

AI Masquerade: Hiding Employee Expansion in the AI Era *

March 16, 2026

Abstract

This paper documents a fundamental regime shift in how the equity market values corporate labor expansion following the release of ChatGPT. Prior to ChatGPT, job creation and stock market valuations moved closely together, but this relationship decouples sharply afterward. Using earnings call transcripts, I show that the stock market previously rewarded hiring signals but penalizes them in the post-ChatGPT period, as investors interpret labor expansion as evidence of lagging AI adoption. In response, firms that are actively hiring either curtail their hiring disclosures or, when they do disclose, simultaneously signal AI adoption as a substitute for workers—a pattern I term the *AI Masquerade*. Yet firms that still disclose hiring intentions are systematically underpriced in the short run, generating predictable long-run excess returns as the market corrects its initial over-penalization.

Keywords: Artificial Intelligence, Labor, Natural Language Processing, Earnings calls, Information Disclosure, Generative AI, ChatGPT, Mispricing

There is a growing debate about how advances in generative AI facilitate automation and reduce the demand for labor (e.g., [Eisfeldt et al. \(2026\)](#); [Sharma \(2026\)](#)). Consistent with this view, [Figure 1](#) shows that the S&P 500 index and aggregate job openings, which had previously moved closely together, began moving in opposite directions following the release of ChatGPT. This divergence raises a central question: how do investors interpret firms' labor expansion in an environment where AI may substitute for workers? Do they view hiring as a positive signal of future growth, or as evidence that firms are relying on labor rather than adopting AI technologies that can replace workers?

This paper studies how the stock market reacts to firms' hiring disclosures in the generative AI era. Using forward-looking hiring statements from earnings call transcripts, I examine how market responses to hiring signals changed before and after the release of ChatGPT. I then investigate the mechanism behind this shift by documenting how market reactions vary with firms' AI adoption disclosures, showing that investors interpret labor expansion as evidence of lagging AI adoption. Finally, I explore three follow-up questions: (1) whether firms strategically reduce hiring disclosures even when they are actively hiring; (2) whether firms that still disclose mask their hiring by signaling AI adoption to substitute for workers—a pattern I term the *AI Masquerade*; and (3) whether the market's penalty reflects rational reassessment or investor overreaction that generates predictable long-run excess returns.

I begin by constructing a measure of managerial hiring signals that captures executives' statements about *near-future* employment expansion, using natural language processing applied to earnings call transcripts of S&P 500 firms. Crucially, the measure captures planned rather than realized hiring, offering a forward-looking proxy for firms' labor demand. I validate this measure by showing that firms with stronger hiring signals exhibit higher employment growth and more job postings in subsequent quarters. Using these signals, I examine how the stock market reacts to firms' forward-looking hiring disclosures around earnings calls. I find that, prior to ChatGPT, hiring signals are positively associated with announcement returns, but this relation reverses in the

post-ChatGPT period. This implies that investors have come to view labor expansion as a negative signal. The market penalty is more pronounced in industries with higher AI substitutability, such as the consumer sector, but attenuates in high-tech sectors where complementarity between AI and skilled labor is stronger.

To understand why investors respond negatively to hiring signals in the generative AI era, I show that they interpret labor expansion as a signal that firms are lagging in AI adoption. I extract signals of AI adoption from earnings call transcripts using the same methodology used to construct the hiring signal. I find that the negative market reaction to hiring signals is significantly mitigated for firms that simultaneously communicate AI adoption. I further distinguish whether executives frame AI as substituting for labor or as complementing it. When AI adoption is described as substituting for labor—such as automating tasks or reducing future hiring needs—the penalty on hiring signals becomes significantly weaker. However, when AI adoption is framed as requiring additional hiring—such as expanding AI-related teams—the market penalty on hiring signals becomes stronger. These findings suggest that investors penalize hiring disclosures because they interpret labor expansion as evidence of lagging AI adoption.

Given that the market penalizes hiring signals, how do firms respond? I find two strategic responses. First, managers reduce hiring disclosures even when they are actively hiring, as if they were not. Second, firms that still choose to disclose mask their hiring intentions by simultaneously signaling that they are adopting AI to substitute for workers (*AI Masquerade*). This behavior resembles greenwashing, where firms adjust disclosures about ESG policies to align investor perceptions with a desired narrative (e.g., [Duchin, Gao, and Xu \(2024\)](#)).

To document the first pattern, I compare hiring signals across firms with different employment growth by classifying firms with employment growth less than or equal to zero and firms with employment growth greater than 10%. This comparison ensures that any reduction in hiring disclosures reflects strategic communication choices rather than an actual slowdown in hiring activity. I find that the gap between these two groups significantly narrows as firms that continue to expand

employment reduce their forward-looking hiring disclosures. In addition, I use job postings data from Lightcast (*formerly Burning Glass*), which more closely align with firms' hiring intentions, to document the same pattern. I find that hiring signals substantially decline in the post-ChatGPT period while job postings increase, suggesting that firms temper hiring-related communication even as actual labor demand stays elevated. This pattern is especially pronounced among firms with high AI-related hiring. I use the methodology in [Babina et al. \(2024\)](#)—which identifies AI-related jobs based on the skills listed in job postings—to define AI-related hiring. Despite a sharp increase in AI job postings in the post-ChatGPT period, these firms significantly reduce their hiring disclosures. Taken together, the evidence suggests that the reduction in hiring signals reflects strategic disclosure adjustment rather than a contraction in labor demand.

To establish the second pattern, I show that firms disclosing hiring intentions while actively hiring are more likely to use AI-substitution language during earnings calls. I find that both hiring signals and job postings are positively and significantly associated with AI-Labor Substitution disclosures, indicating that firms expanding their workforce simultaneously communicate to investors that they plan to use AI to reduce future hiring needs. Strikingly, firms with high AI-related job postings are significantly less likely to invoke such language. This suggests that the substitution narrative is not driven by firms genuinely investing in AI to replace workers, but by broadly hiring firms that strategically deploy AI-substitution rhetoric as a shield against the market's penalty.

I next ask whether the stock market's penalty on hiring disclosures is warranted. I compare the productivity of firms that disclose hiring intentions with those that do not. Firms that communicate hiring plans during earnings calls are more productive than their non-disclosing counterparts, as measured by both return on assets and total factor productivity. Moreover, these differences are larger in the post-ChatGPT period, suggesting that firms selectively disclosing hiring intentions are, if anything, growing faster and becoming more productive over time. While generative AI tools indeed alter the complementarity between labor and capital by facilitating substitution for workers, investors appear to overestimate the extent of this shift, pricing in a larger decline in the

value of labor than what is subsequently realized.

To examine whether this overreaction generates predictable excess returns, I construct a long-short portfolio that goes long firms disclosing positive hiring signals during earnings calls and short all remaining firms. Positions are held for one year with quarterly rebalancing. The strategy generates significant abnormal returns of approximately 6–7 basis points per day in the post-ChatGPT period, with cumulative alpha reaching approximately 10% over the post-ChatGPT sample (6 quarters). No such returns are observed in the pre-ChatGPT period, confirming that the return predictability is specific to the generative AI era. Together, these results imply that investors overestimate the extent to which AI displaces labor, penalize hiring disclosures too severely, and gradually correct this mispricing as firm fundamentals are realized over the following year.

To isolate the source of this mispricing, I compare the return predictability of hiring disclosures with that of firms' realized hiring actions, measured by job postings. I construct an analogous long-short portfolio based on job postings to benchmark the asset pricing implications of firms' hiring actions against their hiring disclosures. While the disclosure-based strategy earns strongly positive abnormal returns in the post-ChatGPT period, the action-based strategy generates no such returns over the same period—if anything, drifting slightly negative. This divergence establishes that the mispricing is not driven by a risk premium associated with labor demand per se, but by a systematic overreaction to the *disclosure* of hiring intentions. The market correctly prices realized hiring activity, yet penalizes firms that communicate plans to expand their workforce, interpreting such disclosures as a signal of vulnerability to AI-driven labor substitution. As underlying firm performance is realized, the market corrects this overreaction, generating the predictable excess returns documented in the talk-based strategy.

