AI Exposure without Labor Data:

Measuring AI's Impact through Firms*

Jiacheng Liu†

I present a new framework for investigating AI's impact by analyzing managerial discussions. This forward-looking approach captures *firm-specific* exposure to AI in near-real time without relying on job postings or resumes. To validate the framework, I show that AI-exposed firms subsequently increase their investments in AI-related labor, suggesting that textual information accurately reflects real resource allocation on average. The framework accommodates a range of topics and can estimate many otherwise hard-to-quantify variables of societal importance. I focus on two foundational agendas: namely, AI's impact on firms and GDP growth. In the cross-section, AI-exposed firms exhibit higher valuations, greater efficiency, and lower leverage relative to firms in the same industry. However, many firm characteristics predate the adoption of AI, suggesting that entrenched firm-level qualities influence the adoption of new technologies. Using a difference-in-differences design, I demonstrate that changes in the cost of capital and leverage are more likely linked to the adoption of AI. Guided by a theoretical model, I document that AI is currently most likely to increase GDP growth by complementing human workers and raising their productivity. Using a range of granular intra-firm data, I identify early signs of both labor displacement and productivity gains. Finally, I showcase the generalizability of this measure by assessing current progress toward AGI. This paper highlights the importance of incorporating firm-level information to better understand and align the economic aspects of AI.

*Acknowledgments: I especially thank my advisor, M. Deniz Yavuz, for his invaluable guidance. This paper would not have been possible without many insightful conversations with him. I am also very grateful to my committee members, Sergey Chernenko, Craig Brown, Peter G. Hansen, and Ha Diep-Nguyen for their continuous support. I thank Ralph Koijen, Kaushik Vasudevan, Mara Faccio, Adem Atmaz, Bryan Kelly, Sebastiano Manzan, Juan-Angel Jimenez-Martin, Kevin Mumford, Michael Woeppel, Siqi Guo, Nathaniel Feige, Zibo Zhou, Valeriia Yatsenko, Robin Luo, Xinyu Cao, Hao Yang (Discussant), Abigail Banan, Tim Seida, Qian Zhu, Qian Wang, Renyuan Luo, Jacob Chanyeol Choi, Jin Kim, Andrew Zhang, Cindy Pan, and Jared DeLisle, as well as participants at the Generative AI in Finance, the Workshop on LLMs and Generative AI for Finance in ACM International Conference on AI in Finance, the Inter-Finance PhD Seminar, Southern Finance Association, and the Krannert PhD Research Symposium for their helpful comments. An earlier version of the paper circulated under the title "Measuring the Economic Impact of AI through Forward-Looking Firm Communications." First Draft: June 2023. All errors are my own.

†Purdue University, jiacheng-liu@purdue.edu. Comments welcome.

1 Introduction

Modern AI, such as chat-based assistants (OpenAI, 2023; Google, 2023), has a wide range of applications (Eloundou et al., 2023). At least partly due to its multifaceted nature, there is substantial uncertainty in estimating its impact (CEA, 2024). A common approach to estimate firms' AI exposure (e.g., Eisfeldt et al. (2023a); Babina et al. (2024); Kogan et al. (2023a); Svanberg et al. (2024)) is to first derive labor exposures to AI—i.e., to estimate how AI influence workers—and then aggregate at the firm level. However, firms' AI exposures can differ from these underlying labor exposures. Indeed, the finance literature notes that managers often behave differently than standard models predict (e.g., Graham (2022); Ben-David and Chinco (2023); Gormsen and Huber (2024)), potentially creating a disconnect between estimates and actual firm-level AI exposure. To bridge this gap, I propose a firm-based approach that directly examines the business-specific impacts of AI through managerial communications.

In this paper, I propose a novel measure of firms' AI exposure using their conference call transcripts. This approach offers several advantages. First, it provides a near-real-time, time-varying, and forward-looking assessment of AI's potential influence. Second, by analyzing the textual context of AI-related discussions, it is possible to capture managerial expectations about the *firm-specific* implications of AI in a wide range of topics from productivity enhancement to AGI progress. Last but not least, managerial communications are widely distributed and highly visible compared to firm-level labor data. Thereby an earnings call-based measure empowers more researchers, especially those in emerging economies, to study the impact of AI.[1] In what follows, I will demonstrate that managerial communications provide rich information content that is hard-to-measure or unobservable via traditionally focus worker resume and job postings.

Is there information about AI in earnings calls? I begin by establishing several stylized facts about firms' adoption of AI. Examining temporal patterns in earnings call discussions, Figure 1 presents four complementary measures that together provide a compre-

[1]The AI exposure will be publicly available with replication package following acceptance of manuscript.

hensive view of how firms' engagement with AI has evolved from 2006 to 2024. Along the extensive margin, the fraction of firms mentioning AI rose from below 5% before 2016 to over 15% by 2024, with a pronounced acceleration following ChatGPT's release. On the intensive margin, the average number of AI mentions per call more than doubled post-ChatGPT to 4.3, and the increasing dispersion suggests heterogeneous strategies. The cumulative adoption curve indicates nonlinear diffusion, reaching about 2,800 firms by 2024, while the 12-month rolling average (Panel D) suggests a structural break rather than a temporary fluctuation. Collectively, these patterns point to a transformative upward shift in firms' engagement with AI technology.

To validate my new AI exposure measure, I benchmark it to established labor-based AI exposure measures, examining both aggregate time series and cross-sectional validity of the new measure. A start point for comparison is whether the firms identified as AI-exposed from earnings calls are also more likely to be identified as AI-exposed from labor data. I provide a series of tests to validate that firms identified as AI-exposed are more likely to be associated with AI-related labor. Without adjusting for the number of jobs, the correlation between my measure and labor-based measures is already high at 0.97. Weighting by the number of jobs in the cross-section increases the correlation to 0.99. Similar results emerge when using the cumulative number of earnings calls mentioning AI. The strong alignment between the measures, especially on a cumulative basis, supports the idea that the new measure captures forward-looking expectations about AI adoption. For instance, while labor market search frictions can delay actual AI job growth, managerial communications can reflect changes in expectations immediately.

To explore this hypothesis further, I conduct an event study examining the cross-sectional relationship between the text-based measure and subsequent labor demand shifts. The findings from the event study suggest that firms' expectations about AI, as expressed in their communications, predict subsequent AI-related labor force dynamics. Firms that mention AI in their filings are more likely to experience increases in AI-related employment in the following years. Using an event study design, firms that first mention AI see a 0.1 percentage point increase in the share of narrow AI jobs in that year, growing

to 0.2 percentage points over the subsequent two years. Though small in absolute terms, these changes represent substantial relative increases given the low baseline share of AI jobs. In absolute terms, these firms add about 52 narrow AI jobs and 154 AI-related jobs overall after their initial AI discussions. Difference-in-differences analyses yield similar patterns, showing that firms discussing AI exhibit an average 0.2 percentage point increase in the share of AI-related jobs compared to firms without such discussions.

As a final validation, I compare the correlation between my measure and labor-based measure at the occupation level over time. The main objective is to see if my earnings call based measure agrees with labor based measure at the occupation level. The main comparison is with Eisfeldt et al. (2023b), who shares the same methodology as Eloundou et al. (2023). Because Eisfeldt et al. (2023b) calculate exposure in early 2023, I choose the earnings calls from 2022 and 2024 to demonstrate Figures 24 and 23 illustrate the correlation between labor-based exposure and earnings call-based exposure for 2022 and 2024, respectively. These figures highlight several key insights about the relationship between the two methodologies. In Figure 24, which focuses on 2022 data, the correlation between labor-based and earnings-based total exposure is 0.818, while the correlation for core exposure is slightly higher at 0.821. These strong correlations indicate that the earnings-based measure aligns closely with labor-based exposure, suggesting that managerial discussions reflect firm-level AI engagement with considerable accuracy. By contrast, Figure 23 presents data from 2024, where the correlations decrease to 0.611 for total exposure and 0.564 for core exposure. This decline likely reflects increasing heterogeneity in AI adoption strategies among firms, as well as a broader diffusion of AI technologies across different sectors. While earnings-based exposure continues to correlate with labor-based measures, the lower correlations suggest that firms may adopt AI in ways that are not fully captured by labor data alone, emphasizing the importance of using complementary methodologies. These findings suggest that while traditional labor-based measures capture structural aspects of AI adoption, the earnings-based measure adds a forward-looking dimension by reflecting managerial expectations and firm-specific strategies. This temporal divergence between the two measures underscores the value

of analyzing managerial discussions to identify early trends and anticipate shifts in AI-related resource allocation.

Having validated and benchmarked the firm-based measure on labor data, I then turn to a range of AI-related outcomes that would be otherwise hard to quantify or even unobservable. To start, I explore the financial implications of AI adoption. Which firm characteristics are impacted by AI? What kind of firm characteristics induce AI adoption? Ideally, this question should be approached from a causal perspective. Nonethless, addressing such causal questions is highly challenging because it is difficult to isolate changes in AI adoption that aren't driven by the firm's internal processes or by omitted factors. Traditionally labor-based approach to AI adoption is especially limited in this regard, due to the lack of within-firm variation in labor data. Indeed, most of the labor-based AI adoption measures start with an occupation exposure, which is relatively stable over time in a given sector. While still not a perfect measure, the earnings call-based measure provides a time-varying measure of AI adoption. Based on the event study results, the first year of AI-related discussion appears to be a good leading proxy for AI adoption. Therefore, I will begin by addressing the more attainable question of identifying characteristics that are associated with AI adoption. Building on this earnings call-based measure of AI exposure, I examine the associations between AI adoption and a range of financial and real outcomes guided by the task-based model. The results suggest that AI adoption is associated with changes in firms' capital structures in an effort to potentially lower their cost of capital. Many other firm characteristics, however, remain persistent, likely reflecting long-standing cross-sectional heterogeneities in the propensity to adopt new technologies. Regarding real outcomes, the evidence points to both labor displacement and productivity gains associated with AI. For example, relative to comparable firms in the same industry, AI-exposed firms tend to have fewer employees after adoption, suggesting labor displacement. Meanwhile, focusing on industries known for earlier AI uptake, pharmaceutical firms associated with AI exposure complete clinical trials more quickly, and internet firms produce more extensively documented software.

Examining financial characteristics, AI-adopting firms appear more efficient, with a

1 percentage point lower labor expense-to-sales ratio than non-adopters, potentially indicating improved labor productivity. They also tend to be more profitable, with a 4 percentage point higher return on equity, and have higher market valuations, as indicated by an 11 percentage point lower book-to-market ratio. However, these firms also have lower leverage and dividend payout ratios. Notably, many of these characteristics predate the recent surge in AI capabilities, consistent with findings in Acemoglu et al. (2023). This pattern suggests that firms adopting AI may have inherent attributes—such as a culture of innovation, superior management practices, or greater risk tolerance—that predispose them to embrace new technologies. These observations underscore the importance of considering potential endogeneity and the direction of causality in the relationship between AI adoption and firm performance.

Motivated by these cross-sectional patterns, I explore whether AI adoption is associated with changes in the perceived cost of capital and discount rates Gormsen and Huber (2024). Under a Diff-in-Differences design, I find that firms that first mention AI see a 0.1 percentage point increase in the perceived cost of capital and a 0.2 percentage point increase in the discount rate. Interstingly, the wedge between the cost of capital and the discount rate narrows after AI-related discussions. Among other hypotheses, I conjecture that firms with substantial AI-related growth options face an interesting capital structure dynamic. The high valuation of AI-related equity may reduce these firms' optimal leverage ratios, as they rely less on debt financing. This shift in capital structure affects the overall cost of capital. Simultaneously, the discount rate may increase due to the heightened perceived risk and expanded growth options associated with AI investments. The observed narrowing of the wedge between cost of capital and discount rates suggests these measures move concordantly following firms' AI-related discussions, consistent with a fundamental revaluation of both the firm's cost of financing and its risk-return profile.

Having explored the financial dimension, I then turn to the real implications of AI. A key objective is to determine whether AI exposure is associated with measurable productivity gains or labor displacement. The literature provides arguments for both. For example, Babina et al. (2024) links AI adoption to increased product innovation, sug-

gesting potential productivity improvements, while Eisfeldt et al. (2023a) examines market valuations and suggests that AI breakthroughs may be associated with labor cost reductions. I focus on two settings—drug development and software creation—where improvements can be clearly measured. Within a differences-in-differences framework, pharmaceutical firms associated with AI exposure complete clinical trials about 158 days (15.8%) faster. AI exposure is also linked to more detailed code documentation, with pull request descriptions increasing by 54 characters (27%) on average. These findings imply operational enhancements in research-intensive industries. At the same time, there is also evidence of labor displacement: compared to similar firms, AI-exposed firms employ fewer workers after adoption. These results are consistent with the characterization of AI as a general-purpose technology (Eloundou et al., 2023), potentially influencing a wide range of economic outcomes. I use a range of granular intra-firm data to further explore the real implications of AI. For productivity gains, I find that AI-adopting firms in the pharmaceutical industry complete clinical trials about 158 days (15.8%) faster. In the IT sector, AI-adopting firms produce more extensively open-sourced software. These intra-firm productivity gains are consistent with the characterization of AI as a general-purpose technology (Eloundou et al., 2023), potentially influencing a wide range of economic outcomes. And motivates a more granular analysis of firm-specific AI adoption.

Reling on the rich information content of the earnings calls, I quantify some key variables in the economics of AI (see, for example, Acemoglu (2024)) by analyzing the qualitative content of firms' AI discussions. Using large language models themselves, I find that over 96% of firms express positive sentiment toward AI's potential, and about 93% adopt an active strategic stance, indicating deliberate efforts to integrate AI technologies. Approximately 49% of AI discussions address labor-related topics, including workforce strategies and skill requirements. Despite the optimistic outlook, firms also acknowledge challenges: nearly 48% indicate that AI could displace workers, while only about 23% mention the creation of new tasks for human employees. This dual recognition underscores the complexity of AI's labor implications and highlights the need for policies that address both opportunities and challenges. By incorporating forward-looking managerial

communications, this paper introduces a new methodology to measure AI exposure and its anticipated economic effects. Complementing existing labor task-based approaches, this communication-based perspective captures strategic intentions not observable solely through labor data, offering a more nuanced understanding of how firms perceive and respond to AI.

This paper contributes to the emerging literature on measuring AI exposures (Webb, 2020; Felten et al., 2018; Brynjolfsson et al., 2018; Eloundou et al., 2023; Babina et al., 2024; Kogan et al., 2023b) and to research on the social and economic aspects of machine learning and AI (e.g., Acemoglu and Restrepo (2018a); Brynjolfsson and Mitchell (2017); Eisfeldt et al. (2023a)). The labor implications of AI have been a focal point of this literature, grounded in task-based models of automation (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018a; Autor et al., 2006; Van Reenen, 2011; Acemoglu and Restrepo, 2022). More recently, Eloundou et al. (2023) highlights that GPT-driven impacts may be pervasive, and Eisfeldt et al. (2023a) examines generative AI's influence on firm value. Using worker-level résumé and job posting data, Babina et al. (2024) finds that AI-investing firms grow faster and innovate more. This paper differs by taking a top-down approach, focusing on firms' own communications rather than relying solely on bottom-up, task-level data. These communication-based measures complement existing task-based ones, as managerial considerations of labor tasks may inform strategic decisions about AI adoption and, conversely, business responses can shape labor demand. The patterns documented here also align with previous worker-level findings (e.g., Peng et al. (2023); Brynjolfsson et al. (2023)). Methodologically, while distinct in research questions, this study builds on a large literature in textual analysis (e.g., Lopez-Lira and Tang (2023); Korinek (2023); Kelly et al. (2021); Cong et al. (2024); Hassan et al. (2019); Loughran and Mcdonald (2014)).

These findings inform our understanding of technological change and labor markets (Acemoglu and Restrepo, 2018a; Autor et al., 2003; Acemoglu and Autor, 2011). Distinguishing between labor-augmenting and labor-displacing technologies is central to interpreting wage inequality, employment dynamics, and productivity trends. By incorporat-

ing firms' expectations, this work provides insights into how contemporary AI technologies are perceived at the firm level and the types of tasks that may be more susceptible to automation or enhancement.

The policy implications are multifaceted. If firms predominantly view AI as labor-augmenting, policies might emphasize workforce development and education programs that enable workers to complement AI-driven technologies. Conversely, the recognition of displacement effects underscores the need for social safety nets and retraining initiatives to support workers transitioning to new roles.

2 Data

2.1 Firm Communication Data

The main input for my new firm's AI exposures is the text content of firm communication. I collected both the earnings call transcripts and SEC filings. Using either SEC filings or earnings calls did not meaningfully change the results. I use firm earnings call transcripts in most of the analysis. For a set of additional robustness tests, I use SEC filings instead.

For firm SEC filings, I obtained all annual and quarterly filings from 2016 to 2022 from the U.S. Securities and Exchange Commission's EDGAR database. These filings provide information about firms' operations and managerial assessments. The SEC describes these reports as "a comprehensive overview of the company's business and financial condition".[2]

The primary dataset for earnings calls is sourced from Capital IQ, providing detailed managerial discussions on firm operations. In the main analysis involving labor data, I focused on the earnings calls dataset because it allows for more precise merging with labor data. When labor data was not used, I included both earnings call and filing data.

[2]See https://www.investor.gov/introduction-investing/investing-basics/glossary/form-10-k

2.2 Clinical Trial Data

I obtained data from the database of clinical trials by the National Library of Medicine (clinicaltrials.gov). This dataset contains information on clinical trials sponsored by the National Institutes of Health (NIH) and other organizations. The data include detailed information about each trial, including the trial phase, enrollment size, and completion date. I link the clinical trial data to the Compustat firms using their names.[3]

2.3 Open Source Software Data

On the productivity of software development, I obtained data from Github, a leading platform for open-source software development. I obtained data on all commits and pull requests to public repositories to until about early 2022. I used the GitHub organization names to link the data to the firms' names within Compustat database.[4]

2.4 Job Postings Data

To study the realized impact of AI, I utilized job posting data from Burning Glass Technologies (BG). This dataset includes over 180 million job postings in the United States from 2007 and 2010 to 2018. BG collects data from more than 40,000 online job boards and company websites, ensuring comprehensive coverage. The data is parsed into a machine-readable format and deduplicated to avoid counting the same vacancy multiple times. Each job posting includes detailed information such as job title, location, occupation, employer name, and required skills.

The BG dataset is advantageous due to its extensive coverage and the rich detail it provides for each job posting. It captures nearly all jobs posted online and represents about 60-70% of all job vacancies in the U.S., whether posted online or offline.

For a firm to be included in the sample, I required the firm to have both labor and earnings call transcripts data from 2011 to 2018, no other filters are applied. For more details on the production of this dataset, see Babina et al. (2024).

[3] For more information about this data, see clinicaltrials.gov
[4] For more information about this data, see Github.com

2.5 Resume Data

To measure firms' AI-related human capital, I utilize comprehensive employment profile data from Cognism. The dataset contains detailed career histories for over 500 million individuals globally. Each profile includes detailed employment records with start and end dates, standardized job titles, employer information, and role descriptions. Cognism uses advanced natural language processing techniques to clean and standardize employment records, including normalizing company names, job titles, and educational credentials.

A key advantage of this dataset is its extensive coverage—approximately two-thirds of the U.S. workforce as of 2018—while maintaining representative distribution across industries and regions. The data are particularly valuable for studying AI-related employment because they capture both technical and non-technical roles involved in AI initiatives, allowing for a comprehensive view of how firms build their AI capabilities. For more details on the production of this dataset, see Babina et al. (2024); Fedyk and Hodson (2023).

2.6 LLM Data Annotation Pipeline

The foundation of the new framework is using LLM to help annotate managerial discussions. However, while LLM-empowered tools are general purpose, in order to apply them for research purposes, I need to design a specific annotation pipeline around the large language model to get useful research outputs. This section describes the main data annotation steps, however, more technical details are in the appendix.

2.6.1 Foundation Model Annotation

I searched the earnings calls to identify mentions of machine learning and artificial intelligence. Specifically, I checked each report for the presence of the terms "machine learning" or "artificial intelligence". A firm was considered to have an AI mention if either term appeared at least once in the call. Once identified, I extracted the entire sentence containing these keywords and compiled a dataset of AI-related discussions from the

filings.

Using this dataset, I produced statistics on the fraction of unique firms with AI mentions. A firm was considered to mention AI in a given year if at least one filing within that year contained an AI mention. I performed the analysis both at the aggregate level for all U.S. firms over time and across different industries using the 2-digit NAICS industry code. Firms were identified by their CIK identifiers. The fraction of firms mentioning AI in a given year was calculated as the number of unique firms with AI mentions divided by the total number of unique firms with filings in that year.

To gain deeper insights and annotate these discussions effectively, I leveraged the advanced language processing capabilities of GPT-series models OpenAI (2023). I used both "gpt-3.5-turbo" and "gpt-4o-mini" to assist in analyzing and annotating the extracted text, allowing for an evaluation of the impact of AI as conveyed in the discussions. This method of using GPT to label text is also used in other studies, such as Eloundou et al. (2023); Eisfeldt et al. (2023a).

In the annotation process, I asked GPT to assess the direction and magnitude of the AI impact for each discussion, as well as other aspects of potential impact. Specifically, I asked GPT to classify each text based on the following aspects: 1. Direction of impact: positive, negative, or other. 2. Significance of impact: low, moderate, high, or other. 3. Primary topic of impact: labor, investment, revenue, competition, M&A, or other. 4. Timeline of impact: happened, current, planning, or other. 5. Aggressiveness of impact: active, passive, or other.

I prompted GPT as follows:

""" System: You will be provided AI-related text from a firm's filing. If the text is not about AI, answer "NA" for all items below. Provide an answer in a list of 5. Your task is to analyze the overall impact of AI and label the following 5 aspects. Think step by step, and choose only one from each aspect.

[ASPECTS: 1. direction of impact: positive, negative, other. 2. significance of impact: low, moderate, high, other. 3. primary topic of impact: labor, investment, revenue, competition, M&A, other. 4. timeline of impact: happened, current, planning, other. 5.

aggressiveness of impact: active, passive, other.]

User: ["We continue to see an opportunity to offer our fraud detection solutions with advanced machine learning capabilities to help customers"]

Assistant: [positive, moderate, revenue, current, passive] """

During the data cleaning step, I excluded some filings from the final sample. Specifically, I excluded discussions that were over 100 words in context to avoid distractions and reduce complexity. I also excluded cases where GPT incorrectly labeled texts, such as selecting more than one topic of impact or failing to generate a high-level analysis of impact.

3 An AI Exposure without Labor Data

3.1 Dynamic Patterns in AI Exposure

I begin by analyzing temporal patterns in firms' AI exposure through their earnings call discussions. Figure 1 presents four complementary measures that together provide a comprehensive view of how firms' engagement with AI has evolved from 2006 to 2024.

Panel A shows the extensive margin of AI exposure—the fraction of firms mentioning AI in their earnings calls. This measure reveals three distinct phases in corporate AI adoption. From 2006 to 2016, AI discussions remained relatively rare, with mention rates consistently below 5%. This low baseline suggests that despite early developments in machine learning, most firms did not view AI as strategically relevant. The 2016 DeepMind Challenge, which demonstrated AI's capability to master complex strategic games, marked the beginning of a second phase characterized by gradual adoption. Mention rates rose steadily to approximately 8% through 2022, indicating growing but still limited corporate engagement with AI technology.

The release of ChatGPT in late 2022 triggered a dramatic third phase, with mention rates more than doubling to exceed 15% by early 2024. This sharp acceleration suggests that generative AI represented a qualitatively different type of technological breakthrough, one that firms viewed as having broader applicability across industries and functions.

The magnitude of this increase—nearly triple the pre-2022 baseline—indicates a fundamental shift in how firms perceive AI's strategic relevance.

Panel B examines the intensive margin, measuring the average number of AI mentions per AI-discussing call. This metric reveals not just whether firms discuss AI, but how substantively they engage with it. The average mentions per call remained relatively stable at 2-3 mentions through 2022, suggesting standardized or perfunctory references to AI initiatives. However, post-ChatGPT, this measure more than doubled to 4.3 mentions per call. More striking is the increased variance in mention frequency, with the 90th percentile reaching 8.2 mentions while the 10th percentile remained at 1.1. This growing dispersion suggests emerging heterogeneity in firms' AI strategies, with some firms pursuing comprehensive AI initiatives while others maintain more limited engagement.

Panel C tracks the cumulative number of unique firms that have discussed AI, providing insight into the technology's diffusion across the corporate landscape. The curve exhibits clear non-linear dynamics, with approximately 2,800 unique firms having discussed AI by early 2024. The slope of this curve steepens notably after each major AI breakthrough—first in 2016 and more dramatically in 2022. This pattern suggests that technological milestones serve as catalysts for broader corporate adoption, potentially triggering strategic responses even among firms that previously viewed AI as peripheral to their operations.

To separate underlying trends from quarterly earnings seasonality, Panel D presents a 12-month rolling average of the fraction of firms discussing AI. This smoothed measure confirms that the post-ChatGPT surge represents a structural break rather than a temporary phenomenon. The trend shows clear acceleration beginning in December 2022 and maintaining momentum through 2023 and early 2024. The persistence of elevated discussion rates, even after controlling for seasonality, suggests that increased corporate attention to AI reflects fundamental strategic repositioning rather than mere rhetorical responses to technological headlines.

Collectively, these measures reveal a transformative shift in corporate engagement with AI technology. The simultaneous increases in both extensive and intensive margins

post-2022 suggest that firms are not only more likely to discuss AI but are engaging with it more substantively. The non-linear patterns in adoption and growing dispersion in engagement intensity point to potential complementarities or network effects in AI adoption decisions. These dynamics have important implications for understanding how transformative technologies diffuse through the corporate sector and for predicting the broader economic impacts of AI advancement.

4 Expected Economic Impact of AI and the Subsequent Labor Dynamics

In this section, we examine whether firms' expectations about AI, as reflected in their earnings call discussions, are predictive of subsequent labor dynamics, particularly in terms of AI-related employment. By analyzing the relationship between AI mentions in earnings calls and changes in AI-related job shares and counts, we assess whether firms' forward-looking communications capture meaningful information about their strategic adoption of AI technologies and the associated adjustments in their workforce.

