure to Generative AI are, on average, more substitutable by ChatGPT, we would expect firms with a higher workforce exposure to Generative AI to have a greater labor-saving opportunity when ChatGPT becomes available. We present two sets of results suggesting that this is the case. First, we show that firms with a higher Generative AI exposure substantially reduce their exposed labor after the release of ChatGPT. This labor-substitution effect, as discussed in the conceptual framework, can reduce costs, increase future cash flows, and enhance firm value. Second, we provide additional evidence of the labor-substitution mechanism by showing that the labor reduction results are even stronger if firms' workforce Generative AI exposure derives solely from their core tasks.

Firms' labor demand and Generative AI exposure The labor-substitution mechanism suggests that firms with a higher Generative AI exposure can save more costs by reducing their reliance on labor to do the tasks in highly exposed occupations after the introduction of ChatGPT. We thus define an occupation as highly exposed to Generative AI if the occupation's Generative AI exposure is within the top tercile of the distribution across all occupations, labeled as *High-GenAI Exposure* occupations. From the Lightcast job posting data, we calculate the number of firms' job postings for High-GenAI Exposure occupations in each month. We further require the firms to be in our stock return test sample, resulting in a final sample of 36,900 firm-month observations from January 2022 to August 2023.

Our first test uses the following difference-in-differences regression specification:

$$\text{High-Exp Job Posting}_{i,t} = \kappa \cdot \text{Post}_t \times \text{GenAI Exp}_i + \text{FEs} + \varepsilon_{i,t}, \quad (10)$$

where $High\text{-Exp Job Posting}_{i,t}$ is firm i 's job posting count for the highly exposed occupations defined above, $GenAI Exp_i$ is firm i 's Generative AI exposure, and FEs includes year-month fixed effects and firm fixed effects which subsume the standalone variables $Post_t$, and $GenAI Exposure_i$.

Column (1) of Table XI confirms that firms with a higher Generative AI exposure reduce their demand for high-GenAI exposure occupations after the release of ChatGPT. The magnitude of the coefficient suggests that the effect is sizable, with a one-standard-deviation increase in firm exposure resulting in a 11% ($= 0.078 \times 1.378$) decline in job postings for the highly exposed occupations.⁵⁴

We further test the labor-substitution mechanism by including the firm-level $ShareSupp_i$ in specification (10) above:

$$\begin{aligned} \text{High-Exp Job Posting}_{i,t} = & \phi \cdot Post_t \times GenAI Exp_i + \zeta \cdot Post_t \times ShareSupp_i \\ & + \psi \cdot Post_t \times GenAI Exp_i \times ShareSupp_i + FEs + \varepsilon_{i,t}, \end{aligned} \quad (11)$$

where $ShareSupp_i$ is the share of firm i 's Generative AI exposure deriving from its workers' supplemental tasks' exposure. Similar to our inference at the occupation level, the coefficient ϕ informs the labor-substitution mechanism, as it captures the labor reactions when a firm's labor exposure to Generative AI derives solely from their core tasks (i.e., $ShareSupp_i = 0$). Column (2) confirms that the reduction in labor demand is more than twice as great (i.e., comparing the coefficients -3.038 and -1.378) for firms where the Generative AI exposure derives solely from their workers' core task exposure.⁵⁵

⁵⁴In Appendix Figure IA.5, we plot the effect of firms' GenAI Exp on their job postings for high-exposure occupations in each quarter from 2021q1 to 2023q3. We observe a strong trend break after the release of ChatGPT, consistent with the above regression results. In Panel A, we observe relatively higher job postings from firms with higher GenAI exposure during the pre-period. Our pre-period in this test coincides with the COVID period, during which tech firms had substantially higher labor demand. Panel B shows that the pre-trend is absorbed by industry heterogeneity when including industry \times quarter fixed effects, while Panel C shows that the results are very similar if we exclude firms from the tech sector. While these findings suggest that our results are unlikely to be fully driven by the rise and fall of tech hiring during and after Covid, we cannot fully rule out that the impact of Covid forms part of the explanation, to the degree that it also impacted non-tech firm hiring of high-exposure occupations. We confirm in Appendix Table IA.15 that our labor demand effect estimates are very similar with and without controlling for industry \times month fixed effects.

⁵⁵In the Internet Appendix IA.14, we examine the demand for mid-exposure and low-exposure which can potentially serve as a placebo test. The table shows that firms' GenAI Exposure predicts a less negative response of job postings for mid-exposure occupations to the release of ChatGPT and an insignificant response of job postings for low-exposure occupations. These results highlight that the job posting effects derive primarily from high-exposure occupations, consistent with our mechanism.

Firms’ employment and Generative AI exposure. We reinforce our labor demand findings by further showing results on firms’ actual employment changes in the highly exposed occupations. While the effect on hiring suggests a shift in firms’ labor demand away from highly exposed occupations, a consistent finding in firms’ employment changes strengthens the evidence for the labor-substitution effect. We use LinkedIn data from Revelio Labs on employees in the highly exposed occupations, aggregated at the firm-by-month level for public firms from January 2022 to December 2023. We estimate the same two specifications in equations (10) and (11) above, and replace the dependent variable with the firm’s employment of highly-exposed occupations in the month, $High\text{-}Exp\ Employment_{i,t}$.

Our findings support our occupation-level and job posting results: Column (3) of Table XI shows that a one-standard-deviation increase in firm Generative AI exposure corresponds to a 0.6% ($= 0.078 \times 0.083$) decline in employment of highly-exposed occupations after the release of ChatGPT. Column (4) shows that the effect is much stronger when firms’ Generative AI exposure derives solely from their workforce’s core tasks (comparing the coefficients -0.414 and -0.083).

In summary, the consistent results across the board of our tests regarding firms’ labor demand and employment changes collectively provide strong support for the cost-saving mechanism: firms with a higher labor exposure to Generative AI can reduce their reliance on the exposed workforce. This labor-substitution mechanism provides an explanation for the predicted and realized changes in future cash flows and changes in firm value following the arrival of Generative AI technology, as documented in the previous sections.

— Insert Table XI about here —

F. Labor-substitution mechanism for firm value and cash flow effects

Here, we present our last set of results to further solidify the labor-substitution mechanism for firm cash flow and value results. We have shown in the section above that the labor-substitution effect is strongest when a firm’s employees’ exposure to Generative AI derives solely from their core tasks. Below, we show that the effects on firm value and future cash flows are also strongest when a firm’s employees’ exposure derives solely from their core tasks.

For our test on firm value, we estimate the event study regression on cumulative abnormal returns in equation (4) while interacting $GenAI\ Exp_i$ with $ShareSupp_i$:

$$CAR_i = \mu \cdot GenAI\ Exp_i + \nu \cdot ShareSupp_i + \omega \cdot GenAI\ Exp_i \times ShareSupp_i + \varepsilon_i. \quad (12)$$

We again focus on the coefficient μ , which informs the market value reaction to the release of ChatGPT for firms where Generative AI exposure derives solely from their workforce’s core tasks, $ShareSupp_i = 0$. Column (1) of Table XII reports that their firm value effects are significantly positive with a magnitude larger than the benchmark case shown in Column (1) of Table V (i.e, comparing the coefficients 0.380 and 0.216 from the two tables).

For our tests regarding future cash flows, we similarly add $ShareSupp_i$ to our analyses of firms’ forecast earnings and realized profitability in Section III. Columns (2) and (3) of Table XII show consistent results that if we focus on firms with their Generative AI exposure derived solely from their workers’ core tasks’ exposure, the effects on forecasts for short-term earnings and long-term growth are both at least 60% larger than the average effects shown in Columns (1) and (2) of Table IX. Column (4) shows a similar result for firms’ realized gross profitability.

In summary, the earlier results on the labor side have already suggested that firms’ exposure to Generative AI primarily captures their workers’ substitutability by the technology, supporting a labor-substitution mechanism for our main results on firm value. Our findings in this section further strengthen this mechanism by showing that all effects are stronger if we restrict firms’ Generative AI exposure to capture solely their workers’ substitutability by the technology.

— Insert Table XII about here —

V. Conclusion

Market equity prices indicate that the arrival and diffusion of large language models and Generative AI represent a major technology shock with important effects on the overall value of firms, leading to winners and losers. This paper uses occupational exposures to Generative AI, along with firm-level measures of occupational composition, to assess the exposure to Generative AI innovations at the firm level for publicly traded U.S. corporations. We find that the release of ChatGPT had a substantial impact on firm value, leading to a difference in firm returns of 0.45% daily. Moreover, we show that this change in valuations can be explained by a mechanism where market participants expect higher cash flows for exposed firms in the future: analysts’ long-run earnings forecasts increase more for exposed firms after the release of ChatGPT, and there are also effects on actual gross profitability that are realized within a year after the release.

To show that firms change their actions in response to their technology exposure, we provide evidence that firms with higher exposure to Generative AI-driven productivity increases are more likely to communicate with their investors about these technologies in their

earnings calls, and differ in their hiring behavior in subsequent months, reducing their job postings for highly exposed occupations and also seeing a decline in employment levels for these roles. We also document a relative decline in wages and hiring for highly exposed occupations at the national level in the months that followed the initial technology shock.

We find that there is substantial heterogeneity in these effects that aligns with a mechanism where it matters *which* tasks within a given occupation are exposed: we argue that if more of a worker's core tasks—as opposed to supplemental tasks—can be automated by a technology, this has more negative consequences for labor demand for those roles, but more positive consequences for firm profitability. All our key outcomes show heterogeneity of the effects of Generative AI exposure that aligns with the mechanism.

These findings raise important issues for policymakers: to the extent that shareholders benefit from increases in firm value while affected workers lose out, the new technology redistributes income in the economy, which may or may not be desirable from a welfare perspective. At the same time, our quantification of the productivity potential of Generative AI for different firms permits policymakers and entrepreneurs to better identify areas of opportunity and targets for disruption as this new technology reshapes the economic landscape. The degree to which this new technology will bring pain or plenty will depend on how firms and regulators can align in realizing the value that is promised by the financial market reaction to its release. **The early labor market effects that we document highlight that there are significant impacts on employment and wages of worker in occupations with substantial core-task exposures to Generative AI.**

REFERENCES

Abis, Simona, and Laura Veldkamp, 2023, The changing economics of knowledge production, *Review of Financial Studies* .

Acemoglu, Daron, and David Autor, 2011, Skills, tasks and technologies: Implications for employment and earnings, in *Handbook of labor economics*, volume 4, 1043–1171 (Elsevier).

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, 2022, Artificial intelligence and jobs: evidence from online vacancies, *Journal of Labor Economics* 40, S293–S340.

Acemoglu, Daron, and Pascual Restrepo, 2018, The race between man and machine: Implications of technology for growth, factor shares, and employment, *American economic review* 108, 1488–1542.

Acemoglu, Daron, and Pascual Restrepo, 2020, Robots and jobs: Evidence from us labor markets, *Journal of political economy* 128, 2188–2244.

Agrawal, Ajay, Joshua S Gans, and Avi Goldfarb, 2019, Artificial intelligence: the ambiguous labor market impact of automating prediction, *Journal of Economic Perspectives* 33, 31–50.

Agrawal, Ajay K, Joshua S Gans, and Avi Goldfarb, 2023, The turing transformation: Artificial intelligence, intelligence augmentation, and skill premiums, Technical report, National Bureau of Economic Research.

Akcigit, Ufuk, and Sina T Ates, 2021, Ten facts on declining business dynamism and lessons from endogenous growth theory, *American Economic Journal: Macroeconomics* 13, 257–298.

Andrews, Isaiah, and Maryam Farboodi, 2025, Do markets believe in transformative ai?, MIT Department of Economics and Sloan School of Business Working Paper.

Autor, David H, Lawrence F Katz, and Melissa S Kearney, 2006, The polarization of the us labor market, *American economic review* 96, 189–194.

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, 2024, Artificial intelligence, firm growth, and product innovation, *Journal of Financial Economics* 151, 103745.

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* .

Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson, 2023, Artificial intelligence and firms' systematic risk, *Available at SSRN* .

Begenau, Juliane, Maryam Farboodi, and Laura Veldkamp, 2018, Big data in finance and the growth of large firms, *Journal of Monetary Economics* 97, 71–87.

Berger, David, Kyle Herkenhoff, and Simon Mongey, 2022, Labor market power, *American Economic Review* 112, 1147–1193.

Bertomeu, Jeremy, Yupeng Lin, Yibin Liu, and Zhenghui Ni, 2025, On the value losses of disruption in technology adoption, Working paper, version August 13, 2025. Earlier version circulated as “Capital Market Consequences of Generative AI: Early Evidence from the Ban of ChatGPT in Italy” .

Bloom, Nicholas, Tarek Alexander Hassan, Aakash Kalyani, Josh Lerner, and Ahmed Tahoun, 2021, The diffusion of disruptive technologies, Technical report, National Bureau of Economic Research.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, Matthew OBrien, and Andrei Shleifer, 2024, Long-Term Expectations and Aggregate Fluctuations 38, 311–347.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen, 2025a, Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence, Technical report, Working paper. Latest version available at <https://digitaleconomy.stanford.edu>

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond, 2025b, Generative AI at work, *The Quarterly Journal of Economics* qjae044.

Caserta, Joe, Holger Harreis, Kayvaun Rowshankish, Nikhil Srinidhi, and Asin Tavakoli, 2023, The data dividend: Fueling generative ai .

Chen, AJ, Miao Ben Zhang, and Zhao Zhang, 2023, Talent market competition and firm growth, *Available at SSRN* .

Cohn, Jonathan B, Zack Liu, and Malcolm I Wardlaw, 2022, Count (and count-like) data in finance, *Journal of Financial Economics* 146, 529–551.

Dell'Acqua, Fabrizio, Edward McFowland III, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, François Candelon, and Karim R. Lakhani, 2023, Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality.

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

EEckhout, Jan, and Laura Veldkamp, 2022, Data and market power, Technical report, National Bureau of Economic Research.

Eisfeldt, Andrea L., Antonio Falato, and Mindy Xiaolan, 2022, Human capitalists, in Martin Eichenbaum, Erik Hurst, and Valery Ramey, eds., *NBER Macroeconomics Annual*, chapter 1, 2 (University of Chicago Press, Chicago).

Eisfeldt, Andrea L, Antonio Falato, and Mindy Z Xiaolan, 2023, Human capitalists, *NBER Macroeconomics Annual* 37, 1–61.

Eisfeldt, Andrea L., and Dimitris Papanikolaou, 2013, Organization capital and the cross-section of expected returns, *The Journal of Finance* 68, 1365 – 1406.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, 2023, Gpts are gpts: An early look at the labor market impact potential of large language models, *arXiv preprint arXiv:2303.10130* .

Fama, Eugene F, and Kenneth R French, 2015, A five-factor asset pricing model, *Journal of financial economics* 116, 1–22.

Farboodi, Maryam, Roxana Mihet, Thomas Philippon, and Laura Veldkamp, 2019, Big data and firm dynamics, in *AEA papers and proceedings*, volume 109, 38–42, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Farboodi, Maryam, and Laura Veldkamp, 2023, Data and markets, *Annual Review of Economics* 15, 23–40.

Frey, Carl Benedikt, and Michael A Osborne, 2017, The future of employment: How susceptible are jobs to computerisation?, *Technological forecasting and social change* 114, 254–280.

Gârleanu, Nicolae, Leonid Kogan, and Stavros Panageas, 2012, Displacement risk and asset returns, *Journal of Financial Economics* 105, 491–510.

Gilardi, Fabrizio, Meysam Alizadeh, and Maël Kubli, 2023, Chatgpt outperforms crowd-workers for text-annotation tasks, *arXiv preprint arXiv:2303.15056* .

Greenwood, Jeremy, 1999, The Third Industrial Revolution: Technology, Productivity, and Income Inequality, *Economic Review* 35, 2–2.

Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, 1997, Long-run implications of investment-specific technological change, *The American economic review* 342–362.

Gutiérrez, Germán, and Thomas Philippon, 2017a, Declining competition and investment in the us, Technical report, National Bureau of Economic Research.

