Artificial Intelligence and Labor Market Signaling *

Jingyi Cui, Gabriel Dias, and Justin Ye †

March 20, 2025

Abstract

We study the effect of generative artificial intelligence on labor market signaling, leveraging the introduction of an AI-powered writing tool on Freelancer.com, an online labor market. The tool generates cover letters customized to job descriptions. Diff-in-diff analyses reveal that the usage of AI assistance increases cover letter quality by 144% – as measured by the textual similarity between cover letters and job descriptions – and raises the probability of receiving a callback by 63%. While AI is only used in a minority of cover letters, we find evidence of a substantial equilibrium response. Following the tool's rollout, the signaling value of cover letters declines sharply, with the correlation between cover letter quality and callback rates and offer rates falling by 40% and 53%, respectively. These results suggest that AI weakens the effectiveness of cover letters as a labor market signal.

*We are grateful to Elliot Ash for inspiration for the project and to Joe Altonji, Steve Berry, Charles Hodgson, John Eric Humphries, Phil Haile, and Katja Seim for helpful comments. We thank Freelancer.com for providing data and insights. This work was supported by funding from the Cowles Foundation for Research in Economics.

†Yale Department of Economics; Contact information: jingyi.cui@yale.edu, gabriel.santamarina@yale.edu, and justin.ye@yale.edu.

1 Introduction

Artificial intelligence is reshaping labor market matching. College applicants and job seekers increasingly use tools like ChatGPT to edit resumes, craft cover letters, and prepare for interviews.[1] The impact of artificial intelligence on matching frictions is complex (Wiles et al. (2023), Cowgill et al. (2024), Gans (2024)). On the one hand, AI could level the playing field by reducing the influence of language skills on matching, allowing workers' true credentials to shine through. Moreover, AI can streamline the job application process, saving applicants time and enabling them to apply for more positions. In these ways, AI could reduce the matching frictions in labor markets. On the other hand, the generative nature of AI may dilute the informational value of application materials. If job seekers use AI to present themselves as stronger candidates than they truly are, employers may struggle to identify high-quality applicants. As a result, AI may exacerbate matching frictions.

Empirical research on the effect of AI on labor market matching is often constrained by an observability challenge. Although anecdotal evidence suggests widespread use of generative AI tools in job and college applications, such usage is often difficult for researchers to detect, let alone link to application outcomes. We address this challenge by studying the roll-out of a platform-native AI cover letter writing tool on Freelancer.com, a major online marketplace that connects freelancers to medium-term jobs in areas including web development, design, and translation.

On April 19, 2023, Freelancer introduced the "AI Bid Writer," a tool that helps workers generate cover letters tailored to employers' job descriptions. This feature was freely available to a large subset of workers based on their membership plans. We analyze the effects of the AI writing tool with data from a random sample of PHP projects posted between March 11 and July 1, 2023. Our data includes project descriptions, applicants' cover letters, experience levels, and other relevant characteristics, as well as employers' interview and hiring decisions, and the reviews and payments received by selected applicants. Crucially, we have data on workers' access to the tool — determined by their membership status — and whether they used it for each application they submitted.

The AI writing tool had a direct impact on a significant fraction of the market. In our data sample, 64% of applications were submitted by workers eligible for the tool. After its introduction, 14% of all bids were written with its assistance. The nature of the roll-out provides a unique opportunity to study both the partial and general equilibrium effects of AI. The latter is particularly important as labor markets adjust to new equilibria, where an increasing number of workers use generative AI for job applications and employers update their beliefs about worker quality in the absence of clear signals indicating AI use.

Our first key finding is that AI improves cover letter quality and raises the per-application prob-

[1]This phenomenon has been widely covered by the media. Recent headlines include "AI is supposed to make applying to jobs easier — but it might be creating another problem" (NBC News, November 17, 2024) and "AI Is Taking Over College Admissions" (The Nation, October 11, 2024).

ability of receiving a callback. To measure cover letter quality, we develop a metric that captures the relevance of cover letters to job descriptions. We validate the measure by showing that, prior to the AI tool's roll-out, the measure strongly predicts hiring outcomes after controlling for other application characteristics. Using a two-way fixed effect specification, we compare changes in outcomes around the introduction of the AI tool between treated workers (those on eligible membership plans throughout the study period) and control workers (those who were never on eligible membership plans). We find that using the AI tool led to a 144% increase in cover letter quality and a 63% increase in the probability of receiving a callback. For comparison, Bertrand and Mullainathan (2004) find that changing the job applicant's name from an African American-sounding name to a white-sounding name increases the callback rate by 50 percent. This suggests that generic cover letters are a barrier to receiving callbacks. AI improves job search outcomes for those with access to the tool.

The impact of AI is heterogeneous. We examine on how its effects vary based on workers' pre-AI ability to write tailored cover letters. This heterogeneity matters because, as we discuss using a model in section 5, the impact of AI depends on whether it complements or substitutes workers' existing job market signals. If higher-quality workers can better articulate their credentials with AI, the technology reduces information frictions. However, if AI helps lower-quality workers more, it could increase information frictions by making better and worse workers appear more alike. Using data on workers' pre-AI cover letter quality, we provide empirical evidence that the AI tool disproportionately benefited workers who previously wrote weaker cover letters, thereby compressing the distribution of cover letter quality among treated workers. Workers at the low end of pre-AI writing ability increased their cover letters by twice as much as workers at the high end of pre-AI writing ability. This finding suggests a general equilibrium effect of generative AI: by diminishing the signaling value of cover letters, it makes it harder for employers to assess worker quality. We study the equilibrium effects formally in the next section.

While AI is only used in a minority of cover letters, we find evidence of a substantial equilibrium response. We analyze how the relationship between cover letter quality and hiring outcomes changed following the AI tool's roll-out. Our findings reveal a persistent decline in the power of cover letter quality in predicting hiring decisions. The correlation between cover letter quality and receiving a callback dropped by 40%, while the correlation between cover letter quality and receiving an offer dropped by 53%. These effects are particularly striking given that only 14% the applications in our sample were written with the tool. Despite its limited direct adoption, the tool's influence permeated the broader market, affecting even those who did not use or have access to it. Employers, faced with greater uncertainty about whether a cover letter genuinely reflects an applicant's ability, adjusted their hiring criteria, relying less on cover letter quality as a screening mechanism.

To elucidate our empirical findings and evaluate alternative AI designs, we develop a simple theoretical model with two worker quality types. The model's predictions align with our results

on both partial and general equilibrium effects and yield two additional insights. First, when AI substitutes existing quality signals, employers are more prone to hiring mistakes. Employers' tendency to make mistakes depends on the intensity of AI adoption in a non-monotonic way. Second, an alternative AI design that complements rather than substitutes quality signals can reduce hiring mistakes. Such a design may involve prompting workers to provide more relevant information on their credentials and present it in a structured, digestible way for employers. While requiring greater investment from platforms and workers, these AI tools could ultimately reduce matching frictions.