4.1 Aggregate Time-Series Analysis

We begin by examining the relationship between firms' AI discussions-based measure of AI exposures and their subsequent AI-related labor dynamics at the aggregate level. Figure 2 plots several key time series measuring AI adoption from 2010 to 2018. The share of AI-related jobs shows a steady upward trend, increasing from 0.12% in 2010 to 0.83% in 2018, with growth accelerating notably after 2015. This pattern reflects the gradual integration of AI capabilities into firms' operations through hiring.

The raw share of firms discussing AI in earnings calls exhibits more dramatic growth, rising from 1.1% in 2010 to 26.5% in 2018. When weighted by number of jobs, the increase is even more pronounced - from 1.2% to 37.7%. This difference between weighted and unweighted measures suggests that larger employers were earlier and more extensive adopters of AI technology, consistent with theories of technological diffusion where larger

firms are better positioned to absorb the fixed costs of new technology implementation.

The cumulative measure of AI discussions provides additional insights into adoption patterns. By 2018, 37.6% of firms had mentioned AI at least once in their earnings calls, indicating widespread recognition of AI's potential impact across the corporate landscape.

The relationship between these different measures is remarkably strong. Even without incorporating employment weights, the correlation between discussion-based and labor-based measures is 0.97. Including employment weights or using the cumulative measure further increases this correlation to 0.99. The stronger fit of the cumulative measure with labor market outcomes supports the interpretation that earnings call discussions capture firms' forward-looking expectations about AI adoption. While actual hiring may lag behind strategic initiatives due to search frictions and other labor market constraints, managerial discussions in earnings calls can immediately reflect changes in firms' technological trajectories and strategic priorities.

This high correlation between discussion patterns and realized labor outcomes, combined with the leading nature of managerial communications, suggests that our text-based measure effectively captures firms' expectations about AI adoption and its labor market implications. The fact that larger firms show earlier and more extensive adoption also aligns with established patterns of technological diffusion, lending further credibility to our measure as a meaningful indicator of AI's expanding role in the economy.

4.2 Cross-Sectional Analysis

We then conduct cross-sectional regressions to explore the association between firms' discussions of AI and their AI-related labor demand. Table 2 presents the regression results where the dependent variables are the share of narrow AI jobs (Columns 1–2) and the share of all AI jobs (Columns 3–4). The key explanatory variable is an indicator variable, AI_exposed, which equals one if a firm mentions AI-related terms in its earnings calls during the sample period and zero otherwise.

In Column 1, we regress the share of narrow AI jobs on AI_exposed without any control variables. The coefficient on AI_exposed is 0.001 with a t-statistic of 3.709, indicating

that firms mentioning AI have, on average, a 0.1 percentage point higher share of narrow AI jobs compared to firms that do not mention AI. This effect is statistically significant at the 1% level. Given that the average share of narrow AI jobs is relatively small, this represents a meaningful increase in AI-related employment for these firms.

In Column 2, we include a set of firm-level control variables: the logarithm of assets (ln Assets), the logarithm of sales (ln Sales), the logarithm of employees (ln Employees), firm age, leverage ratio, and return on assets (ROA). The inclusion of these controls is crucial to account for observable firm characteristics that might influence both the likelihood of discussing AI and the share of AI-related jobs. For instance, larger firms may have more resources to invest in AI technologies and may have different hiring practices. Firm age could capture lifecycle effects, while financial variables like leverage ratio and ROA may reflect firms' capacity to finance new initiatives and their operational efficiency.

After adding these controls, the coefficient on AI_exposed remains virtually unchanged at 0.001, suggesting that the positive association between AI discussions and narrow AI job share is robust to the inclusion of these variables. Among the control variables, return on assets shows a negative and significant coefficient, indicating that more profitable firms may have a slightly lower share of narrow AI jobs. This could be due to these firms already having optimized operations, reducing the immediate need for additional AI personnel, or perhaps they are less pressured to adopt new technologies due to their current profitability.

In Columns 3 and 4, we repeat the analysis with the share of all AI jobs as the dependent variable. The coefficient on AI_exposed increases to 0.003. This suggests that firms discussing AI have a 0.3 percentage point higher share of all AI jobs—a larger effect than for narrow AI jobs. The larger magnitude reflects the broader scope of AI-related roles encompassed in this measure and underscores the pervasive impact of AI discussions on firms' labor demand. Including the control variables in Column 4 does not materially change the coefficient on AI_exposed, which remains significant and of similar magnitude. Return on assets continues to have a negative and significant coefficient, consistent with the earlier finding.

In all regressions, we include industry-fixed effects based on 4-digit NAICS codes to control for unobserved industry-specific factors that could influence both AI discussions and AI employment. By comparing firms within the same industry, we mitigate concerns that industry trends are driving the results. The adjusted R-squared values indicate that the models explain a substantial portion of the variation in AI job shares (approximately 19% for narrow AI and 23% for all AI).

These cross-sectional findings suggest that firms actively engaging in AI discussions are also those with higher shares of AI-related employment, pointing to a meaningful link between firms' communicated expectations about AI and their labor force composition.

Table 1 extends the cross-sectional analysis to examine the relationship between AI discussions and the number of AI jobs, rather than their share. The dependent variables are the count of narrow AI jobs (Columns 1–2) and the count of all AI jobs (Columns 3–4). The key explanatory variable remains the AI_exposed indicator.

In Column 1, without control variables, the coefficient on AI_exposed is 51.723 (t-statistic = 3.709), indicating that firms mentioning AI have, on average, about 52 more narrow AI jobs compared to firms that do not mention AI. This effect is both economically and statistically significant. Given that many firms may have few or no narrow AI jobs, this represents a substantial difference in AI-related employment.

Column 2 includes the same set of firm-level control variables as in Table 2. After adding these controls, the coefficient on AI_exposed increases slightly to 52.723 (t-statistic = 3.716), suggesting that the positive association between AI discussions and narrow AI job count is robust to the inclusion of these variables. Among the control variables, firm size (as measured by $\ln$ Employees) shows a positive and significant coefficient (coefficient = 23.456, t-statistic = 5.678), indicating that larger firms tend to have more narrow AI jobs, which is intuitive given their greater resources and potentially more complex operations.

Columns 3 and 4 repeat the analysis with the count of all AI jobs as the dependent variable. The coefficient on AI_exposed is substantially larger at 153.169 (t-statistic = 7.012) without controls and 154.169 (t-statistic = 7.023) with controls. This indicates

that firms discussing AI have, on average, about 153-154 more AI jobs across all categories—a much larger effect than for narrow AI jobs alone. This larger magnitude reflects the broader scope of AI-related roles encompassed in this measure and underscores the pervasive impact of AI discussions on firms' labor demand across various job categories.

As in Table 2, we include industry-fixed effects based on 4-digit NAICS codes in all regressions to control for unobserved industry-specific factors. The adjusted R-squared values are lower than in Table 1 (approximately 10% for narrow AI and 15% for all AI), which is expected given the greater variability in job counts compared to job shares.

These findings complement those in Table 2 by showing that the positive association between AI discussions and AI-related employment holds not just in relative terms (job shares) but also in absolute terms (job counts). This reinforces the conclusion that firms actively engaging in AI discussions are indeed those with larger AI-related workforces, both proportionally and in absolute numbers.

4.3 Event Study Analysis

To delve deeper into the dynamics of AI-related labor adjustments around the time of AI discussions, we employ an event study methodology. This approach allows us to examine the evolution of AI job shares and counts in the years before and after firms first mention AI in their earnings calls.

Table 3 presents the event study estimates for AI job shares. The dependent variables are the share of narrow AI jobs (Columns 1–2) and the share of all AI jobs (Columns 3–4). The key explanatory variables are indicator variables for each year relative to the first AI mention ($year_rel$), ranging from five years before ($year_rel__-5$) to five years after ($year_rel_5$).

In the pre-event period, the coefficients on the lead terms ($year_rel__-5$ to $year_rel__-1$) are generally small and statistically insignificant. For example, in Column 1, the coefficients on these terms range from -0.000 to 0.000 with t-statistics close to zero. This lack of significant pre-trends supports the validity of the event study design and suggests that the parallel trends assumption holds—that is, firms that will mention AI in the future

are not already experiencing different trends in AI job shares compared to other firms.

In the post-event period, there is a significant increase in AI job shares following the first AI mention. In Column 1, the coefficient on $year_rel_0$ is 0.001 (t-statistic = 4.296), indicating that in the year of the first AI mention, firms experience a 0.1 percentage point increase in the share of narrow AI jobs. This effect grows in subsequent years, reaching 0.002 (t-statistic = 5.066) in $year_rel_1$ and remains significant in later years. The sustained increase suggests that firms continue to expand their narrow AI workforce following their initial engagement with AI discussions.

Similar patterns are observed for all AI job shares in Columns 3 and 4. The coefficients on the post-event indicators are slightly larger, reflecting the broader set of AI-related roles included in this measure. For instance, the coefficient on $year_rel_1$ in Column 3 is 0.002 (t-statistic = 4.953), indicating a significant increase in the share of all AI jobs. These findings highlight that firms' AI discussions are associated with ongoing adjustments in their labor force composition toward AI-related roles.

In these regressions, we include firm-fixed effects to control for unobserved, time-invariant characteristics of firms that could influence both AI discussions and employment patterns. Year-fixed effects are also included to account for common shocks affecting all firms in a given year, such as macroeconomic conditions or technological trends. The inclusion of these fixed effects enhances the identification of the causal impact of AI discussions on AI job shares. The adjusted R-squared values are approximately 50% for narrow AI and 71% for all AI, indicating that the models explain a substantial portion of the variation in AI job shares over time.

Table 4 extends the event study analysis to the number of AI jobs. The dependent variables are the counts of narrow AI jobs (Columns 1–2) and all AI jobs (Columns 3–4). The coefficients on the lead terms are again insignificant, indicating no pre-trends in AI job counts.

In the post-event period, the coefficients become positive and significant. For instance, in Column 1, the coefficient on $year_rel_0$ is 11.364 (t-statistic = 2.013), suggesting that firms add approximately 11 narrow AI jobs in the year they first mention AI. This

number increases in subsequent years, reaching 53.909 (t-statistic = 2.062) in $year_rel_2$. The growing magnitude over time indicates that firms progressively expand their AI workforce following their initial AI discussions. The inclusion of control variables in Columns 2 and 4—similar to those in the cross-sectional analysis—does not materially change the results. The control variables themselves are generally insignificant, indicating that the observed increases in AI job counts are not driven by these firm characteristics.

4.4 Difference-in-Differences Analysis

To further establish the causal relationship between AI discussions and labor dynamics, we implement a difference-in-differences (DiD) approach. This methodology compares the changes in AI job shares and counts for firms that mention AI (treated group) to those that do not (control group), before and after the first AI mention.

Table 5 reports the DiD estimates for AI job shares. The dependent variables are the share of narrow AI jobs (Columns 1–2) and the share of all AI jobs (Columns 3–4). The key variable of interest is post_ai_text, an indicator equal to one for years on or after a firm's first AI mention and zero otherwise.

In Column 1, without control variables, the coefficient on post_ai_text is 0.002 (t-statistic = 7.892), indicating that firms experience a 0.2 percentage point increase in the share of narrow AI jobs after discussing AI, relative to firms that do not mention AI. This effect remains significant and of similar magnitude in Column 2, which includes firm-level controls and industry-fixed effects. Including controls helps address potential confounding factors that might influence both the treatment (AI discussions) and the outcome (AI job shares), enhancing the credibility of the causal interpretation. The control variables are the same as those used in the cross-sectional analysis: $\ln$ Assets, $\ln$ Sales, $\ln$ Employees, firm age, leverage ratio, and ROA.

Similarly, in Columns 3 and 4, the dependent variable is the share of all AI jobs, and the results are consistent with those for narrow AI jobs. The coefficients on post_ai_text are 0.002 (t-statistics = 7.521 and 7.466), indicating a significant increase in all AI job shares following AI discussions. The control variables remain insignificant, and the adjusted R-

squared values are around 50% for narrow AI and 71% for all AI, suggesting a good fit of the models.

In these regressions, we include firm fixed effects to control for unobservable, time-invariant firm characteristics, while year-fixed effects account for time trends common to all firms. Industry fixed effects at the 4-digit NAICS level control for industry-specific trends. The inclusion of these fixed effects and controls strengthens the causal interpretation of the results.

4.5 Robustness Checks and Additional Analyses

To ensure the robustness of the results and address potential concerns, we conduct several additional analyses.

One concern might be that the increase in AI jobs comes at the expense of non-AI jobs, potentially leading to overall job displacement within firms. To examine this possibility, we analyze the effect of AI discussions on non-AI job counts. Table 15 presents the DiD estimates where the dependent variable is the number of non-AI jobs.

The coefficient on post_ai_text is positive but insignificant (coefficient = 236.391, t-statistic = 1.007), suggesting that there is no significant decrease—or increase—in non-AI job counts following AI mentions. This implies that the increase in AI-related employment does not lead to a reduction in other types of jobs within the firm, and firms may be expanding their overall workforce.

We also examine the impact on other categories of jobs, such as robot-related jobs and high-wage jobs, to assess whether AI discussions are associated with shifts in different segments of the workforce. The results (available upon request) indicate that AI discussions are not significantly associated with changes in robot job shares but are negatively associated with high-wage job shares. This may reflect a reallocation of resources toward AI-related roles, which could be more specialized but not necessarily high-wage positions compared to traditional high-wage roles.

We test alternative model specifications, including the addition of interaction terms and lagged dependent variables, and explore the inclusion of other control variables such

as R&D intensity and capital expenditures, which could influence both AI adoption and employment patterns. The main findings remain consistent across these variations, reinforcing the robustness of the results.

Furthermore, we address potential serial correlation and heteroskedasticity by clustering standard errors at the firm level, which provides more reliable inference in the presence of within-firm correlation over time.

As an additional robustness check, we perform placebo tests by randomly assigning AI mention years to firms and re-estimating the difference-in-differences models. The placebo treatments yield insignificant coefficients, supporting the validity of the main findings and suggesting that the results are not driven by chance or spurious correlations.

4.6 Discussion of Results

The collective evidence from the cross-sectional, event study and difference-in-differences analyses indicate a strong and positive association between firms' AI discussions and subsequent increases in AI-related employment. This relationship is evident in both the share and the absolute number of AI jobs and persists after controlling for various firm characteristics and unobserved heterogeneity.

The inclusion of firm, year, and industry fixed effects, along with relevant control variables, helps mitigate concerns about omitted variable bias and enhances the credibility of the causal interpretation. The absence of significant pre-trends in AI job shares and counts before the first AI mention strengthens the validity of the identification strategy.

These findings imply that firms' forward-looking communications about AI are not merely rhetorical but reflect genuine strategic intentions that materialize in tangible labor force adjustments. Firms that actively engage in AI discussions appear to invest in building their AI capabilities by hiring more AI-related personnel, aligning with their stated expectations.

Moreover, the increase in AI-related employment without a corresponding decrease in non-AI jobs suggests that AI adoption may be associated with net job creation within

firms, at least in the short to medium term. This has important implications for understanding the labor market impact of AI, suggesting that concerns about widespread job displacement may be mitigated by firms' expansion of their AI workforce.

These results contribute to the broader literature on technological change and labor markets, highlighting the role of firms' expectations and communications in shaping labor demand. They underscore the importance of considering firms' strategic decisions and forward-looking statements when assessing the economic implications of emerging technologies like AI.

5 Firm characteristics and AI adoption

Having established that firms' discussions of AI are associated with significant increases in AI-related employment, we now examine how various firm characteristics correlate with AI adoption. This analysis explores whether firms that discuss AI exhibit different operational and financial characteristics compared to those that do not.

We examine key performance metrics between AI-exposed and non-AI-exposed firms. Table 7 presents regression results where various firm characteristics are regressed on the indicator,AI_Exposed that equals one only if a firm adopted AI based on the communication, controlling for firm size, age, financial characteristics, and including year and industry fixed effects. Since we are interested in estimating cross-sectional correlations, I do not include firm fixed effects in these regressions as in earlier tables.

In Column 1, the dependent variable is the labor expense to sales ratio. The coefficient on AI_Exposed is -0.01 (t-statistic = -4.622), indicating that AI-exposed firms have a 1 percentage point lower labor expense relative to sales compared to non-exposed firms. This correlation could reflect either that AI adoption leads to labor efficiencies, or that more efficient firms are more likely to adopt AI.

Column 2 focuses on the leverage ratio (debt-to-equity ratio). The coefficient on AI_Exposed is -0.14 (t-statistic = -1.077), but it is not statistically significant.

In Column 3, the dependent variable is inventory turnover. The coefficient on AI_Exposed

is 2.40 (t-statistic = 2.374), suggesting that AI-exposed firms tend to have higher inventory turnover rates. This correlation could reflect either operational improvements from AI or that firms with better inventory management are more likely to adopt AI.

Column 4 examines return on equity (ROE). The coefficient on AI_Exposed is 0.04 (t-statistic = 4.564), indicating that AI-exposed firms have a 4 percentage point higher ROE. This correlation with profitability could reflect either benefits from AI adoption or that more profitable firms are more likely to invest in AI.

In Column 5, we analyze the book-to-market ratio. The coefficient on AI_Exposed is -0.11 (t-statistic = -8.601), indicating that AI-exposed firms tend to have higher market valuations relative to book value. This could reflect either market optimism about AI adoption or that growth firms are more likely to adopt AI.

Column 6 investigates the dividend payout ratio. The coefficient on AI_Exposed is -0.06 (t-statistic = -2.312), suggesting that AI-exposed firms have a 6 percentage point lower dividend payout. This correlation could reflect either reinvestment in AI or that growth-oriented firms are more likely to adopt AI.

Lastly, Column 7 assesses the gross profit margin. The coefficient on AI_Exposed is 0.07 (t-statistic = 3.202), implying that AI-exposed firms tend to have a 7 percentage point higher gross margin. This could reflect either efficiency gains from AI or that firms with higher margins are more likely to invest in AI.

Across all regressions, we control for firm size (using $\ln$ Assets and $\ln$ Sales), firm age, leverage ratio, and return on assets. Year-fixed effects capture macroeconomic conditions and temporal trends affecting all firms, while industry-fixed effects at the 4-digit NAICS level account for industry-specific factors.

The correlations between AI exposure and firm characteristics reveal interesting patterns, though we cannot make causal claims. AI-exposed firms tend to have lower labor costs, higher inventory turnover, better profitability, higher market valuations, lower dividend payouts, and higher gross margins. However, these differences could reflect either the impact of AI adoption or underlying firm characteristics that make certain firms more likely to adopt AI.

These patterns are consistent with multiple interpretations. One possibility is that AI adoption leads to operational and financial improvements. Another is that better-performing firms have more resources and capabilities to invest in AI. A third possibility is that unobserved factors, such as management quality or technological sophistication, drive both AI adoption and firm performance.

We acknowledge significant endogeneity concerns in interpreting these correlations. Selection effects are likely important - firms that adopt AI may systematically differ from non-adopters in ways that also affect performance. While our regressions include extensive controls and fixed effects, unobservable factors could still drive the observed relationships. Future research using quasi-experimental methods could help establish more credible causal evidence.

The analysis provides robust evidence that firms discussing AI in earnings calls differ systematically from other firms across multiple dimensions. While we cannot make strong causal claims, the patterns suggest either that AI adoption is associated with better firm performance or that better-performing firms are more likely to adopt AI. These findings contribute to our understanding of the relationship between technological adoption and firm characteristics while highlighting the challenges in establishing causality in this setting.

These results have implications for interpreting technological change in the economy. The systematic differences between AI-adopting and non-adopting firms suggest that technological diffusion may be influenced by firm characteristics. This raises important questions about whether technology adoption exacerbates or reduces differences between firms. Future research could explore the dynamics of these relationships and their implications for competition, inequality, and economic growth.

6 Qualitative Analysis of AI Discussions by GPT

To gain deeper insights into firms' perspectives on AI and the qualitative aspects of their AI-related communications, we employ GPT-based annotations of AI discussions ex-

tracted from firms' filings and earnings calls. This approach leverages advanced natural language processing techniques to systematically categorize and quantify the sentiment, significance, topics, timelines, and strategic orientations reflected in these discussions. By analyzing these qualitative dimensions, we aim to understand not only whether firms discuss AI but also how they perceive its impact, which areas they emphasize, and what strategic stances they adopt.

6.1 Trends in AI Mentions

Our analysis begins by examining the temporal trends in AI mentions within firms' filings over the period from 2016 to 2022. As illustrated in Figure 1, there is a marked increase in the fraction of firms mentioning AI, rising from approximately 10% in 2016 to over 30% in 2022. This upward trajectory reflects the growing recognition of AI's relevance in the corporate sector and suggests that AI has become an increasingly integral topic in firms' strategic considerations and public disclosures.

6.2 Industry Variation

Exploring the heterogeneity across industries reveals significant cross-sectional variation in the prevalence of AI mentions. Figure 2 displays the fraction of firms mentioning AI in 2021 across various sectors. Industries such as Information, Educational Services, Administrative and Support Services, and Professional, Scientific, and Technical Services exhibit the highest fractions, with over 50% of firms mentioning AI. These sectors are inherently knowledge-intensive and are likely at the forefront of AI adoption due to the potential for AI to enhance service delivery, innovation, and efficiency.

In contrast, industries like Construction and Arts, Entertainment, and Recreation show lower rates of AI mentions, indicating varying degrees of AI adoption and perceived relevance across sectors. Table 2 provides detailed statistics on the growth of AI mentions by industry, including the fraction of firms mentioning AI in 2016 and 2022, as well as the average growth rate and standard deviation. This industry-level analysis underscores that the diffusion of AI is not uniform and is influenced by industry-specific factors such

as the nature of work, competitive pressures, and technological opportunities.

6.3 Sentiment and Significance of AI Impact

Using GPT-based annotations, we categorize the AI discussions along several dimensions to capture the nuances of firms' perspectives. The analysis reveals that an overwhelming majority of firms express a positive sentiment regarding AI's impact. Specifically, 96.18% of the AI mentions are classified as having a positive sentiment, while only 2.98% are neutral, and a mere 0.84% are negative. This widespread optimism suggests that firms perceive AI as a beneficial development with the potential to enhance their operations, products, or services.

Regarding the perceived significance of AI's impact, firms display a range of views. Approximately 42.81% of the mentions indicate that AI has a high impact, 49.24% consider it to have a moderate impact, and 7.95% perceive the impact as low. This distribution implies that while most firms acknowledge the importance of AI, there is variation in the degree to which they believe AI will affect their business. Firms perceiving a high impact may be more likely to make substantial investments in AI technologies and integrate AI into their core strategies.

6.4 Primary Topics of AI Discussions

To understand which areas firms emphasize in their AI-related communications, we analyze the primary topics associated with AI mentions. The GPT-based annotation identifies several key themes. The most frequently mentioned topic, appearing in 49.45% of AI discussions, relates to labor. Firms are actively considering the implications of AI on their workforce, including aspects such as automation of tasks, changes in skill requirements, workforce restructuring, and the need for employee training and development. This focus reflects concerns and strategic thinking about how AI will transform labor dynamics within firms.

Investment emerges as the second most common topic, accounting for 31.74% of mentions. Firms discuss allocating resources toward AI capabilities, including expenditures

on technology acquisition, research and development, and building AI infrastructure. This emphasis indicates that firms recognize the need to invest proactively to harness AI's potential benefits.

Other topics include revenue (6.96%) and competition (6.50%). Discussions around revenue pertain to how AI can enhance sales, create new revenue streams, or improve pricing strategies. Competitive considerations involve how AI adoption may affect firms' market positions, competitive advantages, and responses to rivals' actions. A smaller fraction of firms (5.36%) mention mergers and acquisitions (M&A) activities in relation to AI, suggesting that some firms pursue acquisitions to acquire AI capabilities or to consolidate their position in AI-related markets.

The prominence of labor and investment topics underscores that firms view AI both as a driver of operational transformation and as an area requiring significant resource commitment. This dual focus highlights the strategic importance of AI in shaping firms' future trajectories.

6.5 Timeline of AI Impact

Understanding when firms expect AI to impact their operations provides insights into their strategic planning horizons. The GPT-based annotations categorize the timeline of AI impact as past, current, or future. The largest group, representing 44.65% of mentions, indicates that AI is currently impacting the firm. These firms are likely already implementing AI technologies and experiencing their effects on operations, products, or services. A significant portion, 36.64%, discusses AI's impact in future terms. These firms may be in the planning or early adoption stages, anticipating how AI will affect them and preparing for integration. A smaller group, 18.71%, reflects on past experiences with AI. These firms may have already undergone initial adoption phases and are evaluating the outcomes or lessons learned.

This temporal distribution suggests that while AI is an immediate concern for many firms, others are still gearing up for its anticipated influence. The emphasis on current and future impact indicates that AI is a dynamic area of focus, with firms at different

stages of adoption and integration.

6.6 Strategic Approach toward AI

Assessing firms' strategic orientations toward AI reveals that an overwhelming majority exhibit an active stance. Specifically, 93.24% of firms demonstrate proactive engagement with AI technologies, investing in AI capabilities, and integrating AI into their strategic initiatives. These firms may be developing AI-driven products, optimizing processes with AI, or pursuing AI-related business models. The predominance of an active approach indicates that firms recognize the strategic importance of AI and are taking deliberate steps to capitalize on its potential.

In contrast, a minority of firms (6.76%) adopt a passive approach, possibly reflecting a wait-and-see attitude or limited engagement with AI. These firms may be monitoring AI developments but have not yet committed significant resources or strategic focus to AI adoption. This divergence in strategic approaches highlights varying levels of readiness and commitment to AI across firms.

6.7 Labor-Augmenting and Labor-Displacing Effects

A critical aspect of firms' AI discussions pertains to the expected impact on labor. We analyze whether firms perceive AI as augmenting human labor, displacing it, or creating new tasks. The analysis reveals that a substantial majority of firms indicate that AI will increase worker productivity and augment human capabilities. This view reflects an expectation that AI will enhance employees' effectiveness by automating routine tasks, providing advanced analytics, or enabling more informed decision-making.