In the final set of analyses, I conduct placebo tests to validate the main results. First, I repeat the analysis using physical investment signals extracted from earnings call transcripts. Unlike hiring signals, investment signals do not generate any meaningful patterns around earnings announcements in either the pre- or post-ChatGPT period. This is consistent with investors specifically

viewing labor expansion—rather than capital investment—as a signal of lagging AI adoption, since generative AI primarily substitutes for labor rather than physical capital. Moreover, firms do not reduce investment-related disclosures following the introduction of ChatGPT; if anything, they increase such disclosures even as actual investment declines. This asymmetry further distinguishes the strategic disclosure behavior documented for hiring from a general tendency to suppress forward-looking statements, and confirms that the AI Masquerade is specific to labor-related communication.

Second, I examine whether the negative market reaction to hiring signals is driven by broader macroeconomic conditions rather than the release of generative AI tools. One concern is that the post-ChatGPT period coincides with a broader slowdown in investment and hiring following the COVID-era surge, potentially confounding the main results with recessionary dynamics. To address this, I interact hiring signals with either an NBER recession indicator or the Federal Funds rate in place of the Post-ChatGPT dummy. Neither replicates my main findings, indicating that the stock market does not systematically penalize hiring signals during recessionary periods or periods of tighter monetary policy. This further confirms that the documented penalty is specific to the generative AI era.

The main contribution of this paper is to provide new evidence on how generative AI reshapes the information content of firms' labor market disclosures. I document that the release of ChatGPT fundamentally changed how investors interpret hiring signals. This shift has two important consequences. First, it induces firms to strategically conceal their hiring intentions even when they are actively expanding employment, distorting the information environment in ways that may impede efficient capital allocation. Second, the market's overreaction to hiring disclosures generates predictable long-run abnormal returns. Together, these findings suggest that the rise of generative AI has introduced a new form of disclosure friction, where the fear of being perceived as AI-averse leads firms to suppress value-relevant information about their growth prospects.

This paper contributes to several strands of research. First, it adds to the growing literature on

the interaction between artificial intelligence and finance. Recent work examines how advances in generative AI and machine learning affect firms, financial markets, and economic outcomes, with [Eisfeldt and Schubert \(2025\)](#) reviewing the implications of generative AI for finance. A growing body of work examines the economic and labor-market implications of generative AI ([Brynjolfsson et al., 2023](#); [Chen et al., 2024](#); [Eloundou et al., 2024](#); [Li, 2024](#); [Schubert, 2025](#)), which are closely related to the focus of this paper. For instance, [Azpurua \(2026\)](#) suggests that generative AI reshapes labor demand by reducing hiring for automation-prone tasks while increasing demand for roles where AI augments human workers. [Babina et al. \(2024\)](#) and [Babina et al. \(2025\)](#) study how AI adoption relates to firm growth, innovation, and operating volatility. Other studies explore how generative AI influences financial decision-making and markets, including its effects on corporate valuation ([Eisfeldt et al., 2026](#)), investment behavior and perception biases ([Fedyk et al., 2025](#)), corporate policies ([Jha et al., 2024](#)), asset management ([Sheng et al., 2025](#)), and remote work ([Schubert, 2025](#)). A related set of papers documents how large language models affect financial markets and forecasting, including stock return predictability ([Lopez-Lira and Tang, 2025](#)), macroeconomic links ([Chen et al., 2024](#)), and financial advice ([Lo and Ross, 2024](#)). More broadly, researchers study the role of generative AI and large language models in shaping beliefs, economic agents, and research practices ([Bybee, 2025](#); [Horton et al., 2026](#); [Korinek, 2023](#); [Novy-Marx and Velikov, 2026](#)).

Second, this paper extends the literature that uses textual data and machine learning methods to study financial and economic outcomes. Early work develops tools to extract information from financial text and corporate disclosures ([Loughran and McDonald, 2020](#)), and text-based measures have been used to study industry structure and product differentiation ([Hoberg and Phillips, 2016](#)). More recent research applies machine learning and natural language processing methods in economics and finance ([Ludwig and Mullainathan, 2024](#); [Nagel, 2021](#); [Korinek, 2023](#); [Dell, 2024](#)). Advances in natural language processing, including transformer-based models such as BERT ([Devlin et al., 2019](#)), have enabled researchers to classify and analyze large collections of text ([Caragea et al., 2020](#)). A growing set of papers applies textual analysis and modern language models to

study financial communication, corporate disclosures, and related economic outcomes (Beckmann et al., 2025; Li et al., 2024). This paper contributes to this literature by applying modern language models to corporate earnings call transcripts to analyze managerial hiring disclosures and their implications for stock market reactions.

This paper is also related to the literature on intangible capital, as generative AI is part of firms' intangible capital, reflecting investments in data and organizational capabilities that shape production and decision-making. A large body of work studies the role of intangible assets in firm valuation, investment, and financial structure (Crouzet et al., 2022). For example, Veldkamp and Chung (2024); Goldstein et al. (2021) study the role of data in the economy and finance, highlighting the growing importance of data as an economic asset. Crouzet and Eberly (2023) and Eisfeldt and Papanikolaou (2013) develop a framework linking intangible capital to rents and firm value, while Sun and Xiaolan (2019); Falato et al. (2022); Eisfeldt et al. (2026) study how firms finance intangible capital and examine its implications for corporate finance. This paper contributes to this literature by highlighting how the emergence of generative AI as a new form of intangible capital shapes investors' interpretation of firms' labor expansion and corporate communication.

Lastly, this paper adds to the literature on cross-sectional return predictability based on corporate hiring policies. A large body of work shows that firms that hire more tend to earn lower subsequent stock returns, consistent with lower discount rates (e.g., Belo, Xue, and Zhang (2013), Hou, Xue, and Zhang (2015), Belo et al. (2017)). This paper complements this evidence by showing that the *disclosure* of hiring intentions generates return predictability in the generative AI era. The mispricing is concentrated in what firms say rather than what they do, suggesting that the information environment surrounding corporate labor decisions has become a distinct source of cross-sectional return predictability.

1 Data and Key Variables

1.1 Earnings Calls Transcripts

The primary data for the textual analysis consist of corporate earnings call transcripts for S&P 500 constituent firms from 2002 to 2025, obtained from the Capital IQ (CIQ) Transcripts database via WRDS. The transcript data include the formal presentation section and the subsequent question-and-answer session. To identify management’s disclosures related to future hiring decisions, I implement a natural language processing (NLP) pipeline to extract labor-related text segments from earnings calls. Using a refined dictionary of labor-related keywords—such as headcount, workforce, hiring, and recruitment—the algorithm identifies relevant sentences and captures the surrounding context by extracting a three-sentence window consisting of the preceding sentence, the focal sentence, and the subsequent sentence. I randomly select 1,000 sentence groups and manually label them as either indicating future hiring or otherwise. The key distinction is between planned hiring and employment actions already taken. I cross-validate the labels using ChatGPT and Google Gemini.

Using these manually labeled observations, I train a FinBERT model to classify the sentiment of the extracted labor-related snippets. FinBERT is a pre-trained language model based on the BERT architecture and optimized for financial text, which allows the classification to capture nuanced managerial disclosures regarding future hiring intentions beyond simple keyword-based approaches. For each earnings call, the classified labor-related snippets are aggregated to construct firm-level measures. The primary independent variable, *Hiring Signal*, is defined as the total number of snippets classified as *future hiring* during the call, normalized by the total word count of the earnings call transcript to account for variation in disclosure length.

Figure 2 provides validation for the hiring signals by examining whether firms that disclose stronger hiring intentions subsequently expand employment. In each industry $\times$ year, firms are

sorted into two groups based on the intensity of hiring signals disclosed during earnings calls. Panel A reports realized employment growth. Firms with stronger hiring signals exhibit significantly higher employment growth both contemporaneously and in the subsequent year, indicating that managerial disclosures about future hiring intentions correspond closely to realized employment expansion. I estimate the differences between the two groups using a time-series t -test and find that the differences are 0.051 ($t = 4.48$) for the contemporaneous measure and 0.046 ($t = 5.16$) for the subsequent-year measure.

To more directly capture firms' hiring intentions, I compare firms' job postings from Lightcast in Panel B. Firms that disclose stronger hiring signals post substantially more job openings both in the current year and the following year. The differences are 1526.77 ($t = 4.82$) for the contemporaneous measure and 1509.55 ($t = 4.73$) for the subsequent-year measure. Taken together, these results indicate that hiring signals extracted from earnings call transcripts capture meaningful variation in firms' subsequent hiring intentions and realized hiring activity.