Related Literature This project is primarily related to the literature on signaling in labor markets. Building on the foundational paper by Spence (1978), an extensive body of work has explored how workers can use schooling and other signals to communicate their productivity to future employers. Emerging technologies such as Large Language Models can dramatically change how job applicants signal to employers, which is the focus of a few recent papers. Wiles et al. (2023) study the effect of algorithmic writing assistance on employment through an experiment in a similar labor market. Unlike the generative writing assistance studied here, their intervention is non-generative, focusing mostly on correcting spelling and grammar errors. They show that the tool improved treated workers' employment outcomes without affecting the employers' satisfaction. While we also investigate the role of AI in workers' employment outcomes (*partial equilibrium effects*), we are additionally able to speak to the aggregate market outcomes associated with the introduction of generative AI (*general equilibrium effects*). We find that while the partial equilibrium effects of AI are positive, AI leads to the degradation of cover letters as a useful signal in labor matching.

Relatedly, Cowgill et al. (2024) examine the impact of generative AI on job applications and startup pitches through a survey experiment using participants from Prolific. They find that, on average, AI reduces employers' and investors' screening accuracy. Our study explores similar questions in a real labor market, focusing on the effect of AI on actual hiring decisions. A key difference between the settings is that Cowgill et al. (2024) explicitly inform employers whether applicants used ChatGPT, whereas, in many real-world contexts—including ours—employers typically cannot discern whether an applicant received assistance from AI. As a result, our findings are not directly comparable.

More broadly, our paper contributes to the growing literature on AI's impact on labor markets. Papers including Agrawal et al. (2019), Noy and Zhang (2023), Toner-Rodgers (2024), Brynjolfsson et al. (2025), and Acemoglu (2025) study AI's impact on labor productivity and innovation. Our focus on labor market signaling highlights an alternative mechanism in which AI can profoundly reshape labor markets.

Organization Section 2 provides information on the setting and the data, detailing the policy variation that we leverage. Section 3 uses a two-way fixed effect framework to study AI's impact on cover letter quality and callback rates. Section 4 turns to general equilibrium effects, showing that cover letters become a less informative signal after AI's introduction. Section 5 explores the impact of alternative AI design on matching efficiency. Section 6 concludes.

2 Context and Data

Our analysis employs data from Freelancer.com, which ranks among the world's largest online task-based labor markets both in terms of user base and the volume of posted jobs. Platforms like Freelancer.com, Upwork, and Fiverr serve as digital marketplaces that connect workers and employers worldwide, facilitating the exchange of remotely delivered tasks such as software development, sales and marketing support, and creative work—commonly referred to as "freelancing." Since its inception in 2009, the platform has attracted over 78 million users from 247 countries, regions, and territories, with over $10 billion in gross payments transacted.

2.1 The Matching Process

On Freelancer.com, the matching process proceeds in five main steps. First, firms or individual employers seeking to outsource a specific task post jobs, known as "projects," on the Freelancer.com platform. Second, workers with skills matching the job requirements submit bids on these projects. Bidding for jobs on Freelancer.com is free for workers; however, the number of bids a worker can submit each month is limited. Workers can increase their monthly bidding allowance by subscribing to one of the platform's membership plans (see https://www.freelancer.com/membership/). Each bid typically comprises two components: the proposed payment amount and a written job proposal, which serves as a cover letter outlining the worker's suitability for the task. Third, based on the set of submitted bids and the characteristics of the competing workers, the platform employs a recommendation algorithm (as a function of the volume and strength of a bidder's past reviews, referred to henceforth as the 'bidder review score) to rank the bids for employers. Fourth, employers review the bids, displayed in order of ranking, and select the most appealing one, initiating the work process. Alternatively, employers may opt to leave the job unfilled if none of the bids meet their requirements. During this process, employers can and often reach out to a subset of bidders for a private conversation, which we refer to as a callback. Finally, after the task is completed, both the worker and the employer rate each other's performance by assigning a star rating, which contributes to their respective reputations on the platform. Payments are handled through the platform and are typically divided into milestone payments.

2.2 AI Bid Writing Tool

On April 19, 2023, Freelancer.com launched the AI Bid Writer, a tool designed to streamline the process of proposal writing for freelancers. The feature allows users to generate tailored, AI-written proposals based on the content of job descriptions. By analyzing the text of a job posting, the AI Bid Writer produces a proposal that is contextually relevant and ready for immediate submission, reducing the effort and time required for freelancers to participate in the bidding process. The release of the AI Bid Writer was unannounced beforehand, ensuring that its adoption was unaffected pre-release anticipation. The tool is available exclusively to users subscribed to the Plus or higher membership tiers. Additional information about the feature and its functionality is provided on Freelancer.com's support page: AI Bid Writer Overview (https://www.freelancer.com/support/project/ai-bid-writer).

2.3 TFIDF: a Measure of Textual Alignment

To derive a measure of cover letter quality, we focus on textual alignment, which captures the degree of overlap in keywords between a cover letter and a project description. The underlying assumption is that more tailored cover letters should exhibit greater textual similarity with the job posting, reflecting the freelancer's effort in addressing the specific requirements of the employer. On the contrary, generic cover letters which are re-used across projects would exhibit lower textual similarity.

To quantify this, we utilize Term Frequency-Inverse Document Frequency (TFIDF), a simple text vectorization technique that assigns weights to words based on their importance within a document while accounting for their overall prevalence across the corpus. Specifically, for each project description and its corresponding bid proposal, following a pre-processing step to remove stopwords, we construct TFIDF vectors, where each element in the vector represents the weighted frequency of a word within the text. The TFIDF score for a given word w in document d is computed as:

$$\text{TFIDF}(w,d) = \text{TF}(w,d) \times \text{IDF}(w)$$

where $\text{TF}(w,d)$ represents the term frequency of word w in document d, and $\text{IDF}(w)$ is the inverse document frequency, given by:

$$\text{IDF}(w) = \log\left(\frac{N}{1+n_w}\right)$$

where N is the total number of documents in the corpus, and n_w is the number of documents containing word w.

After computing the TFIDF vectors for both the project description and the bid proposal,

we measure their cosine similarity to obtain a numerical representation of textual alignment. The cosine similarity between two TFIDF vectors, v_p (for the project description) and v_b (for the cover letter), is defined as:

$$\text{Cosine Similarity} = \frac{v_p \cdot v_b}{\|v_p\|\|v_b\|} \tag{1}$$

This measure ranges from 0 to 1, where a value closer to 1 indicates a higher degree of textual similarity between the project description and the cover letter, while a value near 0 suggests little to no overlap.

By constructing this textual alignment measure, we obtain a quantitative metric of how closely a freelancer's cover letter mirrors the employer's job posting, allowing for an analysis of how cover letter quality evolves over time and in response to AI. This measure, while simple, has several advantages. First, it is highly computationally inexpensive. Second, it is largely insensitive to word count and is normalized by document length—hence, variation in the word counts of project descriptions or cover letters do not introduce bias. Third, it is insensitive to word order or sentence structure, and hence explicitly focuses on content similarity while ignoring stylistic variations.[2]

We validate that in the pre-rollout sample, our measure of textual similarity is significantly predictive of hiring. In Table 1, we regress the outcome of whether a bid is selected on the cover letter's textual similarity score, the worker's level of experience (captured by the review score), and the normalized wage (i.e., the bid amount divided by the minimum project budget) using both a logistic and linear specification, controlling for project fixed effects. Using the logistic specification which restricts the sample to only bids on projects which are awarded, the coefficient of 0.625 suggests a one standard deviation increase in textual similarity (0.1459) increases the odds of selection by roughly 9.4%. According to the linear specification, a one standard deviation increase in textual similarity (0.1459) corresponds to an increase in absolute selection probability of 0.058 percentage points, or roughly an 8% increase in the probability of being selected compared to the raw mean.