However, approximately 48.37% of firms acknowledge the potential for AI to displace workers. This recognition suggests concerns about automation leading to job losses or redundancy in certain roles, especially those involving repetitive or easily automatable tasks. Only 22.67% of firms mention that AI will create new tasks for human workers. This lower proportion indicates that while some firms anticipate new roles emerging due to AI, the creation of entirely new job categories is not a dominant theme.

To capture the net effect on labor, we compute an enhancement-displacement score for each firm by subtracting the number of displacement-driven task mentions from the number of enhancement-driven task mentions. The distribution of this score, depicted in Figure 5, is centered around zero with a slight left skew. This suggests that while many firms anticipate a balance between augmentation and displacement, there is a subset of firms expecting significant displacement effects.

6.8 Interpretation and Implications

The GPT-based annotation results provide nuanced insights into firms' perceptions and strategic considerations regarding AI. The overwhelmingly positive sentiment and active strategic approaches suggest that firms view AI as a critical driver of future success and are investing accordingly. The emphasis on labor-related topics highlights the central role of workforce implications in firms' AI strategies.

However, the acknowledgment of potential labor displacement underscores the dual-edged nature of AI adoption. Firms recognize that while AI can enhance productivity, it may also lead to job redundancies, necessitating careful management of workforce transitions. The relatively low expectation of new task creation implies that firms may not fully anticipate the potential for AI to generate new types of employment, which could affect how they prepare their employees for future roles.

From a policy perspective, these findings highlight the importance of supporting firms and workers during the AI adoption process. Initiatives that facilitate reskilling and up-skilling can help mitigate displacement risks and enable workers to capitalize on new opportunities created by AI technologies. Additionally, fostering an environment that encourages innovation while addressing ethical and societal considerations can promote sustainable and inclusive growth in the context of AI advancement.

Overall, the GPT-based annotations enrich our understanding of how firms perceive and plan for AI's impact, complementing the quantitative analysis of labor dynamics. They reveal that firms are not only discussing AI but are actively integrating it into their strategic thinking, with significant implications for labor, investment, and competitive

positioning. These qualitative insights provide a valuable context for interpreting the observed increases in AI-related employment and the associated improvements in firm performance.

7 Real Effect of AI on Firm Productivity

Having established the patterns of AI adoption and discussion across firms, I now examine whether AI adoption is associated with measurable productivity gains. I focus on two specific contexts where productivity can be clearly measured: drug development in pharmaceutical firms and software development efficiency. These settings provide concrete evidence of how AI adoption translates into operational productivity while offering insights that may generalize to other domains where productivity effects are harder to quantify.

However, I wish to emphasize that the key message of the section is to demonstrate that the new measure actually captures some real activity. For detailed estimates of the causality effect of AI, I leave for future work.

7.1 Real Effect of AI on Clinical Trial Duration

I examine whether AI adoption leads to productivity gains in drug development by analyzing clinical trial durations. Using a difference-in-differences design that compares AI-adopting pharmaceutical firms to non-adopters, I find that AI adoption significantly accelerates the drug development process.

The baseline estimates in Table 17 Column 1 indicate that AI adoption reduces trial duration by approximately 158 days (t-statistic = -4.27). This represents a 15.8 percent reduction relative to the sample mean of 1,000 days. The effect remains robust to increasingly stringent specifications. Adding trial phase fixed effects and controls for expanded access programs in Column 2 yields a reduction of 121 days (t-statistic = -2.61). The most demanding specification in Column 3, which includes controls for trial enrollment size and other characteristics, shows a reduction of 102 days (t-statistic = -2.72).

Notably, trials generally took longer in the post period across all firms, as evidenced by the positive coefficient on the post-AI indicator (87.152 days in Column 3). This suggests increasing trial complexity or stricter regulatory requirements over time. However, AI-adopting firms more than offset this trend through productivity gains from AI implementation.

7.2 Real Effect of AI on Software Development Productivity

Next, I investigate AI's impact on software development efficiency by examining pull request documentation. Pull requests are proposed code changes that undergo review before integration into the codebase. More detailed pull request descriptions generally indicate better documentation practices and facilitate code review, potentially improving development productivity.

Using data from GitHub repositories, I find that AI adoption leads to significantly more detailed code documentation. Table 18 Column 1 shows that after AI adoption, pull request descriptions increase by 54 characters (t-statistic = 2.76), representing a 27 percent increase relative to the pre-adoption mean of 200 characters.

The effect strengthens when accounting for developer fixed effects. Column 2 shows an increase of 83 characters (t-statistic = 1.77), suggesting the improvement occurs within developers rather than through composition changes. This magnitude persists in Column 3 after controlling for collaboration patterns through the number of authors per pull request.

These findings reveal an important channel through which AI affects software development: by improving documentation quality rather than just automating code generation. Better documentation can reduce future maintenance costs and knowledge transfer frictions. The results suggest AI tools may generate substantial productivity gains by enhancing development practices beyond direct code production.

8 Conceptual Framework of AI Impact on Firms

To understand the macroeconomic implications of AI, this paper employs a task-based production framework where productivity gains arise from automation and task complementarities. The framework builds on Hulten's theorem to link task-specific cost savings to aggregate total factor productivity (TFP) and GDP growth.[5]

8.1 Task-Based Production Framework

We consider an economy producing a unique final good, Y, which combines a continuum of tasks indexed by $z \in [0, N]$. The production function is given by:

$$Y = B(N) \left(\int_0^N y(z)^{\frac{\sigma-1}{\sigma}} dz \right)^{\frac{\sigma}{\sigma-1}},$$

where $y(z)$ is the output from task z, σ is the elasticity of substitution between tasks, and $B(N)$ captures system-wide effects of task creation. For simplicity, we normalize $B(N) = 1$. Tasks can be performed by either labor or capital, with automation shifting tasks from labor to capital, and task complementarities enhancing productivity in tasks that remain labor-intensive.

AI's productivity effects arise through several key channels. First, AI may automating human tasks, which can be thought as the automation (extensive margin automation) enables capital to substitute for labor in specific tasks, effectively expanding the measure of tasks performed by capital. Second, AI can also increase human worker productivity, which can be thought as task complementarities augment the productivity of labor in the tasks it continues to perform, increasing the marginal product of labor. Third, AI enable automation deepening (intensive margin automation), which effectively increase capital productivity on tasks that were previously automated already. Last but not least, AI can also improve the quality of tasks, which can be thought as creating new tasks for human workers and conceptually can include system-level improvements.

[5]The version of the model analyzed is based on Acemoglu (2024), which itself is based on Acemoglu and Autor (2011); Acemoglu and Restrepo (2018b).

8.2 Aggregate Productivity Gains: Hulten's Theorem

Hulten's theorem provides a direct link between micro-level productivity improvements and aggregate TFP growth. Specifically, the aggregate TFP change can be expressed as:

$$\Delta \ln \text{TFP} = \int_0^N \chi(z)\pi_L(z)\, dz,$$

where $\chi(z) = \frac{p(z)y(z)}{Y}$ represents the GDP share of task z, and $\pi_L(z)$ denotes the productivity improvement or cost savings in task z.

8.3 From TFP to GDP Growth

To translate TFP growth into GDP growth, we incorporate the effects of capital accumulation. Under standard assumptions, the relationship between GDP growth and TFP is given by:

$$\Delta \ln Y = \Delta \ln \text{TFP} + s_K \cdot \Delta \ln K,$$

where s_K is the capital share of GDP, and $\Delta \ln K$ is the proportional increase in the capital stock.

We are most interested in the TFP gains:

$$\Delta \ln \text{TFP} = \text{GDP share of AI-exposed tasks} \times \text{cost savings}$$

8.4 Qualification of Framework assumptions

This section uses the empirical findings from this paper to qualify a set of assumptions that are used to estimate the parameters. To begin the analyses, Figure 15 provides qualifies a set of assumptions that are used to estimate the parameters.

First, most of the impact on firms is via complementarities of AI and human workers, meaning that AI adoption leads to increase in human worker productivity instead of automation.

Second, there are considerable new tasks and system-level improvements being cre-

ated by AI adoption.

Third, on inequality, AI is likly to increase inequality along several dimensions, including skill premium, gender wage gap, and racial wage gap. I find that extensive-margin automation is likely to happen in female workers, while complementarities is likely to happen in male workers tasks. On the seniority side, AI is likely to increase the wage premium due to the exposures (both automation and complementarities) are concentration in the middle tier of seniority. Thus likely to increase the polarization of the labor market.

8.5 Empirical Estimation of Framework Parameters

This section uses the empirical findings from this paper to provide new estimates for key parameters in the task-based framework. These estimates allow for a calibration of the model to quantify the macroeconomic impact of AI adoption.

The baseline assumption is that half of the tasks are exposed to AI, regardless of automation or complementarities. In the paper, I find about 70% of the exposure is via complementarities of AI and human workers with likly 40% productivity gains. This implies a productivity gain of about 0.5 * 0.7 * 0.4 = 14%, due to the complementarities alone.

Another about 20% of the exposure is via new tasks which has about 20% productivity gains. This implies a productivity gain of about 0.5 * 0.2 * 0.2 = 2%, due to the new tasks.

Finally, the remaining 10% of the exposure is via automation deepening, which has about 10% productivity gains. This implies a productivity gain of about 0.5 * 0.1 * 0.1 = 0.5%, due to the automation deepening.

Adding these three together, the overall productivity gain is about 16.5%.

9 Conclusion

This paper examines how AI influences firm behavior. To do so, I develop a novel, near real-time measure of firms' AI exposure using earnings call transcripts. By analyzing

the language managers employ when discussing AI, this approach integrates firm-level expectations and strategic considerations, offering a perspective that complements traditional task- or labor-based measures.

My empirical analysis yields three main insights. First, firms that actively discuss AI subsequently experience increases in AI-related hiring, suggesting that managerial discourse anticipates and aligns with future workforce adjustments. Second, qualitative analyses of these discussions reveal overwhelmingly positive sentiment and proactive strategies for integrating AI, indicating that firms are not only aware of AI's potential but are actively planning to harness it. Finally, I find that increased AI exposure corresponds with improvements in specific real outcomes, such as enhanced productivity in drug discovery.

These findings enrich the broader literature on technology adoption and labor markets by illuminating the crucial role of managerial expectations in guiding firm responses to emerging technologies. By leveraging forward-looking communications, this paper provides new insights into how firms anticipate, adapt to, and capitalize on transformative technologies like AI. Moreover, the results underscore the value of incorporating timely, qualitative measures into the study of economic impacts, especially in fast-paced technological landscapes where traditional data sources may lag behind ongoing developments.

References

Acemoglu, D. (2024). The simple macroeconomics of ai. Technical report, National Bureau of Economic Research.

Acemoglu, D., Anderson, G., Beede, D., Buffington, C., Childress, E., Dinlersoz, E., Foster, L., Goldschlag, N., Haltiwanger, J., Kroff, Z., et al. (2023). Advanced technology adoption: Selection or causal effects? In *AEA Papers and Proceedings*, volume 113, pages 210–214. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 4 of *Handbook of Labor Economics*, chapter 12, pages 1043–1171. Elsevier.

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

Acemoglu, D. and Restrepo, P. (2018b). The race between man and machine: Implications of technology for growth, factor shares, and employment.

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

Autor, D. H., Katz, L. F., and Kearney, M. S. (2006). The polarization of the us labor market. *American economic review*, 96(2):189–194.

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

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

Ben-David, I. and Chinco, A. M. (2023). Modeling managers as eps maximizers. Technical report, National Bureau of Economic Research.

Brynjolfsson, E., Li, D., and Raymond, L. R. (2023). Generative ai at work. Technical report, National Bureau of Economic Research.

Brynjolfsson, E. and Mitchell, T. (2017). What can machine learning do? workforce implications. *Science*, 358(6370):1530–1534.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What Can Machines Learn and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings*, 108:43–47.

CEA (2024). Potential labor market impacts of artificial intelligence: An empirical analysis. Technical report, The White House.

Cong, L. W., Liang, T., and Zhang, X. (2024). Textual Factors: A Scalable, Interpretable, and Data-driven Approach to Analyzing Unstructured Information. *SSRN Electronic Journal*.

Eisfeldt, A. L., Schubert, G., and Zhang, M. B. (2023a). Generative AI and Firm Values. *SSRN Electronic Journal*.

Eisfeldt, A. L., Schubert, G., and Zhang, M. B. (2023b). Generative AI and firm values. NBER Working Paper 31222.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2023). GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *arXiv*.

Fedyk, A. and Hodson, J. (2023). Trading on talent: Human capital and firm performance. *Review of Finance*, 27(5):1659–1698.

Felten, E. W., Raj, M., and Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, 108:54–57.

Google (2023). Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

Gormsen, N. J. and Huber, K. (2024). Firms' perceived cost of capital. Technical report, National Bureau of Economic Research.

Graham, J. R. (2022). Presidential address: Corporate finance and reality. *The Journal of Finance*, 77(4):1975–2049.

Hassan, T. A., Hollander, S., van Lent, L., and Tahoun, A. (2019). Firm-Level Political Risk: Measurement and Effects*. *The Quarterly Journal of Economics*, 134(4):2135–2202.

Kelly, B., Manela, A., and Moreira, A. (2021). Text Selection. *Journal of Business and Economic Statistics*, 39(4):859–879.

Kogan, L., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2023a). Technology and labor displacement: Evidence from linking patents with worker-level data. Technical report, National Bureau of Economic Research.

Kogan, L., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2023b). Technology and labor displacement: Evidence from linking patents with worker-level data. NBER Working Paper 31846.

Korinek, A. (2023). Generative ai for economic research: Use cases and implications for economists. *Journal of Economic Literature*, 61(4):1281–1317.

Lopez-Lira, A. and Tang, Y. (2023). Can chatgpt forecast stock price movements? return predictability and large language models. *SSRN*.

Loughran, T. and Mcdonald, B. (2014). Measuring Readability in Financial Disclosures. *The Journal of Finance*, 69(4):1643–1671.

OpenAI (2023). Gpt-4 technical report. Technical report, OpenAI.

Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590*.

Svanberg, M., Li, W., Fleming, M., Goehring, B., and Thompson, N. (2024). Beyond ai exposure: Which tasks are cost-effective to automate with computer vision? *Available at SSRN 4700751*.

Van Reenen, J. (2011). Wage inequality, technology and trade: 21st century evidence. *Labour economics*, 18(6):730–741.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.

Webb, M. (2020). The impact of artificial intelligence on the labor market. Working paper, Stanford University.

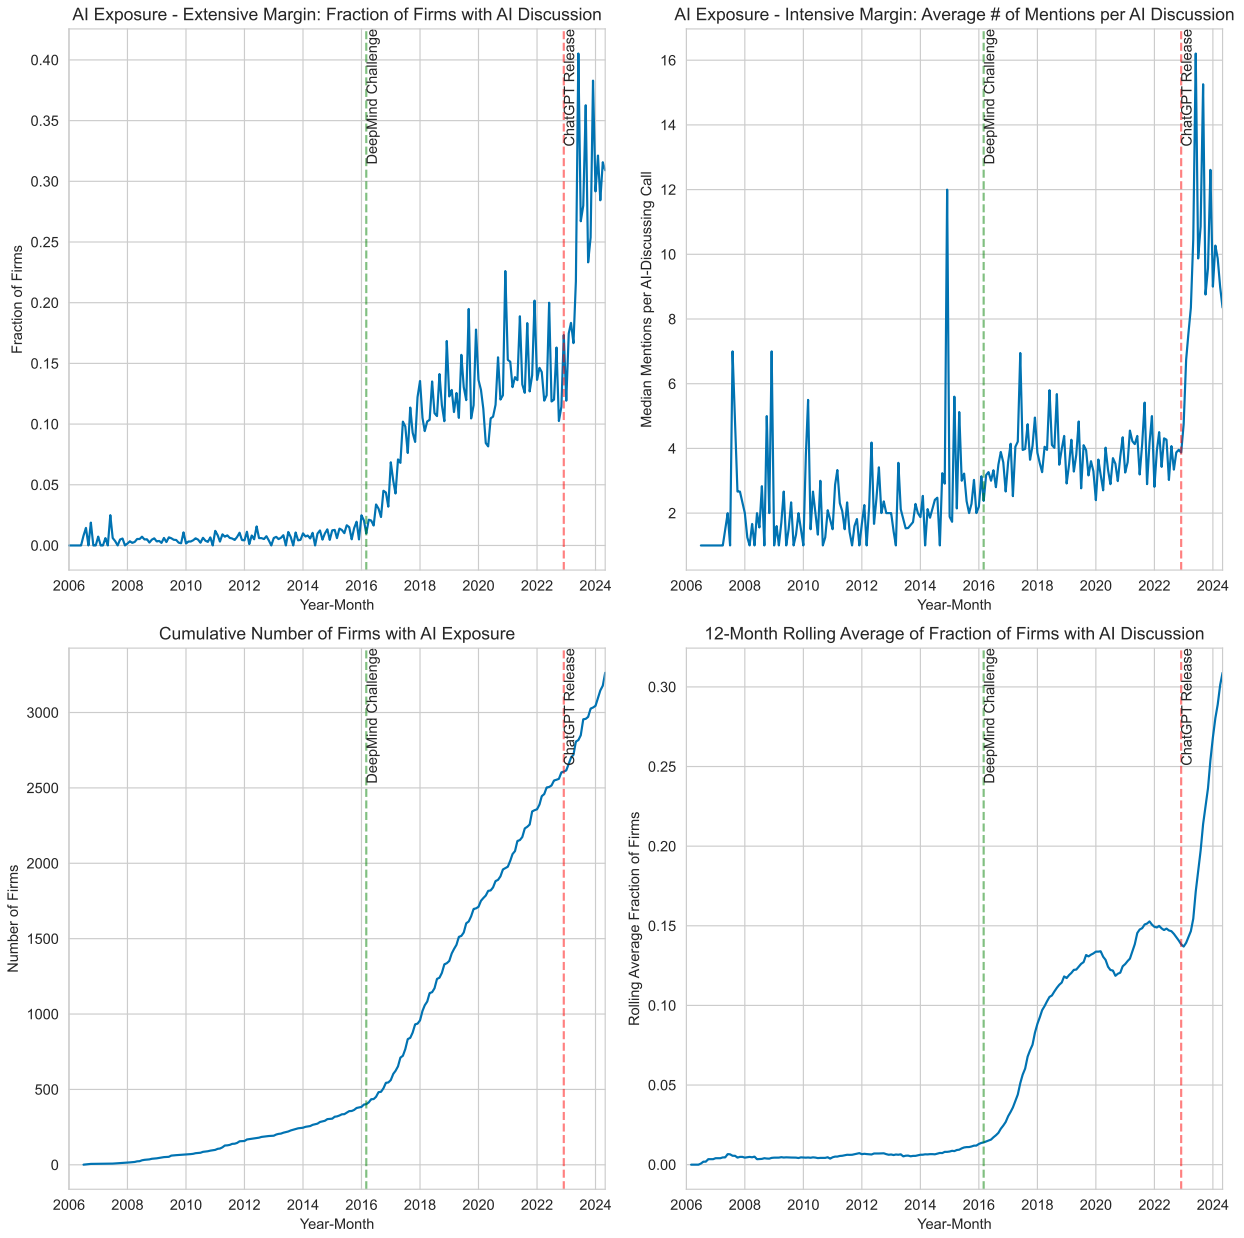


Figure 1: The new AI exposure without labor data

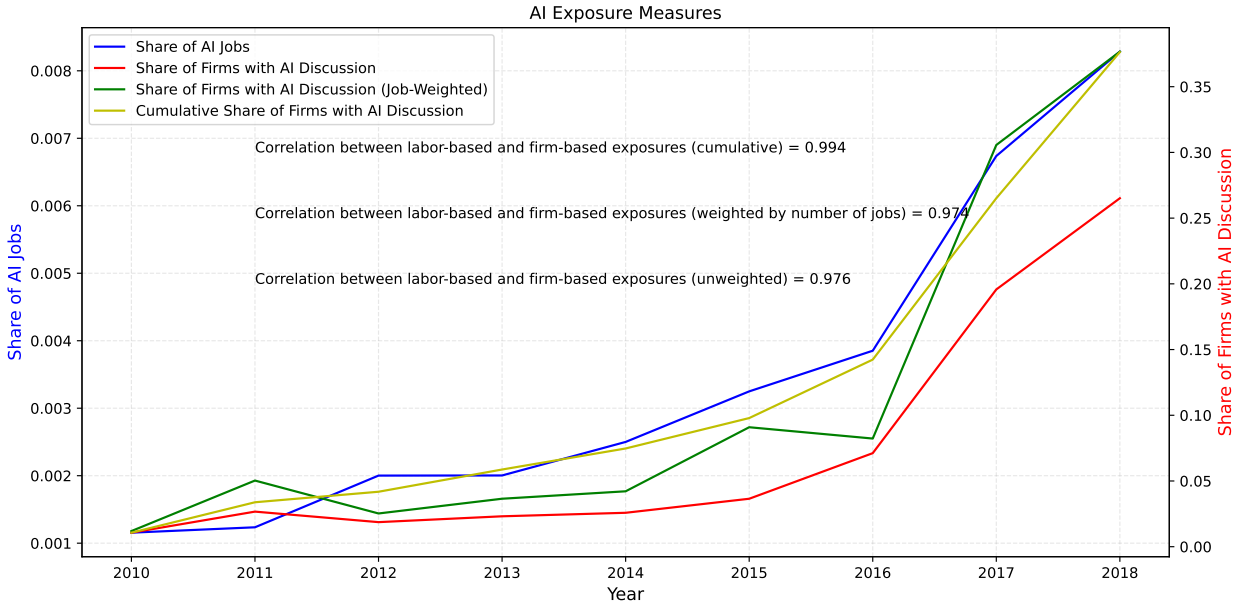


Figure 2: Aggregate shares of AI jobs and firms with AI discussions

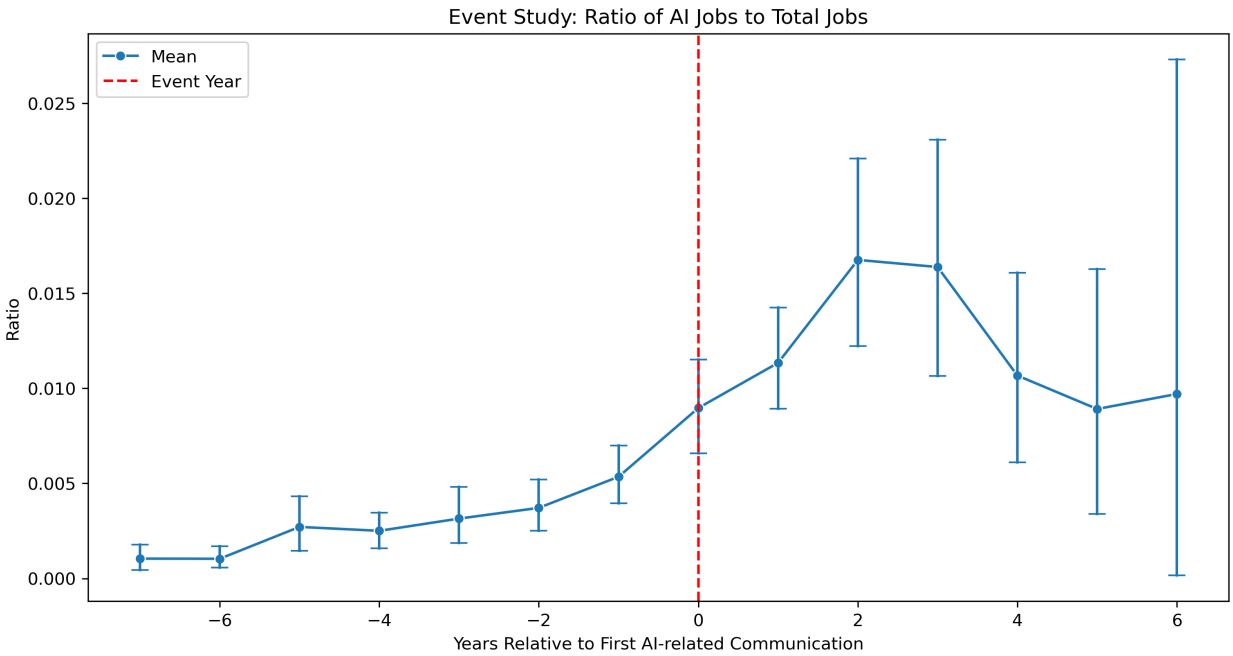


Figure 3: Event study of AI related labor dynamics.

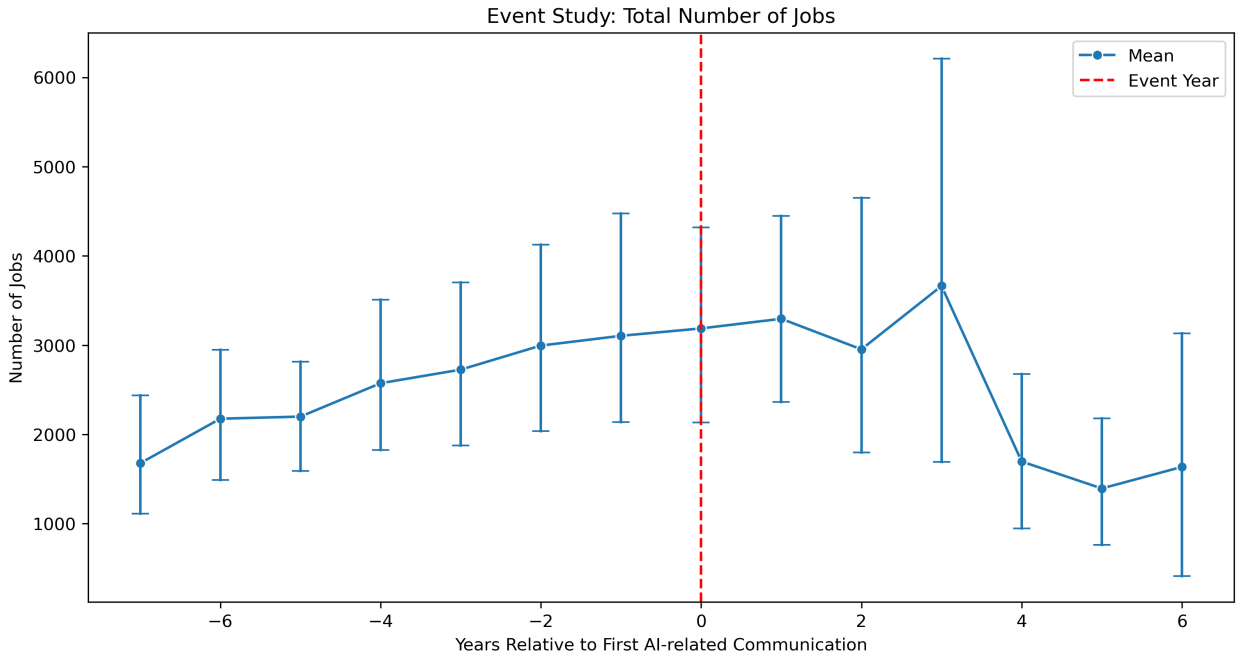


Figure 4: Event study of overall labor dynamics.