Using the same methodology, I also construct signals related to physical capital investment and AI adoption for subsequent analyses. Investment signals identify statements in which executives discuss current or planned increases in physical capital expenditures—such as expanding production capacity, building facilities, or investing in data centers—and serve as a placebo test. AI adoption signals capture snippets in which executives discuss adopting or expanding the use of artificial intelligence technologies and are used to explore the mechanism through which generative AI affects firms' hiring communication. In both cases, the measures are constructed by identifying relevant snippets using targeted keyword dictionaries and normalizing the number of classified snippets by the total word count of the earnings call transcript, following the same procedure used to construct the hiring signal. All signals are aggregated at the earnings call level, with each measure computed for a given transcript (TranscriptID).

1.2 Other Data

I supplement the earnings call data with several additional datasets. First, I obtain firm-level job postings from Lightcast (*formerly Burning Glass*), which aggregates online vacancies from company websites and major job boards. These data provide detailed information on job titles, required skills, and employer identities, allowing the construction of firm-level hiring measures over time. To identify AI-related hiring, I follow the methodology of [Babina et al. \(2024\)](#), which identifies AI-related positions based on the skills listed in job descriptions. Specifically, the AI-relatedness of each skill is measured by its co-occurrence with core AI skills such as machine learning, natural language processing, and computer vision. Job postings are then assigned an AI-relatedness score based on the average AI-relatedness of the required skills, and positions with scores above 0.1 are classified as AI-related. I aggregate these classifications to the firm-quarter level to measure firms' demand for AI-related labor.

I use Compustat to obtain firm-level accounting information, including realized employment growth. In all regressions, I include a standard set of firm-level controls: firm size measured by book assets, Tobin's Q, and the R&D-to-sales ratio. I use CRSP daily stock return data to compute cumulative abnormal returns (CAR) around earnings call dates to examine the stock market's reaction to managerial disclosures.

2 Main Results

2.1 Stock Market Penalizing Hiring Signals

As shown in [Figure 1](#), stock market valuations and aggregate job openings move closely together prior to the release of ChatGPT. Following November 2022, however, the comovement breaks down sharply, and the two series begin to decouple. To assess this pattern more formally, I estimate time-series regressions of the S&P 500 index on the aggregate number of nonfarm job openings. [Table 1](#) reports the results separately for the pre- and post-ChatGPT periods. In the pre-

ChatGPT sample (December 2000–October 2022), job openings are strongly positively associated with stock market valuations, indicating that tighter labor market conditions coincide with higher equity valuations. In contrast, in the post-ChatGPT period (November 2022–December 2025), the relation reverses, suggesting that higher market valuations are instead associated with weaker labor market conditions.

To further assess this regime shift, I estimate firm-level regressions of announcement returns on hiring signals. Specifically, I examine how the stock market reacts to firms’ hiring signals extracted from earnings call transcripts before and after the introduction of ChatGPT. Hiring signals measure the extent to which executives express intentions to increase *future* hiring during earnings calls. Using these hiring signals, I estimate the following specification:

$$CAR_{it} = \alpha + \beta_1 \text{Hiring Signal}_{it} + \beta_2 (\text{Hiring Signal}_{it} \times \text{Post-ChatGPT}_t) + \gamma X_{it} + \varepsilon_{it}. \quad (1)$$

where CAR_{it} is the cumulative abnormal return over the $(-1, +1)$ window around firm i ’s earnings announcement, computed using the Fama–French three-factor model. Post-ChatGPT_t is an indicator equal to one for earnings calls held after November 2022. I include firm-level controls, X_{it} , consisting of book assets, Tobin’s Q , and R&D intensity (R&D-to-sales), as well as industry-by-time and firm fixed effects. Standard errors are clustered at the firm level.

[Table 2](#) reports the results. In the pre-ChatGPT period, the stock market rewards firms’ hiring signals: the coefficient on $\text{Hiring Signal}_{it}$ is positive and statistically significant, indicating that firms communicating stronger hiring intentions earn higher abnormal returns around earnings announcements. This suggests that, prior to November 2022, the market interpreted forward-looking hiring as favorable news about future growth.

In contrast, the interaction term between $\text{Hiring Signal}_{it}$ and Post-ChatGPT_t is negative and statistically significant. The magnitude of the interaction term offsets the baseline effect: the F -statistic for the test of $\beta_1 + \beta_2 = 0$ in column (2) equals 3.27, rejecting the null and confirming that

the sum of β_1 and β_2 is negative. These results indicate that the stock market substantially penalizes hiring signals after the introduction of ChatGPT, consistent with the aggregate time-series evidence in [Figure 1](#).

Next, I analyze whether the post-ChatGPT penalty on hiring signals is heterogeneous across industries. To examine this heterogeneity, I estimate an industry-interacted specification that allows the sensitivity of announcement returns to hiring signals in the post-ChatGPT period to vary across the Fama–French five industries. Specifically, I estimate,

$$CAR_{it} = \sum_{j=1}^4 \beta_j \cdot (1_j \times \text{Hiring Signal}_{it}) + \sum_{j=1}^4 \gamma_j \cdot (1_j \times \text{Hiring Signal}_{it} \times \text{Post-ChatGPT}_t) \\ + \beta \text{Hiring Signal}_{it} + \gamma \cdot (\text{Hiring Signal}_{it} \times \text{Post-ChatGPT}_t) + \eta X_{it} + \alpha_i + \varepsilon_{it}, \quad (2)$$

where 1_j denotes an indicator for industry j under the Fama–French five-industry classification. The “Other” industry ($j = 5$) is omitted, so the coefficients γ_j measure the incremental market reaction to hiring signals in the post-ChatGPT period for each industry relative to this reference sector. The specification includes firm-level controls as well as industry $\times$ time and firm fixed effects.

[Figure 3](#) plots the coefficient estimates of γ_j along with their 90% confidence intervals. The figure reveals substantial cross-industry heterogeneity in how the stock market penalizes hiring signals. The consumer sector exhibits the largest negative coefficient, indicating that hiring signals in this industry are penalized most strongly. The high-technology sector shows a comparatively attenuated penalty, consistent with investors being less inclined to interpret hiring in this sector as detrimental to firm value—a pattern driven by stronger perceived complementarity between AI and skilled labor. The cross-industry pattern in γ_j thus provides direct evidence on where investors believe AI substitution for labor is most likely to occur.

2.2 Mechanism: AI Adoption

Why do investors begin to penalize hiring signals after the introduction of ChatGPT? Investors may interpret labor expansion as a signal that the firm relies on labor-intensive growth rather than adopting artificial intelligence technologies that can substitute for workers. In an environment where generative AI tools can automate tasks previously performed by employees, hiring may be viewed as reliance on traditional labor-intensive production rather than AI-enabled growth, leading investors to react negatively to hiring disclosures because they believe these firms are less capable of leveraging AI to improve productivity. However, if firms simultaneously communicate active AI adoption, hiring may instead be interpreted as complementary to AI technologies rather than as evidence of technological lag—in which case, the negative market reaction to hiring signals should be weaker for firms that also disclose AI adoption.

To examine this mechanism, I construct an *AI Adoption* measure using the same transcript-based methodology used to construct the hiring signal. Specifically, I identify transcript snippets in which executives discuss adopting or expanding the use of artificial intelligence technologies in the post-ChatGPT period.¹ I estimate the following,

$$\begin{aligned} CAR_{it} = & \beta_1 \text{Hiring Signal}_{it} + \beta_2 (\text{Hiring Signal}_{it} \times \text{Post-ChatGPT}_t) \\ & + \beta_3 (\text{Hiring Signal}_{it} \times \text{Post-ChatGPT}_t \times \text{AI Adoption}_{it}) + \beta_4 \text{AI Adoption}_{it} \quad (3) \\ & + \gamma X_{it} + \alpha_i + \varepsilon_{it}. \end{aligned}$$

If investors penalize hiring signals because they interpret them as evidence that firms are less capable of adopting AI, the negative market reaction to hiring signals in the post-ChatGPT period should be attenuated when firms also communicate AI adoption. This prediction implies that the coefficient on the triple interaction term should be positive.