Nonetheless, future iterations of this work may utilize explicit embeddings from large language models to capture a more nuanced measure of writing quality. By incorporating embeddings from models employing transformer-based architectures, future work can move beyond surface-level text similarity and assess semantic closeness, argument coherence, tone, and persuasiveness in bid proposals. Unlike TFIDF, which treats words independently, LLM-based embeddings allow for context-aware representations, capturing how words and phrases contribute meaningfully to the overall proposal. This would enable a richer evaluation of whether a freelancer's bid not only aligns with the job description but also demonstrates clarity, relevance, and persuasive writing quality.

Appendix Table A1 includes examples of textual similarity scores for selected bids on a project.

[2] Wiles et al. (2023) study the impact of an AI tool that primarily corrects spelling and grammar errors in a similar context.

Table 1: Cover Letter Quality and Bid Success

	Offer	
	(1) Logistic	(2) Linear Probability
TFIDF	0.625***	0.004***
	(0.146)	(0.001)
Wage	-0.158***	-0.0007***
	(0.016)	(0.00005)
Experience	2.94***	0.014***
	(0.120)	(0.0005)
Project FE	✓	✓
Outcome mean	0.006	0.006
N	82,061	376,681

Table notes: The table reports regression results where the dependent variable is an indicator for whether the bid is suceessful. Column (1) reports estimates from a logistic regression, while Column (2) reports estimates from a linear probability regression. Standard errors are reported in parentheses and clustered at the project level. TFIDF is a measure of a cover letter's relevance to the project description. Wage is calculated as the wage bid divided by the minimum project budget. Bids that are more than 30 times above the minimum budget are dropped from the sample. Experience is calculated as the ranking of a platform-calculated bidder review score among bids on the project. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

2.4 Data Description and Sample Restrictions

Our analysis focuses on all bids submitted to projects requiring skills categorized under "PHP" during the period period between March 11, 2023, and July 1, 2023. We exclude projects prior to March 11 to avoid distortion from a platform-conducted A/B test of the AI Bid Writer, which ran from February 22, 2023, to March 10, 2023. In constructing our sample, we omit bids from individuals whose membership plans were altered during the period, as well as from those who were not active both before and after the intervention. Consequently, bids corresponding to varying levels of AI access are excluded. Additionally, we exclude bids where the proposed wage exceeds 30 times the minimum project budget, as these outliers likely do not reflect typical bidding behavior. Finally, we omit projects that received only a single bid, as these are likely direct offers.

Table 2 provides descriptive statistics at the bid, bidder, and project levels for the analytic sample, which comprises of 1,170,353 bids submitted by 10,325 unique bidders across a total of 42,611 projects. In the post-release period following April 19, 2023, there are 742,544 bids submitted across 27,586 projects. At the bid level, 0.6% of bids are awarded, and 0.5% are awarded and subsequently accepted by the winning bidder. Ultimately, 0.4% of submitted bids result in a completed project. Wage bids average 3.638 times the minimum project budget, reflecting a

Table 2: Summary Statistics

	Mean	SD	P25	P50	P75
Bid-Level ($N = 1{,}170{,}353$)					
Awarded	0.006	0.078	0	0	0
Accepted	0.005	0.069	0	0	0
Completed	0.004	0.061	0	0	0
Bidder Review Score	0.834	0.136	0.732	0.861	0.947
Normalized Wage Bid	3.638	3.511	1.500	2.222	4.667
Text Similarity Score	0.149	0.169	0.025	0.081	0.219
Bidder Access to AI	0.916	0.278	1	1	1
Post AI ($N = 742{,}544$)					
AI Bid	0.132	0.338	0	0	0
Bidder-Level ($N = 10{,}325$)					
Number of Bids	113.351	410.799	4	8	39
Number of Awards	0.700	2.977	0	0	0
Number of Acceptances	0.589	2.465	0	0	0
Number of Completions	0.064	0.472	2.121	0	0
Avg. Bidder Review Score	0.671	0.140	0.560	0.598	0.762
Avg. Normalized Wage Bid	3.272	1.661	2.156	3.00	4.00
Avg. Text Similarity Score	0.139	0.119	0.046	0.106	0.201
Access to AI	0.360	0.480	0	0	1
Any Use of AI Post-Release	0.206	0.404	0	0	0
Post AI ($N = 10{,}325$)					
Ratio of Bids Using AI	0.085	0.239	0	0	0
Ratio of Bids Using AI Among Users	0.391	0.377	0.038	0.250	0.750
Project-Level ($N = 32{,}231$)					
Number of Bids	32.121	31.380	10	21	44
Awarded	0.335	0.472	0	0	1
Accepted	0.275	0.447	0	0	1
Completed	0.214	0.410	0	0	0
Mean of Bidder Review Score	0.816	0.063	0.781	0.826	0.861
SD of Bidder Review Score	0.125	0.027	0.113	0.127	0.140
Mean of Text Similarity Score	0.156	0.085	0.099	0.145	0.200
SD of Text Similarity Score	0.154	0.068	0.110	0.152	0.196
Post AI ($N = 17{,}204$)					
Any AI Bid	0.825	0.380	1	1	1
Ratio of Bids Using AI	0.130	0.109	0.067	0.122	0.176

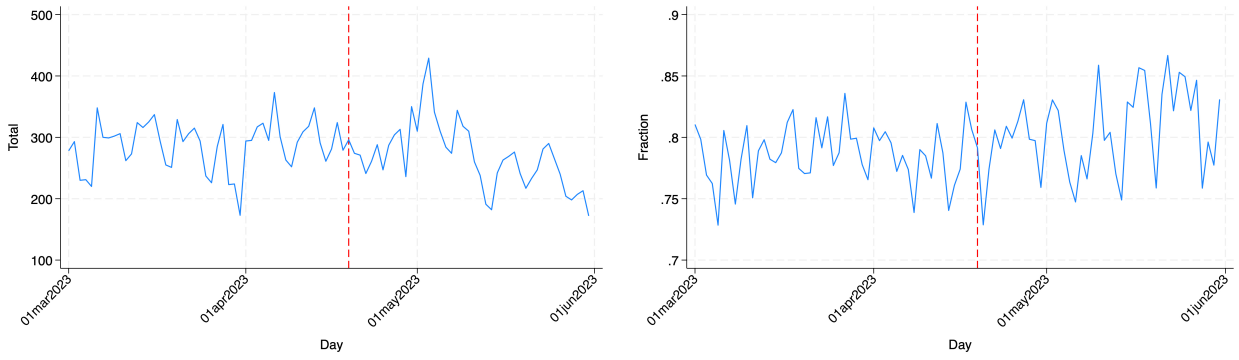
Notes: Summary statistics for three levels of analysis. Bid-Level variables pertain to individual bids. Percentiles (P25, P50, P75) represent the 25th, 50th (median), and 75th percentiles respectively. The Bidder Review Score is a platform calculated score based on the volume and strength of a bidder's past reviews. The normalized wage bid is calculated as the total bid divded by the minimum project budget. The text similarity score is computed as the cosine similarity between the vectorized TF-IDF of the project proposal and the project description.

range of freelancer pricing strategies. Cover letters achieve an average text similarity score of 0.149, though the distribution is positively skewed, with many bids scoring relatively low. Notably, 91.6% of bids in the sample are submitted by bidders with access to the AI Bid Writer tool, and in the post-release period, 13.2% of all bids are written using the tool.