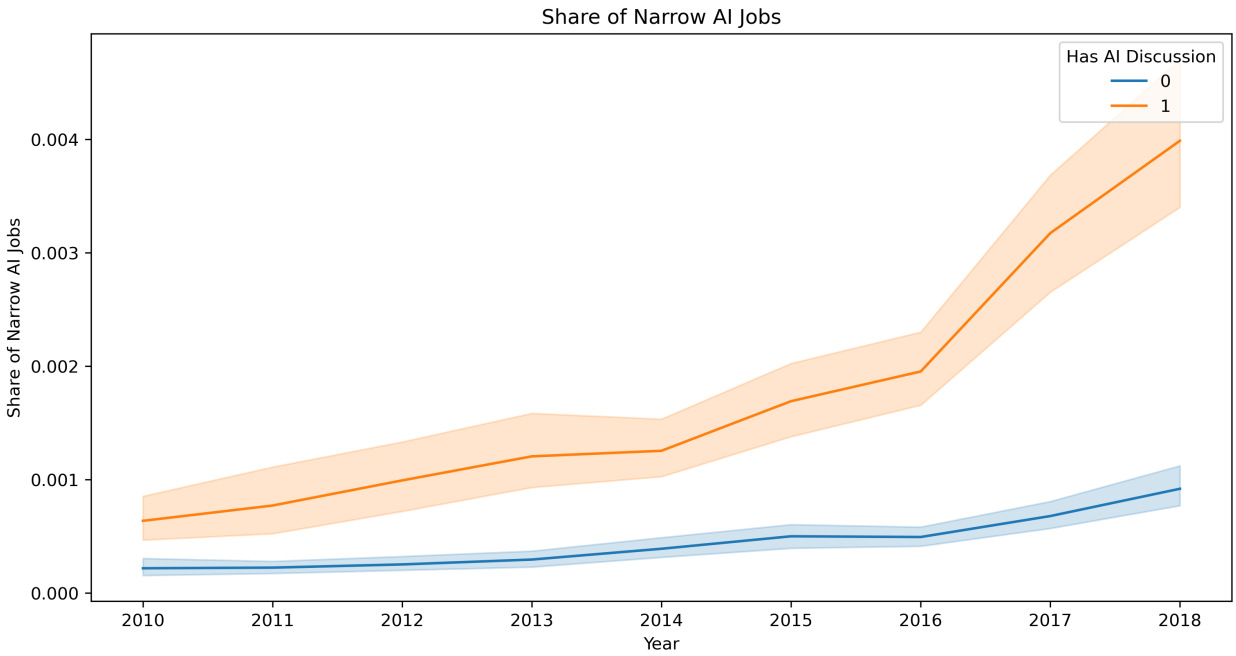
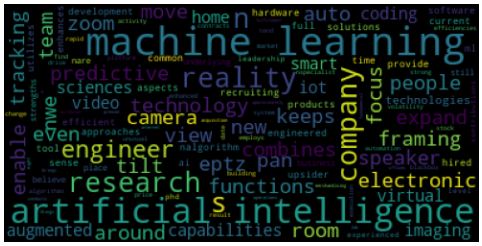


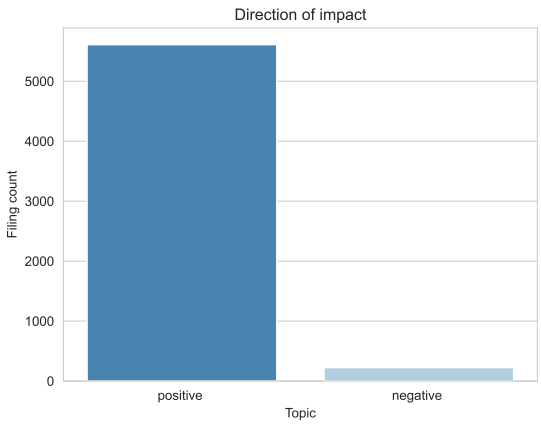
Figure 5: Average within-firm shares of narrow AI jobs



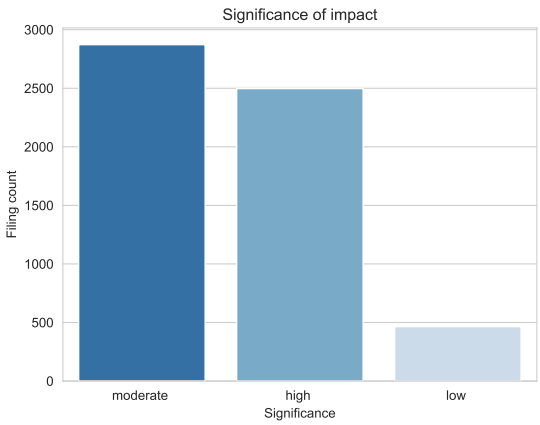
(a) Word cloud of AI discussions in firm filings.

(b) Word cloud of GPT labels.

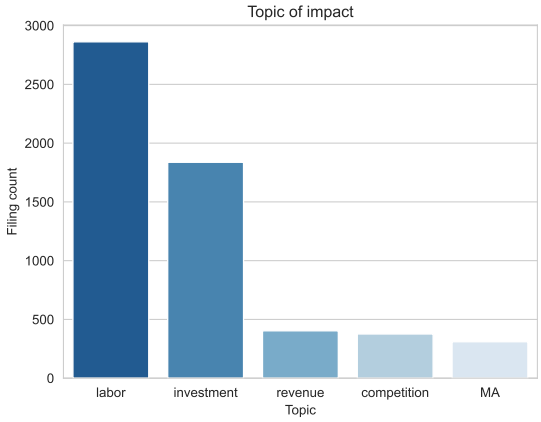
Figure 6: Word clouds of keywords in filings and impact labels.



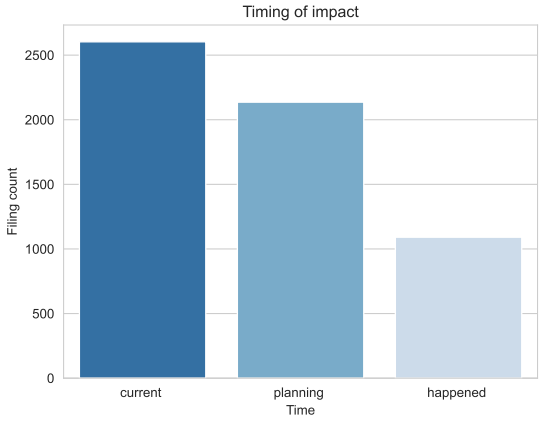
(a) Direction



(b) Significance

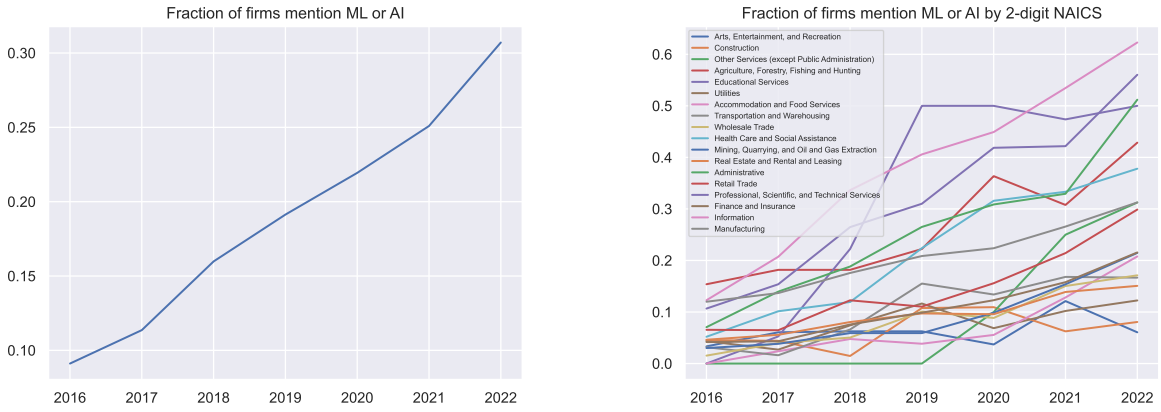


(c) Topic



(d) Time

Figure 7: Number of filings for labels.

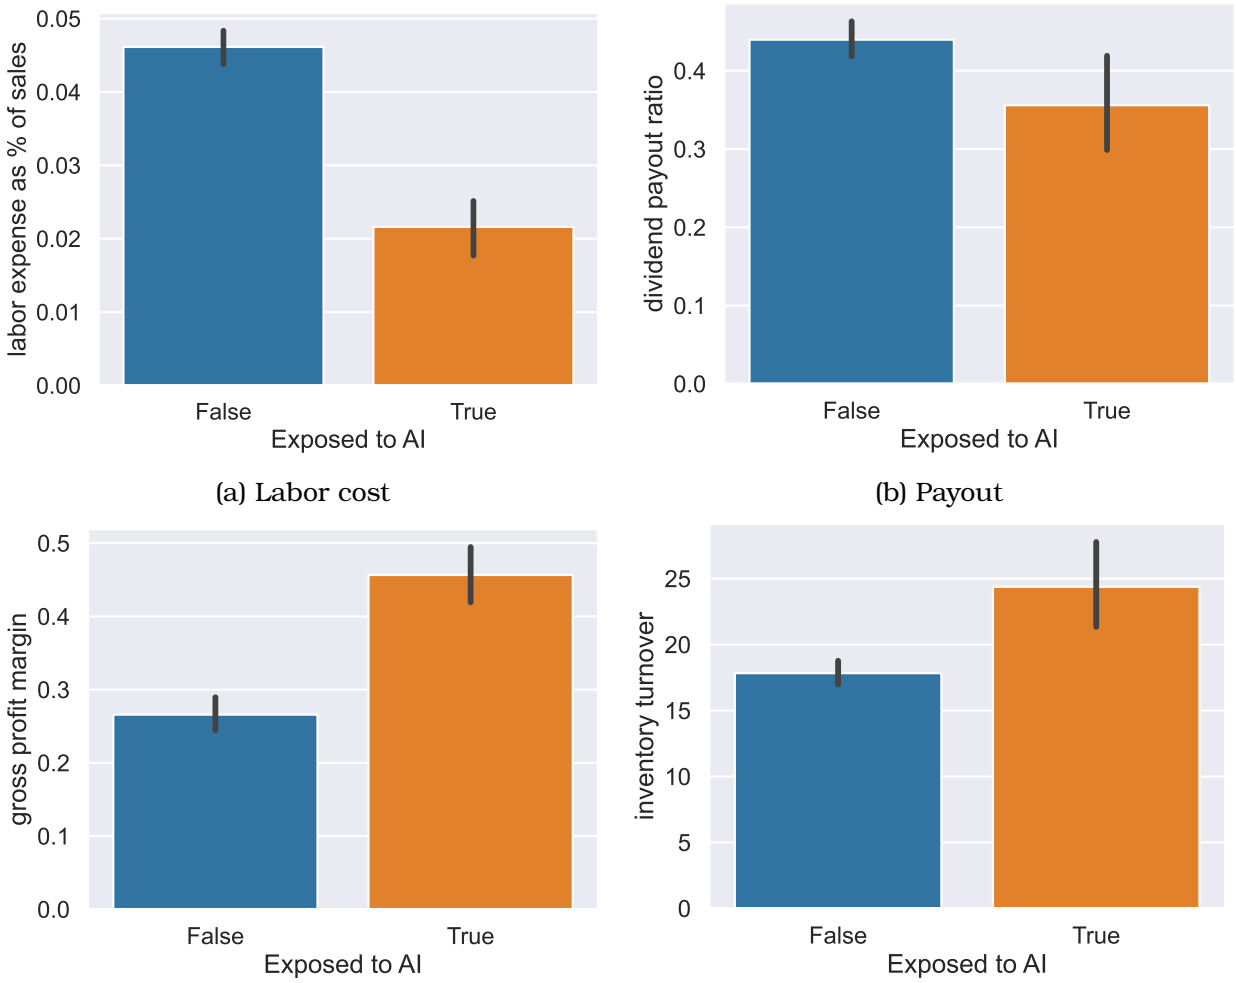


(a) All US

(b) Industries-specific

Figure 8: The fraction of firms mention Machine Learning (ML) or AI in filings.

This figure presents the relationship between firm financial ratios and AI exposure.