Table 3 reports the results. Consistent with the baseline findings, Column (1) shows that the

¹Because explicit AI-adoption language is extremely rare prior to the introduction of ChatGPT and rises sharply afterward (Figure A.1), I focus on the post-ChatGPT sample.

interaction between *Hiring Signal* and *Post-ChatGPT* is negative and statistically significant, confirming that the stock market penalizes hiring signals after the introduction of ChatGPT. The coefficient on the triple interaction is significantly positive, indicating that the post-ChatGPT hiring penalty is substantially weaker for firms that simultaneously communicate AI adoption during earnings calls. These findings support the proposed mechanism: investors interpret hiring more favorably when it is accompanied by signals that the firm is actively adopting AI technologies.

I further distinguish whether executives frame AI as substituting for labor or as complementing it. I classify AI-related disclosures based on whether executives describe AI as reducing future hiring needs or requiring additional hiring. *AI-Labor Substitution* captures statements indicating that AI enables headcount reductions, labor efficiencies, or automation of tasks that reduce future hiring needs. *AI-Labor Complementarity*, in contrast, captures statements indicating that AI adoption requires expanding the workforce, such as hiring additional AI engineers or data scientists.

Column (2) of [Table 3](#) shows that when AI adoption is framed as substituting for labor, the negative market reaction to hiring signals is significantly attenuated. The positive and statistically significant coefficient on the triple interaction term indicates that investors penalize hiring signals less when firms simultaneously communicate that AI adoption will reduce future hiring needs. Column (3) shows that when AI adoption is complementary—framed as requiring additional hiring—the market penalty on hiring signals becomes stronger. The negative and statistically significant triple interaction coefficient indicates that investors react more negatively when firms signal both labor expansion and AI adoption that requires additional hiring. Together, these results confirm that the “AI Shield” is driven specifically by the narrative of labor substitution.

2.3 AI Masquerade: Hiding Employment Growth

The previous subsection shows that the stock market penalizes hiring signals in the post-ChatGPT period. Given this shift in market reaction, I next examine how firms adjust their communication accordingly. If managers anticipate that forward-looking hiring statements are interpreted

unfavorably by investors, they may strategically temper such disclosures even when underlying employment growth remains strong. This possibility forms the basis of the “AI Masquerade” hypothesis.

I first explore how firms undergoing active employment expansion adjust their hiring-related communication after the release of ChatGPT. Within each quarter, firms are sorted based on realized employment growth. Firms with employment growth less than or equal to zero are classified as *No Employment Growth*, while firms with employment growth greater than 10% are classified as *High Employment Growth*. For each quarter, I compute the average hiring signal separately for the two groups and then aggregate these quarterly averages within the pre- and post-ChatGPT periods.

Figure 4 shows that prior to November 2022, firms experiencing high employment growth communicate substantially stronger hiring signals than firms with no employment growth. In the post-ChatGPT period, however, the gap narrows markedly. Although firms with high realized growth continue to expand employment, their forward-looking hiring statements decline relative to the pre-period. The difference-in-differences estimate is significantly negative ($-0.048, t = -2.44$), indicating that firms undergoing active employment expansion substantially reduce their hiring-related disclosures after the introduction of ChatGPT. These results suggest that firms temper hiring communication in response to the market’s adverse reaction to such signals.

I next use job postings data from Lightcast, which more closely align with firms’ hiring intentions, to examine whether this pattern stems from weaker hiring activity or a change in disclosure strategy. Within each quarter, firms are sorted into two groups based on the number of job postings over the subsequent four quarters, and I compute the average hiring signals of these two groups each quarter, aggregating across the pre- and post-ChatGPT periods. Figure 5 plots the average hiring signals (Panel A) and the corresponding number of job postings (Panel B) for these groups in both periods. Panel A shows that prior to ChatGPT, firms with high job postings communicate stronger hiring signals than firms with low job postings. In the post-ChatGPT period, hiring signals

decline for both groups, including firms with high job postings.

Panel B shows that job postings do not decline in the post-ChatGPT period—if anything, the number of postings increases among both groups. Taken together, the reduction in hiring signals is not driven by weaker labor demand. Instead, firms temper hiring-related communication even as their job postings remain elevated, consistent with strategic disclosure adjustment rather than a contraction in employment activity.

The introduction of ChatGPT represents a salient technological shock centered on artificial intelligence, making AI hiring a natural margin along which firms may adjust both real decisions and disclosure. If investors perceive AI as a substitute for labor, the market response to hiring signals may be particularly sensitive when hiring is explicitly AI-related. I therefore examine job postings that are directly related to AI, following the methodology in [Babina et al. \(2024\)](#).

[Figure 6](#) examines this mechanism. Within each quarter, firms are sorted based on subsequent AI-related hiring. Panel A shows that, prior to ChatGPT, firms in the High AI-Hiring group communicate substantially stronger hiring signals than firms with no AI hiring. After ChatGPT, this gap narrows markedly as hiring signals among high AI-hiring firms decline. Panel B shows that AI hiring intensity increases sharply in the post-ChatGPT period. Thus, despite expanding AI-related employment, these firms reduce their communication about near-term hiring. Relative to the patterns based on overall job postings in [Figure 5](#), the decline in hiring signals is even more pronounced among firms that actively expand AI employment, suggesting that firms investing in AI-related labor are particularly likely to conceal their hiring intentions following the shift in market perception.

Finally, I test whether firms that disclose hiring intentions while actively hiring strategically accompany their disclosures with AI-substitution language to offset the market penalty. I use three AI-related signals extracted from earnings calls: (1) *AI Adoption*, which captures the intensity of discussions about adopting or expanding AI technologies; (2) *AI-Labor Substitution*, which cap-

tures statements describing AI as reducing future hiring needs; and (3) *AI-Labor Complementarity*, which captures statements describing AI adoption as requiring additional hiring. Using these measures, I estimate the following in the post-ChatGPT period,

$$\text{AI Signal}_{it} = \beta_1 \text{Hiring Signal}_{it} + \beta_2 \text{Job Posting}_{it} + \gamma X_{it} + \alpha_i + \varepsilon_{it}. \quad (4)$$

where *AI Signal* corresponds to one of the three AI-related measures described above, *Hiring Signal* captures forward-looking hiring intentions expressed during earnings calls, and *Job Posting* measures the number of job postings by the firm in a given quarter.

Table 4 shows that firms that are actively hiring are significantly more likely to frame AI as a substitute for labor. Column (2) shows that both *Hiring Signal* and *Job Posting* are positively and significantly associated with AI-Labor Substitution disclosures—that is, firms disclosing hiring intentions while actively expanding their workforce simultaneously communicate to investors that they plan to use AI to reduce future hiring needs. This is precisely the AI Masquerade: firms growing their headcount strategically couch their disclosures in the language of AI-driven labor substitution to offset the market penalty on hiring. Columns (1) and (3) show that neither AI Adoption broadly nor AI-Labor Complementarity exhibit any significant association with hiring activity, confirming that the strategic framing is specific to substitution language—the narrative that investors most want to hear in the post-ChatGPT period.

Strikingly, AI-Job Posting is negatively and significantly associated with AI-Labor Substitution disclosures, indicating that firms actively expanding their AI-related workforce are in fact less likely to frame AI as a substitute for labor. This implies that the substitution narrative is driven not by genuine AI-for-labor replacement, but by broadly hiring firms deploying it as a strategic shield against the market’s penalty.

3 Asset Pricing Implications

3.1 Stock Market Correction in the Long Run

The previous section shows that the stock market penalizes the disclosure of hiring intentions during earnings calls after the introduction of ChatGPT, which in turn leads firms to reduce such disclosures. A regime shift in the stock market’s reaction to hiring signals would be warranted if firms that disclose hiring signals in the post-ChatGPT period are no longer the growing firms they were previously, or if they are less productive. To examine whether this is the case, I first compare the hiring intensity of firms that disclose positive hiring signals with those that do not. For each quarter, I compute the average job postings for the two groups and then aggregate these quarterly averages separately for the pre- and post-ChatGPT periods.