At the bidder level, the average number of bids submitted during the sample period is 113.351, though this figure masks substantial heterogeneity. A minority of highly active bidders account for the majority of bids, reflecting a skewed distribution of activity. On average, bidders receive 0.700 awards in PHP projects during the sample period. At the time of the AI tool's rollout, 36.0% of bidders had access to the feature, and 20.6% of bidders utilized it post-release. Among those using the tool, 23.6% of their bids in the post-period were generated using AI, indicating significant reliance on the feature among adopters.

At the project level, the average project receives 32.121 bids, underscoring the competitive nature of the platform. Of all projects, 33.5% are awarded, 27.5% are accepted by the winning bidder, and 21.4% are ultimately completed. For projects with bids submitted after the release of the AI tool, 82.5% have at least one AI-generated bid. Across all projects, the average percentage of bids on the project generated by AI is 13.0%, indicating the rapid integration of AI-generated bids into the platform's ecosystem.

Figure 1: Purchases of Memberships with AI Access



(a) Number of Purchases

(b) Fraction of All Purchases

Notes: The left figure plots the purchases and renewals of membership plans that give access to the AI tool. The right figure plots the fraction of purchases and renewals of membership plans that give access to the AI tool out of those of all membership plans.

2.5 Membership Plan Adoption

A potential concern regarding the validity of our design is that the treatment and control groups might have evolved over time, especially around the introduction of the AI tool. In all the

analyses below, as discussed prior, we limit the samples to workers who did not switch between treatment and control groups in the sample period (i.e. workers who always had membership plans that granted them access to the tool when it was introduced and workers who never had such membership plans). Nevertheless, we show evidence that in our data sample, we do not observe an increase in membership plan purchases around the time when AI was added to the plans.

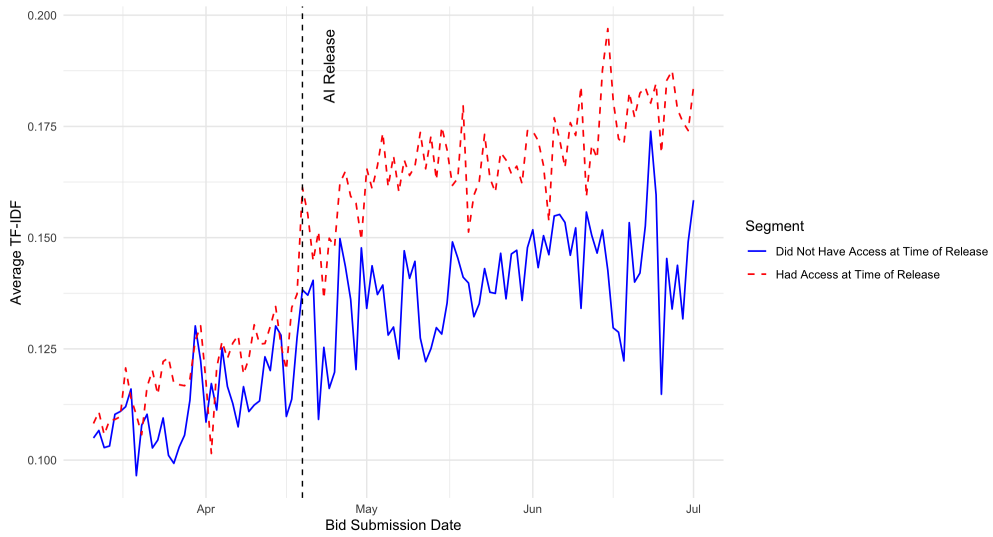
Figure 1a shows that the frequency of purchases and renewals of membership plans that give access to the AI tool is stable in the period studied. Figure 1b shows that the proportion of purchases for AI-qualifying membership plans out of all membership plans is also stable around the time of the tool's rollout. The evidence suggests that the introduction of AI tool did not affect selection into treatment.

3 The Partial Equilibrium Impact of AI

In this section, we study the partial equilibrium effects of the generative AI writing tool. Using a two-way fixed effect specification, we compare the outcomes between workers who did and did not have access to the tool before and after the rollout of the tool. The outcomes we consider are the cover letter quality and the probability of receiving a callback.

3.1 The Impact of AI on Cover Letter Quality

Figure 2: Time Series of Average Cover Letter Quality by Membership Status



Notes: Text similarity is computed as the cosine similarity between TF-IDF vectors of a project proposal and a project description.

We begin by examining the impact of access to the AI writing tool on textual similarity scores

in cover letters. Figure 2 presents a visual comparison of average text similarity scores between bidders with and without access to the AI tool. Before the rollout of the AI tool, the quality of cover letters written by workers with access followed a similar trend to those without access. However, after the rollout, the two lines diverged, with the former experiencing a noticeable and sustained improvement.

To formally estimate the effect, we employ a two-way fixed effects intention-to-treat (ITT) estimator, specified as follows:

$$y_{ijt} = \delta_i + \gamma_t + \beta \text{Post}_t \times \text{Access}_i + x_{ijt} + \epsilon_{ijt}$$

where the outcome y_{ijt} is text similarity score of the cover letter written by worker i at auction j in week t, δ_i captures worker fixed effects, γ_t represents week fixed effects, Post_t is an indicator for whether the bid was submitted post-rollout, Access_i is an indicator for whether the bidder had access to the tool at rollout, and x_{ijt} is a vector of additional controls including the experience level of the worker at the time of bidding (based on their platform-calculated review score) and the normalized wage bid (wage bid divided by the minimum project budget).

While the above recovers an intention-to-treat (ITT) effect of AI access on labor market outcomes, we additionally obtain the local average treatment effect (LATE). The need to estimate LATE arises from imperfect compliance: not all workers with access to the AI writing tool used it, and those who did use it did not necessarily do so for every auction. This non-compliance implies that the ITT estimate is a lower bound of the effect of AI usage, as it averages over cases where the eligible workers used the tool and cases where the eligible workers did not use the tool.

To estimate the LATE, we instrument for AI usage with treatment using the below specification:

$$y_{ijt} = \delta_i + \gamma_t + \beta \text{AI Bid}_{it} + x_{ijt} + \epsilon_{ijt},$$

where the dependent variable, y_{ijt}, again represents textual similarity for bidder i on auction j in week t. The variable AI Bid_{ijt} is an indicator for whether the AI tool was used to write the bid. Observing usage of the AI tool at the application level is a key advantage of our paper. We instrument for AI Bid_{ijt} using $\text{Post}_t \times \text{Access}_i$, the interaction of AI tool access and the post-rollout period. This is a valid instrument because the release of the tool was not pre-announced.