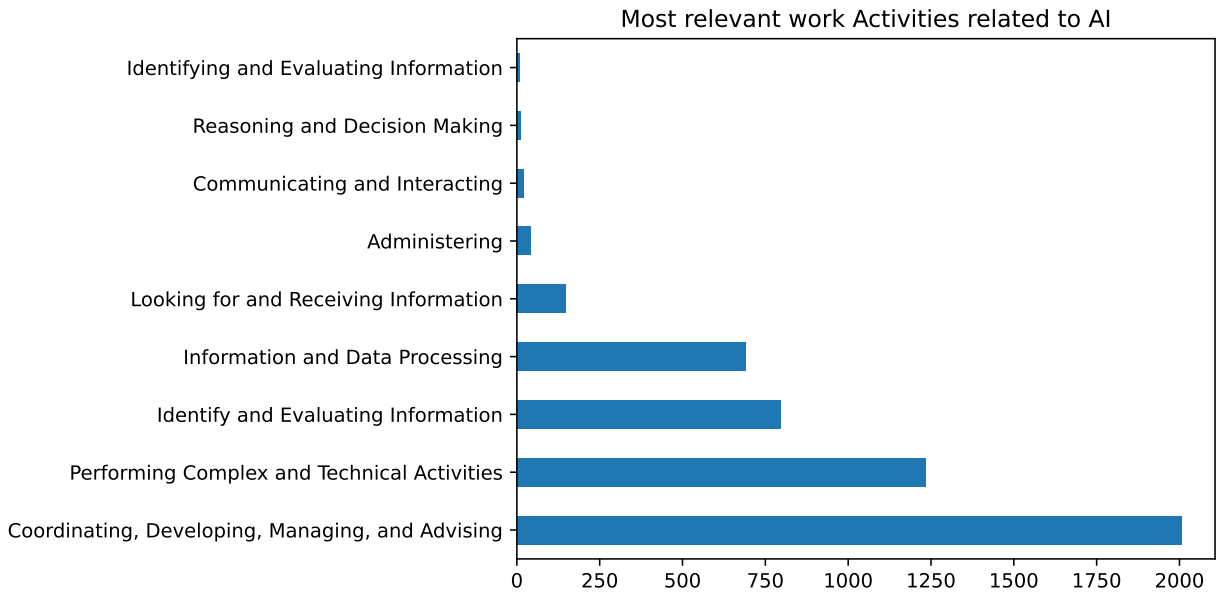
(a) Labor cost

(b) Payout

(c) Profit Margin

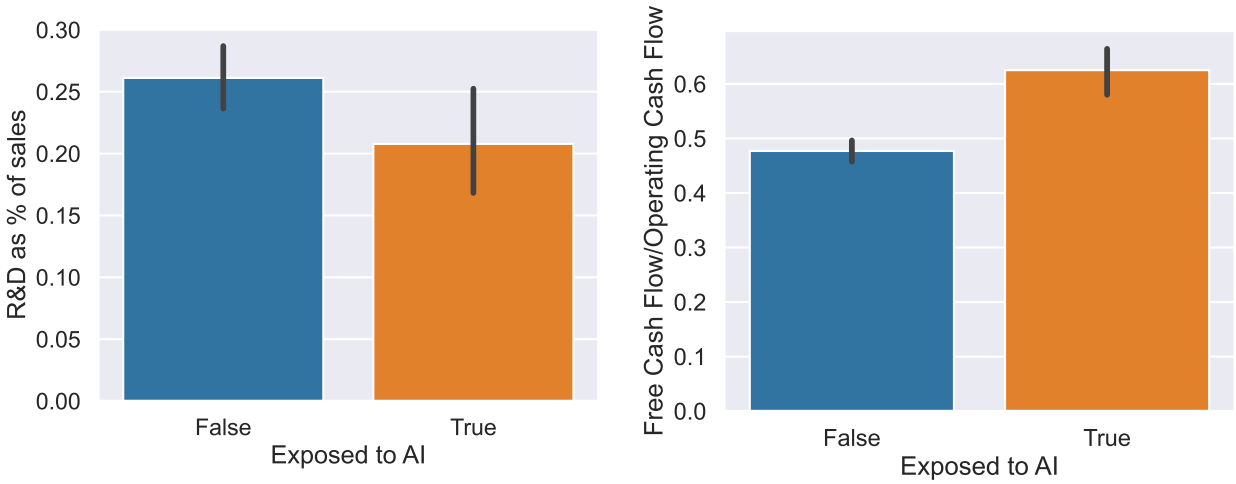
(d) Inventory Turnover

Figure 9: Firm performance and AI



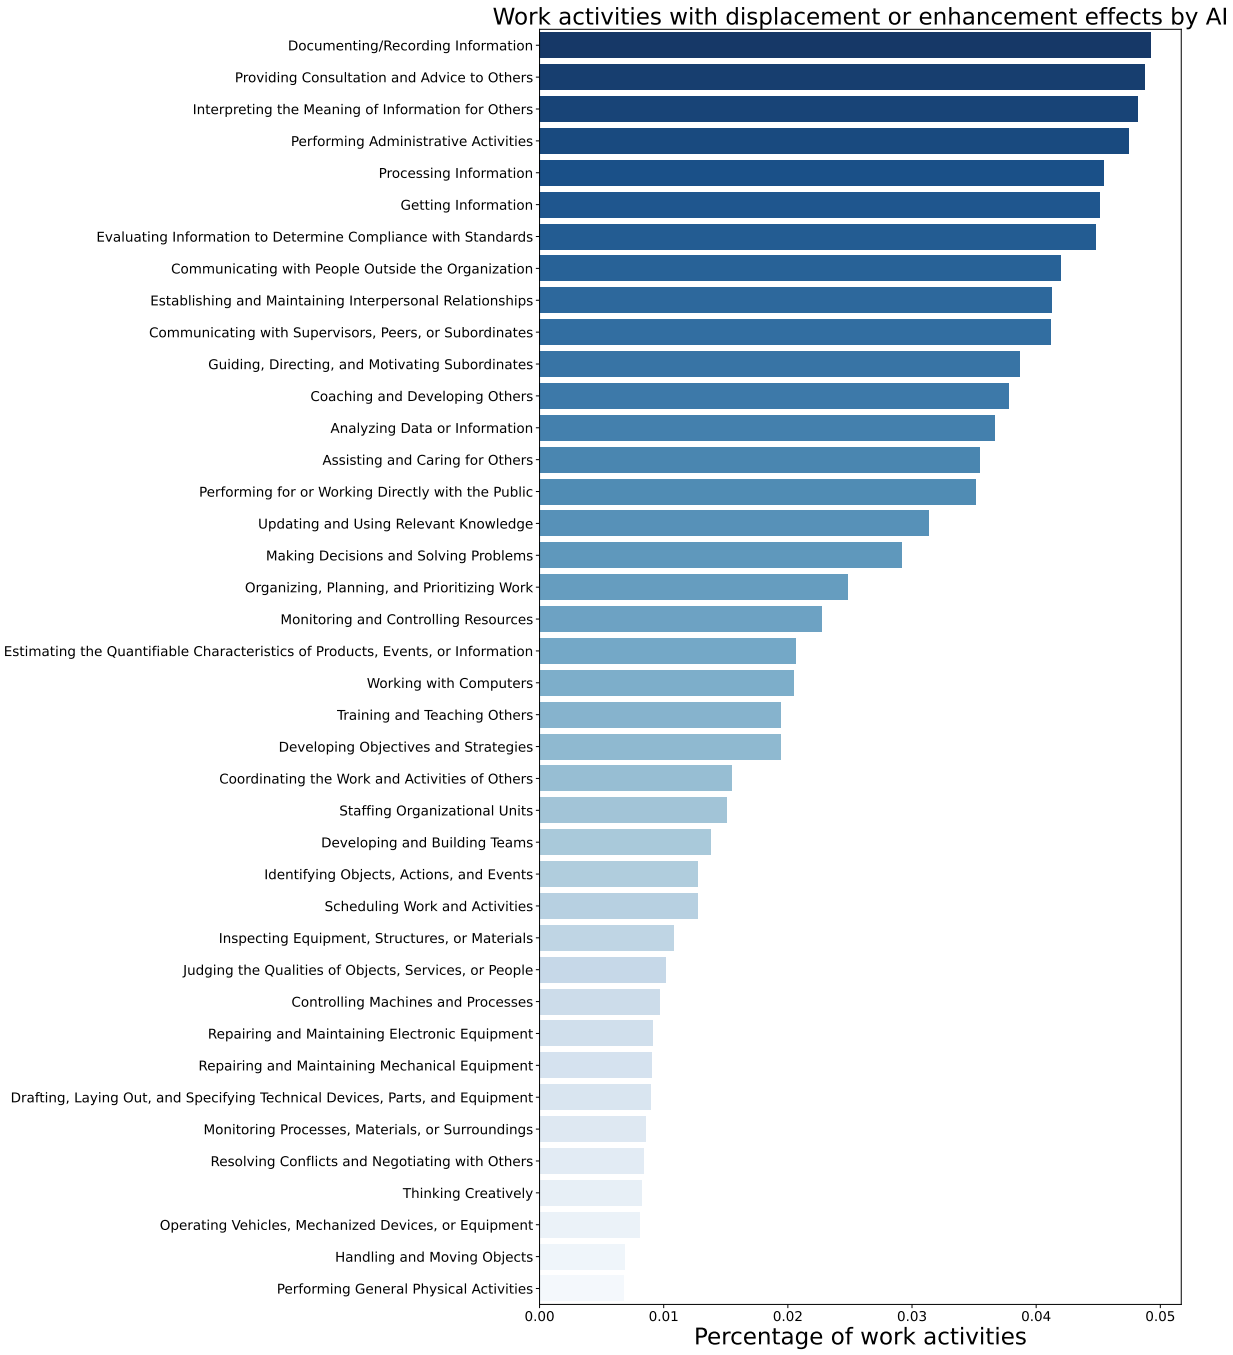
(a) Major Work Content

Figure 10: Most relevant work Activities related to AI

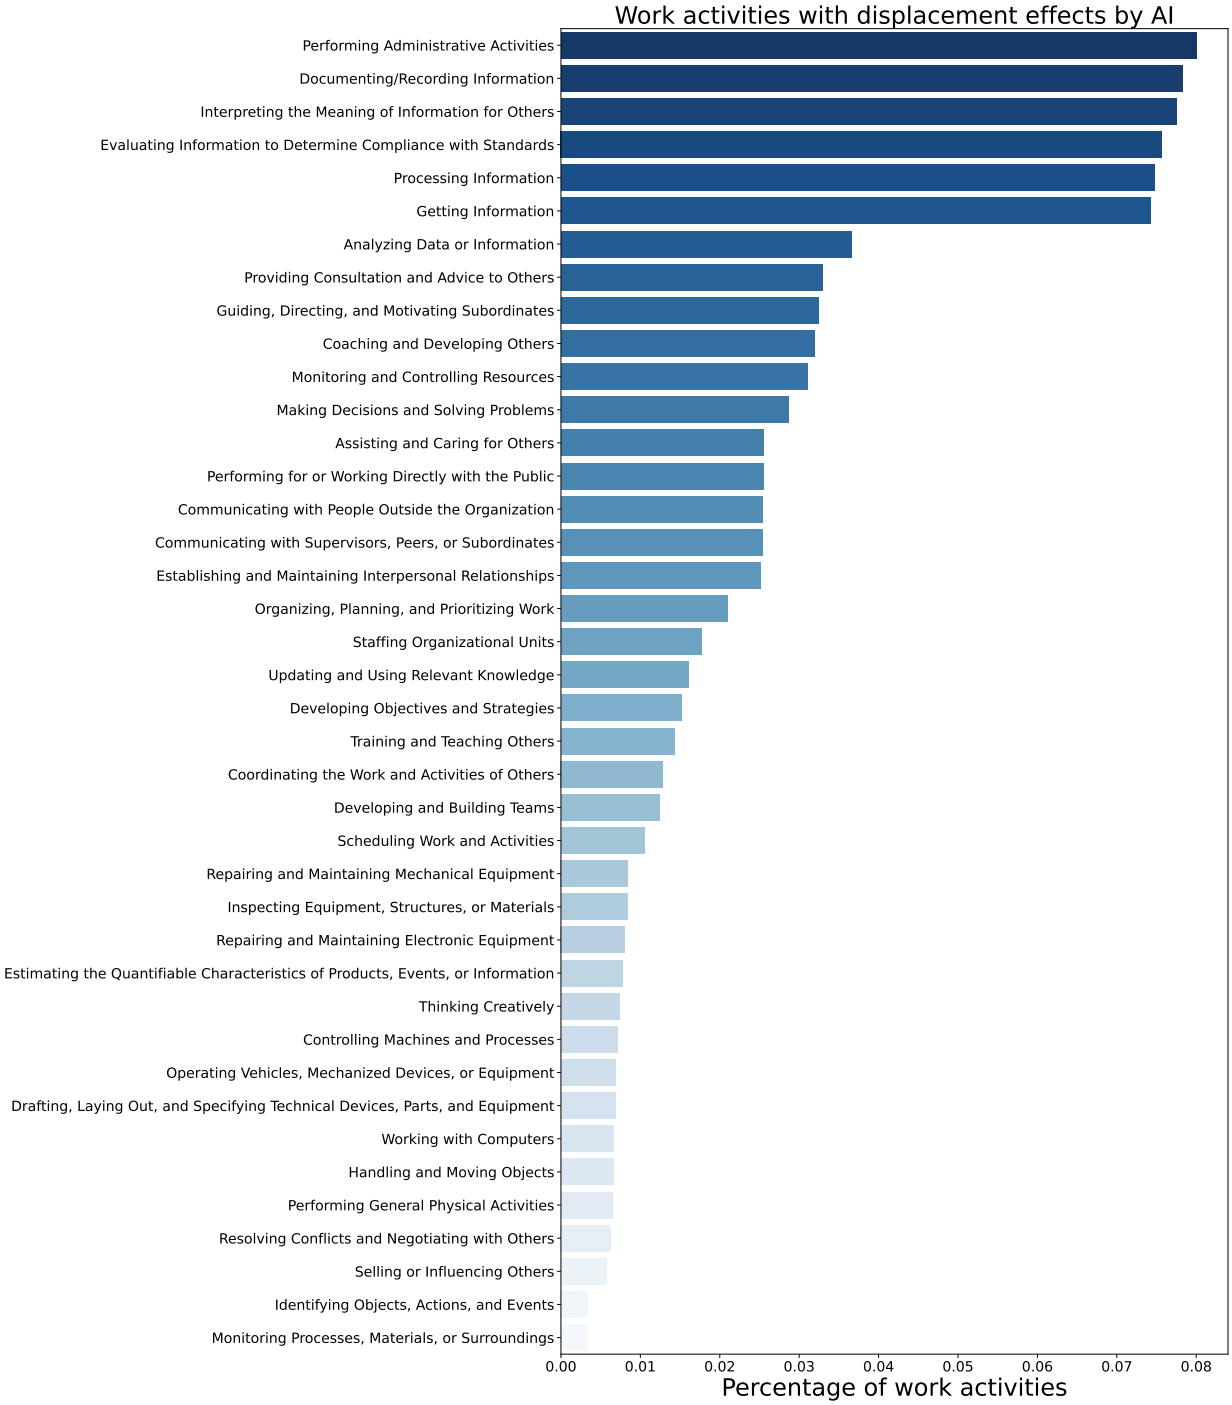


(a) R&D

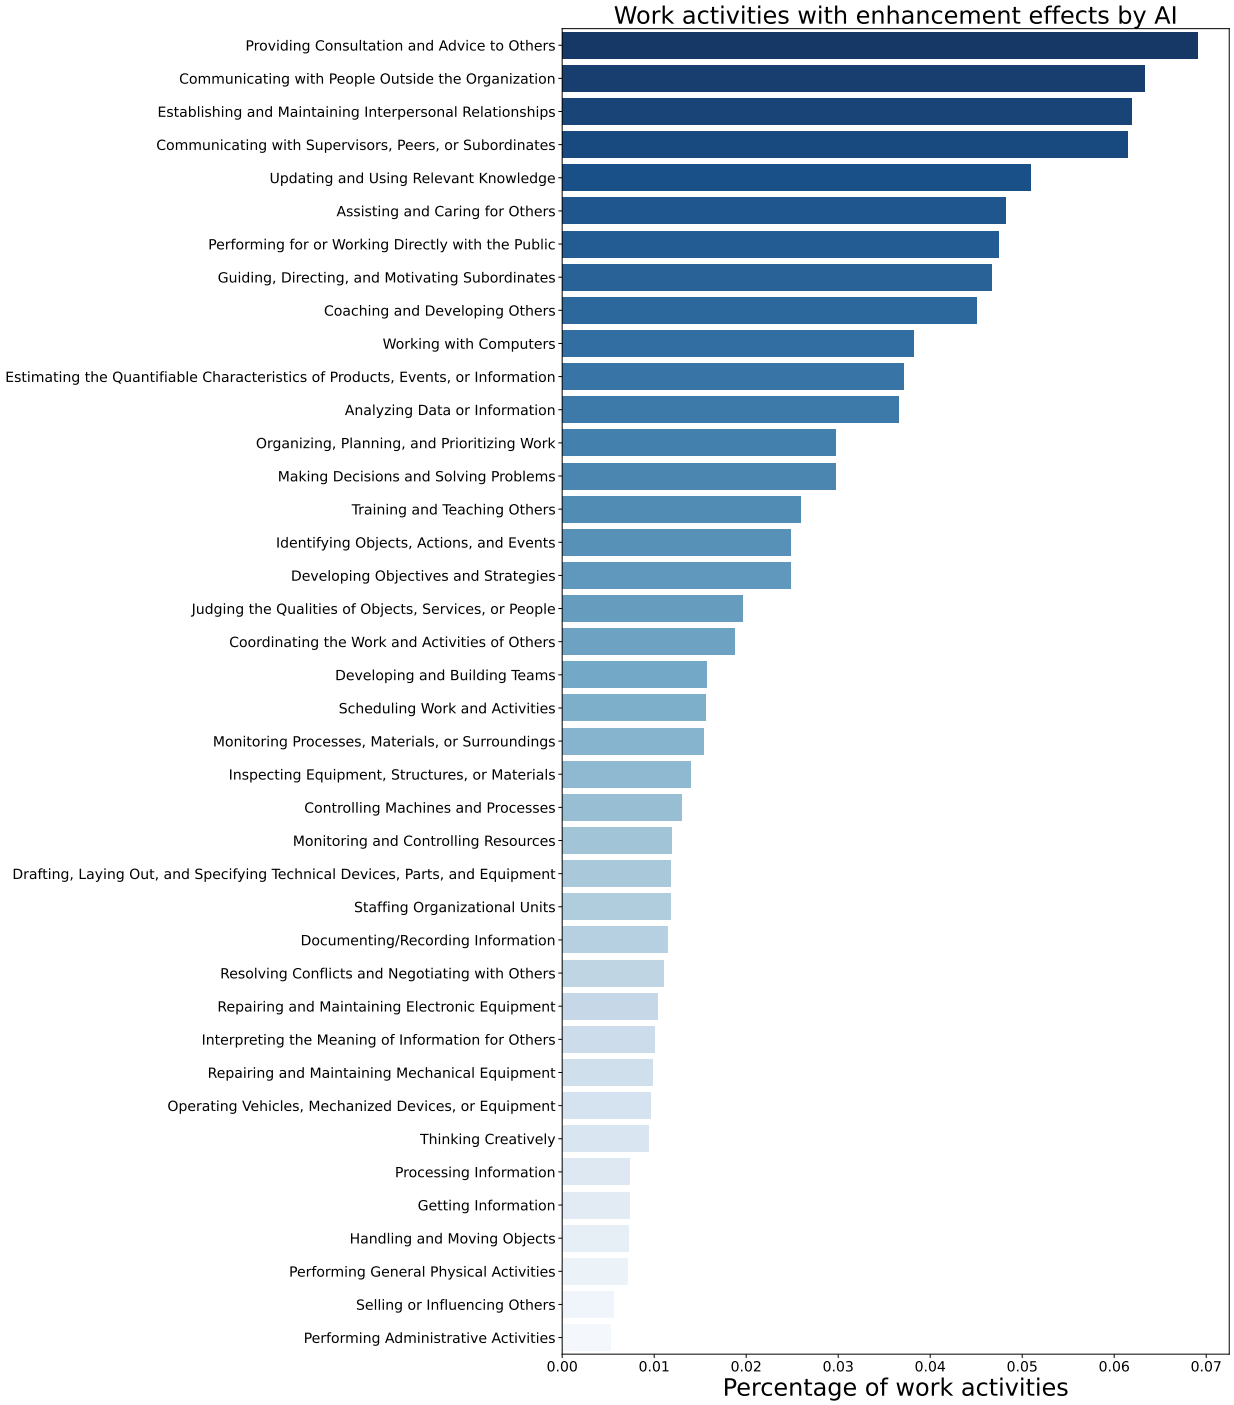
Figure 11: R&D and Free Cash Flow



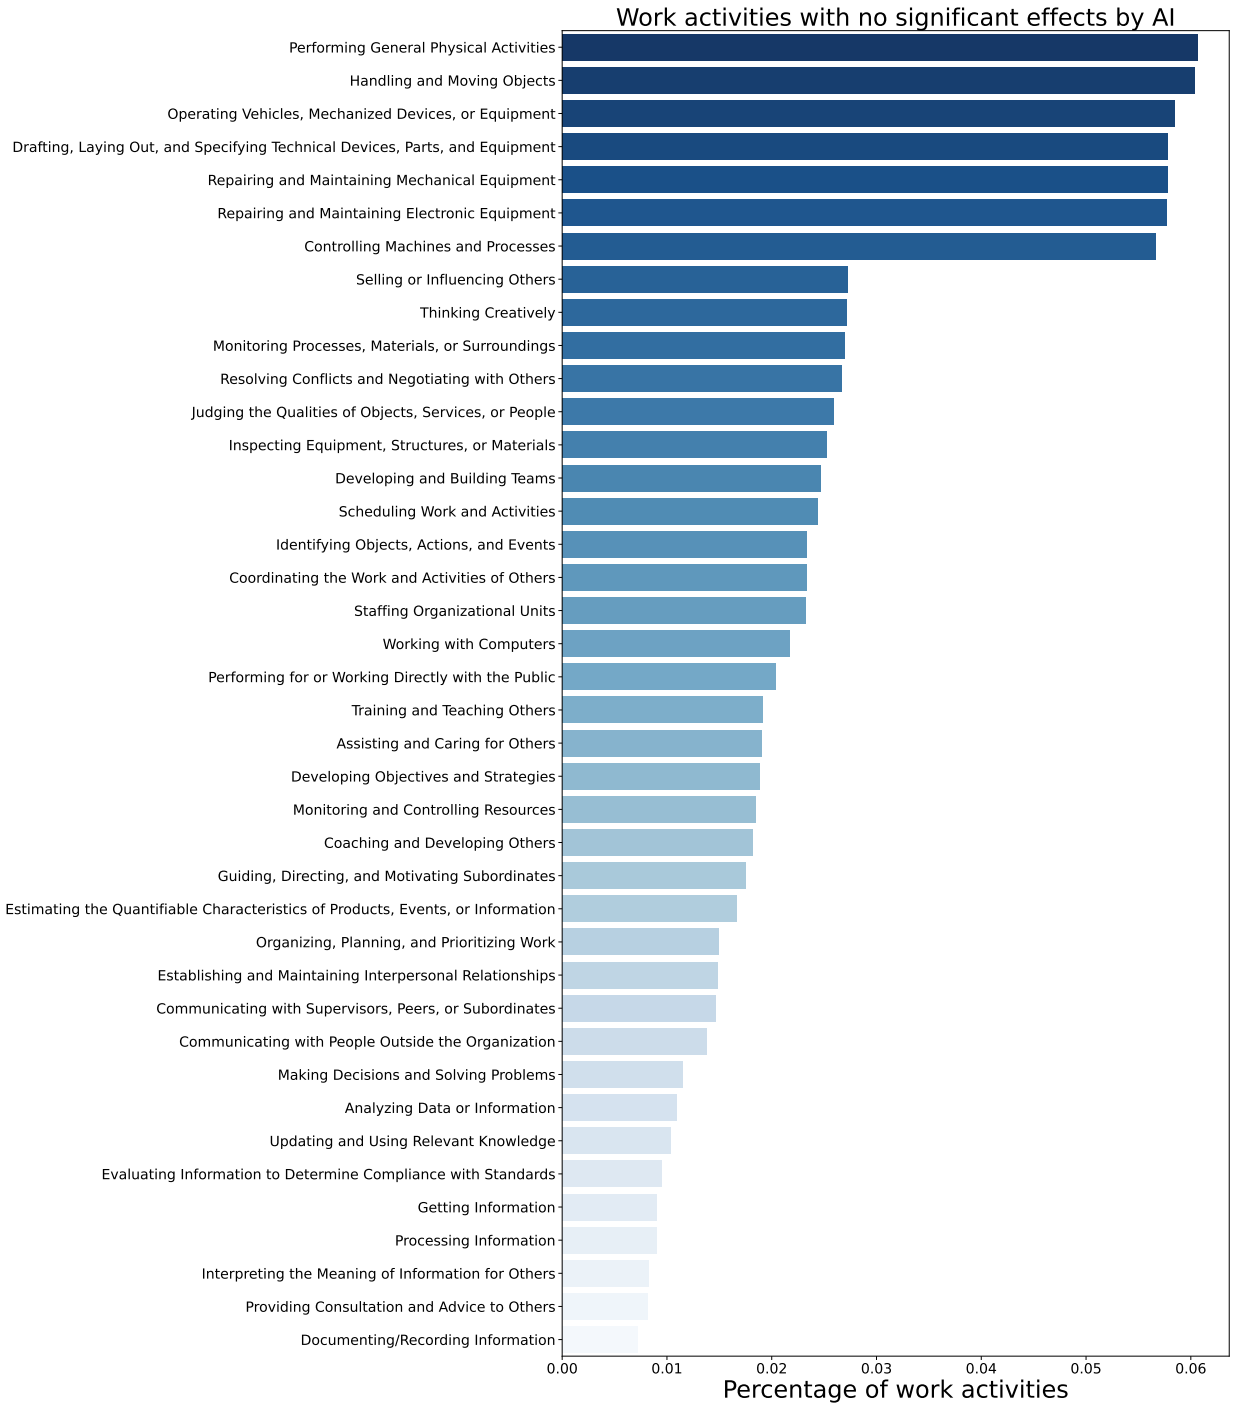
(a) Overall AI exposure on work activities



