AI Shock and Education Premium in Online Labor Markets: Empirical Evidence after the Launch of ChatGPT

Abstract

The rapid diffusion of generative artificial intelligence (AI) has raised important questions about how new technologies reshape labor markets and the distribution of economic opportunities. This study examines the impact of the generative AI shock associated with the public release of ChatGPT on labor outcomes in online freelance markets. Using a worker–month panel dataset from a large global online labor platform covering the period from 2022 to 2025, we implement a difference-in-differences design to compare changes in outcomes between freelancers with different levels of pre-shock AI exposure. Our results show that freelancers with higher pre-existing AI exposure experienced relative declines in earnings, working hours, and contract opportunities following the introduction of generative AI tools. We further document substantial heterogeneity across educational backgrounds: while highly educated freelancers also experienced negative effects, the magnitude of these declines is significantly smaller compared to workers with lower educational attainment. Additional analysis based on university ranking reveals that freelancers from elite institutions are largely insulated from the adverse impact of AI shocks, whereas workers from lower-ranked or unranked universities experience substantial earnings reductions. These findings suggest that generative AI may simultaneously expand technological capabilities while redistributing opportunities across workers, reinforcing the importance of education and institutional signals in digital labor markets.

Keywords: Generative AI; ChatGPT; AI shock; Education premium; Gig economy; Online labor platforms; Upwork; Difference-in-differences

Introduction

The rapid diffusion of large language models (LLMs) has introduced a new technological discontinuity for platform-mediated work. OpenAI released ChatGPT to the public on November 30, 2022, lowering the cost of generating fluent text, code, and other digital outputs for a broad set of non-routine cognitive tasks. Recent advances in large language models (LLMs), such as ChatGPT, significantly reduce the cost of producing text, code, and other digital outputs that were previously performed by knowledge workers. Early empirical evidence suggests that generative AI can substantially increase worker productivity while also reshaping the distribution of labor demand across occupations and tasks (Noy & Zhang, 2023; Brynjolfsson, Li, & Raymond, 2025). At the same time, exposure analyses indicate that a large share of knowledge work activities overlap with capabilities of generative AI systems, raising concerns about potential substitution effects in certain task domains (Felten et al., 2023). Recent studies document both productivity gains and labor market adjustments associated with generative AI (Noy & Zhang, 2023; Brynjolfsson et al., 2025).

These dynamics may be particularly visible in online labor platforms (OLPs) such as Upwork, where digital tasks are traded in highly competitive and transparent markets (Liang et al., 2025). OLPs allow clients to source freelance labor globally, while workers compete based on observable signals such as skills, ratings, and credentials (Kässi & Lehdonvirta, 2018; Wood et al., 2019). Because tasks on these platforms are modular and often digitally delivered, technological changes that alter the production of knowledge work may rapidly affect demand, pricing, and worker outcomes. In addition, platform labor markets are characterized by substantial information asymmetry, making observable signals—such as educational attainment and institutional prestige—important determinants of hiring decisions (Spence, 1973; Weeden, 2002).

The introduction of generative AI may therefore reshape labor market outcomes in two important ways. First, generative AI can augment productivity by enabling workers to complete tasks more efficiently or produce higher-quality outputs (Noy & Zhang, 2023). Second, generative AI may substitute for certain tasks, especially those involving standardized cognitive production, potentially reducing demand for freelancers whose services overlap with AI capabilities (Agrawal et al., 2019). These changes may also alter the role of traditional signals in labor markets. When AI tools lower the cost of generating outputs, clients may rely more heavily on observable credentials—such as education or institutional prestige—to assess worker quality (Rivera, 2012).

Despite growing interest in the labor market consequences of generative AI, empirical evidence on how AI shocks affect worker outcomes and the role of education in moderating these effects remains limited. This study addresses this gap by examining how the public release of ChatGPT in November 2022 reshaped labor market outcomes in an online freelance platform.

Specifically, we address three research questions: (RQ1) How did the introduction of gener-

ative AI affect freelancers' labor market outcomes, including earnings, number of contracts, and labor supply? (RQ2) How do these effects vary across workers with different levels of pre-shock AI exposure? (RQ3) Does education—measured by educational attainment and university prestige—mitigate the labor market impact of AI shocks?

We use a worker–month panel from Upwork spanning January 2022 to December 2025 and implement a difference-in-differences design around the ChatGPT release. Our main treatment is workers' pre-shock AI exposure, measured by whether the worker listed any AI-related skills prior to November 2022. Intuitively, such workers are more likely to operate in task domains that were already AI-adjacent and therefore more directly affected by the sudden democratization of LLM capabilities.

Our exploratory findings are threefold. First, high AI exposure workers experience a statistically and economically significant decline in post-shock earnings and labor supply. Second, the decline is materially smaller among workers with higher educational attainment, implying that the education premium within AI-exposed work expands rather than collapses. Third, credible quality signals—elite university affiliation—are associated with greater resilience, whereas workers from lower-ranked or unranked universities are hit hardest. Together, these patterns suggest that generative AI erodes returns in commoditized segments of knowledge work, while formal education and credible credentials remain valuable as adaptability and quality signals.