Table 3 contains the results from the above specifications. We find that access to the AI tool significantly increased textual similarity scores. Under the ITT specification, the coefficient on the interaction term is 0.0242, which is 17% of the pre-rollout textual similarity standard deviation of 0.1459. Relative to the pre-rollout baseline average textual similarity score of 0.120 among those with access, this corresponds to an overall increase of 20%. We note with the absence of full compliance, the estimate of accessing the tool is a lower bound of the effect of using the AI

Table 3: Effect of AI on Cover Letter Quality

	Cover Letter Quality	
	(1) ITT	(2) LATE
AI	0.0242***	0.1972***
	(0.0032)	(0.0194)
Bidder FE	✓	✓
Week FE	✓	✓
Outcome Std. Dev.	0.1459	0.1459
Outcome Mean	0.120	0.120
N	1,170,353	1,170,353
R^2 (Adj.)	0.3651	0.4184

Table notes: Cover letter quality is measured as text similarity – the cosine similarity between TF-IDF vectors of a project proposal and a project description. Column 1 reports the ITT: the coefficient of Post × Access in the two-way fixed effect specification. Column 2 reports the LATE, where we instrument for the usage of AI using Post × Access. Wage bids and experience levels are controlled for. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

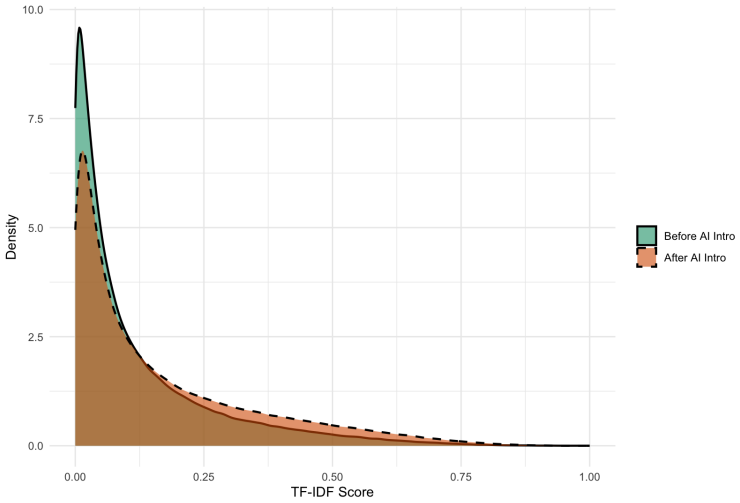
tool, which we discuss now.

Employing the LATE specification in column (2), the coefficient on using AI bids is 0.1972, which is 135% the pre-rollout standard deviation in similarity scores of 0.1459. Relative to the baseline mean similarity score among the treated, these estimates suggest that using the AI Bid Writer is associated with increases of 144% in textual similarity.

Our analysis so far has focused on the average effect of AI on cover letter quality. However, AI could shift the distribution of the cover letter quality in meaningful ways. In Figure 3, we plot the distribution of cover letter quality for workers who had access to the tool before and after the tool's rollout. We note that the post-rollout sample includes both bids that used the tool and bids that did not, despite having access. Notably, we observe a significant decline in the percentage of cover letters exhibiting extremely low textual similarity scores. This makes intuitive sense as now the cost of tailoring is drastically reduced.

Heterogeneity The previous section found that the AI Bid Writer had a significant and positive impact on cover letter quality. The previous section established that the AI Bid Writer significantly enhanced cover letter quality. However, the impact of generative AI assistance may vary across workers, for instance, depending on their prior ability to craft relevant and tailored cover letters. On the one hand, workers who were already skilled at customizing cover letters may be better positioned to leverage the tool's capabilities. On the other hand, those who previously wrote more generic cover letters might experience a greater improvement, as they have more room to grow. As we illustrate with a model in Section 5, understanding whether the AI tool complements or

Figure 3: Distribution of Cover Letter Quality Before and After AI Roll-out



Notes: Text similarity is computed as the cosine similarity between TF-IDF vectors of a project proposal and a project description. The sample is limited to bids by workers who had access to the tool, but is not limited to bids for which the tool was used.

substitutes workers' existing writing abilities is crucial for economic models assessing AI's impact. Our unique data setting allows us to provide direct evidence on the nature of this heterogeneity.

Table 4: Heterogeneous Effect of AI on Cover Letter Quality

	Cover Letter Quality
$\text{Post}_t \times \text{Access}_i$	0.025***
	(0.004)
$\text{Post}_t \times \text{Access}_i \times \text{Pre-AI Ability}_i$	-0.013***
	(0.003)
Bidder FE	✓
Week FE	✓
Outcome Std. Dev.	0.1459
N	1,170,353
R^2 (Adj.)	0.3664

Table notes: The table reports OLS regression results with text similarity (TFIDF) as the outcome. Heterogeneous effects are based on the worker's average TFIDF (ranges from 0 to 1) in the period before the AI tool was introduced. Wage bids and experience levels are controlled for. Standard errors are reported in parentheses and clustered at the bidder level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

To explore whether the impact of the AI writing tool differs based on a worker's pre-AI writing ability, we augment our baseline ITT specification by interacting AI Bid_{it} with a worker's pre-AI

writing ability q_i as shown below:

$$y_{ijt} = \delta_i + \gamma_t + \beta_1 \text{Post}_t \times \text{Access}_i + \beta_2 \text{Post}_t \times \text{Access}_i \times q_i + x_{ijt} + \epsilon_{ijt}.$$

The pre-AI writing ability q_i is measured as the average of the text similarity of worker i's cover letters before the introduction of the AI writing tool. We standardize q_i to have mean 0 and standard deviation 1. We report the results in Table 4.

We find that the AI tool *substitutes* for workers' signaling ability. Workers who were already effective at signaling through cover letters experienced smaller improvements in cover letter quality after gaining access to the AI tool compared to those who previously struggled to do so. The top writer experienced half as much gains as the worst writer. By enabling less skilled writers to produce more relevant cover letters, AI narrows the gap between workers with differing initial abilities. This suggests that the tool may dilute the informational value of cover letters, a hypothesis we formally test in Section 4.

3.2 The Impact of AI on Callbacks

Having established that AI improves the quality of workers' cover letters, we now use the same specifications to study the impact of AI on a key labor market outcome, the probability of receiving a callback. In our setting, as in most labor markets, employers often reach out to a subset of the bidders for a private conversation before making the final hiring decision, which we refer to as a callback. We observe which cover letters resulted in a callback and use that as the outcome variable.

Table 5: The Effect of AI on Receiving a Callback

	Callback	
	(1) ITT	(2) LATE
AI	0.0056*	0.0455**
	(0.0022)	(0.0185)
Bidder FE	✓	✓
Week FE	✓	✓
Outcome Mean	0.072	0.072
N	1,170,353	1,170,353
R^2 (Adj.)	0.0489	0.0477

Table notes: Column 1 reports the ITT: the coefficient of Post × Access in the two-way fixed effect specification. Column 2 reports the LATE, where we instrument for the usage of AI using Post × Access. Experience level and wages bids are controlled for. Standard errors are reported in parentheses and clustered at the bidder level. *** p < 0.01, ** p < 0.05, * p < 0.1.

Table 5 contains the coefficients. We find that the introduction of the AI Bid Writer increases

the probability of receiving a callback for those with access. AI access increases the probability of receiving a callback by 0.56 percentage points and AI usage increases the probability of receiving a callback by 4.55 percentage points. Given the baseline callback rate of 7.21% among the treated, the LATE estimate corresponds to a relative increase in interview probability of 63%.