Gutiérrez, Germán, and Thomas Philippon, 2017b, Investmentless growth: An empirical investigation, *Brookings Papers on Economic Activity* 89.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller, 2025, Artificial intelligence and the labor market, Technical report, National Bureau of Economic Research.

Hansen, Stephen, Peter John Lambert, Nicholas Bloom, Steven J Davis, Raffaella Sadun, and Bledi Taska, 2023, Remote work across jobs, companies, and space, Technical report, National Bureau of Economic Research.

Hoberg, Gerard, and Gordon Phillips, 2016, Text-based network industries and endogenous product differentiation, *Journal of Political Economy* 124, 1423–1465.

Humlum, Anders, and Emilie Vestergaard, 2025, Large language models, small labor market effects, Technical report, National Bureau of Economic Research.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, 2021, Measuring technological innovation over the long run, *American Economic Review: Insights* 3, 303–320.

Kogan, Leonid, and Dimitris Papanikolaou, 2014, Growth opportunities, technology shocks, and asset prices, *The journal of finance* 69, 675–718.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence Schmidt, and Bryan Seegmiller, 2019, Technology-skill complementarity and labor displacement: Evidence from linking two centuries of patents with occupations, *Available at SSRN 3585676* .

Kogan, Leonid, Dimitris Papanikolaou, Lawrence Schmidt, and Jae Song, 2021, Technological Innovation and Labor Income Risk.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence DW Schmidt, and Bryan Seegmiller, 2023, Technology and labor displacement: Evidence from linking patents with worker-level data, Technical report, National Bureau of Economic Research.

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.

Li, Qin, Ben Lourie, Alexander Nekrasov, and Terry Shevlin, 2022, Employee turnover and firm performance: Large-sample archival evidence, *Management Science* 68, 5667–5683.

Lichtinger, Guy, and Seyed Mahdi Hosseini Maasoum, 2025, Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data.

Lopez-Lira, Alejandro, and Yuehua Tang, 2023, Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models, SSRN Scholarly Paper 4412788, Available at SSRN 4412788, Rochester, NY.

Loughran, Tim, and Bill McDonald, 2016, Textual analysis in accounting and finance: A survey, *Journal of Accounting Research* 54, 1187–1230.

Manuelli, Rodolfo E., and Ananth Seshadri, 2014, Frictionless Technology Diffusion: The Case of Tractors, *American Economic Review* 104, 1368–1391.

Mishra, Sagarika, Michael T. Ewing, and Holly B. Cooper, 2022, Artificial intelligence focus and firm performance, *Journal of the Academy of Marketing Science* 50, 1176–1197.

Novy-Marx, Robert, 2013, The other side of value: The gross profitability premium, *Journal of financial economics* 108, 1–28.

Papanikolaou, Dimitris, 2011, Investment shocks and asset prices, *Journal of Political Economy* 119, 639–685.

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor, 2022, Dissecting green returns, *Journal of Financial Economics* 146, 403–424.

Santos Silva, João, and Silvana Tenreyro, 2015, Ppml: Stata module to perform poisson pseudo-maximum likelihood estimation .

Saunders, Adam, and Erik Brynjolfsson, 2016, Valuing information technology related intangible assets, *Mis Quarterly* 40, 83–110.

Schubert, Gregor, 2025, Organizational technology ladders: Remote work and generative ai adoption, *Available at SSRN* .

Schubert, Gregor, Anna Stansbury, and Bledi Taska, 2020, Employer concentration and outside options, Technical report, mimeo, Harvard University.

Seegmiller, Bryan, 2021, Valuing labor market power: The role of productivity advantages, *Available at SSRN 4412667* .

Tuzel, Selale, and Miao Ben Zhang, 2021, Economic Stimulus at the Expense of Routine-Task Jobs, *Journal of Finance* 76, 3347–3399.

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

Veldkamp, Laura, 2023, Valuing data as an asset, *Review of Finance* .

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>.

Yeh, Chen, Claudia Macaluso, and Brad Hershbein, 2022, Monopsony in the us labor market, *American Economic Review* 112, 2099–2138.

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

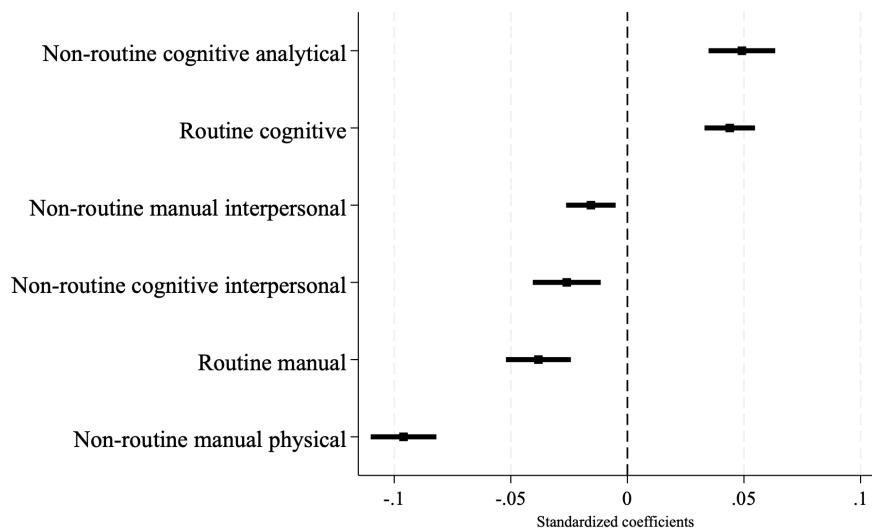


Figure 2. Occupational Generative AI Exposure and Skill Measures. This figure plots the coefficients of regressing occupations' Generative AI exposure (see equation (1)) on six measures of occupations' skills constructed by [Acemoglu and Autor \(2011\)](#). Equation (2) shows the regression specification. Occupation is classified at the SOC 6-digit level. The six skill measures are each standardized to have a mean of zero and a standard deviation of one. See Appendix IA.3.C.4 for details on the construction of these skill measures. The bar around each coefficient shows the 95% confidence interval based on heteroskedasticity-robust standard errors.

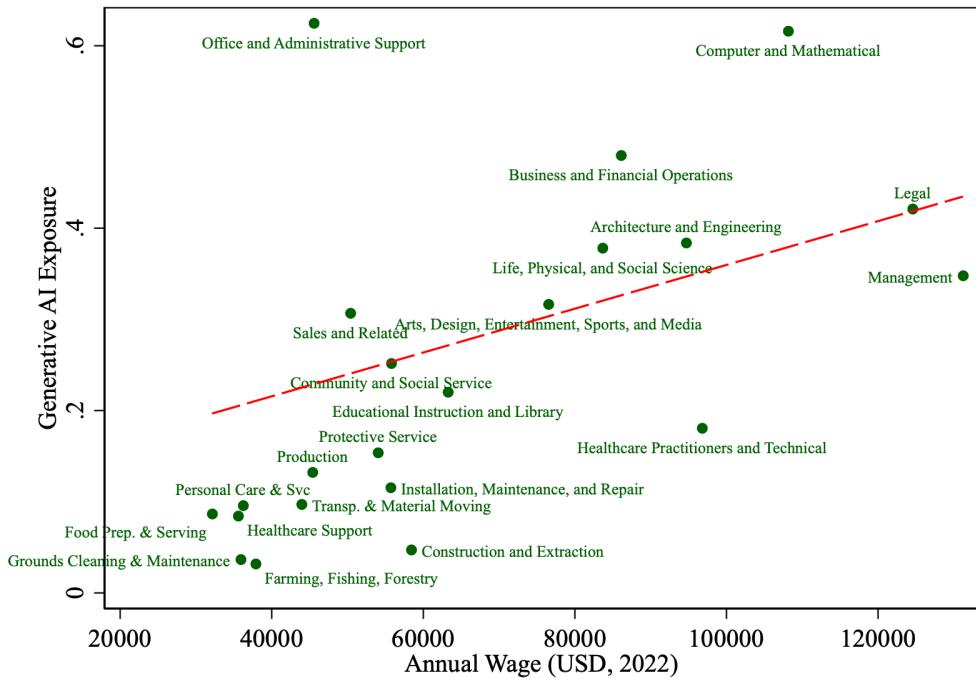


Figure 3. Generative AI Exposure and Wages by Major Occupation Group. The figure plots the relationship between Generative AI exposure and annual wages by SOC 2-digit occupation group. We aggregate the Generative AI exposure from SOC 6-digit to SOC 2-digit using the May 2022 occupational employment distribution from the Bureau of Labor Statistics Occupational Employment and Wage Statistics (OEWS) downloaded at https://www.bls.gov/oes/current/oes_stru.htm. We obtain the average annual wages of the SOC 2-digit occupations from the 2022 OEWS. The dashed red line indicates the employment-weighted linear best fit.

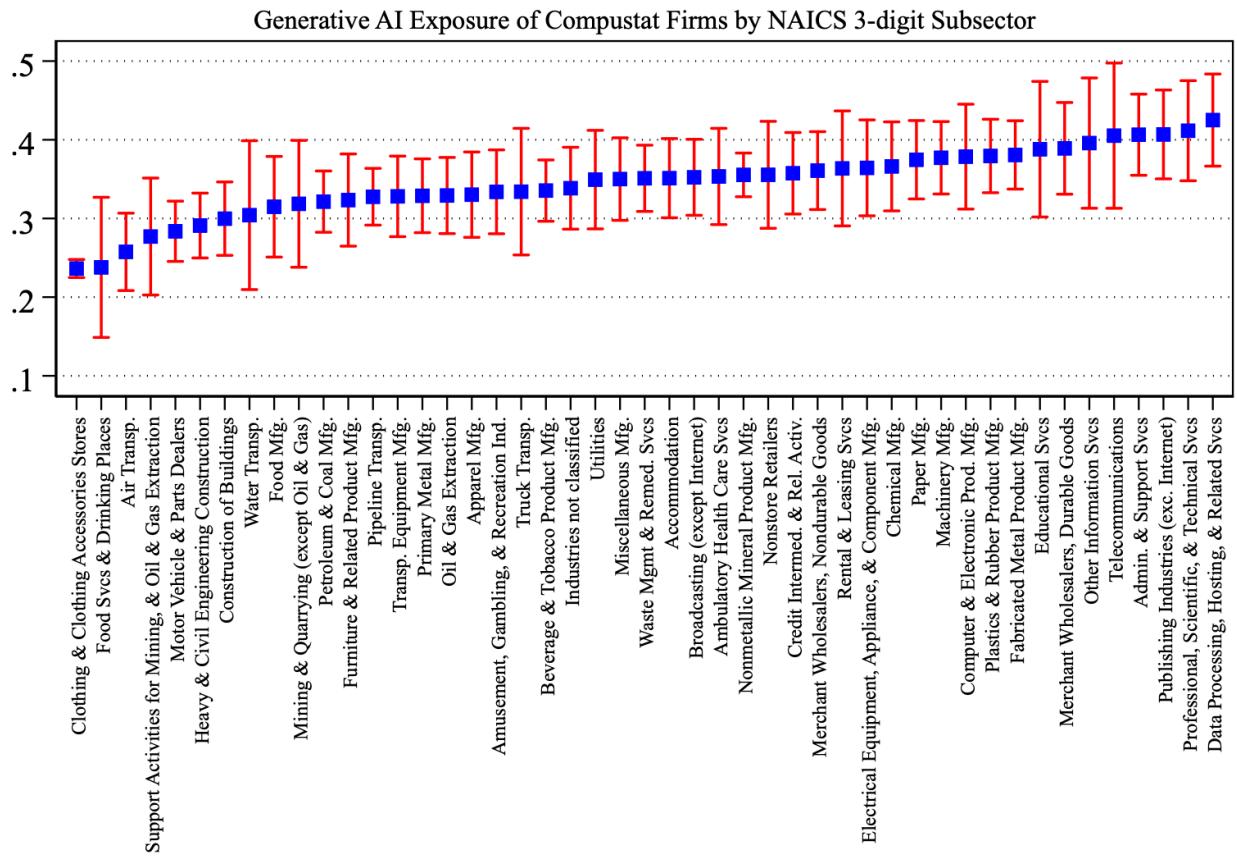


Figure 4. Generative AI Exposure Within and Across Industries. This figure plots the average and the standard deviation of Compustat firms' Generative AI exposure within each NAICS 3-digit industry in 2022. For ease of exposition, we require the industry to have at least 20 firms to be added to this figure. See Appendix Figure IA.2 for the full set of subsectors.

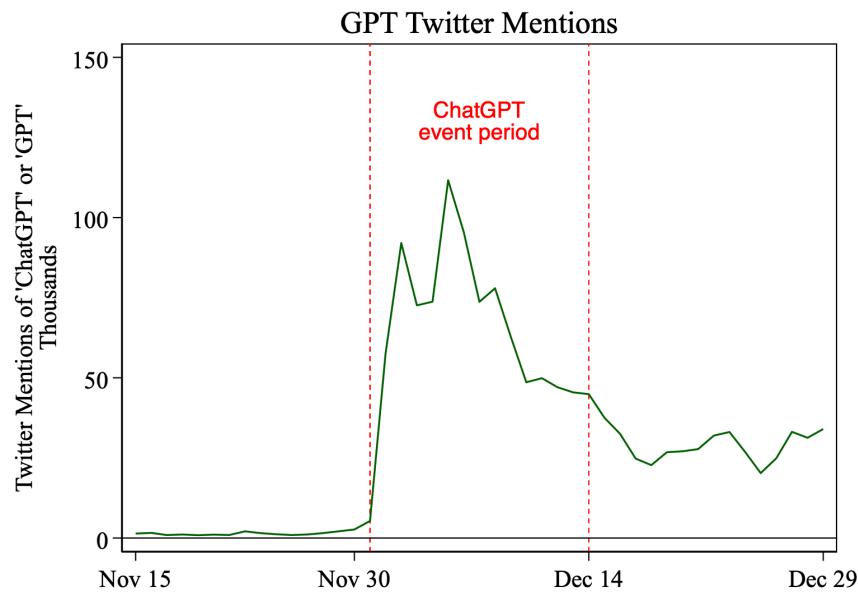
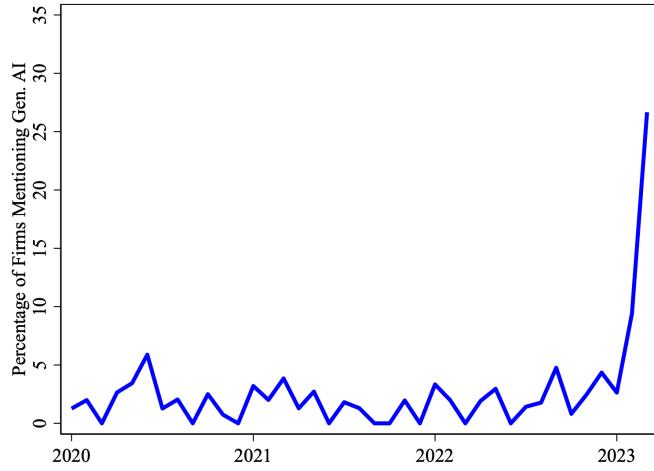


Figure 5. Social Media Attention to the Release of ChatGPT. The figure plots the total count of Twitter mentions of “ChatGPT” or “GPT”, in thousands, on each day from November 14, 2022, to December 29, 2022. The data are from Media Cloud. The red dashed vertical lines indicate the “ChatGPT event period” for our stock market reaction analyses, which includes the two weeks following the release of ChatGPT on November 30, 2022.