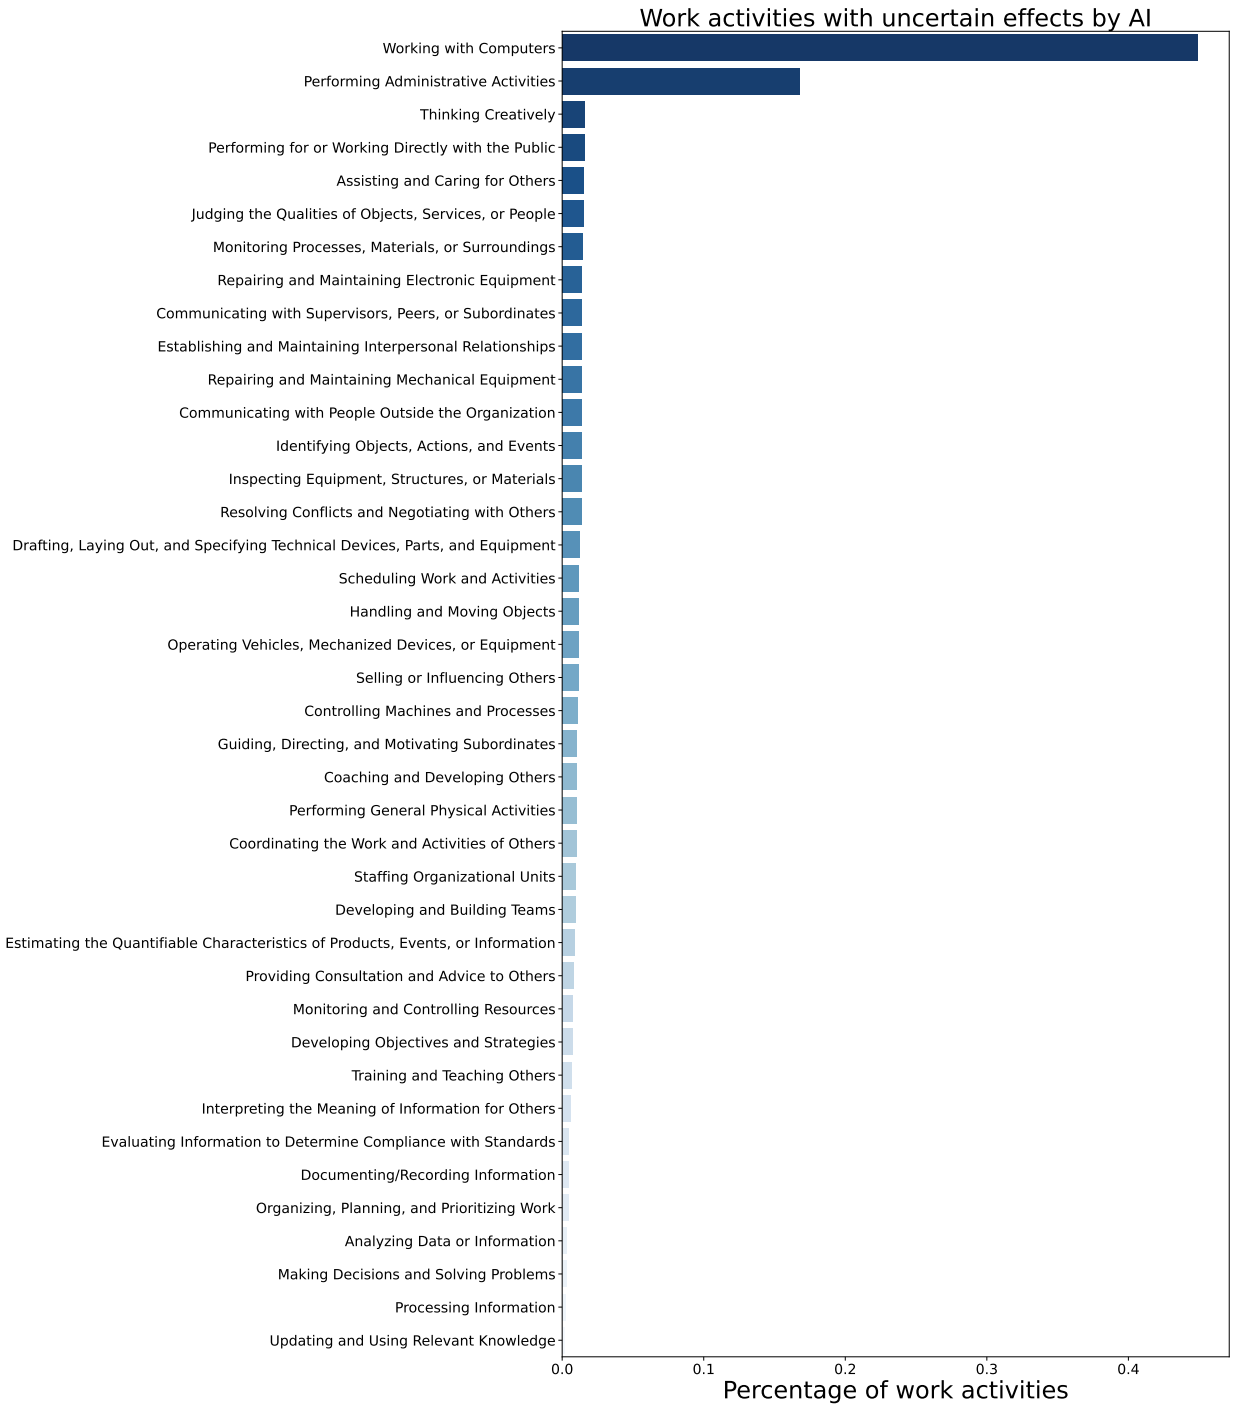
(b) Displacement effect



(c) Enhancement effect



(d) No significant effect



(e) Uncertain effect

Figure 12: AI impact on work activities within firms

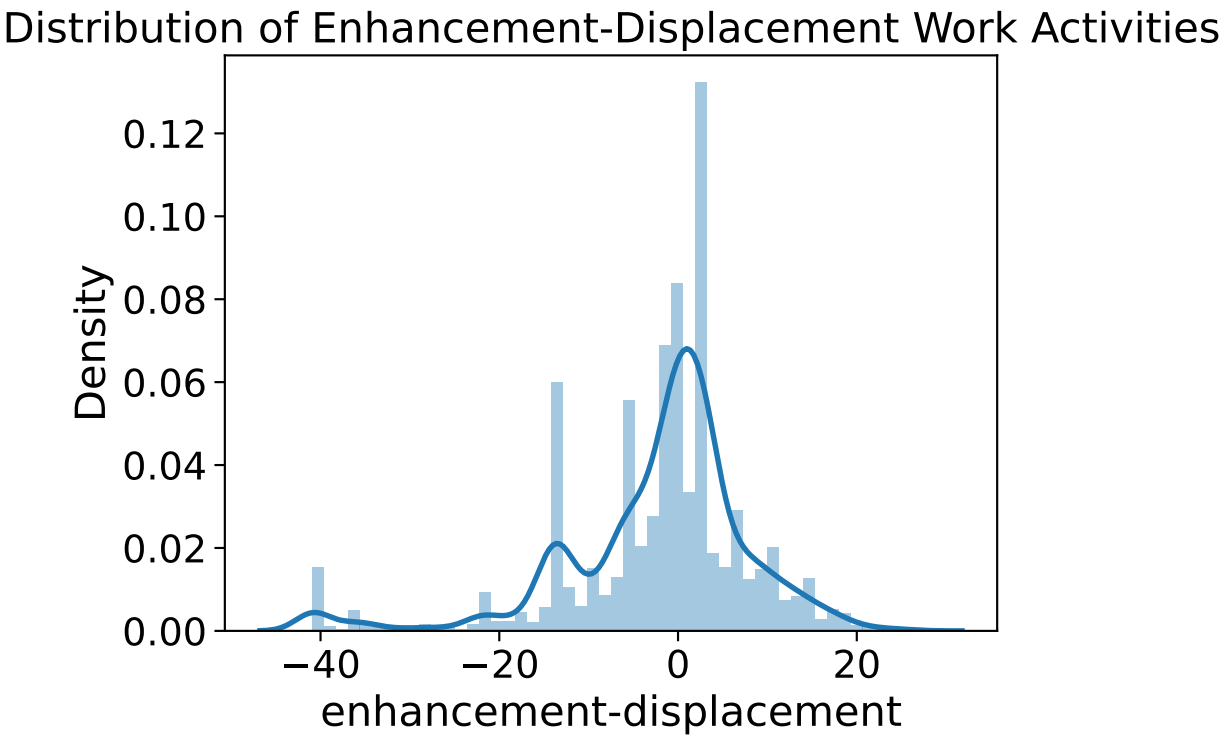


Figure 13: Firm Distribution of Enhancement-Displacement Work Activities

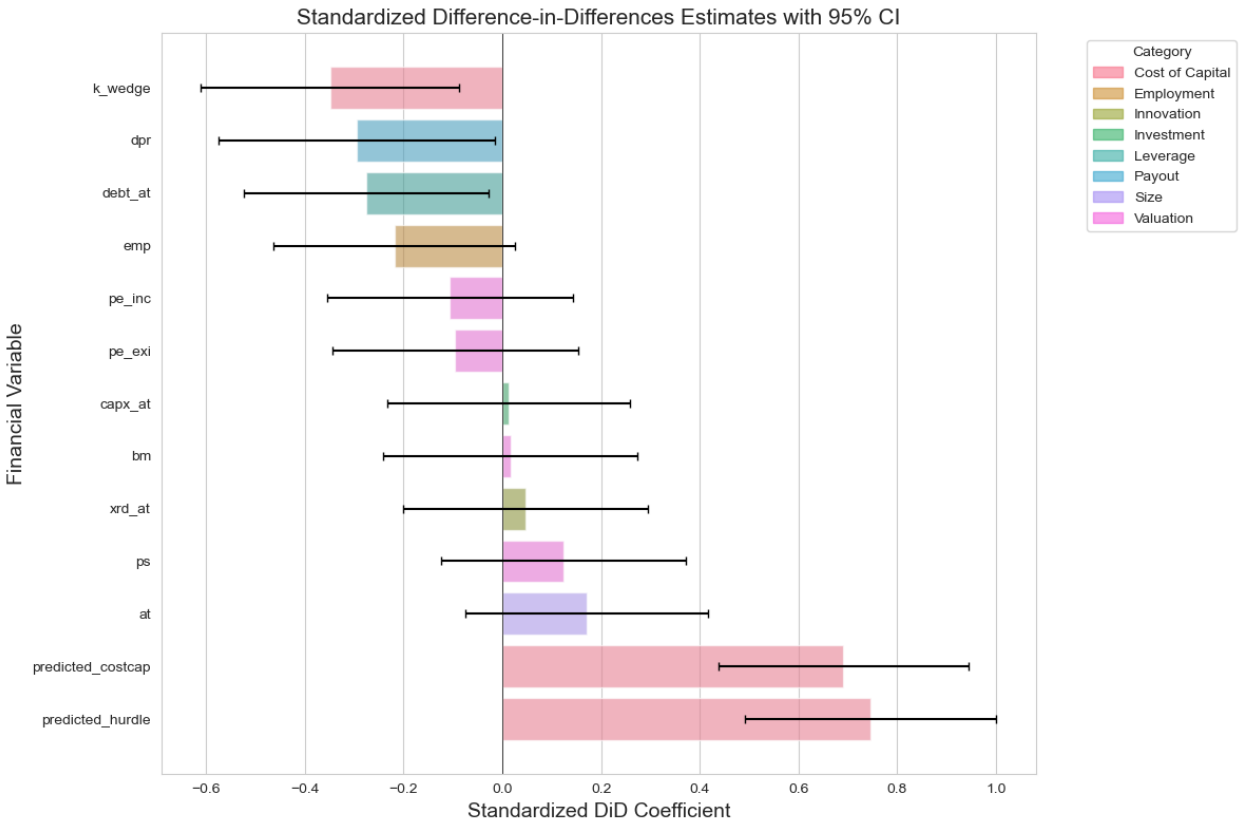


Figure 14: Differences-in-Differences Estimates

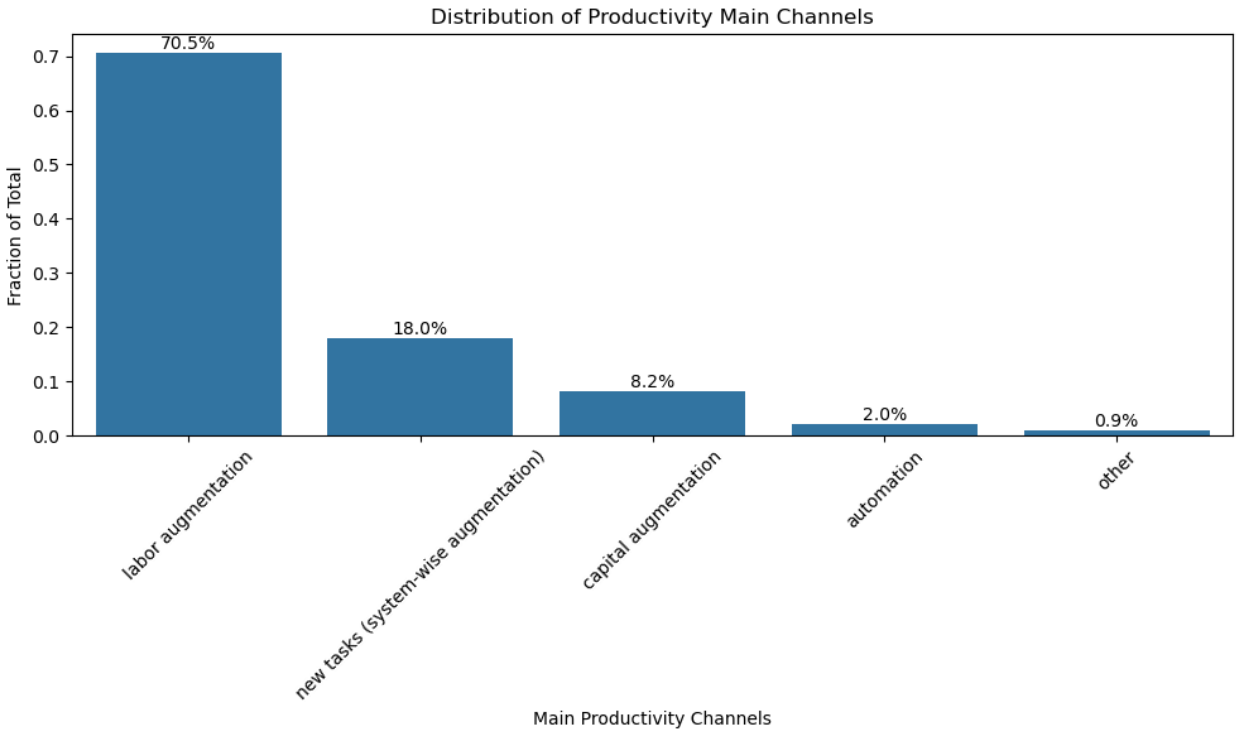


Figure 15: Main Channels of Productivity Gains

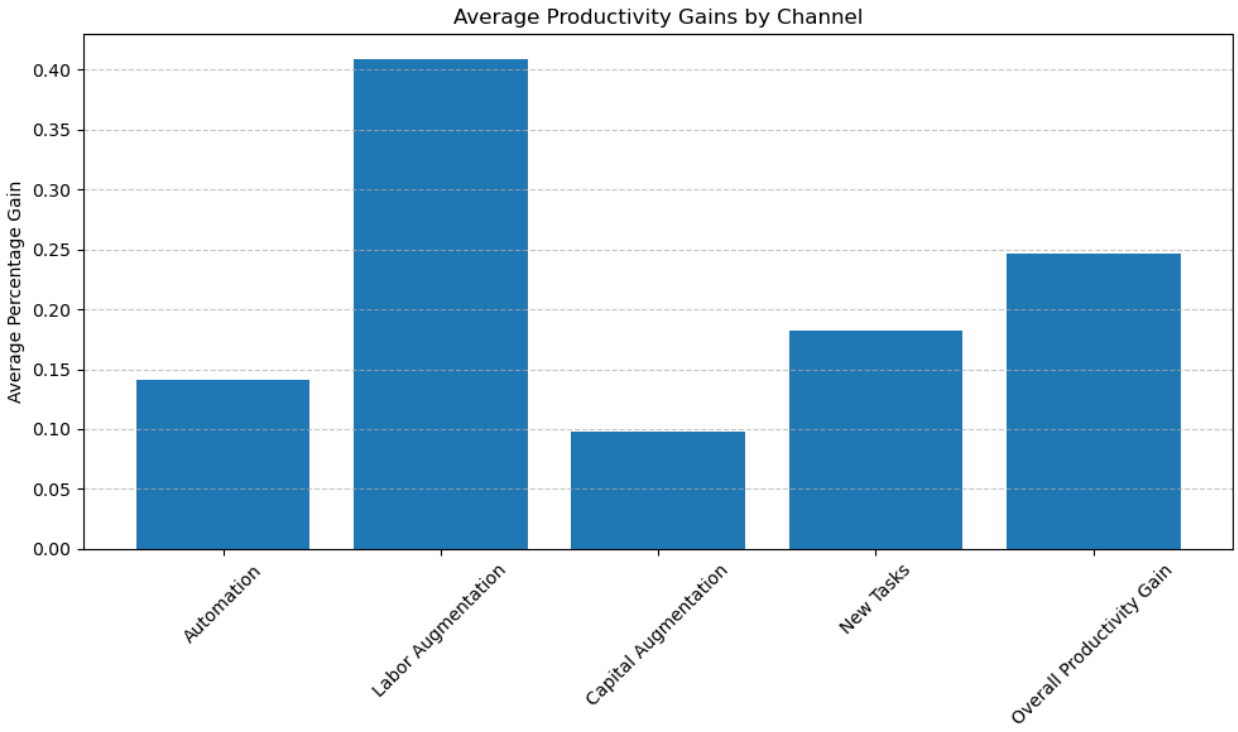


Figure 16: Productivity Impact

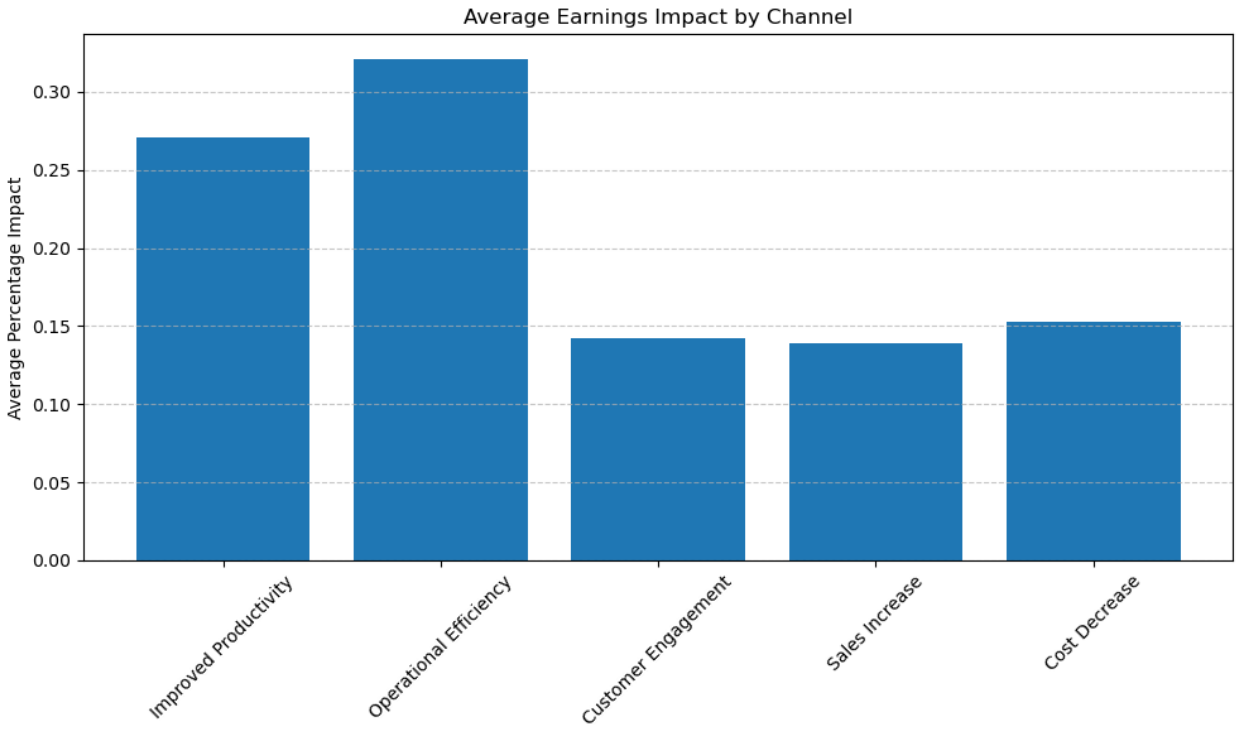


Figure 17: Earnings Impact

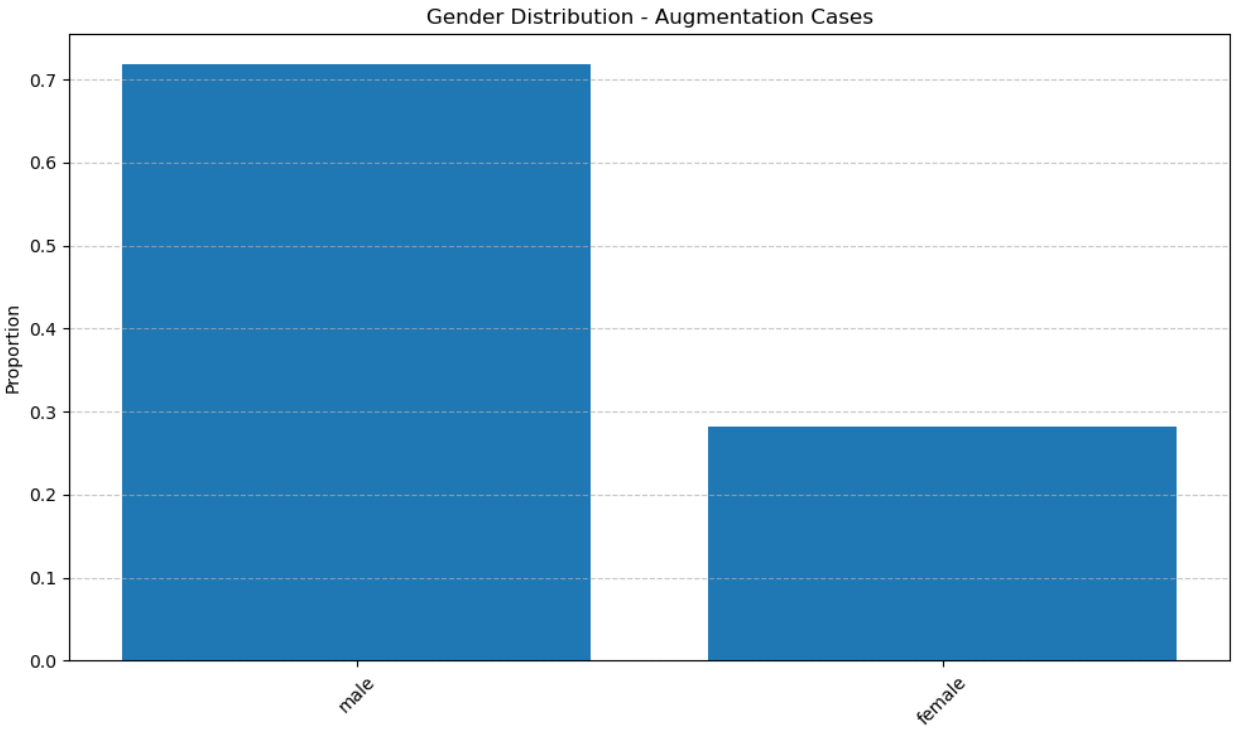


Figure 18: Geneder breakdown of AI-augmented jobs

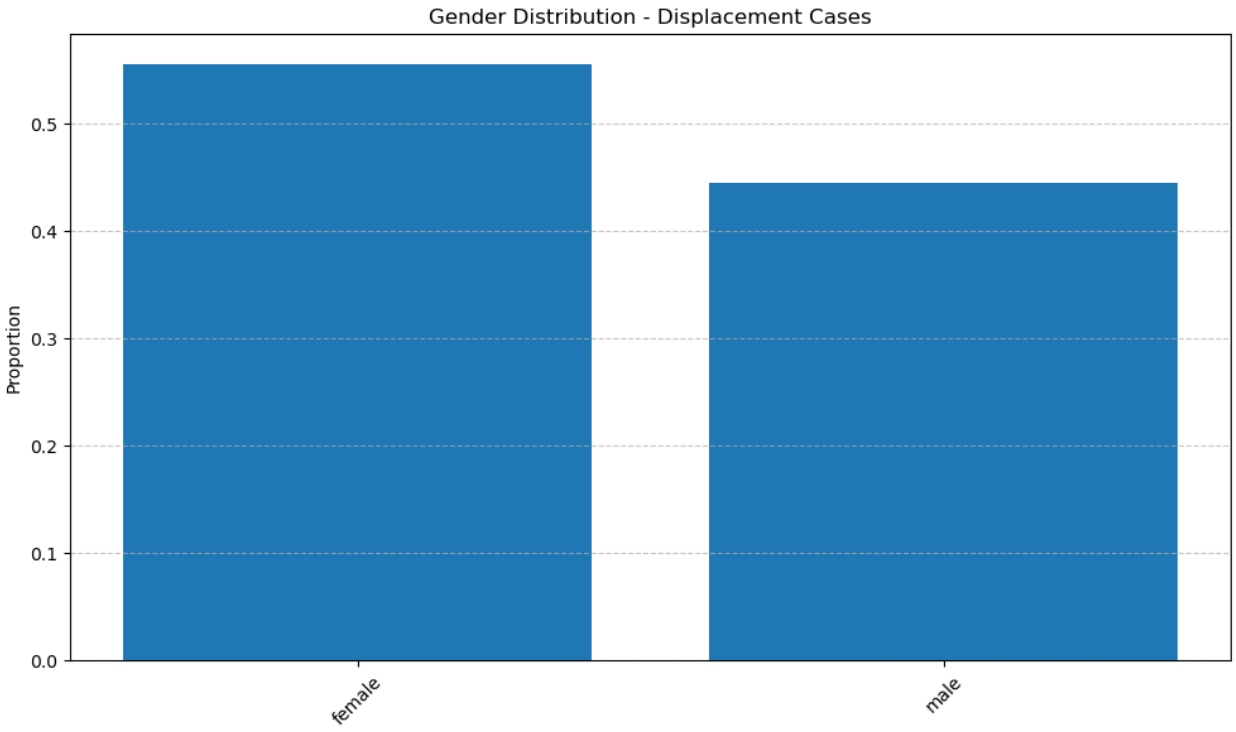


Figure 19: Geneder breakdown of AI-displaced jobs

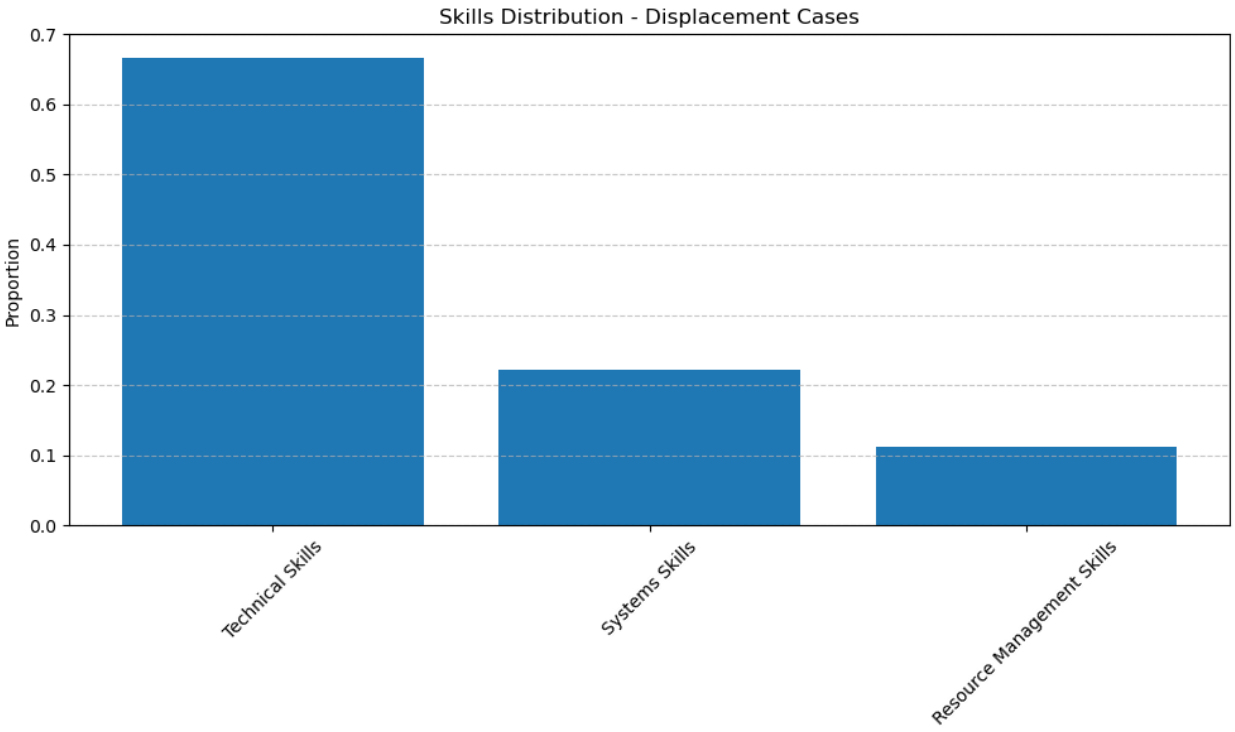


Figure 20: Skills Distribution - Displacement Cases

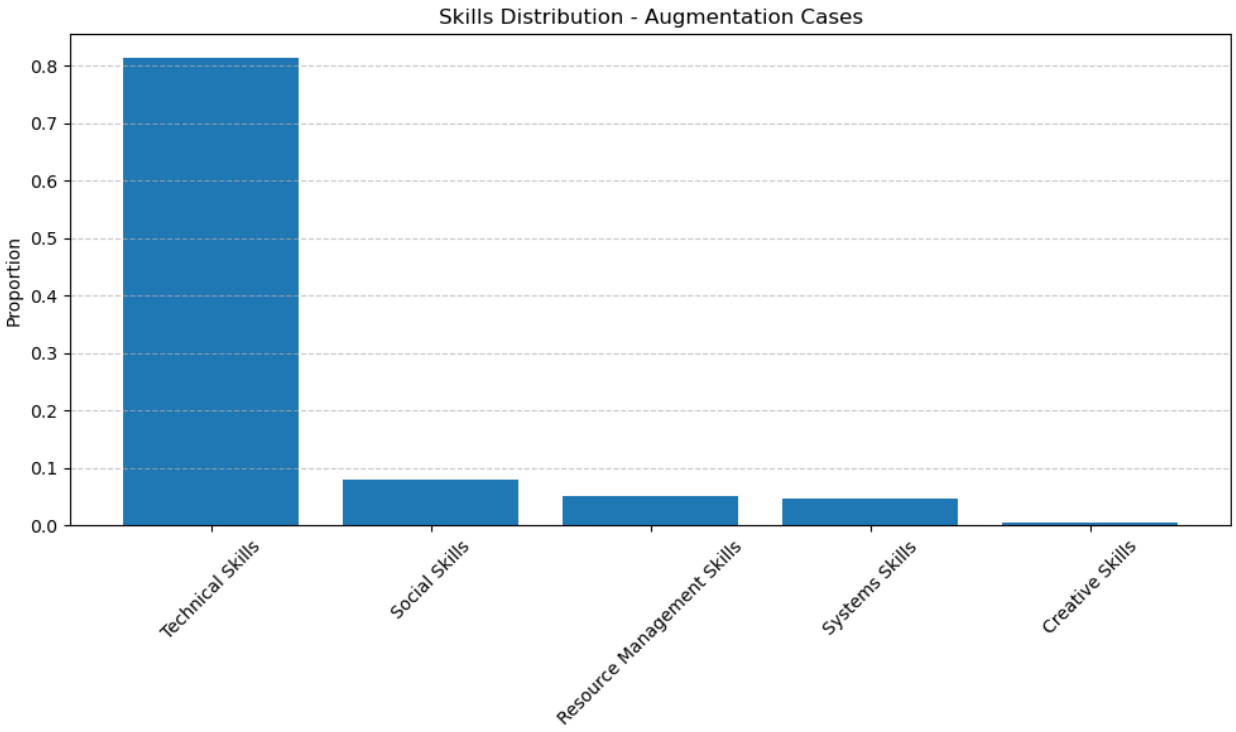


Figure 21: Skills Distribution - Augmentation Cases

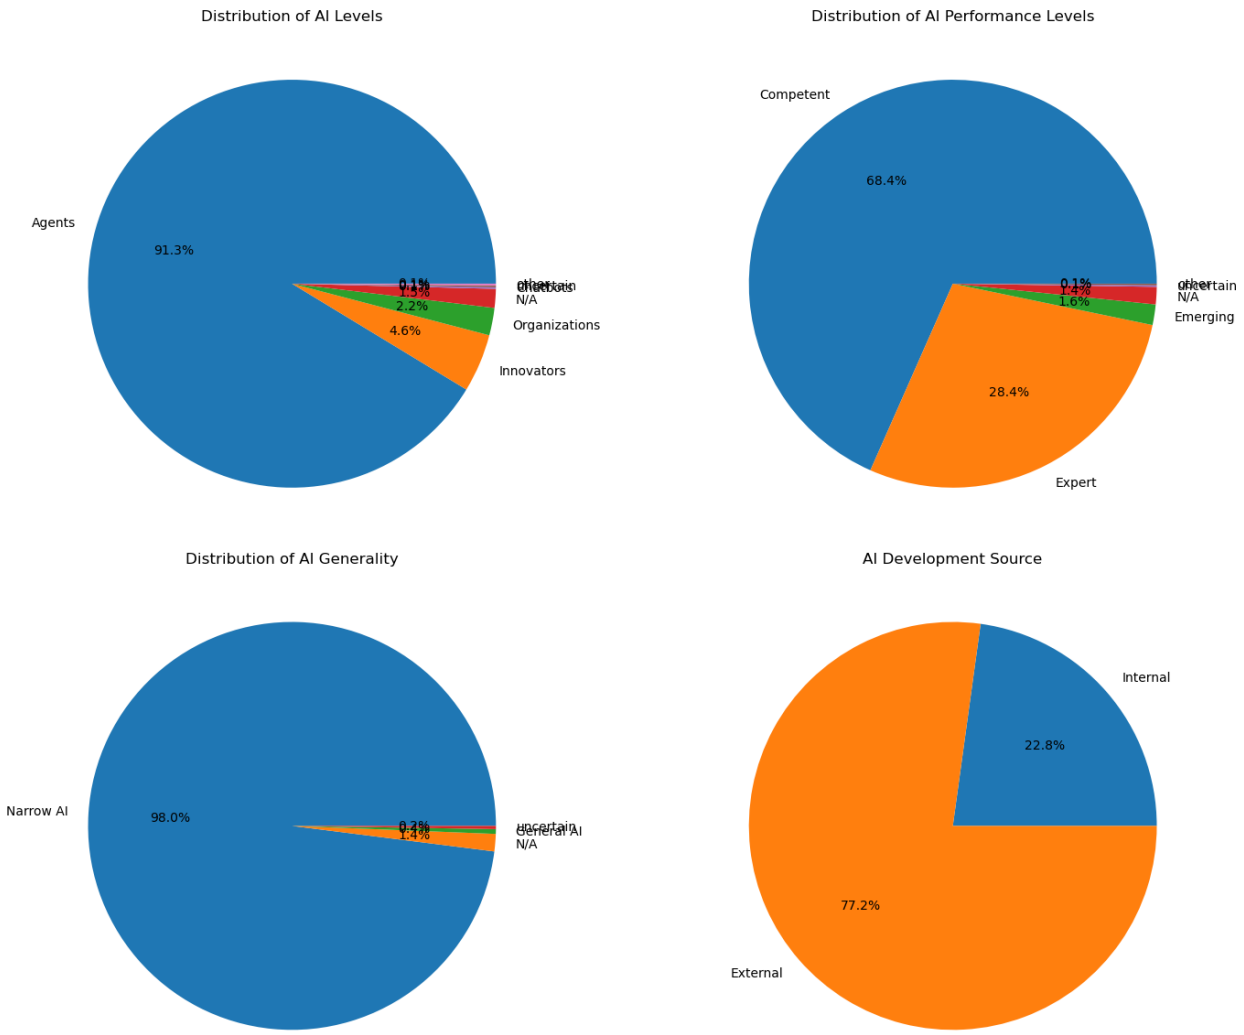


Figure 22: AI Progress and Characteristics

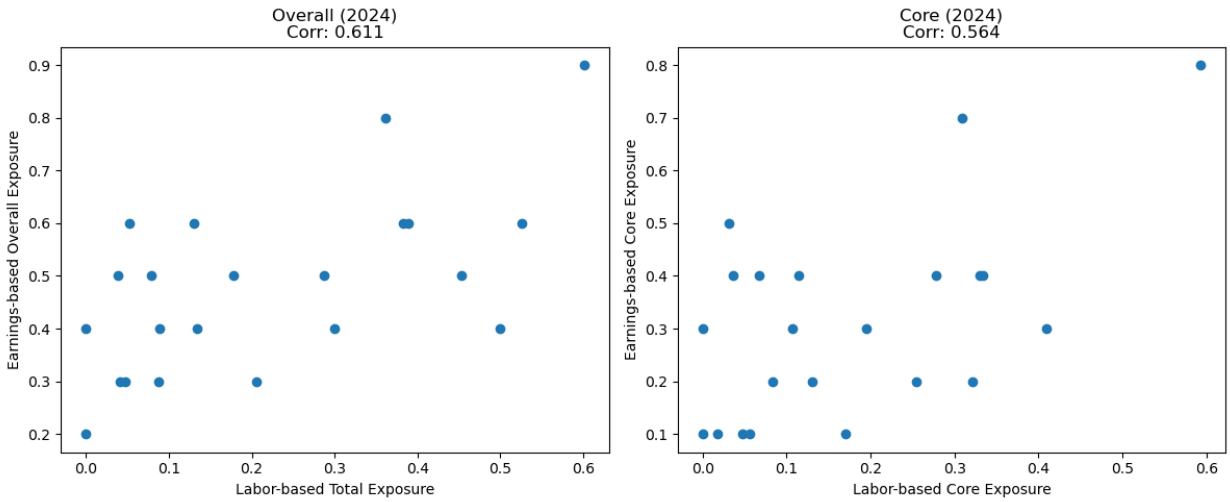


Figure 23: Exposure Correlations - 2024

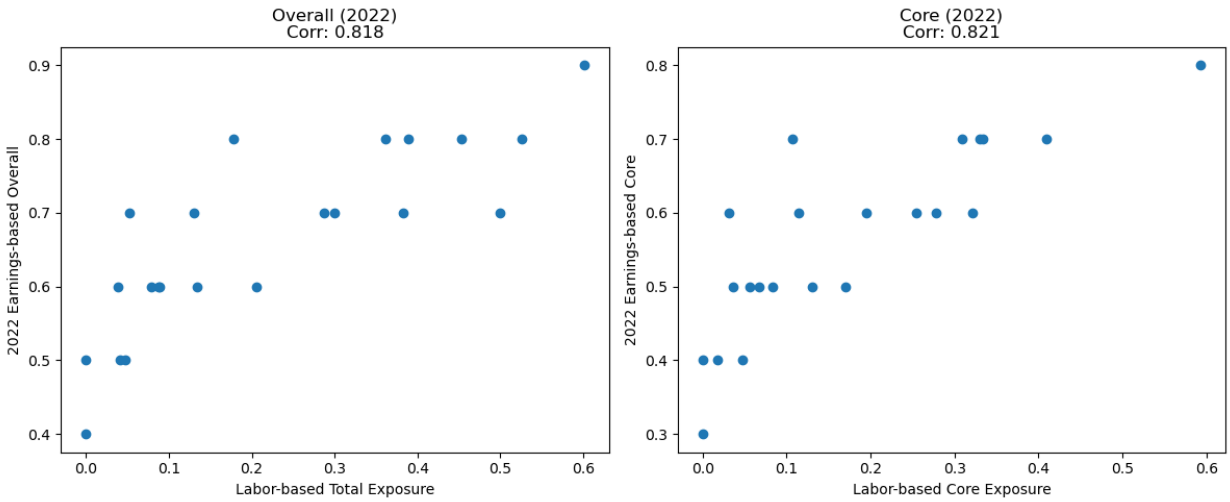


Figure 24: Exposure Correlations - 2022

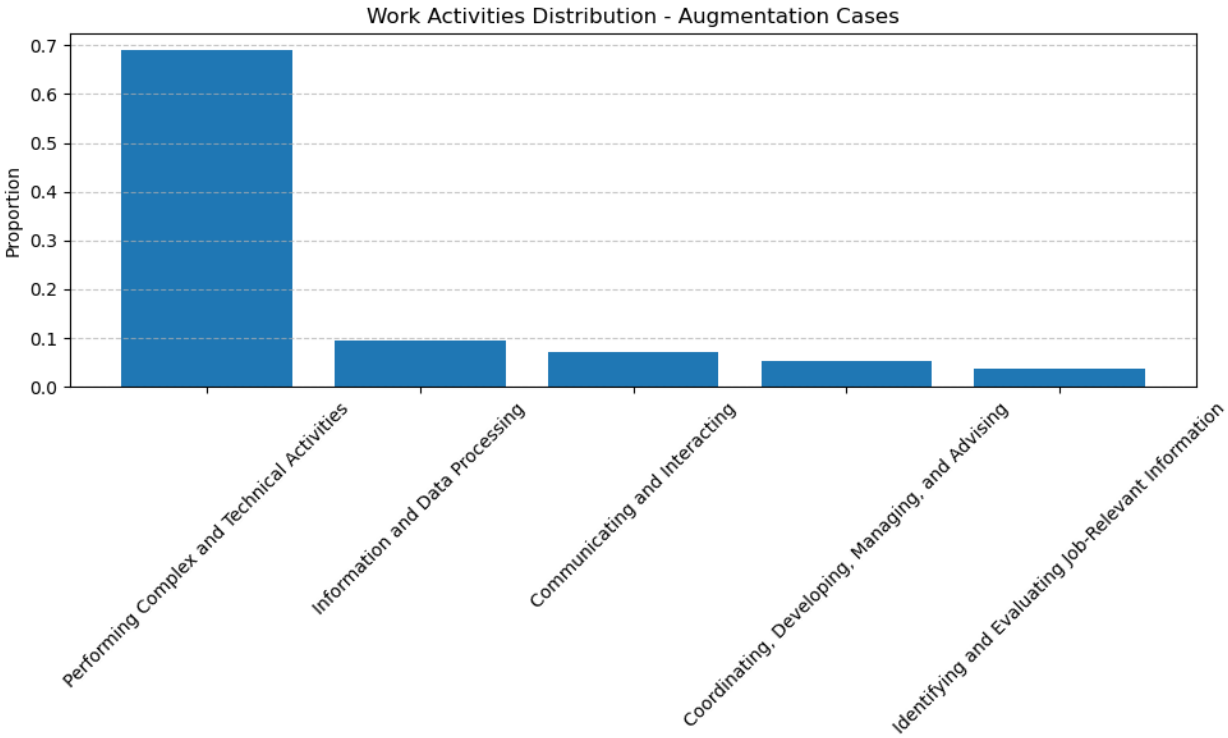


Figure 25: Work Activities Distribution - Augmentation Cases

Table 1: AI Text and Number of Jobs (Narrow AI and All AI): Cross-Sectional

This table presents cross-sectional estimates for the relationship between AI discussion in earnings calls and the number of AI jobs. The dependent variables are the number of narrow AI jobs and the number of all AI jobs. AI text is an indicator equal to one if a firm mentions AI-related terms in its earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio , and return on assets. Standard errors are clustered at the firm-year level.

	(1)	(2)	(3)	(4)
	Narrow AI Jobs	Narrow AI Jobs	All AI Jobs	All AI Jobs
ai_text=1	16.155**	16.158**	16.155**	16.158**
	(2.471)	(2.480)	(2.471)	(2.480)
ln_assets		-0.167		-0.167
		(-0.238)		(-0.238)
ln_sales		1.130		1.130
		(1.377)		(1.377)
ln_employees		0.005		0.005
		(0.011)		(0.011)
firm_age		-0.084		-0.084
		(-1.279)		(-1.279)
leverage_ratio		0.001		0.001
		(0.832)		(0.832)
return_on_assets		-0.065*		-0.065*
		(-2.035)		(-2.035)
Constant	1.982	-4.608	1.982	-4.608
	(1.049)	(-0.316)	(1.049)	(-0.316)
Firm FE	No	No	No	No
Year FE	Yes	Yes	Yes	Yes
4-digit Industry FE	Yes	Yes	Yes	Yes
Observations	10034	10034	10034	10034
Adjusted R-squared	0.114	0.114	0.114	0.114

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2: AI Text and AI Job Shares: Cross-Sectional

This table presents cross-sectional estimates for the relationship between AI discussion in earnings calls and the share of AI jobs. The dependent variables are the share of narrow AI jobs and the share of all AI jobs. AI text is an indicator equal to one if a firm mentions AI-related terms in its earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio , and return on assets. Standard errors are clustered at the firm-year level.

	(1)	(2)	(3)	(4)
	Narrow AI Share	Narrow AI Share	All AI Share	All AI Share
ai_text=1	0.001***	0.001***	0.003***	0.003***
	(3.709)	(3.716)	(7.012)	(7.023)
ln_assets		-0.000		-0.000
		(-0.020)		(-0.290)
ln_sales		-0.000		-0.000
		(-0.996)		(-1.698)
ln_employees		-0.000		-0.000
		(-1.332)		(-1.367)
firm_age		-0.000		-0.000
		(-1.155)		(-0.316)
leverage_ratio		-0.000		0.000
		(-0.466)		(1.564)
return_on_assets		-0.000**		-0.000***
		(-2.456)		(-4.449)
Constant	0.000***	0.001	0.004***	0.005***
	(3.868)	(1.722)	(21.857)	(7.052)
Firm FE	No	No	No	No
Year FE	Yes	Yes	Yes	Yes
4-digit Industry FE	Yes	Yes	Yes	Yes
Observations	10034	10034	10034	10034
Adjusted R-squared	0.188	0.188	0.234	0.234

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: AI Text and AI Job Shares: Event Study

This table presents event study estimates for the effect of AI discussion in earnings calls on AI job shares. The dependent variables are the share of narrow AI jobs and the share of all AI jobs. AI text indicates the first mention of AI-related terms in a firm's earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Post is an indicator equal to one for years on or after a firm's first AI mention. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio , and return on assets . Standard errors are clustered at the firm-year level.

	(1)	(2)	(3)	(4)
	Narrow AI Share	Narrow AI Share	All AI Share	All AI Share
year_rel__5	-0.000	-0.000	-0.000	-0.000
	(-0.994)	(-0.964)	(-0.630)	(-0.625)
year_rel__4	-0.000	-0.000	-0.000	-0.000
	(-0.979)	(-0.968)	(-0.719)	(-0.720)
year_rel__3	-0.000	-0.000	0.000	0.000
	(-0.033)	(-0.036)	(0.283)	(0.279)
year_rel__2	0.000	0.000	0.000	0.000
	(0.375)	(0.357)	(0.768)	(0.746)
year_rel_0	0.001***	0.001***	0.001***	0.001***
	(4.296)	(4.278)	(4.224)	(4.185)
year_rel_1	0.002***	0.002***	0.002***	0.002***
	(5.066)	(4.998)	(4.953)	(4.920)
year_rel_2	0.002***	0.002***	0.003***	0.003***
	(3.907)	(3.882)	(4.380)	(4.352)
year_rel_3	0.002**	0.002**	0.003**	0.003**
	(2.991)	(2.987)	(2.778)	(2.752)
year_rel_4	0.001*	0.001*	0.001	0.001
	(2.008)	(2.005)	(1.805)	(1.799)
year_rel_5	0.001***	0.001***	0.002***	0.002***
	(4.009)	(3.922)	(3.899)	(3.788)
ln_assets		-0.000		-0.000
		(-0.070)		(-0.014)
ln_sales		-0.000		-0.000
		(-0.569)		(-1.230)
ln_employees		-0.000		-0.000
		(-0.816)		(-0.975)
firm_age		-0.000		-0.000
		(-1.154)		(-0.161)
leverage_ratio		0.000		0.000
		(1.325)		(1.053)
return_on_assets		-0.000**		-0.000**
		(-3.019)		(-2.944)
Constant	0.001***	0.001**	0.005***	0.005***
	(13.806)	(3.007)	(63.128)	(15.088)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
4-digit Industry FE	No	Yes	No	Yes
Observations	10034	10034	10034	10034
Adjusted R-squared	0.504	0.493	0.716	0.709

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: AI Text and Number of Jobs (Narrow AI and All AI): Event Study

This table presents event study estimates for the effect of AI discussion in earnings calls on the number of AI jobs. The dependent variables are the number of narrow AI jobs and the number of all AI jobs. AI text indicates the first mention of AI-related terms in a firm's earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Post is an indicator equal to one for years on or after a firm's first AI mention. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio , and return on assets . Standard errors are clustered at the firm-year level.

	(1)	(2)	(3)	(4)
	Narrow AI Jobs	Narrow AI Jobs	All AI Jobs	All AI Jobs
year_rel__5	-10.948	-10.881	-10.948	-10.881
	(-1.553)	(-1.548)	(-1.553)	(-1.548)
year_rel__4	-9.576	-9.519	-9.576	-9.519
	(-1.517)	(-1.478)	(-1.517)	(-1.478)
year_rel__3	-4.916	-4.962	-4.916	-4.962
	(-0.869)	(-0.864)	(-0.869)	(-0.864)
year_rel__2	-1.583	-1.620	-1.583	-1.620
	(-0.302)	(-0.302)	(-0.302)	(-0.302)
year_rel_0	11.364*	11.357*	11.364*	11.357*
	(2.013)	(1.996)	(2.013)	(1.996)
year_rel_1	15.387	15.345	15.387	15.345
	(1.674)	(1.651)	(1.674)	(1.651)
year_rel_2	53.909*	53.852*	53.909*	53.852*
	(2.062)	(2.040)	(2.062)	(2.040)
year_rel_3	77.768	77.830	77.768	77.830
	(1.544)	(1.532)	(1.544)	(1.532)
year_rel_4	32.718*	32.565*	32.718*	32.565*
	(2.017)	(1.996)	(2.017)	(1.996)
year_rel_5	23.283**	23.198**	23.283**	23.198**
	(2.335)	(2.317)	(2.335)	(2.317)
ln_assets		-0.120		-0.120
		(-0.208)		(-0.208)
ln_sales		0.036		0.036
		(0.062)		(0.062)
ln_employees		0.195		0.195
		(0.398)		(0.398)
firm_age		-0.054		-0.054
		(-0.975)		(-0.975)
leverage_ratio		0.000		0.000
		(0.185)		(0.185)
return_on_assets		-0.012		-0.012
		(-0.575)		(-0.575)
Constant	6.072***	7.680	6.072***	7.680
	(7.632)	(0.842)	(7.632)	(0.842)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
4-digit Industry FE	No	Yes	No	Yes
Observations	10034	10034	10034	10034
Adjusted R-squared	0.453	0.441	0.453	0.441

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: AI Text and AI Job Shares: Difference-in-Differences

This table presents difference-in-differences estimates for the effect of AI discussion in earnings calls on AI job shares. The dependent variables are the share of narrow AI jobs and the share of all AI jobs. AI text indicates the first mention of AI-related terms in a firm's earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Post is an indicator equal to one for years on or after a firm's first AI mention. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio , and return on assets . Standard errors are clustered at the firm level.

	(1)	(2)	(3)	(4)
	Narrow AI Share	Narrow AI Share	All AI Share	All AI Share
post $\times ai_text$	0.002***	0.002***	0.002***	0.002***
	(7.892)	(7.824)	(7.521)	(7.466)
ln_assets		-0.000		-0.000
		(-0.173)		(-0.105)
ln_sales		-0.000		-0.000
		(-0.627)		(-1.092)
ln_employees		-0.000		-0.000
		(-0.756)		(-1.105)
firm_age		-0.000		-0.000
		(-1.312)		(-0.169)
leverage_ratio		0.000		0.000
		(1.109)		(1.078)
return_on_assets		-0.000**		-0.000**
		(-2.254)		(-2.276)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
4-digit Industry FE	No	Yes	No	Yes
Observations	10034	10034	10034	10034
Adjusted R-squared	0.505	0.494	0.716	0.710

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: AI Text and Number of Jobs (Narrow AI and All AI): Difference-in-Differences

This table presents difference-in-differences estimates for the effect of AI discussion in earnings calls on the number of AI jobs. The dependent variables are the number of narrow AI jobs and the number of all AI jobs. AI text indicates the first mention of AI-related terms in a firm's earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Post is an indicator equal to one for years on or after a firm's first AI mention. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio , and return on assets . Standard errors are clustered at the firm level.

	(1)	(2)	(3)	(4)
	Narrow AI Jobs	Narrow AI Jobs	All AI Jobs	All AI Jobs
post $\times ai_text$	28.689***	28.675***	28.689***	28.675***
	(3.077)	(3.049)	(3.077)	(3.049)
ln_assets		-0.199		-0.199
		(-0.607)		(-0.607)
ln_sales		-0.060		-0.060
		(-0.162)		(-0.162)
ln_employees		0.271		0.271
		(0.616)		(0.616)
firm_age		-0.051		-0.051
		(-0.984)		(-0.984)
leverage_ratio		0.001		0.001
		(0.516)		(0.516)
return_on_assets		-0.038*		-0.038*
		(-1.907)		(-1.907)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
4-digit Industry FE	No	Yes	No	Yes
Observations	10034	10034	10034	10034
Adjusted R-squared	0.446	0.433	0.446	0.433