This study contributes to the literature in three ways. First, it provides early worker-level evidence on the labor market impact of generative AI in online freelance platforms, a setting where technological shocks may rapidly affect labor demand. Second, it contributes to research on technological change and labor market inequality by examining how AI exposure shapes heterogeneous outcomes across workers. Third, it extends the literature on education and signaling in platform labor markets by showing how traditional credentials interact with new forms of cognitive automation.

Literature Review and Theoretical Background

Online Labor Platforms and Signaling

OLPs organize market exchange through algorithmic matching, search rankings, and reputational feedback (Liang et al., 2023; Liu et al., 2024). Platform governance relies heavily on algorithmic mechanisms such as search rankings, reputation systems, and performance metrics that shape worker visibility and access to opportunities (Kellogg et al., 2020; Möhlmann et al., 2021). Because transactions often occur between previously unacquainted actors, online labor platforms are characterized by substantial information asymmetry between clients and freelancers (Wood et al., 2019). Under such conditions, clients rely on observable signals to evaluate freelancer quality. Signaling theory suggests that educational credentials can serve as indicators of worker ability when productivity is difficult to observe directly

(Spence, 1973). In addition, institutional prestige may function as a stronger signal of baseline capability and occupational closure (Weeden, 2002).

Technological shocks that alter the production of digital tasks may further reshape these signaling dynamics. Research on technological change suggests that new technologies reorganize the task content of work by substituting for certain routine activities while complementing workers who possess higher-level skills or the ability to integrate new tools into complex workflows (Acemoglu & Autor, 2011). Recent evidence on generative AI similarly shows that AI tools can automate components of cognitive production while augmenting workers who effectively incorporate these tools into their work processes (Noy & Zhang, 2023; Brynjolfsson et al., 2025). In platform-based labor markets, such technological changes may therefore redistribute opportunities across workers depending on both their skill sets and the credibility of signals available to clients.

Generative AI and Task Substitution

Generative AI systems based on large language models (LLMs) represent a class of foundation models capable of performing or accelerating a wide range of text- and code-intensive tasks (Epstein et al., 2023; Dhamani & Engler, 2026). Exposure analyses suggest that LLM capabilities overlap substantially with knowledge-work activities and may affect a large share of occupations, particularly those involving cognitive production and information processing (Felten et al., 2023).

Early empirical evidence documents both substitution and augmentation effects. Experimental studies show that access to ChatGPT significantly increases productivity in writing tasks while reducing performance dispersion across workers (Noy & Zhang, 2023). Evidence from firm settings similarly indicates that AI assistance can raise worker productivity, particularly among less experienced employees (Brynjolfsson et al., 2025). These findings highlight the importance of examining heterogeneity in AI exposure and worker resilience in digital labor markets.

Education as Capability and Signal under AI

Education plays an important role in shaping workers' ability to adapt to technological change. Human capital theory suggests that higher levels of education enhance cognitive skills, learning capacity, and the ability to adopt new technologies (Goldin & Katz, 2008; Autor et al., 2003). Workers with higher education may therefore be better positioned to integrate emerging technologies into their workflows and adjust to technological disruptions.

Beyond human capital, education may also function as a signaling mechanism in labor markets characterized by information asymmetry. Educational credentials provide observable signals of worker ability when productivity cannot be directly evaluated (Spence, 1973), while institutional prestige may signal competence and reliability in competitive labor markets (Weeden, 2002).

This signaling role may be particularly important in online labor platforms. Clients hiring freelancers often face uncertainty regarding worker quality, especially when contracting with new or unfamiliar workers. Educational credentials and institutional prestige can therefore function as observable signals that reduce information asymmetry and facilitate hiring decisions.

The introduction of generative AI may further increase the importance of such signals. As AI tools lower the cost of producing standardized outputs, it becomes more difficult for clients to differentiate worker quality based solely on observable task outputs. In this context, credible signals—such as formal education or elite institutional affiliation—may become more valuable indicators of worker capability and trustworthiness.

Data and Empirical Strategy

We use panel data from a large online labour platform covering January 2022 to December 2025, aggregated to the worker-month level. The sample is restricted to active workers who earned positive income prior to the ChatGPT launch (November 2022), ensuring that the panel tracks established participants rather than transient entrants.

Variables

Dependent variables. We examine labor-market outcomes that capture both quantity and price dimensions of online work. Our main dependent variables are monthly earnings, average earnings, monthly hours supplied, and the monthly number of contracts. To mitigate right-skewness and interpret coefficients in semi-elasticity terms, the empirical analysis applies log transformations where appropriate.