Panel A: Share of Firms Mentioning Generative AI



Panel B: Firms' Generative AI Exposure and Their Mentions of Generative AI

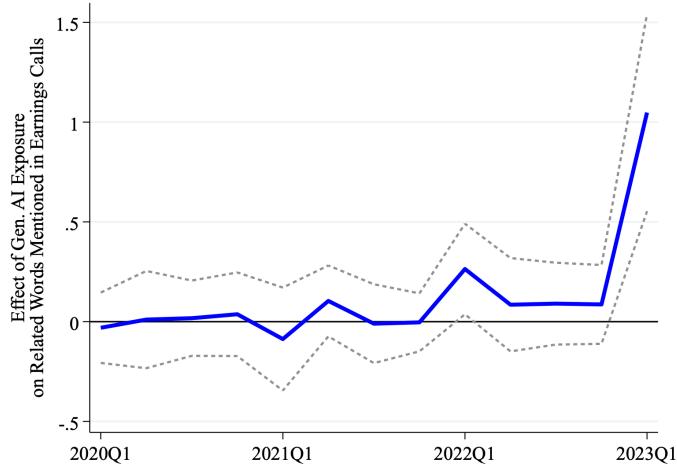


Figure 6. Firm Mentions of Generative AI in Earnings Conference Calls. Panel A plots the percentage of S&P 500 firms mentioning keywords of Generative AI in their quarterly earnings conference call transcripts in each quarter, where the data are manually collected from the *Seeking Alpha* website. We convert each call transcript into a list of lower-case unigrams and bigrams and examine whether it mentions the Generative AI keywords: “lm”, “chatgpt”, “gpt”, “gpt3”, “gpt4”, “generative”, and “language model”. Panel B reports the cross-sectional relationship between firms’ indicator of mentioning Generative AI, $GenAI Mention_{i,t}$, and the labor-based Generative AI exposure, $GenAI Exp_i$, constructed in Section I. We estimate the following regression specification in each fiscal quarter from 2019Q1 to 2023Q1 and plot the point estimates of β_t and the 95% confidence intervals based on heteroskedasticity-robust standard errors: $\mathbb{1}[GenAI Mention]_{i,t} = \alpha_t + \beta_t GenAI Exp_i + \gamma \mathbb{1}[GenAI Mention]_{i,2019} + \varepsilon_{i,t}$. The Appendix Figure IA.3 plots firms’ mentions of generic “engineering” keywords as a placebo test.

Table I
Examples of GPT Scores for Tasks

This table presents examples of the scoring of labor tasks’ exposure to Generative AI by comparing the tasks’ statement to a rubric using the GPT 3.5 Turbo Model. We obtain from the O*Net V27.2 database task statements for each occupation and also an indicator of whether each task is “core” or “supplemental” to the occupation. An occupation has, on average, 22 tasks. We use the GPT model together with a rubric detailing the capabilities of ChatGPT-like tools to classify each task into three categories: if the task can be more efficiently completed with ChatGPT, it receives an exposure score of 1; if the task can be more efficiently completed with ChatGPT only if ChatGPT is further integrated with additional software or applications, it receives a score of 0.5 (Eloundou et al. (2023)); and if the task cannot be completed by ChatGPT, it receives a score of 0. The GPT model also provides an explanation for the classification of each task. Section 1 details the procedure.

Task	Occupation	Type	Score	GPT Explanation for the Score
Adjust sales scripts to better target the needs and interests of specific individuals.	Telemarketers	Core	1	The model can learn from the sales scripts and the needs and interests of specific individuals to adjust the sales scripts.
Write supporting code for Web applications or Web sites.	Web Developers	Core	1	The model can write code based on the specifications provided in the text input.
Develop and maintain media contact lists.	Fundraisers	Supp.	1	The model can learn how to develop and maintain media contact lists during training as part of its general knowledge base, and the task can be completed through text input and output.
Identify opportunities or develop channels for purchase or sale of securities or commodities.	Securities, Commodities, and Financial Services Sales Agents	Core	.5	The model can be used to analyze market trends and provide recommendations for securities or commodities to buy or sell. Additional software could be developed on top of the LLM to retrieve up-to-date information about the market and provide more specific recommendations.
Investigate facts and law of cases and search pertinent sources, such as public records and internet sources, to determine causes of action and to prepare cases.	Paralegals and Legal Assistants	Core	.5	The model can assist in searching pertinent sources and summarizing information, but it may not be able to retrieve up-to-date facts from the internet. Additional software could be developed on top of the LLM to retrieve up-to-date facts and combine them with the LLM capabilities.
Monitor budgets and payroll records, and review financial transactions to ensure that expenditures are authorized and budgeted.	Food Service Managers	Supp.	.5	The model can help with reviewing financial transactions and monitoring budgets by providing summaries and recommendations based on the input data. Additional software could be developed on top of the LLM to retrieve up-to-date financial information and provide more detailed analysis.
Connect heating or air conditioning equipment to fuel, water, or refrigerant source to form complete circuit.	Heating, AC, and Fridge Mechanics and Installers	Core	0	The task requires physical interaction with equipment and materials, which cannot be done by the LLM.
Perform therapeutic procedures, such as injections, immunizations, suturing and wound care, and infection management.	Physician Assistants	Core	0	The task requires physical interaction with patients, which cannot be done by the LLM.
Mentor new faculty.	Business Teachers, Postsecondary	Core	0	The task requires human interaction and decision-making that cannot be automated by the LLM.

Table II
Summary Statistics

This table reports summary statistics of key variations at the occupation level in Panel A and at the firm level in Panel B. *Generative AI Exposure_o* is the SOC 6-digit occupation's Generative AI exposure aggregated from its tasks' Generative AI exposure (see equation (1)). *ShareSupp_o* is the share of the occupation's Generative AI exposure that derives from the occupation's supplemental tasks' exposure to Generative AI as compared to its core tasks' exposure (see equation (6)). For firm-level measures, *Generative AI Exposure* is the firm's Generative AI exposure which is the average of its occupations' Generative AI exposure weighted by the firm's 2022 occupational employment shares from the Revelio Labs database. *Log Size* is the natural logarithm of total assets. *Tobin's Q* is the market value of the firm divided by book assets following [Gutiérrez and Philippon \(2017b\)](#). *ROA* is return on assets measured as EBITDA divided by total assets. *Labor Intensity* is the logarithm of the ratio of employment to the net value of property, plant, & equipment following [Donangelo \(2014\)](#). *Tangibility* is asset tangibility measured as the ratio of PP&E to total assets. *Mkt. Leverage* is the ratio of the firm's total value of debt to the sum of total debt and the market value of equity. See table notes for Tables [VII](#), [IX](#), [VIII](#), and [XI](#) for additional variable definitions.

	Mean	Std. Dev.	p10	p50	p90	Obs.
<i>Panel A: Occupation-level measures</i>						
Generative AI Exposure	0.234	0.209	0.000	0.180	0.526	678
ShareSupp _o	0.227	0.275	0.000	0.141	0.667	576
ShareCore _o	0.773	0.275	0.333	0.859	1.000	576
<i>Panel B: Firm-level measures</i>						
Generative AI Exposure	0.354	0.078	0.268	0.353	0.442	2,518
ShareSupp _f	0.197	0.076	0.127	0.186	0.279	2,515
Cumulative Abnormal Returns	-0.009	0.157	-0.106	-0.010	0.080	2,085
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
Tangibility	0.301	0.266	0.036	0.198	0.758	2,515
Market Leverage	0.239	0.215	0.011	0.179	0.553	2,379
GS Product Exp	0.005	0.069	0.000	0.000	0.000	2,085
GPT10K Product Exp	0.082	0.275	0.000	0.000	0.000	1,910
Count10K Product Exp	0.641	3.891	0.000	0.000	1.000	1,910
BFHH Product Exp	0.002	0.005	0.000	0.000	0.004	1,501
10K Data Assets	0.715	0.773	0.000	1.000	2.000	1,910
AV Data Assets	0.012	0.009	0.005	0.009	0.023	2,043
EPS Forecast	4.270	8.184	-0.720	2.570	10.080	1,728
LTG Forecast	11.214	17.703	-2.815	9.150	26.840	580
Gross Profitability	5.799	6.591	0.637	5.223	13.427	22,275
High-Exposure Job Postings	159	651	0	24	299	36,900
High-Exposure Employment	3936	12065	69	820	7673	31,644

Table III
Firms with Highest and Lowest Generative AI Exposure

This table lists the 15 firms with the highest Generative AI exposure in Panel A and the 15 firms with the lowest exposure in Panel B among the 100 largest US-headquartered publicly traded firms based on their market capitalization as of November 29, 2022. *GenAI Exp* is the firm's labor-based Generative AI exposure defined in Section I. *MktCap* is the firm's market capitalization in \$ billions. *Industry* is classified at the NAICS 3-digit level.

<i>Panel A: Top 15 large firms with the highest Generative AI exposure</i>			
Firm Name	GenAI Exp	MktCap	Industry
International Business Machines Corp	.488	128	Other Information Svcs
Intuit Inc.	.48	110	Publishing Industries (exc. Internet)
QUALCOMM Inc.	.479	123	Computer & Electronic Prod. Mfg.
Fiserv Inc.	.475	64	Data Processing, Hosting, & Related Svcs
NVIDIA Corporation	.468	360	Computer & Electronic Prod. Mfg.
S&P Global Inc	.452	108	Admin. & Support Svcs
Broadcom Inc	.449	234	Computer & Electronic Prod. Mfg.
Verizon Communications Inc	.444	165	Telecommunications
Microsoft Corp	.442	1790	Publishing Industries (exc. Internet)
3M Co	.442	66	Paper Mfg.
Advanced Micro Devices Inc	.441	104	Computer & Electronic Prod. Mfg.
ServiceNow Inc	.434	79	Publishing Industries (exc. Internet)
Adobe Inc	.427	155	Data Processing, Hosting, & Related Svcs
PayPal Holdings Inc	.418	81	Data Processing, Hosting, & Related Svcs
Thermo Fisher Scientific Inc	.411	215	Computer & Electronic Prod. Mfg.

<i>Panel B: Bottom 15 large firms with the lowest Generative AI exposure</i>			
Firm Name	GenAI Exp	MktCap	Industry
Starbucks Corp	.119	114	Food Svcs & Drinking Places
McDonald's Corp	.194	193	Food Svcs & Drinking Places
Target Corp	.235	69	General Merchandise Stores
Walmart Inc	.235	382	General Merchandise Stores
Lowe's Cos Inc	.238	120	Build. Mat., Gard. Equip., Supplies Dealers
TJX Companies Inc (The)	.243	92	Clothing & Clothing Accessories Stores
Costco Wholesale Corp	.252	202	General Merchandise Stores
Union Pacific Corp	.253	127	Rail Transp.
CSX Corp	.256	64	Rail Transp.
United Parcel Service Inc	.256	149	Couriers & Messengers
Home Depot Inc. (The)	.261	321	Build. Mat., Gard. Equip., Supplies Dealers
Norfolk Southern Corp	.272	56	Rail Transp.
Tesla Inc	.283	390	Transp. Equipment Mfg.
Northrop Grumman Corp	.291	84	Computer & Electronic Prod. Mfg.
Mondelez International Inc	.292	91	Food Mfg.

Table IV
Reaction of Portfolios Sorted by Generative AI Exposure

This table reports daily percentage stock returns, as a measure of changes in firm value, during the ChatGPT event period for five value-weighted portfolios sorted on firms' Generative AI exposure. To form the portfolios, we sort NYSE stocks on November 29, 2022, into five quintiles by their Generative AI exposure, and we assign all stocks into the portfolios using the NYSE breakpoints. We aggregate stock returns within each portfolio using the stocks' market capitalization on the prior trading day as the weight. *AMH* is the "Artificial Minus Human" zero net investment portfolio long the portfolio with high exposure and short the portfolio with low exposure. We define the ChatGPT event period as from November 30, 2022, to December 14, 2022, i.e., the two weeks after the release of ChatGPT (see Section II for more details). Panel A reports each portfolio's raw daily return in excess of the risk-free rate. Panels B and C report the CAPM market-adjusted alphas and the Fama-French 5-factor-adjusted alphas, where the factor loadings of each portfolio are estimated using data from the previous six months. Panel D reports the Fama-French 5-factor-adjusted alphas for a sample that excludes firms in the tech sectors, NAICS 51 (information) and NAICS 54 (business services), following [Acemoglu et al. \(2022\)](#). *t*-statistics are shown in parentheses. See Section II.B for more details.

Q1 (H)	Q2	Q3	Q4	Q5 (A)	AMH
<i>A: Excess returns (%)</i>					
-0.169 (-0.62)	0.020 (0.06)	0.141 (0.41)	0.035 (0.11)	0.281 (0.74)	0.450 (3.53)
<i>B: Market-factor-adjusted alpha (%)</i>					
-0.234 (-3.91)	-0.047 (-0.45)	0.074 (0.97)	-0.028 (-0.93)	0.208 (4.54)	0.442 (4.70)
<i>C: FF 5-factor-adjusted alpha (%)</i>					
-0.188 (-2.86)	0.125 (2.31)	-0.017 (-0.35)	-0.026 (-0.82)	0.166 (4.23)	0.354 (3.85)
<i>D: FF 5-factor-adjusted alpha ex. tech sector (%)</i>					
-0.193 (-2.83)	0.131 (2.26)	0.008 (0.22)	0.003 (0.09)	0.192 (2.44)	0.384 (4.77)

Table V
Cumulative Abnormal Returns and Generative AI Exposure

This table reports the relation between firms' cumulative abnormal returns during the ChatGPT event period, $[-1, 10]$ trading days around the release of ChatGPT on November 30, 2022, (CAR $[-1, 10]$) and firms' Generative AI exposure, defined in Section I. To compute the cumulative abnormal return of each stock, we first subtract the daily market return from Kenneth French's website from the daily stock return from Yahoo Finance, and then we cumulate the daily abnormal returns over $[-1, 10]$. See Table II for the definitions of firm characteristic variables. All firm characteristics other than generative AI exposure have been standardized to have a mean of zero and a standard deviation of one. Equation (4) describes the regression specification and the weighting of firms. t -statistics are computed using heteroskedasticity-robust standard errors and reported in parentheses. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

Dep. var.:	CAR[-1,10]							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
GenAI Exp	0.216*** (2.621)	0.226*** (2.672)	0.210** (2.170)	0.214*** (2.747)	0.212** (2.472)	0.208** (2.238)	0.199** (2.374)	0.206** (2.192)
Log Size		-0.010*** (-2.689)					-0.010** (-2.517)	
Tobin's Q			0.003 (0.291)				-0.007 (-0.579)	
ROA				0.009** (1.974)			0.009 (1.356)	
Labor Intensity					0.002 (0.760)		-0.000 (-0.114)	
Tangibility						-0.001 (-0.226)	-0.002 (-0.411)	
Mkt. Leverage							-0.011** (-2.169)	-0.009* (-1.901)
R ²	0.05	0.08	0.06	0.08	0.06	0.06	0.09	0.12
Observations	2085	2078	2078	2073	2054	2069	2073	2047

Table VI
Reaction of Within-Industry Sorted Portfolios

This table reports daily percentage stock returns, as a measure of changes in firm value, during the ChatGPT event period for three value-weighted portfolios sorted on firms' Generative AI exposure within the industry. To form the portfolios, we sort NYSE stocks on November 29, 2022 into three terciles by their Generative AI exposure across all stocks in Panels A1 and B1, within the NAICS 3-digit industry in Panels A2 and B2, and within the [Hoberg and Phillips \(2016\)](#) FIC 50-industry in Panels A3 and B3, and we assign all stocks into the portfolios using the NYSE breakpoints. We aggregate stock returns within each portfolio using the stocks' market capitalization on the prior trading day as the weight. *AMH* is the “Artificial Minus Human” zero net investment portfolio long the portfolio with high exposure and short the portfolio with low exposure. We define the ChatGPT event period as from November 30, 2022, to December 14, 2022, i.e., the two weeks after the release of ChatGPT (see Section II for more details). Panels A1-A3 report the CAPM market-adjusted alphas, and Panels B1-B3 report the Fama-French 5-factor-adjusted alphas. *t*-statistics are shown in parentheses. See Section II.C for more details.