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: The growth of AI mentions. Columns 2016 and 2022 contain fractions of firms that mention AI. Avg. Growth is the mean of yearly percentage changes; Std. Growth is the standard deviation of yearly percentage changes. Fractions above 50% are in bold.

Industry	**2016**	**2022**	**Avg. Growth**	**Std. Growth**
Arts, Entertainment, and Recreation	0.03	0.06	0.369	1.043
Construction	0.05	0.08	0.917	2.670
Other Services (except Public Administration)	0.00	0.31	∞	NaN
Agriculture, Forestry, Fishing and Hunting	0.15	0.43	0.213	0.280
Educational Services	0.00	**0.50**	∞	NaN
Utilities	0.04	0.12	0.369	0.792
Accommodation and Food Services	0.00	0.21	∞	NaN
Transportation and Warehousing	0.03	0.17	0.677	1.341
Wholesale Trade	0.02	0.17	0.594	0.628
Health Care and Social Assistance	0.05	0.38	0.434	0.391
Mining, Quarrying, and Oil and Gas Extraction	0.03	0.21	0.407	0.242
Real Estate and Rental and Leasing	0.05	0.15	0.230	0.188
Administrative	0.07	**0.51**	0.420	0.323
Retail Trade	0.07	0.30	0.328	0.356
Professional, Scientific, and Technical Services	0.11	**0.56**	0.337	0.243
Finance and Insurance	0.04	0.22	0.331	0.236
Information	0.12	**0.62**	0.330	0.255
Manufacturing	0.12	0.31	0.175	0.069

Table 8: Frequency of GPT labels for different aspects of impact

This table presents GPT label frequencies calculated within different aspects of impact. The highest frequencies are bolded.

Aspect	**Label**	**Frequency**
Direction	positive	**0.9618**
	negative	0.0382
Significance	moderate	**0.4924**
	high	0.4281
	low	0.0795
Topic	labor	**0.4945**
	investment	0.3174
	revenue	0.0696
	competition	0.0650
	M&A	0.0536
Time	current	**0.4465**
	planning	0.3664
	happened	0.1871
Aggressiveness	active	**0.9324**
	passive	0.0676

Table 9: GPT label for differentiating labor impacts of AI.

This table presents GPT labels for differentiating labor impacts of AI. The highest frequencies are bolded. Aspects include creating new tasks, displacing workers, and mostly labor-displacing or labor-augmenting. Occupations are classified by their SOC codes. Only top 5 occupations are shown, see appendix for full list.

Aspect	Label	Frequency
create new tasks	no	**0.7476**
	uncertain	0.0257
	yes	0.2267
displace workers	no	0.1364
	uncertain	0.3799
	yes	**0.4837**
increase worker productivity	no	0.0037
	uncertain	0.0331
	yes	**0.9633**
mostly labor-displacing	no	**0.6091**
	uncertain	0.3143
	yes	0.0766
mostly labor-augmenting	no	0.2434
	uncertain	0.0761
	yes	**0.6805**
impacted Occupations (SOC)	Architecture and Engineering Occupations	0.1013
	Business and Financial Operations Occupations	0.1708
	Computer and Mathematical Occupations	**0.3828**
	Healthcare Practitioners and Technical Occupations	0.1758
	Office and Administrative Support Occupations	0.1693

Table 10: Firm performance and AI exposure

This table presents regressions of firm financial ratios on AI exposure, according to $FinRatio_{it} = \beta_0 + \beta_1 \text{AI_exposed}_i + \beta\Gamma + \epsilon_{it}$ where $FinRatio$ denotes the financial ratio of interest. The explanatory variable is a dummy variable, $AI_exposed_i$, that equal to one if a firm has communicated about AI in earnings calls and zeroes otherwise. I also include control variables denoted by the vector Γ including the year and industry fixed effects. Industry are defined by their 2-digit NAICS code. The financial ratios used are as follows: In column (1), the dependent variable is labor expense as a fraction of sales. In column (2), the dependent variable is debt to equity ratio. In column (3), the dependent variable is inventory turnover. In column (4), the dependent variable is return on equity. In column (5), the dependent variable is book to market ratio. In column (6), the dependent variable is dividend payout ratio. Finally, in column (7), the dependent variable is gross profit margin.

	staff_sale	de_ratio	inv_turn	roe	bm	dpr	gpm
AI_exposed	-0.01***	-0.14	2.40**	0.04***	-0.11***	-0.06**	0.07***
	(0.00)	(0.13)	(1.01)	(0.01)	(0.01)	(0.03)	(0.02)
Intercept	-0.01	0.61	-14.74	0.13	1.72***	4.96***	0.10
	(0.10)	(6.27)	(41.46)	(0.46)	(0.41)	(1.13)	(1.15)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	11361	11476	8229	10753	10936	7970	11361
Adjusted R^2	0.40	0.06	0.24	0.07	0.18	0.03	0.07

Note: *p<0.1; **p<0.05; ***p<0.01

Table 11: Firm performance and AI exposure 2005-2015

This table presents regressions of firm financial ratios on AI exposure, according to $FinRatio_{it} = \beta_0 + \beta_1 \text{AI_exposed}_i + \beta\Gamma + \epsilon_{it}$ where $FinRatio$ denotes the financial ratio of interest. The explanatory variable is a dummy variable, $AI_exposed_i$, that is equal to one if a firm has communicated about AI in earnings calls since 2017 and zeroes otherwise. We also include control variables denoted by the vector Γ including the year and industry fixed effects. The industry is defined by the firm's 2-digit NAICS code. The financial ratios used are as follows: In column (1), the dependent variable is labor expense as a fraction of sales. In column (2), the dependent variable is debt to equity ratio. In column (3), the dependent variable is inventory turnover. In column (4), the dependent variable is the return on equity. In column (5), the dependent variable is the book-to-market ratio. In column (6), the dependent variable is the dividend payout ratio. Finally, in column (7), the dependent variable is the gross profit margin.

	staff_sale	de_ratio	inv_turn	roe	bm	dpr	gpm
AI_exposed	-0.00***	-0.27***	-0.58	0.09***	-0.12***	-0.05***	0.08***
	(0.00)	(0.03)	(0.51)	(0.00)	(0.01)	(0.01)	(0.01)
Intercept	0.01	2.37***	8.50	-0.32***	0.68***	0.23*	0.07
	(0.02)	(0.58)	(8.39)	(0.08)	(0.09)	(0.13)	(0.24)
Observations	166,846	170,246	127,735	163,123	164,689	120,533	166,931
Adjusted R^2	0.42	0.32	0.13	0.08	0.09	0.05	0.12

Note: *p<0.1; **p<0.05; ***p<0.01

Table 12: AI Text and Number of Jobs (Non-AI, Robot, and High-Wage): Difference-in-Differences

This table presents difference-in-differences estimates for the effect of AI discussion in earnings calls on AI job shares. The dependent variables are the share of narrow AI jobs and the share of all AI jobs. AI text indicates the first mention of AI-related terms in a firm's earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Post is an indicator equal to one for years on or after a firm's first AI mention. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio , and return on assets . Standard errors are clustered at the firm level.

	(1)	(2)	(3)	(4)	(5)	(6)
	Non-AI Jobs	Non-AI Jobs	Robot Jobs	Robot Jobs	High-Wage Jobs	High-Wage Jobs
post_ai_text	236.391	236.740	3.878*	3.880*	40.832	41.093
	(1.007)	(1.000)	(1.850)	(1.833)	(0.830)	(0.827)
ln_assets		12.632		-0.130		-1.825
		(0.545)		(-1.189)		(-0.561)
ln_sales		-47.286**		0.070		-2.199
		(-2.002)		(0.543)		(-0.569)
ln_employees		16.062		-0.009		2.645
		(0.721)		(-0.123)		(0.687)
firm_age		-3.199*		-0.001		0.006
		(-1.835)		(-0.081)		(0.033)
leverage_ratio		0.030		0.000		0.002
		(0.679)		(0.996)		(0.345)
return_on_assets		0.184		-0.009**		-0.151
		(0.111)		(-1.999)		(-1.073)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
4-digit Industry FE	No	Yes	No	Yes	No	Yes
Observations	10034	10034	10034	10034	10034	10034
Adjusted R-squared	0.738	0.732	0.229	0.211	0.760	0.755

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 13: AI Text and Number of Jobs (Non-AI, Robot, and High-Wage): Event Study

This table presents event study estimates for the effect of AI discussion in earnings calls on robot and high-wage job shares. The dependent variables are the share of robot jobs and the share of high-wage jobs. AI text indicates the first mention of AI-related terms in a firm's earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Post is an indicator equal to one for years on or after a firm's first AI mention. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio, and return on assets. Standard errors are clustered at the firm level.

	(1)	(2)	(3)	(4)	(5)	(6)
	Non-AI Jobs	Non-AI Jobs	Robot Jobs	Robot Jobs	High-Wage Jobs	High-Wage Jobs
year_rel__5	-336.527**	-339.513*	-1.431	-1.407	-32.436	-32.593
	(-2.465)	(-2.221)	(-1.352)	(-1.339)	(-1.271)	(-1.174)
year_rel__4	98.287	100.157	-0.901	-0.901	-4.430	-4.554
	(0.488)	(0.480)	(-1.038)	(-1.021)	(-0.148)	(-0.150)
year_rel__3	339.098	335.373	-0.253	-0.254	21.382	21.135
	(1.571)	(1.523)	(-0.332)	(-0.328)	(0.604)	(0.585)
year_rel__2	201.820	196.274	-0.382	-0.382	-4.131	-4.366
	(0.639)	(0.612)	(-0.359)	(-0.353)	(-0.175)	(-0.185)
year_rel_0	332.829	335.189	0.678	0.675	44.994	45.391
	(1.134)	(1.138)	(0.874)	(0.847)	(1.183)	(1.171)
year_rel_1	169.112	164.065	1.769	1.773	-9.265	-9.216
	(0.705)	(0.671)	(1.064)	(1.049)	(-0.282)	(-0.273)
year_rel_2	296.030	296.798	10.608	10.597	56.521	56.529
	(0.861)	(0.849)	(1.140)	(1.128)	(0.865)	(0.853)
year_rel_3	466.422	459.198	14.213	14.236	115.856	115.038
	(0.890)	(0.873)	(1.480)	(1.467)	(1.048)	(1.026)
year_rel_4	311.079	303.823	4.996	4.992	53.130	52.650
	(0.870)	(0.839)	(1.796)	(1.776)	(0.807)	(0.789)
year_rel_5	77.408	67.138	4.746	4.761	1.796	1.689
	(0.332)	(0.287)	(1.740)	(1.714)	(0.053)	(0.049)
ln_assets		15.121		-0.115		-1.598
		(0.712)		(-0.863)		(-0.431)
ln_sales		-47.379		0.090		-2.182
		(-1.790)		(0.653)		(-0.526)
ln_employees		16.798		-0.028		2.613
		(0.743)		(-0.210)		(0.815)
firm_age		-3.136*		-0.002		0.002
		(-1.900)		(-0.184)		(0.010)
leverage_ratio		0.026		0.000		0.002
		(0.607)		(0.822)		(0.310)
return_on_assets		0.440		-0.004		-0.077
		(0.344)		(-0.778)		(-0.420)
Constant	1948.963***	2258.565***	1.537***	2.009	171.348***	197.979***
	(50.655)	(5.295)	(6.091)	(1.041)	(63.948)	(3.674)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
4-digit Industry FE	No	Yes	No	Yes	No	Yes
Observations	10034	10034	10034	10034	10034	10034
Adjusted R-squared	0.738	0.732	0.235	0.218	0.760	0.755

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 14: AI Text and Number of Jobs (Non-AI, Robot, and High-Wage): Cross-Sectional

This table presents cross-sectional estimates for the relationship between AI discussion in earnings calls and the number of non-AI, robot, and high-wage jobs. The dependent variables are the number of non-AI jobs, the number of robot jobs, and the number of high-wage jobs. AI text is an indicator equal to one if a firm mentions AI-related terms in its earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio , and return on assets. Standard errors are clustered at the firm-year level.