[Figure 7](#) presents the results. Panel A compares overall job postings, while Panel B focuses on AI-related hiring, following the definition in [Babina et al. \(2024\)](#). Panel A shows that, relative to the pre-ChatGPT period, firms that disclose hiring signals actively increase their hiring intensity. This pattern has two implications. First, firms that communicate hiring intentions are indeed the firms expanding employment. Second, in the post-ChatGPT environment, firms that disclose hiring intentions appear more selective and, if anything, are growing even faster. Panel B further shows that AI hiring increases in the post-ChatGPT period for both groups and remains markedly higher among firms that disclose hiring signals, suggesting that firms actively expand AI employment regardless of whether they disclose such intentions during earnings calls.

Next, I compare the productivity of firms that disclose hiring intentions with those that do not, using two measures: (1) an accounting-based measure, return on assets (ROA), and (2) a structurally estimated measure of total factor productivity following [Olley and Pakes \(1996\)](#). For each quarter, I compute the average productivity of the two groups and then aggregate these quarterly averages across the pre- and post-ChatGPT periods.

Figure 8 shows that firms with hiring signals are more productive in both periods. The productivity gap between the two groups widens in the post-ChatGPT period. Panel A shows that the difference-in-differences estimate for ROA is 0.023 ($t = 3.07$), indicating that firms disclosing hiring signals experience a larger improvement in profitability relative to non-disclosing firms. This pattern also holds for the Olley and Pakes (1996)’s productivity measure in Panel B, although the difference-in-differences estimate is not statistically significant. Collectively, Figure 7 and Figure 8 indicate that firms disclosing hiring intentions are growing faster and are more productive—differences that, if anything, strengthen in the post-ChatGPT period.

This evidence implies that the stock market’s penalty on hiring signals is consistent with investor overreaction rather than rational reassessment. While generative AI tools do alter the complementarity between labor and capital by facilitating substitution for workers, investors appear to overestimate the extent of this shift, pricing in a larger decline in the value of labor than what is subsequently realized. Such overreaction generates predictable excess returns as the market gradually corrects this mispricing. To examine this, I construct a long-short portfolio based on firms’ hiring disclosures. At the end of each quarter, firms that disclose positive hiring signals are placed in the long portfolio, while all remaining firms are placed in the short portfolio. Positions are held for one year with quarterly rebalancing.

Figure 9 plots the cumulative performance of this portfolio beginning in 2023Q4, incorporating hiring signals observed after the introduction of ChatGPT in 2022Q4. The solid navy line shows cumulative raw returns, and the dashed red line shows cumulative abnormal returns (alpha) from the Fama-French three-factor model. Following portfolio formation in 2023Q4, the strategy generates steadily increasing returns over time, with both the raw cumulative return and the three-factor alpha exhibiting a clear upward trend. The strategy generates no excess returns in the pre-ChatGPT period, confirming that the return predictability is specific to the generative AI era and that the initial market penalty on hiring disclosures is not fully justified by fundamentals.

Table 5 reports the portfolio performance in a regression setting. The strategy earns positive

and statistically significant returns of approximately 6–7 basis points per day in the post-ChatGPT period. In column (1), the average return is 6.8 basis points per day ($t = 3.21$), and columns (2) and (3) show that these returns remain statistically significant after controlling for the Fama–French three factors. Notably, firms that disclose hiring signals are predominantly growth firms in the pre-ChatGPT period but resemble value firms in the post-ChatGPT period. Together, these results indicate that investors initially penalize hiring signals but subsequently revise their beliefs as underlying firm performance is realized.

3.2 Talk vs Action

The previous subsection establishes that the stock market initially penalizes firms that disclose hiring intentions during earnings calls, even though such firms are genuinely expanding employment and becoming more productive. This overreaction generates predictable long-run abnormal returns as investors gradually correct their misperception over the following year. A natural question is whether this mispricing is unique to what executives *say*, or whether it also extends to what firms *do*. If investors systematically misprice labor demand in the generative AI era, one would expect a portfolio based on firms’ realized hiring actions to generate similar abnormal returns. This subsection examines this hypothesis by comparing the long-run return predictability of hiring disclosures with that of hiring actions.

To construct the action-based portfolio, I sort firms into four portfolios at the end of each quarter based on the number of job postings, a direct measure of firms’ realized hiring activity. A long-short portfolio is then constructed by going long firms with high job postings and short firms with low job postings. As in the talk-based strategy, positions are held for one year with quarterly rebalancing.

[Figure 10](#) presents the cumulative abnormal returns for both strategies. The solid navy line reproduces the talk-based portfolio from [Figure 9](#), which generates steadily increasing abnormal returns following portfolio formation in 2023Q4, reaching approximately 8–10% by the end of the

sample. The dashed red line shows the corresponding performance of the action-based portfolio. In stark contrast, the action-based strategy generates no abnormal returns over the same period—if anything, drifting slightly negative. This divergence is striking: firms with high job postings are not systematically mispriced, yet firms that *disclose* hiring intentions earn significant positive abnormal returns over the following year.

The contrast between *Talk* and *Action* isolates the source of mispricing. Realized hiring activity is publicly observable and the market prices it correctly. It is the *disclosure* of hiring intentions during earnings calls that triggers investor overreaction in the post-ChatGPT period—investors penalize firms for communicating plans to expand their workforce, interpreting such disclosures as a signal of vulnerability to AI-driven labor substitution. As firm fundamentals are realized over the subsequent year, the market corrects this overreaction, generating the abnormal returns documented in [Figure 9](#). Together, these results indicate that the return predictability of hiring signals is driven not by a risk premium associated with labor demand, but by a systematic overreaction to the disclosure of hiring intentions in an era of rapid AI adoption.

4 Robustness: Placebo Tests

4.1 Investment Signal

As the first placebo test, I repeat the analysis using managerial investment signals extracted from earnings call transcripts. The investment signal is defined as the number of transcript snippets in which executives state that the firm is actively investing or planning to increase investment in physical capital, including references to capital expenditures, expansion of production capacity, or infrastructure investments such as building data centers. I normalize the number of such snippets by the total transcript word count so that the measure is comparable across firms and earnings calls.

Using this investment signal, I first examine whether the stock market also penalizes investment

signals in the post-ChatGPT period as follows,

$$CAR_{it} = \alpha + \beta_1 \text{Investment Signal}_{it} + \beta_2 (\text{Investment Signal}_{it} \times \text{Post-ChatGPT}_t) + \gamma X_{it} + \varepsilon_{it}. \quad (5)$$

Table 6, Panel A shows that this is not the case. Investment signals in both the pre- and post-ChatGPT periods are unassociated with cumulative abnormal returns around earnings announcements. While investment signals also capture future growth opportunities, investment typically requires substantial capital expenditures, and a large literature documents that firms undertaking high levels of investment tend to earn lower subsequent returns, as investors interpret such behavior as management taking advantage of a low cost of capital to finance new projects (e.g., Fama and French (2006); Fama and French (2015); Liu, Whited, and Zhang (2009)). Thus, although investment signals may convey growth prospects, they are not expected to generate excess returns.

I also examine whether firms reduce investment signals in the post-ChatGPT period. Within each quarter, firms are sorted into two groups based on investment, measured by annual growth in physical capital (net property, plant, and equipment). Figure 11, Panel A shows that investment signals increase among high-investment firms while declining among low-investment firms. The difference-in-differences estimate is 1.200 ($t = 7.41$), indicating that firms actively investing in physical capital communicate stronger investment intentions in the post-ChatGPT period.

This widening gap in investment signals, however, is not driven by a corresponding increase in actual investment. Figure 11, Panel B shows that the investment gap in fact narrows, contrary to the pattern in Panel A. High-investment firms substantially reduce investment relative to the pre-period, with a difference-in-differences estimate of -0.155 ($t = -4.98$). These results indicate that firms do not reduce investment-related disclosures following the introduction of ChatGPT—if anything, they increase such disclosures even as actual investment declines.

4.2 Macroeconomic Conditions

The post-ChatGPT period coincides with a broader slowdown in investment and hiring following the COVID-era surge, as shown in [Figure 1](#). One may therefore argue that the main results are driven by broader macroeconomic conditions rather than the release of generative AI tools.