Q1 (H)	Q2	Q3 (A)	AMH
<i>A1: Market-factor-adjusted alpha (%)</i>			
-0.174 (-2.87)	0.069 (1.36)	0.124 (2.86)	0.298 (3.70)
<i>A2: NAICS 3-digit ind.-neutral market-factor-adjusted alpha (%)</i>			
-0.165 (-3.98)	0.029 (0.50)	0.145 (9.43)	0.311 (6.06)
<i>A3: FIC 50 ind.-neutral market-factor-adjusted alpha (%)</i>			
-0.183 (-3.45)	0.045 (1.82)	0.123 (2.04)	0.306 (2.97)
<i>B1: FF 5-factor-adjusted alpha (%)</i>			
-0.059 (-1.12)	0.012 (0.32)	0.105 (2.41)	0.164 (2.03)
<i>B2: NAICS 3-digit ind.-neutral FF 5-factor-adjusted alpha (%)</i>			
-0.146 (-3.78)	0.063 (0.86)	0.156 (5.32)	0.302 (6.31)
<i>B3: FIC 50 ind.-neutral FF 5-factor-adjusted alpha (%)</i>			
-0.106 (-2.48)	0.081 (3.61)	0.121 (2.22)	0.226 (2.57)

Table VII
Product Exposure versus Labor Exposure to Generative AI

This table reports the relation between firms' cumulative abnormal returns during the ChatGPT event period, $[-1, 10]$ trading days around the release of ChatGPT on November 30, 2022, (CAR $[-1, 10]$) and firms' labor-based Generative AI exposure (our main measure defined in Section 1) while controlling for various proxies for firms' product exposure to AI. To compute the cumulative abnormal return of each stock, we first subtract the daily market return from Kenneth French's website from the daily stock return from Yahoo Finance, and then we cumulate the daily abnormal returns over $[-1, 10]$. The definitions of the four proxies of firms' product exposure to AI are described in detail in Section II.C.3. *GPT10 Product Exp* is based on GPT 3.5 Turbo assessment of firms' annual reports' description of their products. *Count10K Product Exp* is based on counting AI-related keywords in firms' annual reports. *GS Product Exp* is based on Goldman Sachs' classification of firms as "near-term beneficiaries of AI". *BFHH Product Exp* is based on firms' share of AI-skilled workers constructed by Babina et al. (2024). All product exposure measures are standardized to have a mean of zero and a standard deviation of one, except *GS Product Exp* and *GPT10K Product Exposure*, which are binary variables. All regressions include NAICS 3-digit fixed effects. Equation (5) describes the regression specification and the weighting of firms. *t*-statistics are computed using heteroskedasticity-robust standard errors and reported in parentheses. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

Dep. var.:	CAR[-1,10]					
	(1)	(2)	(3)	(4)	(5)	(6)
GenAI Exp	0.288*** (2.770)	0.291*** (2.835)	0.187** (2.246)	0.264*** (3.074)	0.200** (2.223)	0.175** (2.029)
<i>GPT10K Product Exp</i>		0.011 (0.614)				-0.016 (-0.845)
<i>Count10K Product Exp</i>			0.015*** (5.521)			0.014 (1.608)
<i>GS Product Exp</i>				0.041** (2.293)		0.013 (0.437)
<i>BFHH Product Exp</i>					0.012*** (5.012)	0.001 (0.164)
R-squared	0.36	0.37	0.40	0.38	0.38	0.40
Observations	2080	1905	1905	2080	1493	1374
NAICS 3-dig. FE	X	X	X	X	X	X

Table VIII
Generative AI Exposure and Firm Data Assets

This table reports coefficients from regressions of cumulative abnormal returns of firms during the ChatGPT event period, $[-1, 10]$ trading days around the release of ChatGPT on November 30, 2022, on firms' Generative AI exposure interacted with proxies for their data assets. To compute the cumulative abnormal return of each stock, we first subtract the daily market return from Kenneth French's website from the daily stock return from Yahoo Finance, and then we cumulate the daily abnormal returns over the $[-1, 10]$ period. Section II.D details the construction of the two proxies for firms' data assets. *10K Data Assets* is a score of 0, 1, 2, or 3, indicating the firm is not likely, slightly likely, moderately likely, or highly likely to have data that can be used as an input into large language model analytics based on the assessment of firms' annual reports using the GPT 3.5 Turbo Model. *AV Data Assets* is constructed following [Abis and Veldkamp \(2023\)](#) and measures the intensity of data management skills in the firm's workforce. All regressions include NAICS 3-digit fixed effects and weight firms by their market capitalization as of November 29, 2022. *t*-statistics are computed using heteroskedasticity-robust standard errors and reported in parentheses. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

Dep. var.:	CAR[-1,10]	
	(1)	(2)
GenAI Exp	-0.035 (-0.403)	-0.085 (-0.992)
GenAI Exp \times <i>10K Data Assets</i>	0.216** (2.144)	
GenAI Exp \times <i>AV Data Assets</i>		12.122** (2.339)
<i>10K Data Assets</i>	-0.073* (-1.816)	
<i>AV Data Assets</i>		-3.815 (-1.644)
R ²	0.11	0.12
Observations	1,910	2,043
NAICS 3-dig. FE	X	X

Table IX
Mechanism: Future Cash Flow Impact and Generative AI Exposure

This table reports the impact of the release of ChatGPT on analyst forecasts of firms' future earnings per share (EPS) and on firms' reported quarterly profitability. Analyst forecasts are updated monthly for each firm in the I/B/E/S data. For each firm in each month, we obtain the median analyst forecasts of the firm's EPS in the fiscal year ending in December 2023 and the firm's long-term annual percentage growth rate in earnings (LTG). The dependent variables in Columns (1) and (2) are the firm's median forecasts of the corresponding measures made in October 2022 and January 2023. The dependent variable in Column (3) is the firm's quarterly gross profitability defined as $(\text{Revenue} - \text{COGS})/\text{Assets}$ following [Novy-Marx \(2013\)](#) in percentage using the Compustat quarterly data from 2022Q1 to 2024Q3. The independent variable *Post* is an indicator for whether the time period is after the release of ChatGPT (i.e., after November 2022). *GenAI Exp* is a continuous measure of GenAI exposure at the firm level. See Section [III](#) for more details. All regressions include both firm-level and time-period fixed effects. *t*-statistics in parentheses are computed using standard errors clustered at the firm level. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

<i>Data:</i>	I/B/E/S Analyst Forecasts		Compustat
<i>Measure:</i>	EPS	LTG	Gross Profit
	(1)	(2)	(3)
Post \times GenAI Exp	1.142** (2.325)	9.398* (1.849)	1.711** (2.213)
Observations	1,728	580	22,275
<i>Addl. controls & fixed effects</i>			
Time FEs	X	X	X
Firm FEs	X	X	X

Table X
Validation: Occupational Hiring and Wage Impact and Generative AI Exposure

This table reports the impact of the release of ChatGPT on job postings and wages of occupations. In Columns (1)-(3), the sample includes the total number of job postings at the occupation-month level, and in Columns (4)-(6) at the firm-occupation-month level in the U.S. from January 2022 to August 2023, aggregated from the Lightcast (formerly Burning Glass) job posting database. In Columns (7)-(9), the sample consists of the worker-level hourly wage rate from the Census Current Population Survey (CPS) in monthly panels from January 2022 to October 2023. *GenAI Exp* is the occupation-level Generative AI exposure. *ShareSupp* and *ShareCore* are the occupation's share of Generative AI exposure deriving from the occupation's supplemental and core tasks' exposure to Generative AI. Worker demographic controls in Columns (7)-(9) include the worker's gender, age, age squared, work experience, experience squared, race (white versus not), and years of education. Individual observations in Column, (7)-(9) are weighted by the CPS sampling weight. The job posting panels include many zeros, which makes it impractical to log-transform the dependent variable. Instead, we estimate all regressions using the Poisson model with fixed effects following [Cohn et al. \(2022\)](#), which means that coefficients can be interpreted as semi-elasticities, i.e., in terms of percentage changes in the dependent variable. *t*-statistics in parentheses are computed using standard errors double clustered at the month level and the occupation level in Columns (1)-(3) and (7)-(9), and triple clustered at the month, occupation, and firm levels in Columns (4)-(6). *** and ** denote significance at 1%, 5%, and 10% levels, respectively. See Section IV.D for more details.

Panel Unit:	Dep var.:	Job Postings						Hourly Wage	
		Occ. × Month			Occ. × Firm × Month			Worker × Month	(9)
		(1)	(2)	(3)	(4)	(5)	(6)		
Post × GenAI Exp _o	-0.387*** (-4.406)	-0.672*** (-5.369)	0.387** (2.074)	-0.657*** (-4.354)	-1.014*** (-5.718)	0.735*** (2.057)	-0.023 (-1.158)	-0.065*** (-2.377)	0.087** (2.185)
Post × ShareSupp _o		-0.258*** (-3.543)			-0.303** (-2.525)		-0.011 (-0.905)		
Post × GenAI Exp _o × ShareSupp _o			1.059*** (3.842)		1.749*** (3.597)		0.152*** (2.725)		
Post × ShareCore _o				0.258*** (3.544)		0.203*** (2.494)		0.011 (0.905)	
Post × GenAI Exp _o × ShareCore _o				-1.059*** (-3.845)		-1.749*** (-3.600)		-0.152*** (-2.725)	
Observations	15,540	12,920	12,920	32,496,147	27,001,499	27,001,499	73,604	66,841	66,841
Addl. controls & fixed effects									
Month FE	X	X	X	X	X	X	X	X	X
Occupation FE	X	X	X	X	X	X	X	X	X
Firm × Month FE									
Worker demographics							X	X	X

Table XI
Labor-Substitution Channel: Firm Hiring and Employment

This table reports the impact of the release of ChatGPT on job postings and employment in high-exposure occupations by firms. In Columns (1)-(2) the sample includes the high-exposure (top tercile) job postings at the firm-month level from January 2022 to August 2023 aggregated from the Lightcast (formerly Burning Glass) job posting database. In Columns (3)-(4), the sample consists of total employment in high-exposure (top tercile) occupations at publicly traded firms for each month, computed from LinkedIn data for publicly traded U.S. firms from January 2022 to December 2023. *GenAI Exp* is the firm-level Generative AI exposure, which is an employment-weighted average across occupations within the firm. *ShareSupp* is the firm's share of Generative AI exposure deriving from the supplemental tasks' exposure to Generative AI. The job posting panels include many zeros, which makes it impractical to log-transform the dependent variable. Instead, we estimate all regressions using the Poisson model with fixed effects following [Cohn et al. \(2022\)](#), which means that coefficients can be interpreted as semi-elasticities, i.e., in terms of percentage changes in the dependent variable. *t*-statistics in parentheses are computed using standard errors double-clustered at the month level and the firm level. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

Dep var.:	High-Exposure Job Postings		High-Exposure Employment	
	(1)	(2)	(3)	(4)
Post \times GenAI Exp	-1.378*** (-3.747)	-3.038*** (-2.991)	-0.083* (-1.810)	-0.414*** (-3.169)
Post \times ShareSupp		-2.302** (-2.029)		-0.665*** (-3.074)
Post \times GenAI Exp \times ShareSupp		8.100** (2.037)		1.500** (2.356)
Observations	36,900	36,880	31,644	31,644
<i>Addl. controls & fixed effects</i>				
Month FE	X	X	X	X
Firm FE	X	X	X	X

Table XII
Labor-Substitution Channel: Firm Value and Future Cash Flows

This table reports the heterogeneous impact of the release of ChatGPT on firm value and future cash flows based on the firm's workforce's supplemental versus core task exposure to Generative AI. The heterogeneous exposure is captured by an interaction between firms' generative AI exposure (*GenAI Exp*) and firms' share of Generative AI exposure deriving from their workers' supplemental tasks' exposure to Generative AI (*ShareSupp*). Column (1) shows the results on firms' cumulative abnormal returns (CARs) during the release period. See details about the specification in Table V. Columns (2) and (3) show the results on monthly analyst forecasts of firms' 2023 year-end earnings per share (EPS) and the long-term annual percentage growth rate of EPS (over the next business cycle) made from May 2022 to May 2023 using the I/B/E/S data. Column (4) shows the results on firms' quarterly gross profitability. In Columns (2)-(4), *Post* is an indicator for whether the time period is after the release of ChatGPT (i.e., after November 2022). See details about the specification and sample in Table IX. *t*-statistics, reported in parentheses, are computed using heteroskedasticity-robust standard errors in Column (1) and using standard errors clustered by firm in Columns (2)-(4). ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

Measure:	CAR[-1,10]	EPS	LTG	Gross Profit
	(1)	(2)	(3)	(4)
GenAI Exp	0.380*** (2.579)			
ShareSupp	0.272 (1.447)			
GenAI Exp × ShareSupp	-1.937*** (-3.226)			
Post × GenAI Exp		2.563** (2.431)	21.077** (1.967)	6.087*** (3.369)
Post × ShareSupp		1.599** (2.487)	9.121 (0.754)	6.786** (2.434)
Post × GenAI Exp × ShareSupp		-4.892* (-1.795)	-87.229** (-2.055)	-24.678*** (-3.193)
Observations	2,084	11,664	4,125	22,267
<i>Addl. controls & fixed effects</i>				
Time FEs	X	X	X	X
Firm FEs	X	X	X	X

Internet Appendix for
“Generative AI and Firm Values”

Andrea L. Eisfeldt

Gregor Schubert

Bledi Taska

Miao Ben Zhang

IA.1 Appendix Figures

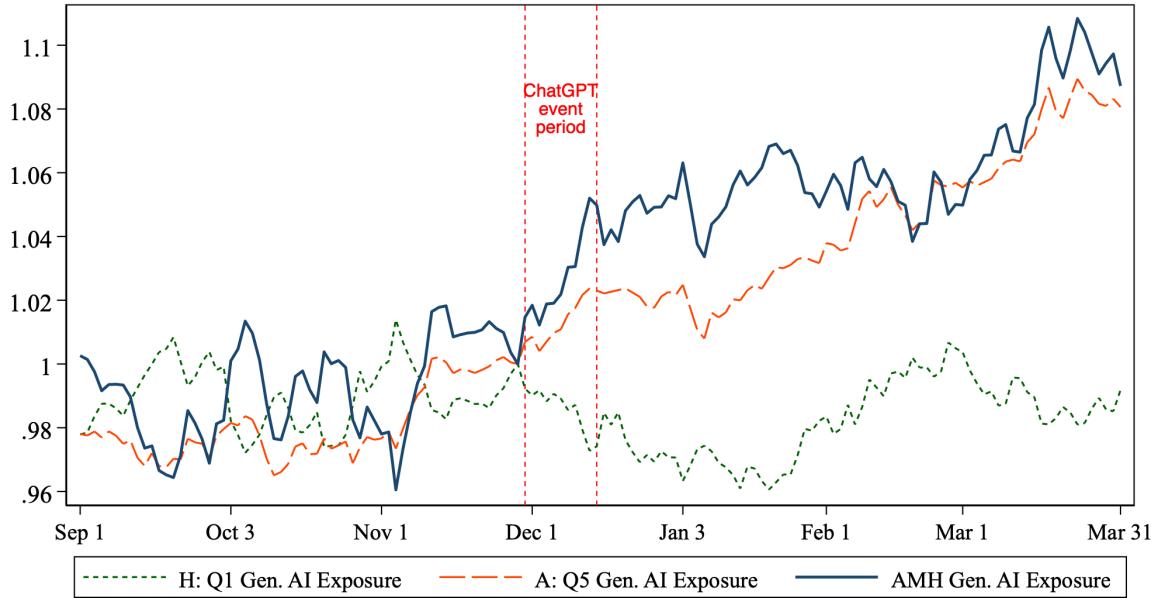


Figure IA.1. Cumulative abnormal returns by Generative AI exposure: Extended pre-period The figure plots the cumulative abnormal returns (CARs) of value-weighted quintile portfolios sorted by firms' labor-based Generative AI exposure. The graph shows the CARs of the lowest-exposure quintile portfolio, "Human" (H), the highest-exposure quintile portfolio, "Artificial" (A), and the zero investment portfolio that goes long A and shorts H, "Artificial-minus-Human" (AMH). Market-adjusted daily abnormal returns are cumulated from November 29, 2022, the day before the release of ChatGPT, and are based on factor exposures computed over the 4-month period preceding the period shown in the graph. Daily stock returns are from Yahoo Finance. The dashed vertical lines indicate the "ChatGPT event period" from November 30, 2022, to December 14, 2022. See details of the definition of firms' Generative AI exposure in Section I and the construction of the portfolios, the calculation of portfolios' CARs, and the determination of the ChatGPT event period in Section II.

Figure IA.2. Generative AI exposure across and within subsectors: complete list. This figure plots the average and the standard deviation of Compustat firms' Generative AI exposure within each NAICS 3-digit subsector.