	(1)	(2)	(3)	(4)	(5)	(6)
	Non-AI Jobs	Non-AI Jobs	Robot Jobs	Robot Jobs	High-Wage Jobs	High-Wage Jobs
ai_text=0	0.000	0.000	0.000	0.000	0.000	0.000
	(.)	(.)	(.)	(.)	(.)	(.)
ai_text=1	1402.090***	1403.350***	1.374*	1.375*	134.570**	134.599**
	(3.640)	(3.643)	(1.866)	(1.876)	(2.674)	(2.677)
ln_assets		74.780*		-0.045		0.957
		(1.967)		(-0.566)		(0.165)
ln_sales		4.278		0.213		6.090
		(0.083)		(1.390)		(0.890)
ln_employees		5.800		-0.048		-2.180
		(0.089)		(-0.323)		(-0.598)
firm_age		-0.163		-0.005		0.340
		(-0.045)		(-0.621)		(0.677)
leverage_ratio		0.067		0.000		-0.008
		(0.660)		(1.383)		(-0.569)
return_on_assets		-1.529		-0.011**		-0.513
		(-0.409)		(-2.308)		(-1.695)
Constant	1477.771***	661.637	1.358***	0.155	125.360***	58.084
	(9.772)	(0.781)	(6.518)	(0.084)	(5.168)	(0.483)
Firm FE	No	No	No	No	No	No
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
4-digit Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10034	10034	10034	10034	10034	10034
Adjusted R-squared	0.096	0.095	0.148	0.148	0.145	0.145

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 15: AI Text and Robot and High-Wage Job Shares: Difference-in-Differences

This table presents difference-in-differences estimates for the effect of AI discussion in earnings calls on robot and high-wage job shares. The dependent variables are the share of robot jobs and the share of high-wage jobs. AI text indicates the first mention of AI-related terms in a firm's earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Post is an indicator equal to one for years on or after a firm's first AI mention. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio, and return on assets. Standard errors are clustered at the firm level.

	(1)	(2)	(3)	(4)
	Robot Share	Robot Share	High-Wage Share	High-Wage Share
post_ai_text	0.000	0.000	-0.005*	-0.005*
	(0.961)	(0.940)	(-1.715)	(-1.667)
ln_assets		0.000		-0.001
		(1.281)		(-0.889)
ln_sales		0.000		-0.000
		(0.110)		(-0.262)
ln_employees		-0.000		-0.000
		(-0.029)		(-0.604)
firm_age		-0.000		0.000**
		(-0.659)		(1.995)
leverage_ratio		0.000		-0.000
		(0.657)		(-0.412)
return_on_assets		0.000		-0.000*
		(0.925)		(-1.921)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
4-digit Industry FE	No	Yes	No	Yes
Observations	10034	10034	10034	10034
Adjusted R-squared	0.636	0.628	0.649	0.641

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 16: AI Text and Robot and High-Wage Job Shares: Event Study

This table presents event study estimates for the effect of AI discussion in earnings calls on robot and high-wage job shares. The dependent variables are the share of robot jobs and the share of high-wage jobs. AI text indicates the first mention of AI-related terms in a firm's earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Post is an indicator equal to one for years on or after a firm's first AI mention. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio , and return on assets . Standard errors are clustered at the firm-year level.

	(1)	(2)	(3)	(4)
	Robot Share	Robot Share	High-Wage Share	High-Wage Share
year_rel__5	-0.000	-0.000	0.008	0.008
	(-0.514)	(.)	(1.725)	(1.703)
year_rel__4	-0.000	-0.000	0.008	0.008
	(-1.201)	(.)	(1.260)	(1.221)
year_rel__3	-0.000	-0.000	0.002	0.002
	(-0.255)	(.)	(0.408)	(0.412)
year_rel__2	-0.000	-0.000	-0.006	-0.006
	(-1.230)	(.)	(-1.295)	(-1.291)
year_rel_0	0.000	0.000	-0.003	-0.003
	(0.577)	(.)	(-0.807)	(-0.781)
year_rel_1	-0.000	-0.000	-0.002	-0.002
	(-0.469)	(.)	(-0.524)	(-0.474)
year_rel_2	0.000	0.000	-0.013***	-0.013***
	(0.549)	(.)	(-5.385)	(-4.891)
year_rel_3	0.000	0.000	-0.004	-0.004
	(1.091)	(.)	(-0.473)	(-0.489)
year_rel_4	0.000	0.000	-0.006	-0.005
	(0.376)	(.)	(-0.799)	(-0.767)
year_rel_5	0.000	0.000	-0.003	-0.003
	(0.419)	(.)	(-1.219)	(-0.955)
ln_assets		0.000		-0.001
		(.)		(-0.933)
ln_sales		0.000		-0.000
		(.)		(-0.212)
ln_employees		-0.000		-0.000
		(.)		(-0.671)
firm_age		-0.000		0.000*
		(.)		(2.058)
leverage_ratio		0.000		-0.000
		(.)		(-0.229)
return_on_assets		0.000		-0.000
		(.)		(-1.606)
Constant	0.005***	0.004	0.079***	0.086***
	(151.397)	(.)	(107.116)	(8.254)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
4-digit Industry FE	No	Yes	No	Yes
Observations	10034	10034	10034	10034
Adjusted R-squared	0.636	0.628	0.650	0.642

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 17: Effect of AI Adoption on Clinical Trial Duration

This table presents difference-in-differences estimates for the effect of AI discussion in earnings calls on clinical trial duration. The dependent variable is the duration of clinical trials in days. AI text indicates the first mention of AI-related terms in a firm's earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Post is an indicator equal to one for years on or after a firm's first AI mention. Controls include enrollment size of the trial. Standard errors are clustered at the firm level.

	Trial Duration (Days)		
	(1)	(2)	(3)
post_ai=1 × ai_firm=1	-157.653***	-121.296*	-101.672*
	(36.913)	(46.512)	(37.305)
post_ai=1	104.622***	102.161**	87.152**
	(33.050)	(28.635)	(20.574)
ai_firm=1	38.205**	-14.378	-21.271
	(14.941)	(50.040)	(48.576)
enrollment			0.118**
			(0.033)
Phase Fixed Effects	No	Yes	Yes
Expanded Access	No	Yes	No
Controls	No	No	Yes
Observations	11476.000	9091.000	9035.000
Adjusted R2	0.002	0.050	0.062

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 18: Effect of AI Adoption on Pull Request Length

This table presents difference-in-differences estimates for the effect of AI discussion in earnings calls on pull request length. The dependent variable is the length of pull requests in lines of code. AI text indicates the first mention of AI-related terms in a firm's earnings calls during the sample period. The sample covers 2010-2018, and firms are required to have data for all years in this period. Post is an indicator equal to one for years on or after a firm's first AI mention. Controls include log(assets), log(sales), log(employees), firm age, leverage ratio , and return on assets . Standard errors are clustered at the firm level.

	Pull Request Length		
	(1)	(2)	(3)
post_ai=1 × ai_firm=1			
post_ai=1	54.053***	82.513*	83.483*
	(19.601)	(46.525)	(46.783)
ai_firm=1			
n_authors			-0.175
			(0.519)
Author Fixed Effects	No	Yes	Yes
Repository Fixed Effects	No	No	No
Controls	No	No	Yes
Observations	7650.000	6651.000	6651.000
Adjusted R2	0.001	0.533	0.533

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A Distribution of Company Types and Countries

To understand which types of firms are discussing AI, I examine the distribution of AI exposure across company types and countries. Figure 26 presents this distribution through two complementary perspectives.

Panel A shows the distribution of company types in the sample. Public companies dominate the distribution, accounting for 82.3% of firms discussing AI in earnings calls. This concentration likely reflects both the nature of my earnings call data, which primarily covers public firms, and the greater technological and financial resources available to public companies for AI adoption. Investment firms, including both public (4.2%) and private (3.1%) investment companies, constitute the second largest category. This significant presence of investment firms suggests that the financial sector is particularly attuned to AI's potential impact, both as adopters of the technology and as allocators of capital to AI-intensive companies. The remaining share is split among various specialized entities such as mutual funds (2.8%) and other financial institutions (7.6%).

Panel B reveals the geographic distribution of the sample, focusing on the top 10 countries by company share. The United States dominates the sample, representing 73.4% of all firms discussing AI. This concentration reflects both the size of the U.S. market and its leadership in AI technology development. The United Kingdom and Canada follow as distant second and third, with 5.8% and 5.2% respectively. China, despite its significant AI investments, accounts for only 3.4% of the sample, likely due to the English-language nature of my earnings call data. The remaining top 10 countries are primarily developed economies, including Japan (2.7%), Germany (2.1%), and France (1.8%), suggesting that AI adoptions are more prevalent in advanced economies.

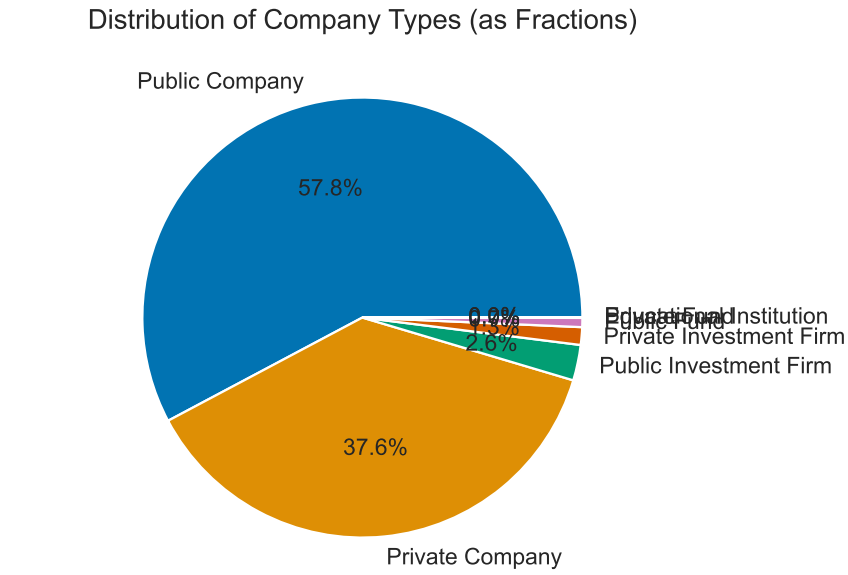


Figure 26: Distribution of Company Types and Countries

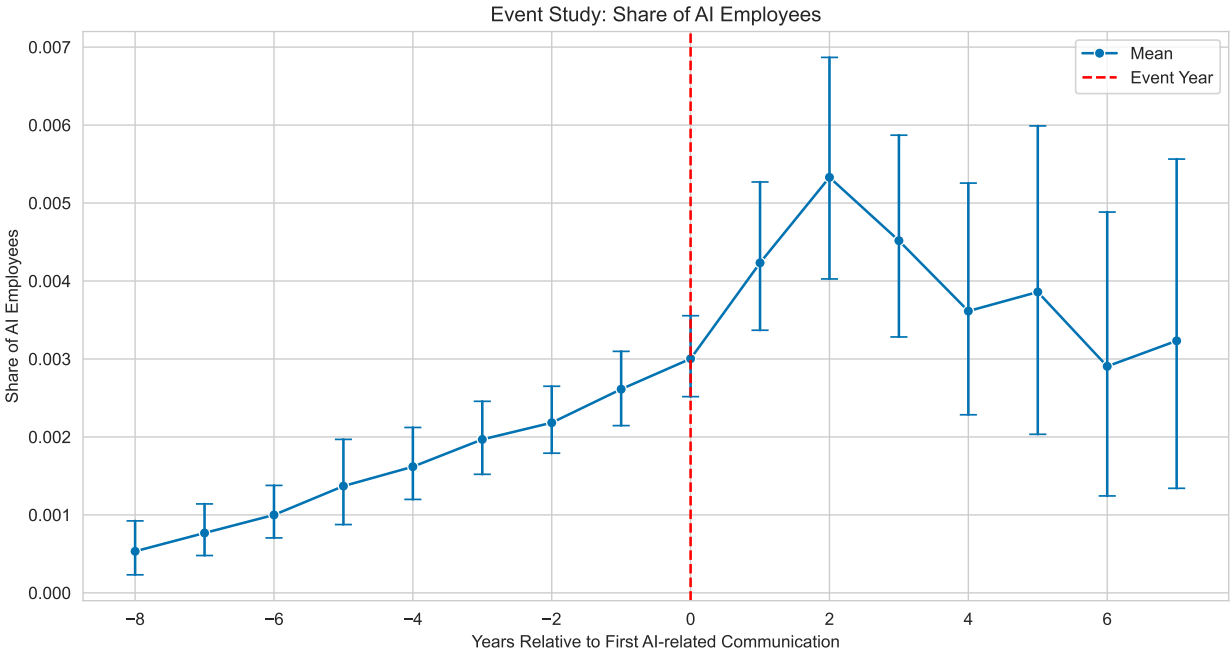


Figure 27: Event study of resume-based share of AI workers

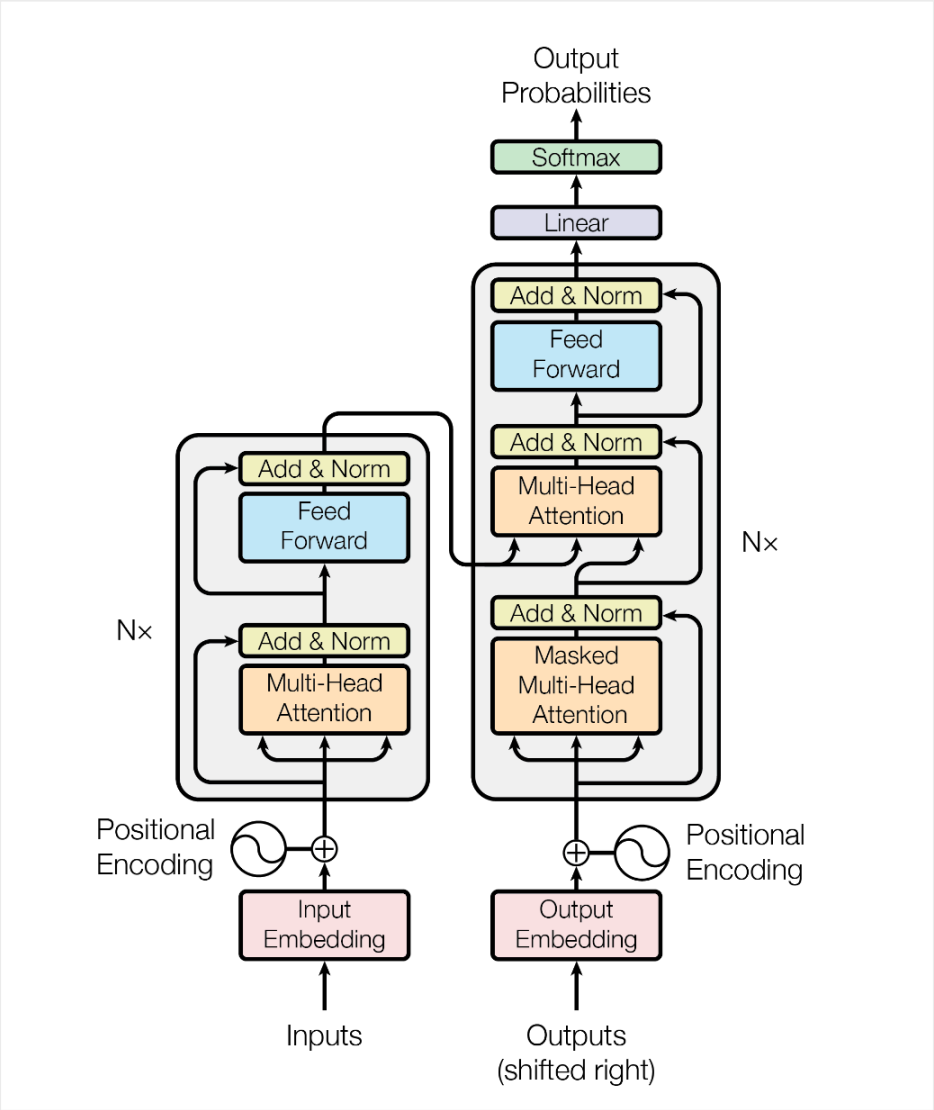


Figure 28: The Transformer - model architecture from Vaswani et al. (2017)

B Using Share of AI-related Employees Resume

A potential concern is that AI-related job postings do not necessarily imply that the firm has hired AI workers. To address this concern, I examine the share of AI knowledge workers and AI employees using the resume data of workers. This data directly measures the share of workers of a firm that has AI knowledge or AI skills thus alleviates the concern that AI related workers has unique labor market frictions. For example, if AI workers are harder to acquire, we would expect that AI-related job postings would be upward biased in comparison to the actual share of AI workers. These frictions are unlikely to be present in resume data.

Figure 28 shows the share of AI workers, defined as the share of workers with skills and backgrounds indicating AI expertise. The analysis reveals several key patterns. First, there is a gradual pre-trend in AI knowledge workers from years -8 to 0, with the share increasing from about 0.15% to 0.35% of total employment. This suggests firms slowly build AI capabilities even before public discussions. Indeed, we expect that some firms are only ex-post discussing AI in earnings calls because they have already hired AI workers. However, following the first AI mention, there is a sharp acceleration in AI knowledge worker hiring. The share jumps from 0.35% at the time of first mention to a peak of approximately 0.55% two years after, representing a more than 50% increase relative to the event year. This post-event increase is both economically and statistically significant, with confidence intervals indicating clear separation from pre-event levels. The effect appears persistent but shows some moderation in later years, with the share stabilizing around 0.45% by year 4. The pattern suggests that firms' public AI discussions coincide with, or slightly precede, substantial investments in AI human capital, though the pre-trend indicates the need for careful interpretation of the causal relationship between AI discussions and workforce composition.

C Appendix: Technical Details of AI Used for Textual Analysis

This section provides technical details of the AI used for textual analysis. I first provide a technical overview of the architecture of GPT-4. Then I describe the entire pipeline of data annotation. The code and data will be made publicly available at my website.

C.1 Backgrounds on the Architecture of GPT-4

In this appendix, I provide an overview of the AI used in this paper, GPT-4 OpenAI (2023), a state-of-the-art language model developed by OpenAI. The transformer architecture, introduced by Vaswani et al. (2017), has become a cornerstone in natural language processing due to its ability to capture long-range dependencies and parallelize training processes.

The transformer model relies on a sequence-to-sequence architecture that replaces recurrent and convolutional layers with self-attention mechanisms. This design allows the model to weigh the importance of different parts of the input data effectively.

At the heart of the transformer is the self-attention mechanism, which computes a representation of the input sequence by relating different positions within the sequence. For a given input sequence $X = [x_1, x_2, \ldots, x_n]$, the self-attention operation is defined as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V$$

This mechanism allows the model to focus on relevant parts of the input when generating each part of the output.

To capture information from different representation subspaces, the transformer employs multi-head attention:

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \ldots, \text{head}_h)W^O$$

where each head is computed as:

$$\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V)$$

and W^O is an output projection matrix. This allows the model to attend to information from different perspectives simultaneously.

Following the attention layers, the transformer applies a fully connected feedforward network to each position separately and identically:

$$\text{FFN}(x) = \text{ReLU}(xW_1 + b_1)W_2 + b_2$$

where W_1, W_2 are weight matrices and b_1, b_2 are biases.

Since the transformer architecture does not inherently capture sequential order, positional encodings are added to the input embeddings:

$$\text{PE}_{(pos,2i)} = \sin\left(\frac{pos}{10000^{2i/d_{\text{model}}}}\right)$$

$$\text{PE}_{(pos,2i+1)} = \cos\left(\frac{pos}{10000^{2i/d_{\text{model}}}}\right)$$

where pos is the position and i is the dimension index. These encodings inject information about the position of each token in the sequence.

To facilitate training and improve gradient flow, the transformer incorporates residual connections and layer normalization:

$$\text{LayerNorm}(x + \text{Sublayer}(x))$$

This ensures that each sublayer can learn modifications to the identity function, which aids in training deep architectures.

GPT-4 scales the transformer architecture significantly, increasing the number of layers, model dimensionality, and training data volume. While specific details about GPT-4's architecture are proprietary, it is generally believed to involve a larger number of transformer blocks to capture complex patterns and training on extensive and diverse datasets

to enhance generalization.

GPT-4 is trained using the autoregressive language modeling objective, maximizing the likelihood of the next token x_t given the previous tokens $x_{<t}$:

$$\mathcal{L} = -\sum_{t} \log P(x_t \mid x_{<t})$$

This objective enables the model to generate coherent and contextually relevant text. For more details of training, see OpenAI (2023).

C.2 Pipeline of Data Annotation

This section describes the pipeline of data annotation. At the core of the pipeline is the use of GPT-series models to annotate the text.

C.2.1 Pre-processing

The pre-processing step is to extract the text of interest from the raw data. For earnings calls, I used the entire transcript of the earnings call, without differentiating between whether the speaker is from the firm or the analyst. In an robustness check, I also used the transcript of the CEOs, and found similar results.

The first step of the pip eline is to identify mentions of machine learning and artificial intelligence. Specifically, I checked each report for the presence of the terms "machine learning" or "artificial intelligence". A firm was considered to have an AI mention if either term appeared at least once in the call. Once identified, I extracted the entire sentence containing these keywords and compiled a dataset of AI-related discussions from the filings. Using this dataset, I produced statistics on the fraction of unique firms with AI mentions. A firm was considered to mention AI in a given year if at least one filing within that year contained an AI mention. I performed the analysis both at the aggregate level for all U.S. firms over time and across different industries using the 2-digit NAICS industry code. Firms were identified by their CIK identifiers. The fraction of firms mentioning AI in a given year was calculated as the number of unique firms with AI mentions divided

by the total number of unique firms with filings in that year.

C.2.2 Foundation Model Annotation

To gain deeper insights and annotate these discussions effectively, I leveraged the advanced language processing capabilities of GPT-series models OpenAI (2023). I used both "gpt-3.5-turbo" and "gpt-4o-mini" to assist in analyzing and annotating the extracted text, allowing for an evaluation of the impact of AI as conveyed in the discussions. This method of using GPT to label text is also used in other studies, such as Eloundou et al. (2023); Eisfeldt et al. (2023a). I used gpt 3.5 turbo to annotate the text in this step.

In the annotation process, I asked GPT to assess the direction and magnitude of the AI impact for each discussion, as well as other aspects of potential impact. Specifically, I asked GPT to classify each text based on the following aspects: 1. Direction of impact: positive, negative, or other. 2. Significance of impact: low, moderate, high, or other. 3. Primary topic of impact: labor, investment, revenue, competition, M&A, or other. 4. Timeline of impact: happened, current, planning, or other. 5. Aggressiveness of impact: active, passive, or other.

Once I obtained these high-level annotations, I used gpt 4o mini to dive deeper into the economic aspects of the impact, including: 1. The main channel of labor impact for this use of AI. 2. Is there an automation effect by this use of AI? 3. In the automation channel of this labor impact, what occupations are most exposed to AI? 4. What is the most likely percentage of cost savings effect in this automation channel? 5. The main channel of productivity impact for this use of AI. 6. Is there an automation effect by this use of AI? 7. In the automation channel of this productivity impact, what occupations are most exposed to AI? 8. What is the most likely percentage of cost savings effect in this automation channel?

I required the output to be in JSON format so that it can be used for statistical analysis. I set the temperature to 0 to ensure the output is deterministic and thus replicable, I also set the maximum tokens to 4000 to ensure the output is not too long. No other special instructions or hyperparameters were given to the model.

D Appendix: AGI Progress

Figure 22 provides a snapshot of how AI systems are characterized in terms of capability, performance, and development source. Notably, the bottom-left chart highlights that an overwhelming majority of current AI deployments—98%—are classified as "narrow AI," meaning they excel at specific tasks but lack the broader, humanlike reasoning that underpins artificial general intelligence (AGI). Only a small fraction of the observed AI references (under 2% in total) fall under AGI or "general AI," underscoring how far most practical applications remain from fully autonomous, human-level cognition.

At the same time, the charts also reveal that AI development tends to be outsourced externally (roughly 77% of projects), suggesting that many firms rely on partnerships or third-party services rather than in-house development. This pattern may reflect both the complexity of advanced AI research and the limited internal resources within organizations. Taken together, the data show that progress in AGI remains in its early stages—both in adoption rates and real-world usage—while most firms continue to leverage narrower, specialized AI solutions developed by external providers. This relatively small footprint of AGI highlights the ongoing research challenges and high uncertainty surrounding the leap from task-specific systems to general intelligence.