To examine this possibility, I test whether the post-ChatGPT effect simply captures recessionary conditions by interacting hiring signals with either an NBER recession indicator or the Federal Funds rate in place of the Post-ChatGPT dummy,

$$\begin{aligned} CAR_{it} = & \alpha + \beta_1 \text{Hiring Signal}_{it} \\ & + \beta_2 (\text{Hiring Signal}_{it} \times \text{NBER Recession or Federal Funds Rate}_t) + \gamma X_{it} + \varepsilon_{it}. \end{aligned} \tag{6}$$

[Table 6](#), Panel B shows that neither interaction produces a negative market reaction. The coefficients on both interaction terms are small and statistically insignificant, indicating that the stock market does not penalize hiring signals during recessionary periods or periods of tighter monetary policy. The negative reaction to hiring signals documented in the main results is therefore specific to the post-ChatGPT period and is not attributable to macroeconomic conditions.

5 Conclusion

This paper documents a regime shift in the stock market’s reaction to firms’ hiring disclosures following the introduction of ChatGPT. Prior to November 2022, hiring signals communicated during earnings calls were interpreted as favorable information about future growth and were associated with positive abnormal returns. After the introduction of generative AI tools, however, the market reaction reverses and hiring signals are penalized. This shift suggests that investors increasingly view labor expansion as potentially detrimental to firm value in an environment where AI may substitute for workers.

Despite this negative market reaction, the evidence indicates that firms disclosing hiring intentions are in fact the firms experiencing stronger growth. Firms with hiring signals exhibit higher hiring intensity, expand AI-related employment more aggressively, and are more productive than firms without such signals. These patterns become even stronger in the post-ChatGPT period. At the same time, firms appear to strategically adjust their communication by reducing hiring-related disclosures even when their underlying labor demand remains strong. This behavior is consistent with the idea that managers anticipate unfavorable investor reactions and therefore temper their forward-looking hiring statements.

Finally, I show that the stock market's penalty on hiring signals generates predictable excess returns. A long-short portfolio that goes long firms with positive hiring signals and short firms without such signals earns economically large and statistically significant returns after the introduction of ChatGPT. The gradual outperformance of hiring-signal firms suggests that the market initially overreacts to hiring disclosures and subsequently corrects this mispricing as firm fundamentals become realized. Taken together, the findings highlight how technological shocks can alter investor beliefs about labor and lead firms to strategically adjust their communication with capital markets. Future research could estimate how generative AI changes the optimal mix between labor and capital by estimating the decline in complementarity between workers and software tools or physical capital. Developing structural or empirical measures of this shift would help quantify the extent to which the market's reaction to hiring signals reflects a rational reassessment of production technology versus an overreaction to uncertainty surrounding the economic impact of generative AI.

References

- Azpurua, A. E. (2026). Research: How AI is changing the labor market. *Harvard Business Review*.
- Babina, T., A. Fedyk, A. He, and J. Hodson (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151, 103745.
- Babina, T., A. Fedyk, A. He, and J. Hodson (2025). Artificial intelligence makes firm operating performance less volatile. *AEA Papers and Proceedings* 115, 35–39.
- Beckmann, L., H. Beckmeyer, I. Filippou, S. Menze, and G. Zhou (2025). Unusual financial communication: ChatGPT, earnings calls, and financial markets. *Working Paper*.
- Belo, F., J. Li, X. Lin, and X. Zhao (2017). Labor-force heterogeneity and asset prices: The importance of skilled labor. *Review of Financial Studies* 30(10), 3669–3709.
- Belo, F., C. Xue, and L. Zhang (2013). A supply approach to valuation. *Review of Financial Studies* 26(12), 3029–3067.
- Brynjolfsson, E., D. Li, and L. R. Raymond (2023). Generative ai at work. *NBER Working Paper*.
- Bybee, J. L. (2025). The ghost in the machine: Generating beliefs with large language models. *Working Paper*.
- Caragea, D., M. Chen, T. Cojoianu, M. Dobri, K. Glandt, and G. Mihaila (2020). Identifying FinTech innovations using BERT. *2020 IEEE International Conference on Big Data*, 1117–1126.
- Chen, J., G. Tang, G. Zhou, and W. Zhu (2024). ChatGPT, stock market predictability and links to the macroeconomy. *Working Paper*.
- Chen, W. X., S. Srinivasan, and S. Zakerinia (2024). Displacement or complementarity?: The labor market impact of generative AI. *Working Paper*.

- Crouzet, N. and J. Eberly (2023). Rents and intangible capital: A Q+ framework. *Journal of Finance* 78(4), 1873–1916.
- Crouzet, N., J. C. Eberly, A. L. Eisfeldt, and D. Papanikolaou (2022). The economics of intangible capital. *Journal of Economic Perspectives* 36(3), 29–52.
- Dell, M. (2024). Deep learning for economists. *Working Paper*.
- Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Working Paper*.
- Duchin, R., J. Gao, and Q. Xu (2024). Sustainability or greenwashing: Evidence from the asset market for industrial pollution. *Journal of Finance* 80(2), 699–754.
- Eisfeldt, A. and D. Papanikolaou (2013). Organization capital and the cross-section of expected returns. *Journal of Finance* 68(4), 1365–1406.
- Eisfeldt, A. L., A. Falato, D. Lee, and M. Z. Xiaolan (2026). Equity pay beyond the C-suite. *Working Paper*.
- Eisfeldt, A. L. and G. Schubert (2025). Generative ai and finance. *Annual Review of Financial Economics* 17, 363–393.
- Eisfeldt, A. L., G. Schubert, B. Taska, and M. B. Zhang (2026). Generative AI and firm values. *Journal of Finance Forthcoming*.
- Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Working Paper*.
- Falato, A., D. Kadyrzhanova, J. Sim, and R. Steri (2022). Rising intangible capital, shrinking debt capacity, and the U.S. corporate savings glut. *Journal of Finance* 77(5), 2799–2852.
- Fama, E. F. and K. R. French (2006). Profitability, investment, and average returns. *Journal of Financial Economics* 82(3), 491–518.

- Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of Financial Economics* 116(1), 1–22.
- Fedyk, A., A. Kakhbod, P. Li, and U. Malmendier (2025). AI and perception biases in investments: An experimental study. *Working Paper*.
- Goldstein, I., C. S. Spatt, and M. Ye (2021). Big data in finance. *Review of Financial Studies* 34(7), 3213–3225.
- Hoberg, G. and G. Phillips (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy* 124(5), 1423–1465.
- Horton, J. J., A. Filippas, and B. S. Manning (2026). Large language models as simulated economic agents: What can we learn from homo silicus? *Working Paper*.
- Hou, K., C. Xue, and L. Zhang (2015). Digesting anomalies: An investment approach. *Review of Financial Studies* 28(3), 650–705.
- Jha, M., J. Qian, M. Weber, and B. Yang (2024). ChatGPT and corporate policies. *Working Paper*.
- Korinek, A. (2023). Generative AI for economic research: Use cases and implications for economists. *Journal of Economic Literature* 61(4), 1281–1317.
- Li, J. (2024). Is generative AI an existential threat to human creatives?: Insights from financial economics. *Working Paper*.
- Li, K., F. Mai, R. Shen, C. Yang, and T. Zhang (2024). Dissecting corporate culture using generative AI: Insights from analyst reports. *Working Paper*.
- Liu, L. X., T. M. Whited, and L. Zhang (2009). Investment-based expected stock returns. *Journal of Political Economy* 117(6), 1105–1139.
- Lo, A. W. and J. Ross (2024). Can ChatGPT plan your retirement?: Generative AI and financial advice. *Working Paper*.

- Lopez-Lira, A. and Y. Tang (2025). Can ChatGPT forecast stock price movements?: Return predictability and large language models. *Working Paper*.
- Loughran, T. and B. McDonald (2020). Textual analysis in finance. *Annual Review of Financial Economics* 12, 357–375.
- Ludwig, J. and S. Mullainathan (2024). Machine learning as a tool for hypothesis generation. *Quarterly Journal of Economics* 139(2), 751–827.
- Nagel, S. (2021). Machine learning in asset pricing. *Princeton University Press*.
- Novy-Marx, R. and M. Velikov (2026). Artificial intelligence-powered (finance) scholarship. *Journal of Economic Literature* 64(1), 5–37.
- Olley, G. S. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64(6), 1263–1297.
- Schubert, G. (2025). Organizational technology ladders: Remote work and generative AI adoption. *Working Paper*.
- Sharma, R. (2026). Why people still matter in the AI era. *Financial Times*.
- Sheng, J., Z. Sun, B. Yang, and A. Zhang (2025). Generative AI and asset management. *Working Paper*.
- Sun, Q. and M. Z. Xiaolan (2019). Financing intangible capital. *Journal of Financial Economics* 133(3), 564–588.
- Veldkamp, L. and C. Chung (2024). Data and the aggregate economy. *Journal of Economic Literature* 62(2), 458–484.