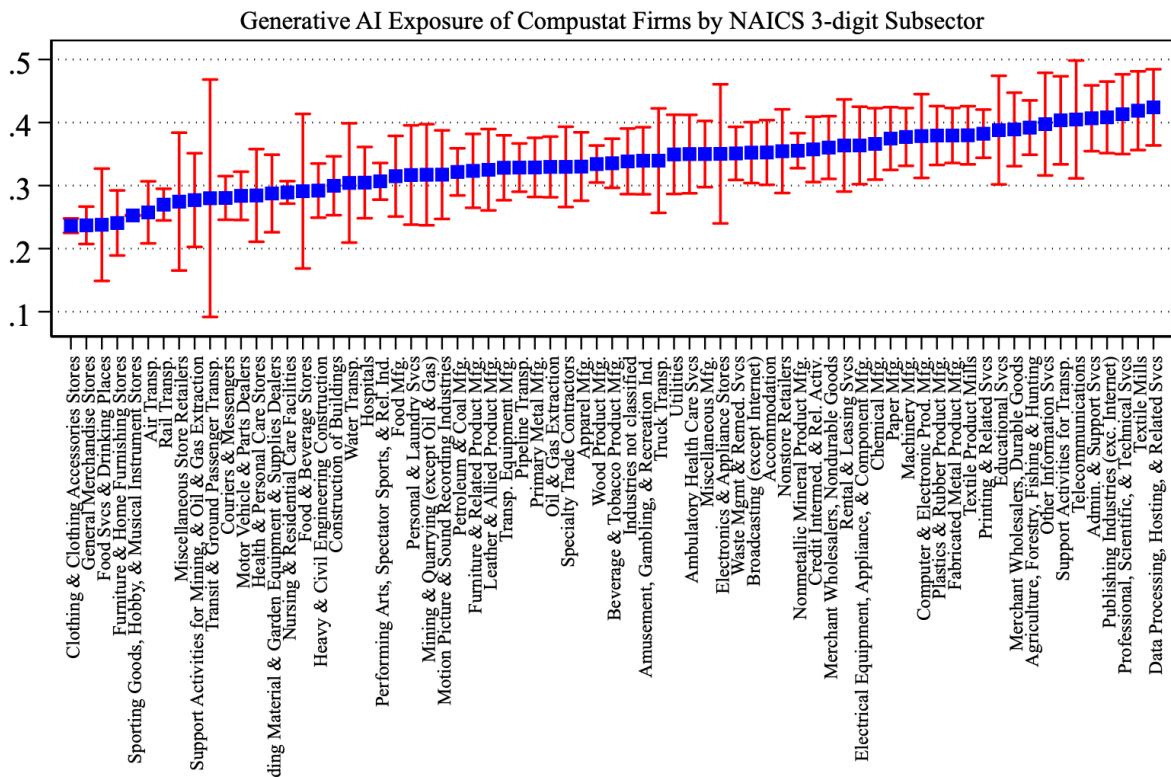
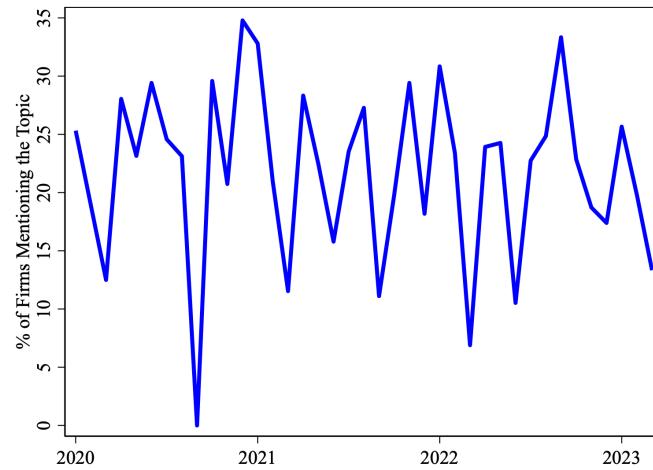


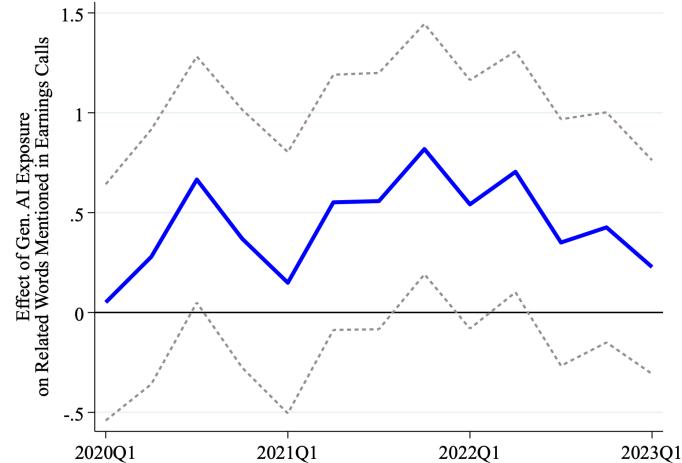
Figure IA.3. Firm-level Generative AI exposure and “engineering” topic mentions in company earnings conference calls. These graphs use data on the share of S&P 500 firms’ earnings calls that mention a particular topic. The quarterly earnings conference call transcripts for S&P 500 firms are manually collected from the *Seeking Alpha* website. Each call transcript is converted into a list of lower-case unigrams and bigrams. Panel A shows the share of earnings calls mentioning keywords about engineering: “engineer” and “engineering”. Panel B then shows the result of estimating regression specifications of the form

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

for cross-sections in each fiscal quarter from 2019 Q1 to 2023 Q1. The graph shows the estimates $\hat{\beta}_t^X$ that represent the effect of higher Generative AI exposure on the likelihood that a firm mentions “engineering” as a topic in that quarter’s earnings call, controlling for whether the firm already mentioned the topic in any 2019 earnings call. Dotted lines show 95% confidence intervals based on heteroskedasticity-robust standard errors.



(A) Share of earnings calls mentioning Engineering



(B) Effect of Gen. AI exposure on firm mentioning Engineering

Figure IA.4. Price reaction of IT producers versus the AMH portfolio: This plot shows the release event period cumulative abnormal returns adjusted for FF5 factor exposure the AMH portfolio and a long-short portfolio of IT producers versus Non-IT producers, where IT producers are defined based on [Saunders and Brynjolfsson \(2016\)](#) as all firms with primary industry codes in computer and electronic product manufacturing (NAICS 334), software publishing (NAICS 5112), data processing, hosting, and related services (NAICS 518), and computer system design and related services (NAICS 5415). The figure shows that the impact of the release of ChatGPT on IT producers occurred in one day, consistent with the industry effect being simpler than, and distinct to, the labor effects. Daily abnormal returns are cumulated from November 29, 2022, the day before the release of ChatGPT, and are based on factor exposures computed over the 6-month period preceding the period shown in the graph. Daily stock returns are from Yahoo Finance. The dashed vertical lines indicate the “ChatGPT event period” from November 30, 2022, to December 14, 2022.

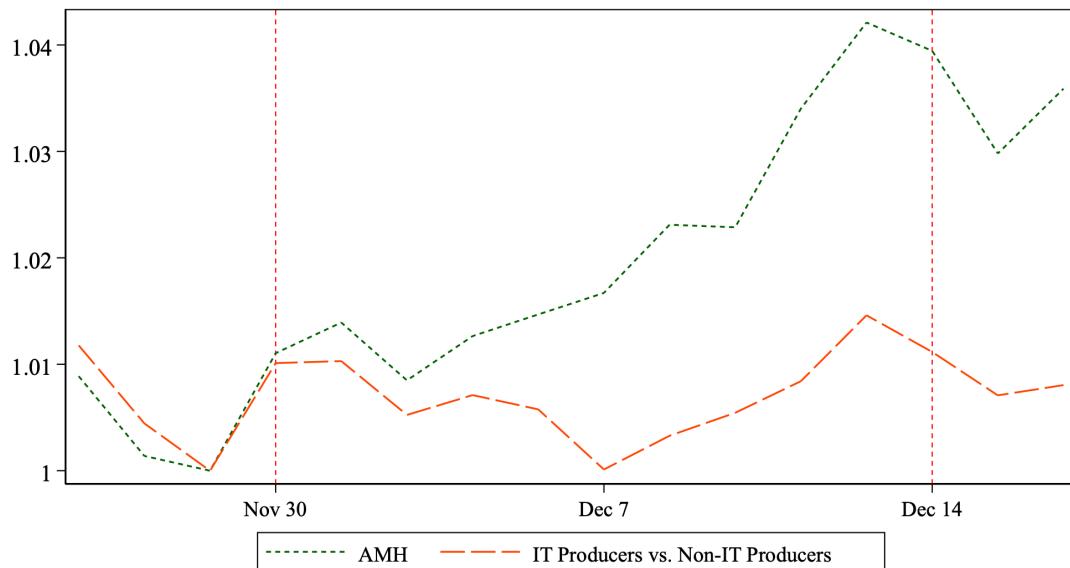
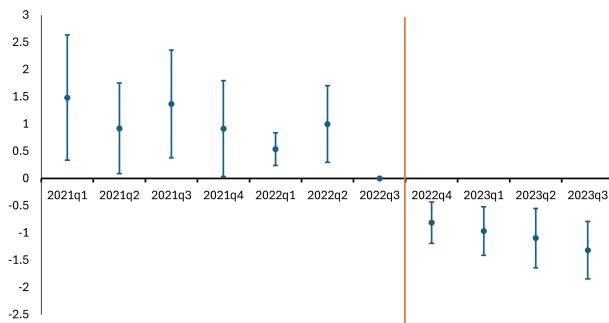
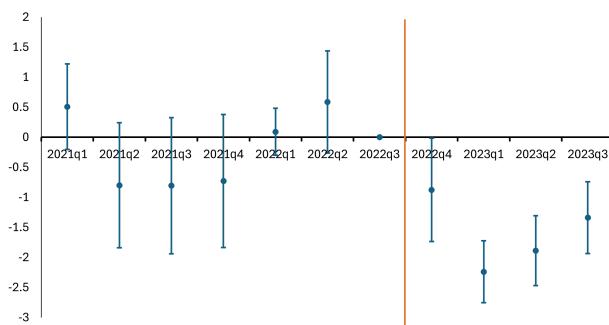


Figure IA.5. Firms' Generative AI exposure and the hiring effects: This figure plots the dynamic effects of firm-level Generative AI exposure on firm-level job posting for high-GenAI-exposure occupations before and after the release of ChatGPT. In Panel A, the regression specification is the following: $y_{i,q} = \sum_{q \in [2021q1, 2023q3]} \beta_q \cdot 1_q \times \text{GenAI Exp}_i + \alpha_i + \alpha_q + \epsilon_{i,q}$, where α_i indicates firm fixed effects and α_q indicates quarter fixed effects. The figure plots the coefficients β_q , which capture the relative size of the effect of firm-level Generative AI exposure on the outcome variable in each quarter *relative* to 2022Q3. Panel B replaces the quarter fixed effects with industry-quarter fixed effects, where industry is defined at the NAICS 3-digit level. Panel C further excludes firms from the tech sectors (see Table IV).

Panel A: With Quarter Fixed Effects



Panel B: With Industry-Quarter Fixed Effects



Panel C: With Industry-Quarter Fixed Effects (Exclude Tech Sectors)

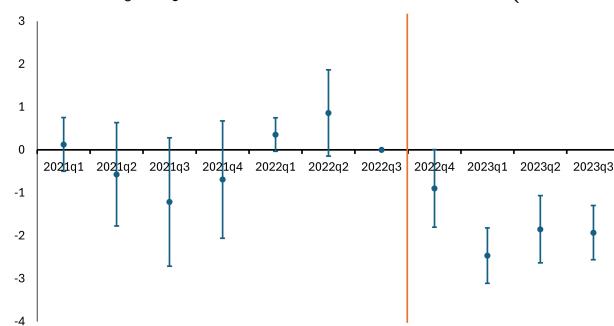
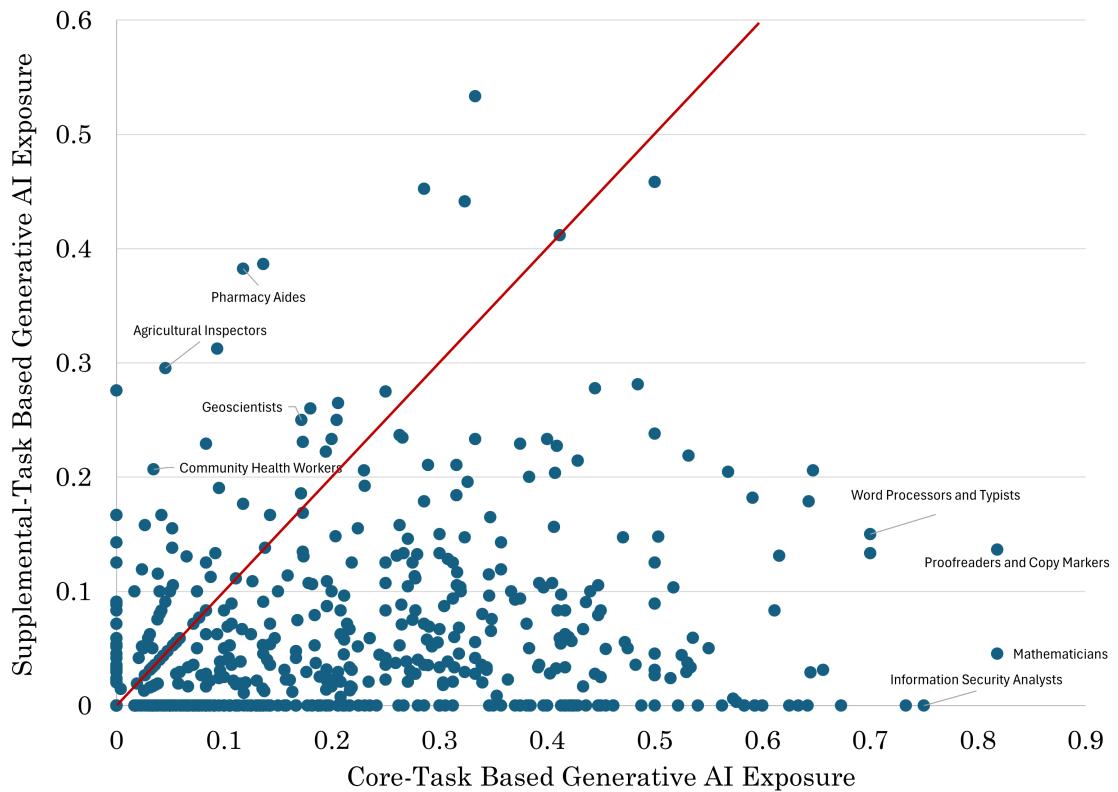


Figure IA.6. Occupations' core-task and supplemental-task exposures to Generative AI: The plot shows the distribution of occupations' core-task exposure to Generative AI in x -axis and supplemental-task exposure in y -axis. Each dot represents a SOC 6-digit occupation. An occupation's supplemental-task exposure to Generative AI is computed as the occupation's GenAI Exposure (equation (1)) multiplied by its' ShareSupp (equation (6)). The red line is the 45 degree line.



IA.2 Appendix Tables

Table IA.1
Highest and Lowest Generative AI Exposure Score Occupations

See Section I.B for details. Note that 15% (102 of 678) of the occupations that exist in our Revelio Labs data have zero exposure, so only a subset of the lowest exposure occupations is shown.