We consider the 63% increase in the probability of receiving a callback to be substantial, yet within the range of effects found in studies examining the determinants of callback rates. For comparison, Bertrand and Mullainathan (2004) find that changing the job applicant's name from an African American-sounding name to a white-sounding name increases the callback rate by 50%. Our results suggest that generic cover letters are a barrier to receiving callbacks. Tools that help workers better customize the cover letters through just one click help workers get a foot in the door. At least in partial equilibrium, workers benefit from the AI writing tool because it gets them closer to landing the job. In the next section, we show that there are downsides to the AI writing tool when we consider general equilibriume effects.

4 The General Equilibrium Impact of AI

As previously shown, the AI writing tool had a larger impact on workers who previously wrote more generic cover letters. In other words, the AI writing tool substitutes from workers' existing ability. Therefore, AI potentially weakens the relationship between workers' innate ability and the cover letters that they produce, making cover letters a less informative signal. We now formally test this hypothesis by examining how the correlation between cover letter text similarity and hiring decisions changes after the release of the tool.

Our approach follows the specification:

$$y_{ijt} = \gamma_t + \beta_1 q_{ijt} + \beta_2 \text{POST}_t \times q_{ijt} + x_{ijt} + \epsilon_{ijt},$$

where y_{ijt} is the outcome (either selection or interview) of the application by worker i for project j in week t, POST_t is an indicator for post-rollout, q_{ijt} is the text similarity of the cover letter, and x_{ijt} includes the wage bid and worker's experience level.

As Table 6 shows, before the introduction of the AI tool, the textual similarity of the cover letter strongly predicts the worker's chances of receiving an interview and receiving an offer. However, after the tool was introduced, that predictive power decreases. The correlation between cover letter quality and callback rates decreases by 40% (the change is significant at the 5% level but not at the 1% level). The correlation between cover letter quality and whether the worker receives an offer decreases by 53%.

Table 6: Signaling Power of Cover Letter: Pre vs. Post AI

	Interview	Offer
TFIDF	0.005***	0.034***
	(0.0009)	(0.001)
POST × TFIDF	-0.003**	-0.016***
	(0.001)	(0.003)
Wage	-0.0002***	-0.002***
	(0.00002)	(0.00009)
Experience	0.011***	0.091***
	(0.0003)	(0.001)
Week FE	✓	✓
Outcome Mean	0.060	0.006
N	1,265,934	1,265,934

Table notes: Wage bids and experience levels are controlled for. Standard errors are reported in parentheses and clustered at the bidder level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

We estimate an alternative specification

$$y_{ijt} = \gamma_t + \sum_{\tau=-\kappa_1}^{\kappa_2} \beta_\tau 1\{t = t^* + \tau\} q_{ijt} + x_{ijt} + \epsilon_{ijt},$$

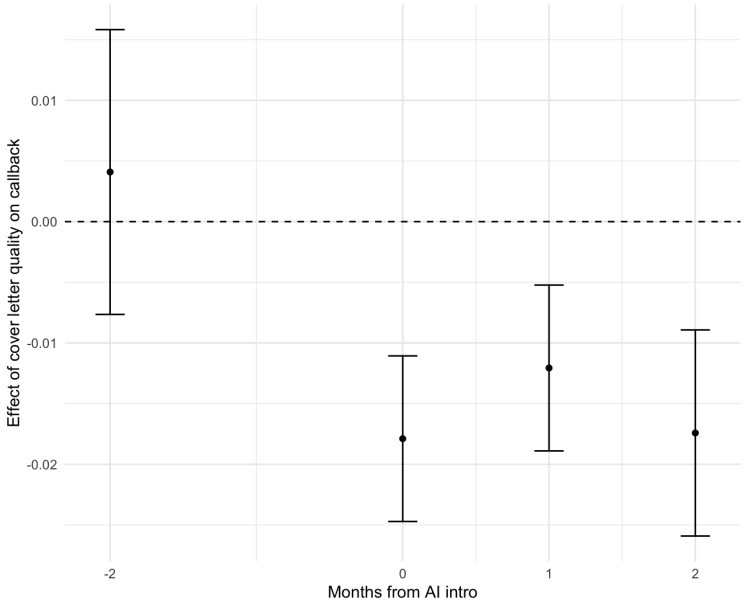
where t is month, y_{ijt} is whether the cover letter by worker i for job j in month t received a callback, q_{ijt} is the cover letter quality, and x_{ijt} includes the usual controls. Figure 4 presents the coefficients on month indicators times textual similarity scores, which are significantly negative in several months following the rollout of AI Bid Writer, suggesting that having a higher TF-IDF is less advantageous post AI-rollout.

Together, these findings suggest a weakening of the signaling value of well-tailored cover letters. Despite only a minority of bids utilizing AI-generated proposals, the tool's availability disrupts the traditional link between cover letter quality and intrinsic bidder ability, reducing employers' reliance on textual similarity as a key hiring signal.

5 Model

In this section, we present a simple model that captures some of the forces at play when AI is made available to a subset of job applicants. The goal is to understand how the presence of this technology affects the signaling power of cover letter quality, and ultimately match outcomes.

Figure 4: Post AI Roll-out, Cover Letters Become Less Predictive of Callback



Notes: Interaction term event study coefficients on Month × Textual Similarity with interview as the outcome. Wages, experience levels, and month FEs are controlled for.

5.1 Set-up

This model consists of workers (i) and firms (j). Workers are either high-productivity (with probability ρ_q) or low-productivity. The quality of the cover letter, denoted by h_{ij}, is a function of workers' productivity $q_i \in \{0,1\}$ and their access to the AI tool $A_i \in \{0,1\}$:

$$h_{ij} = g(q_i, A_i) + \epsilon_{ij},$$

where $\epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$. Workers have private information about their productivity and AI access. Firms observe the cover letter quality and use it to infer worker productivity but do not know whether AI was used. However, they do know that AI access is independent of worker quality. Thus, cover letters serve as a signaling device for hiring decisions. The randomness in cover letter quality and firms' inability to observe AI use introduce uncertainty in assessing worker productivity.

$g(q,A)$	$A=0$	$A=1$
$q=1$	$g(1,0)$	$g(1,1)$
$q=0$	$g(0,0)$	$g(0,1)$

Guided by our empirical results, we assume $g(q,1) \geq g(q,0), \forall q$, meaning AI does not worsen cover letter quality for either type of workers. We also assume $g(1,A) \geq g(0,A), \forall A$, ensuring that

high-productivity workers produce cover letters at least as strong as those of low-productivity workers. This assumption aligns with our finding that firms value cover letter quality in hiring decisions, treating it as a signal of workers' private information of productivity.