Figure 1: **Job Openings and the S&P 500: A Disconnect After the Release of ChatGPT.**

This figure illustrates how the stock market index and job openings were closely aligned before November 2022, but became disconnected following the release of ChatGPT. I plot the number of nonfarm job openings (in thousands) and the S&P 500 index.




Figure 2: Hiring Signals and Hiring Activity.

This figure examines whether hiring signals disclosed during earnings calls correspond to actual hiring activity. Hiring Signal measures the degree to which executives state intentions to increase future hiring during the earnings call. Specifically, it is defined as the number of near-future hiring-related snippets normalized by total transcript words. In each industry $\times$ year, firms are sorted into two groups based on the intensity of hiring signals. For each group, I compute the average outcome within each year and then average these annual values across the sample period so that each year receives equal weight. Panel A reports employment growth, and Panel B reports the number of job postings, using both contemporaneous (t) and subsequent-year ($t + 1$) values. In Panel A, the differences between the two groups are 0.051 ($t = 4.48$) for the contemporaneous measure and 0.046 ($t = 5.16$) for the subsequent-year measure. In Panel B, the differences are 1526.77 ($t = 4.82$) for the contemporaneous measure and 1509.55 ($t = 4.73$) for the subsequent-year measure. The sample is based on S&P 500 firms.




Figure 3: The Industry-Specific Market Penalty on Hiring Signals Post-ChatGPT.

This figure plots the coefficients γ_j from the following specification,

$$CAR_{it} = \sum_{j=1}^5 \gamma_j \cdot (1_j \times \text{Hiring Signal}_{it} \times \text{Post-ChatGPT}_t) + \sum_{j=1}^5 \beta_j \cdot (1_j \times \text{Hiring Signal}_{it}) + \eta \cdot X_{it} + \alpha + \varepsilon_{it},$$

where CAR is the cumulative abnormal return over the $(-1,+1)$ window around the earnings calls using the Fama-French three-factor model and is expressed in basis points. 1_j denotes an indicator for each of the Fama-French 5 industries. Hiring Signal measures the degree to which executives state intentions to increase future hiring during the earnings call, standardized across firms. The specification includes firm fixed effects and industry $\times$ time fixed effects. The coefficient estimates are relative to the omitted reference group, "Other." Standard errors are clustered at the firm level, and 90% confidence intervals are shown.




Figure 4: Hiding Hiring Intentions: Hiring Signals and Realized Employment Growth.

This figure illustrates a relationship between managerial labor market signaling and employment growth rates following the release of ChatGPT. The sample is divided into pre- and post-ChatGPT periods (November 2022). Within each quarter, firms are sorted based on employment growth rates. Firms with employment growth less than or equal to 0% are classified as *No Employment Growth*, while firms with employment growth greater than 10% are classified as *High Employment Growth*. The figure reports the average hiring signals of these two groups. Hiring signal is defined as the number of near-future hiring-related snippets normalized by total transcript words and is expressed in basis points. The difference-in-differences estimate is -0.045 ($t = -2.32$). The sample is based on S&P 500 firms.




Figure 5: Hiding Hiring Intentions: Hiring Signals and Job Posting.

This figure shows that firms reduce hiring signals despite increasing job postings following the release of ChatGPT. The sample is divided into pre- and post-ChatGPT periods (November 2022). Within each quarter, firms are sorted into two portfolios based on the number of job postings over the subsequent four quarters. The figure reports the average hiring signals (Panel A) and the number of job posts (Panel B) of these two groups. Hiring signal is defined as the number of near-future hiring-related snippets normalized by total transcript words and is expressed in basis points. The sample is based on S&P 500 firms.



Figure 6: Hiding Hiring Intentions: Hiring Signals and AI Job Posts.

This figure shows that firms reduce hiring signals despite increasing AI-job postings following the release of ChatGPT. The sample is divided into pre- and post-ChatGPT periods (November 2022). Within each quarter, firms are sorted into two groups based on AI-related hiring over the subsequent four quarters, following the methodology in [Babina et al. \(2024\)](#). *No AI-Hiring* group consists of firms with zero AI hiring during this period, and *High AI-Hiring* group consists of firms in the top 25% of AI hiring. The figure reports the average hiring signals (Panel A) for both groups and the number of AI-related job posts (Panel B) for firms in the High AI-Hiring group. Hiring signal is defined as the number of near-future hiring-related snippets normalized by total transcript words and is expressed in basis points. The sample is based on S&P 500 firms.




Figure 7: Hiring Intensity by Signal Disclosures.

This figure compares the number of job postings for firms that disclose positive hiring signals during earnings calls and those that do not. For each quarter, I compute the average number of job postings over the subsequent four quarters for the two groups and then aggregate these averages across quarters separately for the pre- and post-ChatGPT periods. Panel A shows overall job postings. Panel B focuses on AI-related hiring, defined using AI job postings following the methodology of [Babina et al. \(2024\)](#). The sample is based on S&P 500 firms.




Figure 8: Hiring Signals and Productivity.

This figure compares subsequent firm productivity for firms with lower and higher hiring signals from earnings calls, separately for the pre- and post-ChatGPT periods (November 2022). The sample is divided into pre- and post-ChatGPT periods (November 2022). Within each quarter, firms are sorted into two portfolios based on whether they exhibit positive hiring signals. The figure reports the average productivity in subsequent years for firms in each portfolio. Panel A uses accounting-based productivity measured by return on assets in percentage points while Panel B uses total factor productivity following [Olley and Pakes \(1996\)](#). Difference-in-differences estimates are 0.943 ($t = 3.07$) in Panel A and 0.185 ($t = 1.27$) in Panel B. The sample is based on S&P 500 firms.



Figure 9: Long-Short Portfolio Performance Based on Hiring Signals After ChatGPT.

The figure plots the cumulative performance of a long-short portfolio constructed from firm-level hiring signals. At the end of each quarter, firms with positive hiring signals are placed in the long portfolio, while all remaining firms are placed in the short portfolio. Positions are held for one year, resulting in overlapping portfolios with quarterly rebalancing. The solid navy line shows cumulative raw returns, and the dashed red line shows cumulative abnormal returns (alpha) from the Fama-French three-factor model. Portfolio formation begins at the start of 2023Q4, so that the strategy incorporates hiring signals observed after the introduction of ChatGPT in 2022Q4. The vertical dashed line indicates the portfolio formation date. The sample is based on S&P 500 firms.



Figure 10: **Talk vs Action.**

The figure compares the asset pricing implications of what executives say during earnings calls and what their firms actually do. At the end of each quarter, firms are sorted into four portfolios based on the number of job postings. A long-short portfolio is then constructed by going long firms with high job postings and short firms with low job postings, to capture the long-run return predictability of firms' hiring actions. Positions are held for one year, resulting in overlapping portfolios with quarterly rebalancing. The solid navy line shows the cumulative abnormal returns (alpha) from the Fama-French three-factor model for hiring signals based on executives' statements during earnings calls (*Talk*), as in [Figure 9](#). The dashed line shows the corresponding alpha based on firms' hiring actions measured by job postings (*Action*). The sample is based on S&P 500 firms.




Figure 11: Placebo Test: Investment Signal.

This figure presents a placebo test examining managerial investment disclosures around the release of ChatGPT. The sample is divided into pre- and post-ChatGPT periods (November 2022). Within each quarter, firms are sorted into two groups based on investment, measured by annual growth in physical capital (net property, plant, and equipment). Panel A reports the average investment signal for these two groups (low- and high-investment firms). Investment signal is defined as the number of snippets in which executives state that the firm is actively investing or planning to invest in physical capital, normalized by total transcript words and expressed in basis points. Panel B reports the average investment for these two groups. The difference-in-differences estimate is 1.200 ($t = 7.41$) in Panel A and -0.155 ($t = -4.98$) in Panel B. The sample is based on S&P 500 firms.