SOC Code	Occupation Title	GenAI Exp	ShareSupp
41-9041	Telemarketers	0.96	0.48
43-9081	Proofreaders and copy markers	0.95	0.14
43-3031	Bookkeeping, accounting, and auditing clerks	0.87	0.62
15-2021	Mathematicians	0.86	0.05
15-1251	Computer programmers	0.85	0.24
43-9022	Word processors and typists	0.85	0.18
43-3011	Bill and account collectors	0.83	0.16
27-3091	Interpreters and translators	0.82	0.50
43-9111	Statistical assistants	0.82	0.22
15-1254	Web developers	0.81	0.06
43-6011	Executive secretaries and executive administrative assistants	0.77	0.24
43-3051	Payroll and timekeeping clerks	0.77	0.26
43-6014	Secretaries and administrative assistants, except legal, medical, and executive	0.77	0.37
43-5061	Production, planning, and expediting clerks	0.76	0.58
15-1212	Information security analysts	0.75	0.00
43-6013	Medical secretaries and administrative assistants	0.75	0.29
27-3043	Writers and authors	0.75	0.2
43-4021	Correspondence clerks	0.74	0.61
43-9061	Office clerks, general	0.74	0.32
41-3091	Sales representatives of services, except advertising, insurance, financial services, and travel	0.73	0.00
:	:	:	:
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 IA.2
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. *Subsector* is defined at the NAICS 3-digit level.

Company Name	Gen. AI exposure	MktCap	Subsector
International Business Machines Corp	488	128	Other Information Svcs
Intuit Inc.	.48	110	Publishing Industries (exc. Internet)
QUALCOMM Inc.	.479	123	Computer & Electronic Prod. Mfg.
Fiserv Inc.	.475	64	Data Processing, Hosting, & Related Svcs
NVIDIA Corporation	.468	360	Computer & Electronic Prod. Mfg.
S&P Global Inc	.452	108	Admin. & Support Svcs
Broadcom Inc	.449	234	Computer & Electronic Prod. Mfg.
Verizon Communications Inc	.444	165	Telecommunications
Microsoft Corp	.442	1790	Publishing Industries (exc. Internet)
3M Co	.442	66	Paper Mfg.
Advanced Micro Devices Inc	.441	104	Computer & Electronic Prod. Mfg.
ServiceNow Inc	.434	79	Publishing Industries (exc. Internet)
Adobe Inc	.427	155	Data Processing, Hosting, & Related Svcs
PayPal Holdings Inc	.418	81	Data Processing, Hosting, & Related Svcs
Thermo Fisher Scientific Inc	.411	215	Computer & Electronic Prod. Mfg.
Intuitive Surgical Inc	.404	93	Computer & Electronic Prod. Mfg.
Automatic Data Processing Inc	.398	99	Data Processing, Hosting, & Related Svcs
Comcast Corp	.396	148	Broadcasting (except Internet)
Vertex Pharmaceuticals Inc	.395	74	Chemical Mfg.
Analog Devices Inc	.392	84	Computer & Electronic Prod. Mfg.
AbbVie Inc	.391	286	Chemical Mfg.
Regeneron Pharmaceuticals Inc	.39	79	Chemical Mfg.
Gilead Sciences Inc	.388	107	Chemical Mfg.
Micron Technology Inc.	.388	55	Computer & Electronic Prod. Mfg.
Intel Corp	.386	109	Computer & Electronic Prod. Mfg.
Bristol-Myers Squibb Co	.385	151	Chemical Mfg.
Illinois Tool Works Inc.	.382	67	Machinery Mfg.
Netflix Inc	.381	131	Rental & Leasing Svcs
Meta Platforms Inc	.381	315	Other Information Svcs
Lam Research Corp	.38	58	Machinery Mfg.
SALESFORCE INC	.379	130	Publishing Industries (exc. Internet)
General Dynamics Corp	.378	68	Transp. Equipment Mfg.
Abbott Laboratories	.376	191	Computer & Electronic Prod. Mfg.
AT&T Inc	.375	131	Telecommunications
Applied Materials Inc	.374	82	Machinery Mfg.
Booking Holdings Inc	.373	76	Other Information Svcs
General Electric Co	.373	91	Industries not classified
Merck & Co Inc	.372	282	Chemical Mfg.
T-Mobile US Inc	.371	173	Telecommunications
Johnson & Johnson	.371	462	Chemical Mfg.
Honeywell International Inc	.368	143	Transp. Equipment Mfg.
Alphabet Inc	.366	1134	Data Processing, Hosting, & Related Svcs
Amgen Inc	.365	140	Chemical Mfg.
Eli Lilly and Co	.364	329	Chemical Mfg.
Apple Inc	.364	2072	Computer & Electronic Prod. Mfg.
Philip Morris International Inc	.364	157	Beverage & Tobacco Product Mfg.
DEERE & COMPANY	.364	128	Machinery Mfg.
Texas Instruments Inc	.363	150	Computer & Electronic Prod. Mfg.
Emerson Electric Co.	.363	57	Computer & Electronic Prod. Mfg.
Caterpillar Inc	.358	124	Machinery Mfg.
CVS Health Corp	.356	121	Ambulatory Health Care Svcs
Cisco Systems Inc	.355	196	Computer & Electronic Prod. Mfg.
Zoetis Inc	.355	68	Chemical Mfg.
Pfizer Inc	.352	288	Chemical Mfg.
Southern Co (The)	.351	78	Utilities
Danaher Corp	.35	193	Computer & Electronic Prod. Mfg.
Procter & Gamble Co (The)	.342	363	Chemical Mfg.
Raytheon Technologies Corp	.339	148	Transp. Equipment Mfg.
Colgate-Palmolive Co	.337	65	Chemical Mfg.
Becton Dickinson and Co	.331	72	Miscellaneous Mfg.
NextEra Energy Inc	.329	166	Utilities
Walt Disney Co (The)	.328	158	Broadcasting (except Internet)
Altria Group Inc	.327	82	Beverage & Tobacco Product Mfg.
Air Products and Chemicals Inc.	.326	68	Chemical Mfg.
Waste Management Inc.	.325	64	Waste Mgmt & Remed. Svcs
Duke Energy Corp	.322	79	Utilities
EOG Resources Inc.	.322	76	Oil & Gas Extraction
Exxon Mobil Corp	.32	450	Petroleum & Coal Mfg.
Estee Lauder Cos Inc (The)	.32	89	Chemical Mfg.
Amazon.com Inc	.317	860	Nonstore Retailers
Stryker Corp	.317	93	Miscellaneous Mfg.
Schlumberger Ltd	.316	76	Support Activities for Mining, & Oil & Gas Extraction
Conocophillips	.316	144	Oil & Gas Extraction
HCA Healthcare Inc	.312	67	Hospitals
Occidental Petroleum Corp	.307	57	Oil & Gas Extraction
Coca-Cola Co (The)	.306	275	Beverage & Tobacco Product Mfg.
Boston Scientific Corp	.305	66	Computer & Electronic Prod. Mfg.
Freeport-McMoRan Inc	.305	54	Mining & Quarrying (except Oil & Gas)
PepsiCo Inc	.303	249	Beverage & Tobacco Product Mfg.
Chevron Corp	.301	344	Petroleum & Coal Mfg.
Berkshire Hathaway Inc	.3	676	Industries not classified
Lockheed Martin Corp	.299	124	Transp. Equipment Mfg.
Boeing Co	.298	114	Transp. Equipment Mfg.
Sherwin-Williams Co (The)	.296	61	Chemical Mfg.
Activision Blizzard Inc	.295	60	Publishing Industries (exc. Internet)
Mondelez International Inc	.292	91	Food Mfg.
Northrop Grumman Corp	.291	84	Computer & Electronic Prod. Mfg.
Tesla Inc	.283	390	Transp. Equipment Mfg.
Norfolk Southern Corp	.272	56	Rail Transp.
Home Depot Inc. (The)	.261	321	Build. Mat., Gard. Equip., Supplies Dealers
United Parcel Service Inc	.256	149	Couriers & Messengers
CSX Corp	.256	64	Rail Transp.
Union Pacific Corp	.253	127	Rail Transp.
Costco Wholesale Corp	.252	202	General Merchandise Stores
TJX Companies Inc (The)	.243	92	Clothing & Clothing Accessories Stores
Lowe's Cos Inc	.238	120	Build. Mat., Gard. Equip., Supplies Dealers
Walmart Inc	.235	382	General Merchandise Stores
Target Corp	.235	69	General Merchandise Stores
McDonald's Corp	.194	222	Food Svcs & Drinking Places
Starbucks Corp	.119	110	Food Svcs & Drinking Places

Table IA.3
Generative AI Exposure Scores by NAICS 3-digit Industry

NAICS Subsector	Industry Title	Exposure Score
518	Data Processing, Hosting, and Related Svcs	.42
313	Textile Mills	.42
541	Professional, Scientific, and Technical Svcs	.41
511	Publishing Industries (exc. Internet)	.41
561	Admin. and Support Svcs	.41
517	Telecommunications	.4
488	Support Activities for Transp.	.4
519	Other Information Svcs	.4
111	Agriculture, Forestry, Fishing and Hunting	.39
423	Merchant Wholesalers, Durable Goods	.39
611	Educational Svcs	.39
323	Printing and Related Svcs	.38
314	Textile Product Mills	.38
332	Fabricated Metal Product Mfg.	.38
326	Plastics and Rubber Product Mfg.	.38
334	Computer and Electronic Prod. Mfg.	.38
333	Machinery Mfg.	.38
322	Paper Mfg.	.37
325	Chemical Mfg.	.37
335	Electrical Equipment, Appliance, and Component Mfg.	.36
532	Rental and Leasing Svcs	.36
424	Merchant Wholesalers, Nondurable Goods	.36
522	Credit Intermed. and Rel. Activ.	.36
327	Nonmetallic Mineral Product Mfg.	.36
454	Nonstore Retailers	.35
721	Accommodation	.35
515	Broadcasting (except Internet)	.35
562	Waste Mgmt and Remed. Svcs	.35
443	Electronics and Appliance Stores	.35
339	Miscellaneous Mfg.	.35
621	Ambulatory Health Care Svcs	.35
221	Utilities	.35
484	Truck Transp.	.34
713	Amusement, Gambling, and Recreation Ind.	.34
999	Industries not classified	.34
312	Beverage and Tobacco Product Mfg.	.34
321	Wood Product Mfg.	.33
315	Apparel Mfg.	.33
238	Specialty Trade Contractors	.33
211	Oil and Gas Extraction	.33
331	Primary Metal Mfg.	.33
486	Pipeline Transp.	.33
336	Transp. Equipment Mfg.	.33
316	Leather and Allied Product Mfg.	.33
337	Furniture and Related Product Mfg.	.32
324	Petroleum and Coal Mfg.	.32
512	Motion Picture and Sound Recording Industries	.32
212	Mining and Quarrying (except Oil and Gas)	.32
812	Personal and Laundry Svcs	.32
311	Food Mfg.	.31
711	Performing Arts, Spectator Sports, and Rel. Ind.	.31
622	Hospitals	.3
483	Water Transp.	.3
236	Construction of Buildings	.3
237	Heavy and Civil Engineering Construction	.29
445	Food and Beverage Stores	.29
623	Nursing and Residential Care Facilities	.29
444	Building Material and Garden Equipment and Supplies Dealers	.29
446	Health and Personal Care Stores	.28
441	Motor Vehicle and Parts Dealers	.28
492	Couriers and Messengers	.28
485	Transit and Ground Passenger Transp.	.28
213	Support Activities for Mining, and Oil and Gas Extraction	.28
453	Miscellaneous Store Retailers	.27
482	Rail Transp.	.27
481	Air Transp.	.26
451	Sporting Goods, Hobby, and Musical Instrument Stores	.25
442	Furniture and Home Furnishing Stores	.24
722	Food Svcs and Drinking Places	.24
452	General Merchandise Stores	.24
448	Clothing and Clothing Accessories Stores	.24

Table IA.4
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 II for variable definitions. Panel B controls for fixed effects at the NAICS 3-digit level. All variables are winsorized at the top and bottom 1%. Standard errors are clustered at the industry level and reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% level, respectively.

Panel A: Across All Firms						
	(1)	(2)	(3)	(4)	(5)	(6)
Log Size	-2.65** (-2.212)					
Tobin's Q		3.08*** (3.854)				
ROA			-21.52* (-1.938)			
Labor Intensity				7.89*** (3.395)		
Tangibility					-89.93*** (-4.290)	
Mkt. Leverage						-61.87*** (-4.796)
Observations	2,517	2,380	2,513	2,387	2,515	2,379
R-squared	0.007	0.013	0.006	0.038	0.108	0.033
Panel B: Within-Industry						
	(1)	(2)	(3)	(4)	(5)	(6)
Log Size	-0.73 (-0.715)					
Tobin's Q		0.52 (0.792)				
ROA			-7.30 (-0.722)			
Labor Intensity				5.59*** (3.588)		
Tangibility					-61.79** (-2.464)	
Mkt. Leverage						-22.96*** (-3.237)
Observations	2,517	2,380	2,513	2,387	2,515	2,379
R-squared	0.305	0.317	0.303	0.336	0.321	0.320

Table IA.5**Reaction of Portfolios Sorted by Generative AI Exposure: Shorter Event Period**

This table reports daily percentage stock returns, as a measure of changes in firm value, during the ChatGPT event period for five value-weighted portfolios sorted on firms' Generative AI exposure. To form the portfolios, we sort NYSE stocks on November 29, 2022 into five quintiles by their Generative AI exposure, and we assign all stocks into the portfolios using the NYSE breakpoints. We aggregate stock returns within each portfolio using the stocks' market capitalization on the prior trading day as the weight. *AMH* is the "Artificial Minus Human" zero net investment portfolio long the portfolio with high exposure and short the portfolio with low exposure. We define the shorter ChatGPT event period here as from November 29, 2022 to December 7, 2022, i.e., one day before to 1 week after the release of ChatGPT (see Section II for more details). Panel A reports each portfolio's raw daily return in excess of the risk free rate. Panels B and C report the CAPM market-adjusted alphas and the Fama-French 5-factor-adjusted alphas, where the factor loadings of each portfolio are estimated using data from the previous six months. Panel D reports the Fama-French 5-factor-adjusted alphas for a sample that excludes firms in the tech sectors, NAICS 51 (information) and NAICS 54 (business services), following [Acemoglu et al. \(2022\)](#). *t*-statistics are shown in parentheses. See Section II.B for more details.

Q1 (H)	Q2	Q3	Q4	Q5 (A)	AMH
<i>A: Excess returns (%)</i>					
-0.246 (-0.49)	-0.253 (-0.52)	-0.077 (-0.15)	-0.087 (-0.17)	0.002 (0.00)	0.248 (1.89)
<i>B: Market-factor-adjusted alpha (%)</i>					
-0.117 (-2.40)	-0.120 (-1.93)	0.056 (0.69)	0.038 (1.12)	0.148 (2.83)	0.265 (2.96)
<i>C: FF 5-factor-adjusted alpha (%)</i>					
-0.080 (-1.87)	0.073 (0.72)	-0.042 (-0.83)	0.013 (0.31)	0.138 (3.02)	0.218 (2.76)
<i>D: FF 5-factor-adjusted alpha ex. tech sector (%)</i>					
-0.085 (-1.86)	0.082 (0.79)	-0.035 (-0.79)	0.057 (1.54)	0.191 (2.83)	0.276 (3.59)

Table IA.6**Reaction of Portfolios Sorted by Generative AI Exposure: Longer Event Period**

This table reports daily percentage stock returns, as a measure of changes in firm value, during the ChatGPT event period for five value-weighted portfolios sorted on firms' Generative AI exposure. To form the portfolios, we sort NYSE stocks on November 29, 2022 into five quintiles by their Generative AI exposure, and we assign all stocks into the portfolios using the NYSE breakpoints. We aggregate stock returns within each portfolio using the stocks' market capitalization on the prior trading day as the weight. *AMH* is the "Artificial Minus Human" zero net investment portfolio long the portfolio with high exposure and short the portfolio with low exposure. We define the *longer* ChatGPT event period here as from November 29, 2022 to December 21, 2022, i.e., one day before to 3 weeks after the release of ChatGPT (see Section II for more details). Panel A reports each portfolio's raw daily return in excess of the risk free rate. Panels B and C report the CAPM market-adjusted alphas and the Fama-French 5-factor-adjusted alphas, where the factor loadings of each portfolio are estimated using data from the previous six months. Panel D reports the Fama-French 5-factor-adjusted alphas for a sample that excludes firms in the tech sectors, NAICS 51 (information) and NAICS 54 (business services), following [Acemoglu et al. \(2022\)](#). *t*-statistics are shown in parentheses. See Section II.B for more details.