This simple framework captures various ways AI can influence cover letter quality. We define AI as having decreasing returns if it improves the cover letter of low-productivity workers more than those of high-productivity workers:

$$g(1,1)-g(1,0)<g(0,1)-g(0,0)$$

or the other way around in the case of increasing returns. Finally, this function nests from commonly used parametric specifications in the literature in which the effect of AI is assumed to be additive,

$$g(q,A)=\beta_q q+\beta_A A,$$

or multiplicative,

$$g(q,A)=(\bar{\alpha}+\alpha q)*(\bar{\beta}+\beta A).^{3}$$

Expected productivity As mentioned above, cover letter quality acts as a signal of worker productivity. Given this, to understand the effect of AI on hiring outcomes we need to understand how cover letter quality translates into expected productivity. Let ρ_A be the probability of access to the tool (in this model, access to AI is random and independent of worker quality). Let $\psi(h;g(q,A),\sigma_\epsilon^2)$ be the pdf of a Normal distribution with mean $g(q,A)$ and variance σ_ϵ^2 evaluated at h. To ease notation we define $\psi_{q,A}:=\psi(h;g(q,A),\sigma_\epsilon^2)$. Bayes's rule implies

$$\mathbb{E}[q_i|h_{ij}]=\frac{\rho_q(\rho_A\psi_{1,1}+(1-\rho_A)\psi_{1,0})}{\rho_q(\rho_A\psi_{1,1}+(1-\rho_A)\psi_{1,0})+(1-\rho_q)(\rho_A\psi_{0,1}+(1-\rho_A)\psi_{0,0})}.$$

Hiring decisions We assume that the firm's utility is positive if and only if it works with a worker whose productivity surpasses a firm-specific threshold

$$u_{ij}=1\{q_i-T_j>0\}.$$

Therefore, the firm hires a worker if and only if

$$\mathbb{E}[q_i|h_{ij}]>T_j.$$

[3] Cowgill et al. (2024), for example, adopts an additive specification, but with continuous types.

For the numerical exercises below, we consider the case where the threshold equals the average quality $T_j = \rho_q$.

5.2 Implications

Partial equilibrium The first implication of the model is that, abstracting from the independent shock ϵ, access to the AI tool increases the probability of being hired. Specifically, workers who receive the treatment, i.e., who use AI in their cover letter, have a higher chance of being hired compared to workers of the same type who do not receive the treatment. This result follows from the fact that cover letter quality is monotonic in A ($g(q,1) > g(q,0), \forall q$) and that $\mathbb{E}[q|g(1,1)] - \mathbb{E}[q|g(1,0)] > 0, \quad \mathbb{E}[q|g(0,1)] - \mathbb{E}[q|g(0,0)] > 0$. When we account for the presence of the exogenous shock, this conclusion still holds under certain conditions. A sufficient condition that guarantees $\frac{\partial \mathbb{E}[q|h]}{\partial h} > 0$ is $g(1,0) \geq g(0,1)$.

General equilibrium: expected productivity We next investigate how the effect of cover letter quality on hiring decisions changes with the advent of AI. We first consider the case of AI with decreasing returns.

As Figure 5a illustrates, the correlation between expected quality and cover letter quality — represented by the slope of the curves — decreases in AI adoption rate ρ_A.[4] Intuitively, higher adoption makes it more difficult to differentiate between high and low productivity workers through cover letter quality.

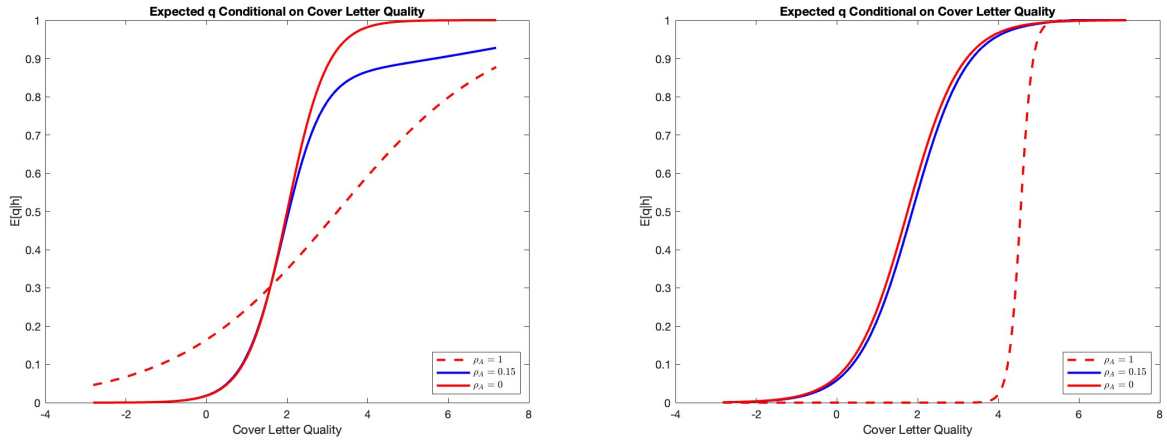
When considering the case of increasing returns, the conclusion reverses: higher adoption facilitates worker differentiation. As ρ_A increases from 0 to 1, the curves steepen, meaning that changes in cover letter quality lead to larger changes in $\mathbb{E}[q_i|h_{ij}]$, as illustrated in Figure 5b.

General equilibrium: firms' welfare We conclude by examining the impact of AI on firms' utility. Under our assumption, firms gain utility only when they hire high-productivity workers. However, due to information frictions, firms are prone to hiring mistakes — they do not always hire high-productivity workers, nor do they consistently reject low-productivity ones. Intuitively, firm welfare improves as the likelihood of correctly hiring high-quality workers and rejecting low-quality ones increases. We analyze how firm welfare varies with the rate of AI adoption and the nature of AI technology through numerical simulations.

As Figure 6a and Figure 6b illustrate, when AI exhibits decreasing returns, firms' likelihood of hiring both high-productivity and low-productivity workers is not monotonic in the share of AI adoption. At the limit, firms are more likely to hire high-productivity workers when all

[4] We use $\rho_A = 0.15$ as an intermediate case because that is approximately the fraction of bids made using AI in the context we study.

Figure 5: Expected Productivity Conditional on Cover Letter Quality



(a) Decreasing Returns **(b)** Increasing Returns

Notes: These numeric examples were made using $\rho_q = 0.5$ and $\sigma_\epsilon^2 = 1$. The decreasing returns case was computed with $g(1,1)=3.5$, $g(1,0)=3$, $g(0,1)=3$, $g(0,0)=1$ and the increasing returns was computed with $g(1,1)=8$, $g(1,0)=2.5$, $g(0,1)=1.1$, $g(0,0)=1$.

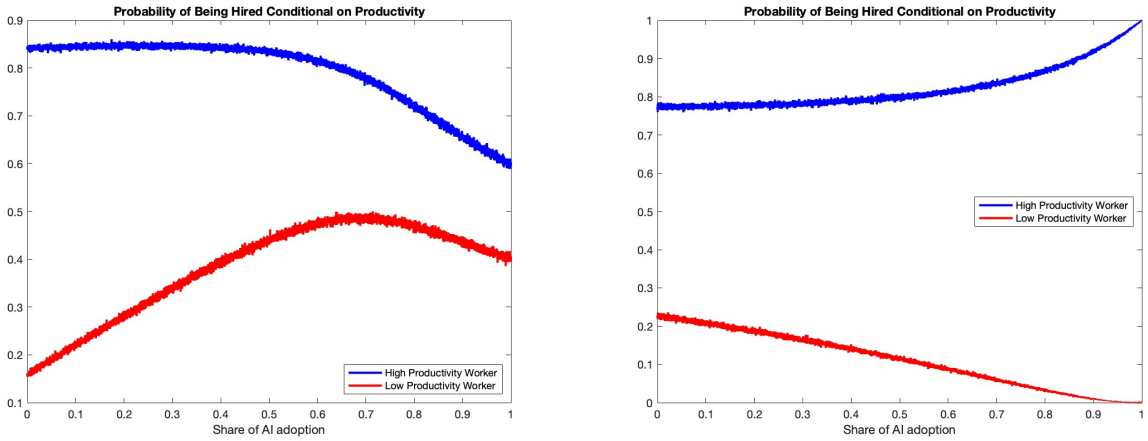
workers adopt AI than when none do. However, when AI exhibits increasing returns, firms benefit monotonically as AI adoption rises.