Table 1: Job Openings and Stock Valuations

This table examines a contemporaneous relationship between labor market conditions and stock market outcomes. I report time-series regressions of the S&P 500 index on the aggregate number of nonfarm job openings. t statistics are based on Newey-West standard errors with three-month lags.

	S&P 500 Index	
	(1)	(2)
Job Openings	0.430*** (25.778)	-0.598*** (-9.064)
Observations	263	25
R^2	0.899	0.803
Sample Period	Dec 2000 – Oct 2022	Nov 2022 – Dec 2025
t statistics in parentheses		
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$		

Table 2: Market Reaction to Hiring Signals: Pre- vs Post-ChatGPT

This table reports regressions of cumulative abnormal returns (CARs) around earnings calls on firms' hiring signals extracted from earnings call transcripts, before and after the introduction of ChatGPT. CAR is measured over the (-1,+1) window using the Fama-French three-factor model and is expressed in basis points. Hiring Signal measures the degree to which executives state intentions to increase future hiring during the earnings call, standardized across firms. Post-ChatGPT is a dummy equal to one for observations from December 2022 onward. The sample consists of S&P 500 firms. t statistics are based on standard errors clustered at the firm level.

	CAR	
	(1)	(2)
Hiring Signal	22.100*** (3.181)	18.766*** (2.810)
Hiring Signal $\times$ Post-ChatGPT	-57.081** (-2.038)	-66.563** (-2.427)
Observations	49270	49264
R^2	0.266	0.292
Industry $\times$ Date FE	Yes	Yes
Firm FE	No	Yes
Sample Period	All	All

t statistics in parentheses

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

Table 3: **Mechanism: AI Adoption Mitigating the Hiring Signal Penalty**

This table examines whether the stock market's penalty on hiring signals after the introduction of ChatGPT is attenuated when firms also communicate AI adoption during earnings calls. CAR is measured over the (-1,+1) window using the Fama-French three-factor model and is expressed in basis points. Hiring Signal measures the degree to which executives state intentions to increase future hiring during the earnings call, standardized across firms. AI Adoption captures the intensity with which executives discuss adopting or expanding the use of artificial intelligence technologies during the earnings call in the post-ChatGPT period. AI-Labor Substitution captures disclosures in which AI is described as lowering future hiring requirements, whereas AI-Labor Complementarity captures disclosures in which AI adoption is described as requiring additional hiring in the post-ChatGPT period. Post-ChatGPT is a dummy equal to one for observations from December 2022 onward. The sample consists of S&P 500 firms. t statistics are based on standard errors clustered at the firm level.

	CAR		
	(1)	(2)	(3)
Hiring Signal	18.910*** (2.831)	18.843*** (2.818)	18.834*** (2.820)
Hiring Signal $\times$ Post-ChatGPT	-80.777*** (-2.805)	-79.578*** (-2.916)	-64.809** (-2.282)
Hiring Signal $\times$ AI Adoption $\times$ Post-ChatGPT	4.054*** (2.801)		
AI Adoption	6.236* (1.674)		
Hiring Signal $\times$ AI-Labor Substitution $\times$ Post-ChatGPT		6.022*** (3.510)	
AI-Labor Substitution		1.064 (0.172)	
Hiring Signal $\times$ AI-Labor Complementarity $\times$ Post-ChatGPT			-2.883** (-2.003)
AI-Labor Complementarity			12.930* (1.964)
Observations	49264	49264	49264
R^2	0.292	0.292	0.292
Industry $\times$ Date FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes

t statistics in parentheses

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

Table 4: AI Masquerade: Hiding Employee Expansion with AI

This table examines whether firms that disclose hiring intentions in the post-ChatGPT period also mention AI adoption more frequently to avoid the market penalty. AI Adoption captures the intensity with which executives discuss adopting or expanding the use of artificial intelligence technologies during the earnings call in the post-ChatGPT period. AI-Labor Substitution captures disclosures in which AI is described as lowering future hiring requirements, whereas AI-Labor Complementarity captures disclosures in which AI adoption is described as requiring additional hiring in the post-ChatGPT period. Hiring Signal measures the degree to which executives state intentions to increase future hiring during the earnings call, standardized across firms. (AI-)Job Posting is the number of (AI-related) job postings by the firm in a given quarter, measured in thousands. The sample consists of S&P 500 firms. t statistics are based on standard errors clustered at the firm level.

	AI Adoption	AI-Labor Substitution	AI-Labor Complementarity
	(1)	(2)	(3)
Hiring Signal	-0.132 (-0.515)	0.319** (2.288)	0.112 (0.332)
Job Posting	52.806 (1.079)	272.079*** (2.642)	14.474 (0.551)
AI-Job Posting	1.734 (0.813)	-2.369** (-2.224)	-0.171 (-0.160)
Observations	2791	2791	2791
R^2	0.604	0.299	0.336
Industry $\times$ Date FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes

t statistics in parentheses

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

Table 5: **Long-Short Portfolio Performance Based on Hiring Signals**

This table reports time-series regressions of daily returns on a long-short portfolio formed based on firm-level hiring signals. At the end of each quarter, firms with positive hiring signals are placed in the long portfolio, while all remaining firms are placed in the short portfolio. Portfolios are held for one year, resulting in overlapping positions with quarterly rebalancing. The sample period is from January 2022 to December 2024. t -statistics are based on Newey-West standard errors with 1 quarter (63 days) lag.

	Hiring Singals Long-Short (%)		
	(1)	(2)	(3)
Post 2023 4Q	0.068*** (3.214)	0.060*** (3.013)	0.063*** (3.057)
$R^m - R^f$		3.594* (1.932)	4.413*** (3.634)
SMB		-0.714 (-0.340)	-1.812 (-0.660)
HML		-16.584*** (-6.457)	-21.148*** (-10.005)
$(R^m - R^f) \times \text{Post 2023 4Q}$			-5.845 (-0.941)
$\text{SMB} \times \text{Post 2023 4Q}$			-2.613 (-0.655)
$\text{HML} \times \text{Post 2023 4Q}$			19.304*** (3.778)
Intercept	-0.047*** (-3.057)	-0.042*** (-3.909)	-0.041*** (-3.703)
Observations	753	753	753
R^2	0.008	0.230	0.285

t statistics in parentheses

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

Table 6: **Placebo**

This table reports two placebo tests on market reactions to earnings-call disclosures. Panel A uses capital investment signals extracted from earnings call transcripts as a placebo measure and examines their relation with cumulative abnormal returns (CARs) around earnings calls. Panel B examines whether the market reaction to hiring signals differs during NBER recession periods or periods of higher Federal Funds rates by interacting hiring signals with a recession indicator or the Federal Funds rate. The dependent variable is CAR measured over the (-1,+1) window around the earnings call using the Fama-French three-factor model and expressed in basis points. The sample consists of S&P 500 firms. t statistics are based on standard errors clustered at the firm level.

Panel A. Investment Signal

	CAR	
	(1)	(2)
Investment Signal	0.000 (0.816)	0.001 (1.141)
Investment Signal $\times$ Post-ChatGPT	0.001 (0.485)	0.001 (0.653)
Observations	41002	40998
R^2	0.236	0.257
Industry $\times$ Year FE	No	Yes
Firm FE	Yes	Yes

t statistics in parentheses

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

Panel B. NBER Recession & Fed Funds Rate

	CAR			
	(1)	(2)	(3)	(4)
Hiring Signal	19.297*** (2.882)	16.062** (2.457)	20.059** (2.429)	16.469** (2.127)
Hiring Signal $\times$ NBER Recession	8.698 (0.295)	2.250 (0.080)		
Hiring Signal $\times$ Fed Funds Rate			-0.187 (-0.047)	-0.273 (-0.068)
Observations	49270	49264	49270	49264
R^2	0.266	0.292	0.266	0.292
Industry $\times$ Year FE	Yes	Yes	Yes	Yes
Firm FE	No	Yes	No	Yes

t statistics in parentheses

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

Appendix: Additional Results




Figure A.1: **AI Adoption Signal.**

This figure presents the aggregate frequency of AI adoption signals within my sample period. AI Adoption captures the intensity with which executives discuss adopting or expanding the use of artificial intelligence technologies during the earnings call. The vertical red dashed line marks the release of ChatGPT in 2022Q4, after which the incidence of AI-related signaling in earnings call transcripts exhibits a sharp, significant increase. The sample is based on S&P 500 firms.