Q1 (H)	Q2	Q3	Q4	Q5 (A)	AMH
<i>A: Excess returns (%)</i>					
-0.289 (-1.34)	-0.128 (-0.56)	-0.152 (-0.57)	-0.091 (-0.38)	-0.027 (-0.09)	0.261 (2.38)
<i>B: Market-factor-adjusted alpha (%)</i>					
-0.151 (-2.82)	0.015 (0.18)	-0.010 (-0.17)	0.042 (0.90)	0.128 (2.34)	0.279 (3.09)
<i>C: FF 5-factor-adjusted alpha (%)</i>					
-0.132 (-2.55)	0.095 (2.17)	-0.053 (-2.17)	0.033 (0.81)	0.124 (3.25)	0.255 (3.53)
<i>D: FF 5-factor-adjusted alpha ex. tech sector (%)</i>					
-0.137 (-2.59)	0.111 (2.44)	-0.041 (-1.23)	0.034 (0.78)	0.132 (1.78)	0.269 (2.75)

Table IA.7**Reaction of Portfolios Sorted by Generative AI Exposure: Non-event Period**

This table reports daily percentage stock returns, as a measure of changes in firm value, in the days surrounding, but outside of, the ChatGPT event period for five value-weighted portfolios sorted on firms' Generative AI exposure. To form the portfolios, we sort NYSE stocks on November 29, 2022 into five quintiles by their Generative AI exposure, and we assign all stocks into the portfolios using the NYSE breakpoints. We aggregate stock returns within each portfolio using the stocks' market capitalization on the prior trading day as the weight. *AMH* is the "Artificial Minus Human" zero net investment portfolio long the portfolio with high exposure and short the portfolio with low exposure. We define the ChatGPT event period as from November 30, 2022 to December 14, 2022, i.e., the two weeks after the release of ChatGPT (see Section II for more details). The table then shows placebo event study returns for the ten trading days before (Nov. 14- Nov. 29, 2022) and after (Dec 15 - Dec 29, 2022) the ChatGPT-release event period. Panel A reports each portfolio's raw daily return in excess of the risk free rate. Panels B and C report the CAPM market-adjusted alphas and the Fama-French 5-factor-adjusted alphas, where the factor loadings of each portfolio are estimated using data from the previous six months. Panel D reports the Fama-French 5-factor-adjusted alphas for a sample that excludes firms in the tech sectors, NAICS 51 (information) and NAICS 54 (business services), following [Acemoglu et al. \(2022\)](#). *t*-statistics are shown in parentheses. See Section II.B for more details.

Q1 (H)	Q2	Q3	Q4	Q5 (A)	AMH
<i>A: Excess returns (%)</i>					
-0.140 (-0.67)	-0.177 (-0.92)	-0.290 (-1.09)	-0.110 (-0.56)	-0.232 (-0.93)	-0.092 (-0.88)
<i>B: Market-factor-adjusted alpha (%)</i>					
0.062 (0.65)	0.032 (0.50)	-0.081 (-1.08)	0.087 (2.09)	-0.004 (-0.14)	-0.066 (-0.72)
<i>C: FF 5-factor-adjusted alpha (%)</i>					
0.069 (0.73)	0.049 (1.12)	-0.129 (-3.12)	0.071 (2.06)	0.019 (0.71)	-0.050 (-0.51)
<i>D: FF 5-factor-adjusted alpha ex. tech sector (%)</i>					
0.065 (0.68)	0.047 (0.94)	-0.151 (-3.78)	0.041 (1.24)	0.023 (0.60)	-0.042 (-0.50)

Table IA.8**Within-industry Portfolios Sorted by Generative AI Exposure: Non-event Period**

This table reports daily percentage stock returns, as a measure of changes in firm value, in the days surrounding, but outside of, the ChatGPT event period for three value-weighted portfolios sorted on firms' Generative AI exposure within the industry. To form the portfolios, we sort NYSE stocks on November 29, 2022 into three terciles by their Generative AI exposure across all stocks in Panels A1 and B1, within the NAICS 3-digit industry in Panels A2 and B2, and within the [Hoberg and Phillips \(2016\)](#) FIC 50-industry in Panels A3 and B3, and we assign all stocks into the portfolios using the NYSE breakpoints. We aggregate stock returns within each portfolio using the stocks' market capitalization on the prior trading day as the weight. *AMH* is the "Artificial Minus Human" zero net investment portfolio long the portfolio with high exposure and short the portfolio with low exposure. We define the ChatGPT event period as from November 30, 2022 to December 14, 2022, i.e., the two weeks after the release of ChatGPT (see Section II for more details). The table then shows placebo event study returns for the ten trading days before (Nov. 14- Nov. 29, 2022) and after (Dec 15 - Dec 29, 2022) the ChatGPT-release event period. Panels A1-A3 report the CAPM market-adjusted alphas, and Panels B1-B3 report the Fama-French 5-factor-adjusted alphas. *t*-statistics are shown in parentheses. See Section II.C for more details.

Q1 (H)	Q2	Q3 (A)	AMH
<i>A1: Market-factor-adjusted alpha (%)</i>			
0.035 (0.53)	-0.019 (-0.39)	0.029 (1.25)	-0.006 (-0.08)
<i>A2: NAICS 3-digit ind.-neutral market-factor-adjusted alpha (%)</i>			
-0.013 (-0.35)	0.069 (1.51)	0.020 (0.80)	0.033 (0.78)
<i>A3: FIC 50 ind.-neutral market-factor-adjusted alpha (%)</i>			
-0.005 (-0.09)	0.048 (2.11)	0.052 (1.91)	0.057 (0.85)
<i>B1: FF 5-factor-adjusted alpha (%)</i>			
0.062 (1.33)	-0.060 (-2.30)	0.037 (1.58)	-0.024 (-0.41)
<i>B2: NAICS 3-digit ind.-neutral FF 5-factor-adjusted alpha (%)</i>			
0.009 (0.32)	0.089 (2.80)	0.022 (0.83)	0.014 (0.30)
<i>B3: FIC 50 ind.-neutral FF 5-factor-adjusted alpha (%)</i>			
0.015 (0.24)	0.021 (0.77)	0.070 (1.89)	0.054 (0.83)

Table IA.9
Product AI Exposure Proxies and Reported AI Use

These regressions show the relationship between different proxies for industries' product exposure to AI and the reported AI use by firms in the industry. The AI use in production by industry is measured as the average rate of "yes" responses to the question *"In the last two weeks, did this business use Artificial Intelligence (AI) in producing goods or services?"* by 3-digit NAICS subsector in the U.S. Census Bureau Business Trends and Outlook Survey from October 23, 2023. The proxies for product exposure to AI are aggregated from firm-level measures as employment-weighted means by NAICS subsector. The definitions of the four proxies for firms' product exposure to AI are described in detail in Section II.C.3. *GPT10 Product Exp* is based on GPT 3.5 Turbo assessment of firms' annual reports' description of their products. *Count10K Product Exp* is based on counting AI-related keywords in firms' annual reports. *GS Product Exp* is based on Goldman Sachs' classification of firms as "near-term beneficiaries of AI". *BFHH Product Exp* is based on firms' share of AI-skilled workers constructed by [Babina et al. \(2024\)](#). All product exposure measures are standardized to have a mean of zero and a standard deviation of one, except *GS Product Exp* and *GPT10K Product Exposure*, which are binary variables. The last column includes all four measures at the same time. t-statistics based on heteroskedasticity-robust standard errors are shown in parentheses. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

Dep. var.:	AI Use in Last 2 Wks %, Oct. 2023 (BTOS Survey)				
	(1)	(2)	(3)	(4)	(5)
GPT10K Product Exp	18.919*** (3.609)				19.460*** (2.807)
Count10K Product Exp		1.262** (2.229)			-0.506 (-0.947)
GS Product Exp			11.795* (1.916)		-0.790 (-0.216)
BFHH Product Exp				1128.449** (2.631)	588.859 (1.459)
R-squared	0.50	0.33	0.18	0.35	0.55
Observations	57	57	57	57	57

Table IA.10
Cumulative Abnormal Returns and Generative AI Exposure: Other Abnormal Return Models

This table reports the relation between firms' cumulative abnormal returns during the ChatGPT event period, $[-1, 10]$ trading days around the release of ChatGPT on November 30, 2022, (CAR $[-1, 10]$) and firms' labor-based Generative AI exposure defined in Section I. To compute the baseline cumulative abnormal return of each stock, we first subtract the daily market return from Kenneth French's website from the daily stock return from Yahoo Finance, and then we cumulate the daily abnormal returns over $[-1, 10]$. Columns (1)-(2) show abnormal returns as defined in the baseline model where returns are adjusted by risk free rates. Columns (3)-(4) compute abnormal returns relative to the CAPM model, and columns (5)-(6) relative to the Fama French 5-factor model. These abnormal returns are computed by first estimating factor loadings for each firm separately either for the market factor or the Fama French 5-Factor model using data for the 6 months from May 18, 2022 - Nov 14, 2022 (ending two weeks before the GPT release), and then computing abnormal returns for later periods as the difference between the raw excess returns and the returns predicted by the factor loadings. Where indicated, regressions include NAICS 3-digit fixed effects. Equation (5) describes the regression specification and the weighting of firms. *t*-statistics are computed using heteroskedasticity-robust standard errors and reported in parentheses. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

Dep. var.:		CAR[-1,10]					
Abnormal returns:		Baseline		CAPM		FF5	
		(1)	(2)	(3)	(4)	(5)	(6)
GenAI Exp		0.216*** (2.621)	0.288*** (2.770)	0.229*** (2.817)	0.284*** (2.693)	0.158** (1.961)	0.278** (2.195)
R-squared		0.05	0.36	0.06	0.35	0.03	0.26
Observations		2085	2080	2085	2080	2085	2080
NAICS 3-dig. FE		X		X		X	

Table IA.11
Correlations between Firm-Level Generative AI exposure, AI Product Market Exposure, and Data Assets

This table reports the correlation between the firm-level Generative AI exposure measure defined in Section I.D, the different product market exposure (“PE”) measures defined in Section II.C.3, and the data asset (“DA”) measures defined in Section II.D.

	GenAI Exp	GPT10K PE	Count10K PE	GS PE	BFHH PE	10K DA	AV DA
GenAI Exp	1.00						
GPT10K PE	0.16	1.00					
Count10K PE	0.12	0.28	1.00				
GS PE	0.05	0.16	0.21	1.00			
BFHH PE	0.15	0.18	0.42	0.20	1.00		
10K DA	0.21	0.38	0.17	0.14	0.13	1.00	
AV DA	0.58	0.36	0.28	0.19	0.31	0.43	1.00

Table IA.12
Generative AI Exposure and Firm Data Assets and Firm Size

This table reports coefficients from regressions of cumulative abnormal returns of firms during the ChatGPT event period, $[-1, 10]$ trading days around the release of ChatGPT on November 30, 2022, on firms' Generative AI exposure interacted with proxies for their data assets and also interacted with proxies for firm size in the form of quintiles of total firm employment in Compustat as of October 2022. The largest firm size quintile is omitted as the reference category. To compute the cumulative abnormal return of each stock, we first subtract the daily market return from Kenneth French's website from the daily stock return from Yahoo Finance, and then we cumulate the daily abnormal returns over the $[-1, 10]$ period. Section II.D details the construction of the two proxies for firms' data assets. *10K Data Assets* is a score of 0, 1, 2, or 3, indicating the firm is not likely, slightly likely, moderate likely, or highly likely to have data that can be used as an input into large language model analytics based on the assessment of firms' annual reports using the GPT 3.5 Turbo Model. *AV Data Assets* is constructed following [Abis and Veldkamp \(2023\)](#) and measures the intensity of data management skills in the firm's workforce. All regressions include NAICS 3-digit fixed effects and employment quintile fixed effects, and weight firms by their market capitalization as of November 29, 2022. *t*-statistics are computed using heteroskedasticity-robust standard errors and reported in parentheses. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

Dep. var.:	CAR[-1,10]		
	(1)	(2)	(3)
GenAI Exp	0.231** (2.564)	-0.018 (-0.202)	-0.073 (-0.803)
GenAI Exp \times 10K Data Assets		0.216** (2.159)	
GenAI Exp \times AV Data Assets			12.594** (2.424)
10K Data Assets		-0.073* (-1.846)	
AV Data Assets			-4.052* (-1.736)
GenAI Exp \times 1[Quint1 Total Emp.]	0.107 (0.579)	0.311 (1.500)	0.323* (1.670)
GenAI Exp \times 1[Quint2 Total Emp.]	-0.274 (-0.800)	-0.295 (-0.814)	-0.146 (-0.383)
GenAI Exp \times 1[Quint3 Total Emp.]	0.015 (0.124)	0.073 (0.606)	0.079 (0.642)
GenAI Exp \times 1[Quint4 Total Emp.]	-0.182 (-1.536)	-0.196 (-1.630)	-0.208* (-1.706)
R-squared	0.06	0.11	0.13
Observations	2,069	1,899	2,028
NAICS 3-dig. FE	X	X	X
Quint. of Employment FE	X	X	X

Table IA.13
Future Revenue Impact and Generative AI Exposure

This table reports the impact of the release of ChatGPT on analyst forecasts of firms' future revenues and on firms' reported quarterly revenues. Analyst forecasts are updated at a monthly frequency for each firm in the I/B/E/S data. For each firm in each month, we obtain the median analyst forecasts of the firm's revenues in the fiscal year ending in December 2023 in billions and the firm's long-term annual percentage growth rate in revenues (LTG (Revenues)). The dependent variables in Columns (1) and (2) are the firm's median forecasts of the corresponding measures made in October 2022 and January 2023. The dependent variable in Column (3) is the firm's quarterly revenues in billions in the Compustat quarterly data from 2022Q1 to 2024Q3. The independent variable *Post* is an indicator for whether the time period is after the release of ChatGPT (i.e., after November 2022). *GenAI Exp* is a continuous measure of GenAI exposure at the firm level. See Section III for more details. All regressions include both firm-level and time-period fixed effects. *t*-statistics in parentheses are computed using standard errors clustered at the firm level. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

<i>Data:</i>	I/B/E/S Analyst Forecasts		Compustat
<i>Measure:</i>	Revenues 2023	LTG (Revenues)	Revenues
	(1)	(2)	(3)
Post \times GenAI Exp	0.508 (0.807)	0.825 (0.323)	0.171 (0.953)
Observations	1,718	372	22,279
<i>Addl. controls & fixed effects</i>			
Time FEs	X	X	X
Firm FEs	X	X	X

Table IA.14
Firm Hiring of Occupations by Exposure

This table reports the impact of the release of ChatGPT on firms' job postings for occupations with different exposures to Generative AI. The sample includes the high/medium/low-exposure (top/medium/bottom tercile) job postings at the firm-month level from January 2022 to August 2023 aggregated from the Lightcast (formerly Burning Glass) job posting database. *ShareSupp* is the firm's share of Generative AI exposure deriving from the supplemental tasks' exposure to Generative AI. The job posting panels include many zeros, which makes it impractical to log-transform the dependent variable. Instead, we estimate all regressions using the Poisson model with fixed effects following [Cohn et al. \(2022\)](#) which means that coefficients can be interpreted as semi-elasticities, i.e., in terms of percentage changes in the dependent variable. *t*-statistics in parentheses are computed using standard errors double clustered at the month level and the firm level. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

Dep var.:	Job Postings for High-Exposure Occ		Job Postings for Medium-Exposure Occ		Job Postings for Low-Exposure Occ	
	(1)	(2)	(3)	(4)	(5)	(6)
Post \times GenAI Exp	-1.378*** (-3.747)	-3.038*** (-2.991)	-1.007*** (-2.783)	-0.611 (-0.596)	-0.534 (-1.427)	-0.745 (-0.791)
Post \times ShareSupp		-2.302** (-2.029)		0.463 (0.399)		-0.149 (-0.159)
Post \times GenAI Exp \times ShareSupp		8.100** (2.037)		-0.694 (-0.160)		1.926 (0.635)
Observations	36,900	36,880	34,720	34,700	30,100	30,080
Month FE	X	X	X	X	X	X
Firm FE	X	X	X	X	X	X

Table IA.15
Firm Hiring Effects Controlling for Industry Heterogeneity

This table reports the impact of the release of ChatGPT on firms' job postings for high-exposure occupations with and without controlling for industry heterogeneity. The sample includes the high-exposure (top tercile) job postings at the firm-month level from January 2022 to August 2023 aggregated from the Lightcast (formerly Burning Glass) job posting database. *GenAI Exp* is the firm-level Generative AI exposure, which is an employment-weighted average across occupations within the firm. *ShareSupp* is the firm's share of Generative AI exposure deriving from the supplemental tasks' exposure to Generative AI. The job posting panels include many zeros, which makes it impractical to log-transform the dependent variable. Instead, we estimate all regressions using the Poisson model with fixed effects following [Cohn et al. \(2022\)](#) which means that coefficients can be interpreted as semi-elasticities, i.e., in terms of percentage changes in the dependent variable. *t*-statistics in parentheses are computed using standard errors double clustered at the month level and the firm level. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

	(1)	(2)	(3)	(4)
Post × GenAI Exp	-1.378*** (-3.747)	-1.337*** (-2.836)	-3.038*** (-2.991)	-2.641*** (-3.019)
Post × GenAI Exp × ShareSupp			8.100** (2.037)	7.315** (2.079)
Post × ShareSupp			-2.302** (-2.029)	-2.108** (-2.021)
Firm FE	X	X	X	X
Month FE	X		X	
Industry-Month FE		X		X
Observations	36,900	36,600	36,880	36,580

Table IA.16
Cumulative Abnormal Returns and Generative AI Exposure with Interactions

This table reports the relation between firms' cumulative abnormal returns during the ChatGPT event period, $[-1, 10]$ trading days around the release of ChatGPT on November 30, 2022, (CAR $[-1, 10]$) and firms' Generative AI exposure, defined in Section I. To compute the cumulative abnormal return of each stock, we first subtract the daily market return from Kenneth French's website from the daily stock return from Yahoo Finance, and then we cumulate the daily abnormal returns over $[-1, 10]$. See Table II for the definitions of firm characteristic variables. All firm characteristics other than generative AI exposure have been standardized to have a mean of zero and a standard deviation of one. Equation (4) describes the regression specification and the weighting of firms. t -statistics are computed using heteroskedasticity-robust standard errors and reported in parentheses. ***, **, and * denote significance at 1%, 5%, and 10% levels, respectively.