6 Conclusion

This study provides empirical evidence on the labor market consequences of AI-assisted job applications. While AI tools enable freelancers to generate more polished and tailored applications with less effort, our findings suggest that they fundamentally alter how employers interpret application quality. The widespread adoption of AI-assisted writing reduces the informational value of written proposals, making them less effective as a hiring signal. As a result, employers may struggle to distinguish between high- and low-quality candidates based on application materials alone, potentially reshaping hiring dynamics in online labor markets.

To harness AI's potential for improving labor market screening, we argue that AI should be designed to complement rather than substitute existing worker signals. Instead of merely generating persuasive text, AI could be leveraged to extract and highlight more relevant aspects of a worker's background—such as past work experience, skill assessments, and verified credentials—presenting this information in a structured and transparent manner. Such an approach would preserve the role of applications as a credible signal while reducing frictions in the hiring process, ultimately improving match quality between employers and workers.

Figure 6: Probability of Being Hired by Type



(a) Decreasing Returns

(b) Increasing Returns

Notes: These numeric examples were made using $\rho_q = 0.5$ and $\sigma_\epsilon^2 = 1$. The decreasing returns case was computed with $g(1,1)=3.5$, $g(1,0)=3$, $g(0,1)=3$, $g(0,0)=1$ and the increasing returns was computed with $g(1,1)=8$, $g(1,0)=2.5$, $g(0,1)=1.1$, $g(0,0)=1$.

References

Acemoglu, Daron, "The simple macroeconomics of AI," *Economic Policy*, 2025, *40* (121), 13–58.

Agrawal, Ajay, Joshua S Gans, and Avi Goldfarb, "Artificial intelligence: the ambiguous labor market impact of automating prediction," *Journal of Economic Perspectives*, 2019, *33* (2), 31–50.

Agrawal, Ajay K, Joshua S Gans, and Avi Goldfarb, "Ai adoption and system-wide change," Technical Report, National Bureau of Economic Research 2021.

Altonji, Joseph G and Charles R Pierret, "Employer learning and statistical discrimination," *The quarterly journal of economics*, 2001, *116* (1), 313–350.

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination," *American economic review*, 2004, *94* (4), 991–1013.

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

Cowgill, Bo, Pablo Hernandez-Lagos, and Nataliya Langburd Wright, "Does AI Cheapen Talk? Theory and Evidence From Global Entrepreneurship and Hiring," *Theory and Evidence*

From Global Entrepreneurship and Hiring (July 26, 2024). Columbia Business School Research Paper, 2024, (4702114).

Gans, Joshua S, "How will Generative AI impact communication?," *Economics Letters*, 2024, *242*, 111872.

Horton, John and Shoshana Vasserman, "Job-seekers send too many applications: Experimental evidence and a partial solution," *Manuscript, MIT*, 2021.

Horton, John J, "Buyer uncertainty about seller capacity: Causes, consequences, and a partial solution," *Management Science*, 2019, *65* (8), 3518–3540.

_ , Ramesh Johari, and Philipp Kircher, "Sorting through Cheap Talk: Theory and Evidence from a Labor Market," Technical Report, National Bureau of Economic Research 2024.

Hui, Xiang, Oren Reshef, and Luofeng Zhou, "The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market," *Organization Science*, 2024, *35* (6), 1977–1989.

Noy, Shakked and Whitney Zhang, "Experimental evidence on the productivity effects of generative artificial intelligence," *Science*, 2023, *381* (6654), 187–192.

Otis, Nicholas, Rowan Clarke, Solene Delecourt, David Holtz, and Rembrand Koning, "The uneven impact of generative AI on entrepreneurial performance," 2024.

Spence, Michael, "Job market signaling," in "Uncertainty in economics," Elsevier, 1978, pp. 281–306.

Toner-Rodgers, Aidan, "Artificial intelligence, scientific discovery, and product innovation," *arXiv preprint arXiv:2412.17866*, 2024.

Wiles, Emma, Zanele T Munyikwa, and John J Horton, "Algorithmic writing assistance on jobseekers' resumes increases hires," Technical Report, National Bureau of Economic Research 2023.

Table A1: Example Similarity Scores

"Brief: I need to add a horizontal scrolling section to a wordpress website. User is scrolling vertical (up & down) while content is moving left and right.

Solution: Plain HTML/CSS/JS is ok. Even better if we can use class/ID on the WPBakery Page Builder row to activate. Has to work on the demo site.

I have some examples that can be copied (CSS and JS methods)

Demo site: [LINK REDACTED]"

Proposal	**Similarity Score**	**AI Bid Writer**
"Hello! I understand that you need to add a horizontal scrolling section to your WordPress website, but the user needs to be able to scroll up and down while content is moving left and right. I also check the attached url with reference and I can make this perfectly. My name is Sandeep and I am a Full Stack Developer with over 7+ years of experience in the industry. I specialize in larger web applications and software development, as well as larger website builds. My solution is to use class/ID on the WPBakery Page Builder row to activate the section. This method works on the demo site and doesn't require any modifications to the site's code. If you are interested in my services please feel free to take a look at my portfolio or contact me directly for more information. Thank you for your consideration!"	0.541	Yes
"Hello, I have experience working on projects similar to yours, so I'm confident that I can deliver an effective solution that meets your needs. I understand that you need to add a horizontal scrolling section to your wordpress website so that the user can scroll vertically (up & down) while the content moves left and right. I can help with this project by using plain HTML/CSS/JS methods that work on the demo site. If you're interested in moving forward with this project, please let me know and we can set up an initial interview/call to discuss further. Thank you for considering me!"	0.488	Yes
Hello I am a senior Web and Mobile app developer and I have worked on many projects similar to this one. I understand that you need to add a horizontal scrolling section to your wordpress website, but instead of just moving from left to right, the content needs to move from left to right while the user is still able to scroll up and down through the section. My skillset allows me to tackle projects like this and bring the user an exciting experience. I am confident that my expertise can help bring this project to success. If given the chance, I would love to work on this site with you! Please don't hesitate to reach out if you have any questions or would like more information about me or my skillset. Thank you for your consideration!	0.354	Yes
"Hi there, you're looking for professional web developer to add a horizontal scrolling section to a wordpress website. . As I've expertise in this area, so I can show you satisfied result in your timeline. I'd like to discuss more on chat so we can get clear idea each. Thanks"	0.325	No
Hi, I can help u to add Vertical Scrolling function to ur WP Website I checked the demo site and found that It's easy task for my team because we had 10 + years experience in it Plz chat now bro !	0.339	No
"Hi Joshua M.! Because of my considerable experience in website development, I believe I would be a wonderful fit for this position. I have examined every detail, and everything appears doable, but if you could start the chat, I have a few questions. Here's what I can bring to your project: 1) Top-notch experience 2) Knowledge of best practices for data security 3) 24/7 communication: I will keep you updated with project's status and consistent updates 4) Someone who cares about helping you succeed and bring more value to your organization I am looking forward to hearing more about your exciting project!	0.001	No