Dep. var.:	CAR[-1,10]						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GenAI Exp	0.216*** (2.621)	0.252*** (2.768)	0.158*** (2.937)	0.208*** (2.724)	0.269*** (2.952)	0.205** (2.219)	0.187*** (2.694)
GenAI Exp \times Log Size		0.098 (1.549)					
GenAI Exp \times Tobin's Q			0.282*** (3.312)				
GenAI Exp \times ROA				0.093* (1.722)			
GenAI Exp \times Labor Intensity					0.174** (-2.489)		
GenAI Exp \times Tangibility						0.079 (1.210)	
GenAI Exp \times Mkt. Leverage							-0.228*** (-2.814)
R-squared	0.05	0.09	0.16	0.08	0.08	0.06	0.13
Observations	2085	2078	2078	2073	2054	2069	2073

IA.3 Methodology Notes

C.1 GPT prompt for task-based exposure scoring

The following prompt structure was based on the rubric language by [Eloundou et al. \(2023\)](#), as well as insights by [Willison \(2023\)](#) and [Underwood \(2023\)](#) about how to best structure API calls for GPT classification. We conducted this categorization on March 29, 2023. To rule out misclassification due to the model’s confusion of ChatGPT’s known capabilities with ChatGPT’s in-development but not well-known capabilities such as image recognition, we follow [Eloundou et al. \(2023\)](#) and add a fourth category (E3) in addition to the main categories used in our analysis, indicating that the task is not exposed to ChatGPT, either directly or by software integration, but can be completed more efficiently if image capabilities of ChatGPT were accessible. Only 0.6% of the tasks are categorized as E3. We consider tasks that are exposed to image recognition to be not exposed throughout our study. Here are the instruction prompts submitted before asking GPT 3.5 Turbo to classify each task statement (using the version as of March 29th, 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” parameter 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.”

C.2 Consistency of Generative AI scoring

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

Table IA.C.1
Exposure score variation across GPT scoring runs

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

C.3 Validation of Generative AI scoring

How reliable is the approach of classifying task statements using an LLM rather than human research assistants? The benefits of LLM-scoring in terms of costs and speed are clear (the classification task was accomplished in less than two days using an LLM while it would have taken a human research assistant at least 200 hours based on our estimates, and that would be without producing explanations for the assigned score). How does the labeling by an LLM compare in terms of “quality”? Note that for the tasks in question it is not the correct question whether an LLM would produce fully correct labels, but whether a human research assistant that could feasibly be hired for the task would have a better ability to produce correct labels.

The ability of human annotators is particularly relevant as the scoring of $> 19K$ tasks that span the entirety of U.S. occupations requires an exceptionally experienced annotator to fully be able to comprehend and contextualize the description of *any* possible occupation’s activities. We consider this as part of the advantages of using a state-of-the-art LLM to do this task: the breadth of occupations represented in the training texts used for LLMs is likely to far exceed the occupational contexts that any human annotator could interpret with confidence. As a result, the LLM could reasonably be considered more of an “expert” for this task than the human research assistants who would otherwise be employed in such a task (usually undergraduate or graduate students).

One validation of how an LLM’s scoring of tasks as exposed to Generative AI technologies compares to human labeling of the same tasks is already provided by [Eloundou et al. \(2023\)](#), who found that, depending on the exact measure of exposure one is considering, human labels agreed with the GPT-produced labels in 65.6-82.1% of labels (see Table 2 in that paper). While, the authors of that paper included researchers at OpenAI and therefore had access to “experienced human annotators who have reviewed GPT-3, GPT-3.5 and GPT-4 outputs as part of OpenAI’s alignment work,” we do not have access to similarly skilled annotators.

To provide further validation of this method, we recruited two research assistants (one with a graduate degree and one without a college degree) to label a random sample of tasks with the same instructions as were provided to the LLM, and also had one of the authors of this paper label the same set of tasks as an “experienced” annotator. We find agreement between these labelers that looks as shown in Appendix Table [IA.C.2](#) (which is similar to the validation table shown in [Eloundou et al. \(2023\)](#)).

Table IA.C.2
Agreement between different annotators

This table shows agreement between annotators of the same random sample of 100 tasks from the full set of task statements on O*NET. Agreement is defined as the % of tasks that receive the same score for two different score definitions: (1) Column 1 only considers whether there was agreement on whether the “direct exposure” label E1 applied or not. (2) Column 2 shows agreement on whether any exposure (direct OR indirect) was assigned to the task.

Agreement measure (%):	E1 only	E1 or E2
GPT vs. Experienced	81	86
GPT vs. RA # 1	80	72
GPT vs. RA # 2	78	81
Experienced vs. RA # 1	79	74
Experienced vs. RA # 2	75	83
RA # 1 vs. RA # 2	82	73

As the table shows, we find very similar agreements between the LLM and our human annotators as [Eloundou et al. \(2023\)](#), ranging from 72% to 86% agreement. If we just focus on whether the annotator determined *direct* Generative AI exposure, we find 78-81% agreement.

Variation across human labelers. We can also compare the human annotators to one another. As Table IA.C.2 shows, there is substantial variation in labels *between* human annotators. Our human annotators have similar levels of agreement with one another as they do with the GPT labels, ranging from 73 to 83% agreement across the different exposure measures, and 79-82% when determining if there is *any* Generative AI exposure. While this level of idiosyncratic variation across human annotators might be in part due to the fact that we do not have access to highly skilled occupational analysts, it is likely reflective of the annotator skill level that would normally be available for this type of labeling task (albeit at much higher expense of time, money, and research assistant welfare) if LLMs could not be used. Comparing this agreement across human labelers to the consistency between different LLM runs with random variation in prompting shown in Table IA.C.1, this suggests that LLMs may be less variable in their output than human annotators would be and might thus enable better replicability of results.

C.4 [Acemoglu and Autor \(2011\)](#) measures of skill requirements of occupations

For the analysis shown in Figure 2, we need measures of the skill requirements of different occupations. We draw on standard measures from the literature that allow for comparability

of our results to the characteristics of previous waves of automation.

Acemoglu and Autor (2011) construct six skill measures for each SOC occupation based on O*Net measures in the following steps. First, they assign detailed skill requirements from the O*Net's database to each of the six aggregated skill measures, using the following O*Net measurements for each occupation:

Non-routine cognitive: analytical skill:

4.A.2.a.4 Analyzing data/information

4.A.2.b.2 Thinking creatively

4.A.4.a.1 Interpreting information for others

Non-routine cognitive: interpersonal skill:

4.A.4.a.4 Establishing and maintaining personal relationships

4.A.4.b.4 Guiding, directing and motivating subordinates

4.A.4.b.5 Coaching/developing others

Routine cognitive skill:

4.C.3.b.7 Importance of repeating the same tasks

4.C.3.b.4 Importance of being exact or accurate

4.C.3.b.8 Structured v. Unstructured work (taking the reverse value)

Routine manual skill:

4.C.3.d.3 Pace determined by speed of equipment

4.A.3.a.3 Controlling machines and processes

4.C.2.d.1.i Spend time making repetitive motions

Non-routine manual: physical skill

4.A.3.a.4 Operating vehicles, mechanized devices, or equipment

4.C.2.d.1.g Spend time using hands to handle, control or feel objects, tools or controls

1.A.2.a.2 Manual dexterity

1.A.1.f.1 Spatial orientation

2.B.1.a Social perceptiveness

Second, they obtain an importance scale of each detailed skill requirement for each SOC8 occupation, and standardize the importance scale of each detailed skill to have a mean of 0 and a standard deviation of 1 across occupations.

Third, they compute an occupation's six skill measures as the average of the standardized importance scales of their corresponding detailed skill measures.

C.5 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 the NYSE as of the sorting date. All portfolios are formed based on equal weighted sorts on November 29, 2022, and weights for computing portfolio returns are adjusted based on daily returns to mimic passive buy-and-hold exposure.

Industry-neutral portfolios are computed by first forming within-industry equal-weighted tercile portfolios, and then averaging portfolio returns for the same terciles across industries. Returns for within-industry terciles and for all global (not industry-neutral) portfolio quintile sorts are value-weighted, while across-industry averages are industry market-cap. weighted. *AMH* is the "Artificial Minus Human" zero net investment portfolio long highest exposure quantile (*A*) stocks and short lowest exposure quantile (*H*) stocks. The data set for estimating portfolio returns consists of daily stock returns from Yahoo Finance and the Fama-French factors, including the market factor and risk free returns are obtained from Ken French's website.

We define the ChatGPT event period as from November 30, 2022 to December 14, 2022, i.e., the two weeks after the release of ChatGPT (see Section II for more details). CAPM market-adjusted alphas and the Fama-French 5-factor-adjusted alphas are based on the factor loadings of each portfolio estimated using data from the six months preceding Nov. 15, 2022. Abnormal returns are computed as the difference between raw portfolio returns and the product of factor loadings and the factor returns on each day of the event period. Alphas for the event period are then computed as the intercept in a regression of abnormal returns over the event period on a constant, with Newey-West standard errors with five lags.

C.6 Product Market AI Exposure Measures

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

Company annual report data. The 10-K annual reports filed by companies at the SEC’s EDGAR system are obtained in pre-cleaned text files from Bill McDonald’s “Software Repository for Accounting and Finance” website,⁵⁶ based on the work in [Loughran and McDonald \(2016\)](#). We then use regular expressions to break up the text into the different “items” contained within the report.⁵⁷, focusing only on reports filed in 2022, and on the “Business” section of each report.

AI-related business description keywords. For our first product market AI exposure measures we follow [Hoberg and Phillips \(2016\)](#) and use the “Business” description section of a firm’s 10-K annual report to infer information about its product markets. We construct a simple measure of AI relevance for the firm’s business by tokenizing the text and counting the number of mentions of “AI” or “artificial intelligence”. We use the total count of mentions as a proxy for a firm’s products either using AI or depending on the use of it by other actors in the value chain. We do this by first lower-casing and tokenizing the text in the “Business” section, constructing lists of all possible unigrams and bigrams contained in the text, and then counting the occurrences of “ai” and “artificial intelligence” in the text.

GPT assessment of business product exposure to AI. The method of counting keywords potentially discards relevant information contained in the full text and the context and interdependence of the AI-related concepts discussed in it. Therefore, we also use a method that uses a large language model to evaluate the text of the business description in the firm’s annual report, asking it to determine whether there would be direct positive product market impact of a Generative AI boom on the described company. This allows us to keep the text in its original form and take into account the context within which AI or related topics are discussed. We develop a rubric that provides guidance to the model in the form of asking it to consider whether the firm’s products might be involved in enabling or scaling AI technologies, or might benefit from a direct incorporation of the new AI capabilities (the full text of the prompt is shown below). The model is then given two examples of scores applied to company business descriptions and given a new business description and asked to apply a binary label of whether the firm is “directly product market exposed to AI” or not, and also to provide an explanation for the score, which allows for an audit of the model’s

⁵⁶See <https://sraf.nd.edu/>

⁵⁷We built on code provided by Yu Zhu at <https://yuzhu.run/how-to-parse-10x/> and coding support from ChatGPT in this analysis.

reasoning. To economize on computing resources and API costs, we only do this analysis for annual reports in our sample that belong to firms for which we have previously computed a task-based AI exposure score, which have stock price data in our sample, and for which the “Business” section contains the word “and” at least three times (which removes annual reports that omitted the “Business” section or where our text extraction parsed a snippet that is too small). We submit the prompt for scoring together with a “Business” section to OpenAI’s GPT 3.5 Turbo API for evaluation. We are limited by the context window to evaluating only the first 3000 words of the business description for each firm, which is rarely binding and allows for ample business description for almost all companies - as the beginning of the section tends to provide the general overview of the company. The result of this procedure is binary product market AI exposure scores for $\sim 2.2K$ companies. Moreover, in early trial runs we audited the explanations provided by the model to ensure that the prompt leads to scoring that closely corresponds to a human scorer’s interpretation of the business impact.

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

Resume AI skill share from [Babina et al. \(2024\)](#): we use a measure of the share of workers at a firm that have AI skills on their resumes constructed by [Babina et al. \(2024\)](#), who show that this measure is a good predictor of AI-related product innovations and R&D spending during the pre-ChatGPT wave of AI advancements. We use both the last available data point on the stock of AI skills (as a share of firm employment) as proxies for a firm’s level of investment in using AI-related tools pre-ChatGPT. We use the replication data available from the authors and follow their methodology to construct these regression variables based on the Cognism resume AI skill share data provided. The last year of available skill share data is 2021 for the majority of firms in the sample, and we only keep skill share level data if the last available data is from no earlier than 2018.

C.7 Product Market AI Exposure: GPT Prompt for Scoring

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

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

Prompt preceding the “Business” text submission”:

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

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

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

C.8 Measuring Firm “Data Value” for Generative AI

In order to measure the “data value” of a company that might contribute to its ability to productively deploy Generative AI tools and their analytics capabilities, we develop a

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

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

GPT assessment of business “data value” for LLMs. Again, we also use a method that uses a large language model to evaluate the text of the business description in the firm’s annual report. We develop a rubric that provides guidance to the model in the form of asking it to first consider subcategories of data relevance. In particular, the LLM is asked to first assess whether a firm’s business description suggests access to relevant data based on its coverage of 6 categories: the general nature of the company’s business, the scale and reach of the firm, data collection mechanisms, data utilization, data infrastructure & management, and data regulation and privacy (full prompt shown below). The model is asked to assign a score of 0, 1, 2, or 3 (no, little, moderate, or high relevance) in each category. Only then is the model asked to also provide an overall score for the degree to which a firm is likely to have data that can be used as an input into LLM analytics (low, moderate, or high data value).

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

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

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

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

C.9 Data Value Assessment: GPT Prompt for Scoring

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

```

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

```

Prompt preceding the “Business” text submission”:

```

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

```