Independent variables. Our focal explanatory variables operationalize an AI shock and workers' pre-shock AI-related skill endowments. We define *Post* as an indicator equal to 1 for months from November 2022 onward and 0 otherwise. AI exposure is constructed based on freelancers' historical work experience prior to the introduction of generative AI. Specifically, we examine all jobs completed by each worker before the release of ChatGPT and identify whether these jobs involved AI-related technical skills. For each completed job, we inspect the associated skill tags and classify them as AI-related if they fall within artificial intelligence, computer science, or information technology domains, including machine learning, data science, natural language processing, deep learning, and related computational techniques. A worker is classified as having high AI exposure if at least one pre-ChatGPT job involves AI-related technical skills. Otherwise, the worker is classified as having low AI exposure. This measure captures whether a freelancer was already operating in task domains closely related to AI technologies prior to the generative AI shock.

Moderating variables. To study the education premium under AI shocks, we incorporate education and pedigree measures. Education is captured by the worker's highest degree category in month t. We construct a binary indicator in which *High Education* includes

workers with a master's or doctoral degree, while *Low Education* includes those with a bachelor's degree, associate degree, or high school diploma. We use the most recent available ranking for each institution to classify university prestige. University ranking is measured using the *Times Higher Education* (THE) World University Rankings dataset (2016–2026), obtained from a publicly available ranking repository. Freelancers' reported universities are matched with the ranking database and categorized into four groups: Top 100, Rank 101–500, Rank 501+, and Unranked institutions. University ranking is treated as time-invariant for each worker and matched using the worker's highest degree.

Control variables. Control variables include worker i's tenure on the platform (log-transformed), overall service quality metrics (e.g., ratings or service scores), profile richness (e.g., number of listed skills), and platform recognition (e.g., Top Rated status) in month t.

Table 1: Summary Descriptives

Category	Variable	Obs.	Mean	Std. dev.	Min	Max
Dependent Variable	$TotalHours_{it}$	465,048	19.7335	865.4473	0	399256.7
	$TotalEarnings_{it}$	465,048	708.5475	34829.54	0	15224775
	$NumOfWork_{it}$	465,048	0.9847693	53.62812	0	12425
Independent Variable	$HighAI_i$	465,048	0.2066152	0.4048774	0	1
	$PostGPT_t$	465,048	0.1283566	0.4616075	0	10.8
Moderating Variable	$HighestDegree_{it}$	465,048	Not suitable			
	$UniversityRank_{it}$	110,121	844.0694	487.2669	0	2160
Control Variable	$avgservicescore_{it}$	465,048	4.709885	0.9633503	0	5
	$avgskills_{it}$	465,048	4.733456	0.9630515	0	5
	$toprated_{it}$	465,048	0.5636087	0.495938	0	1
	$tenure_{it}$	465,048	455.7974	319.0526	0	1899

Empirical Strategy

We estimate the causal impact of the AI shock using a difference-in-differences (DiD) design centered on the launch of ChatGPT (November 2022). The baseline specification can be written as

$$Y_{it} = \beta\Big(Post_t \times HighAIExposure_i\Big) + \gamma X_{it} + \mu_i + \tau_t + \varepsilon_{it}, \tag{1}$$

where Y_{it} denotes an outcome for worker i in month t; $Post_t$ indicates the post-ChatGPT period; $HighAIExposure_i$ captures pre-shock AI exposure; X_{it} is a vector of time-varying controls; μ_i are worker fixed effects; and τ_t are month fixed effects. The coefficient of interest is β, which captures the differential post-shock change for workers with higher pre-shock AI exposure.

To quantify how AI shocks reshape the education premium, we extend Equation (1) by interacting $Post_t \times HighAIExposure_i$ with education measures (highest degree and university-

rank groups). Inference is based on heteroskedasticity-robust standard errors clustered at the worker level.

Potential mechanisms. Education may influence workers' resilience to technological shocks through two mechanisms. First, higher education may enhance cognitive flexibility and the ability to integrate new technologies into complex workflows. Second, education—particularly institutional prestige—may serve as a signaling mechanism in markets with information asymmetry. When generative AI lowers the cost of producing standardized outputs, credible signals of worker quality may become even more valuable for clients making hiring decisions.

Preliminary Results

Preliminary Results for Comparison

Table 2 reports the baseline difference-in-differences estimates of the impact of the AI shock on freelancers' labor market outcomes. The key variable of interest is the interaction term *Post* × *High AI Exposure*, which captures the differential post-shock effect for workers with higher pre-shock AI exposure. The results show that the interaction term is negative and statistically significant across all outcome variables, including total earnings ($\beta = -0.2411$, $p < 0.01$), average earnings ($\beta = -0.2091$, $p < 0.01$), total hours ($\beta = -0.1039$, $p < 0.01$), and the number of contracts ($\beta = -0.0458$, $p < 0.01$). These estimates indicate that freelancers with higher pre-shock AI exposure experienced significant relative declines in labor market outcomes after the introduction of ChatGPT, compared with freelancers with lower AI exposure. This pattern is consistent with a task substitution mechanism, where generative AI reduces the demand for certain tasks previously performed by freelancers with AI-related skills.

Two mechanisms may explain the observed pattern. First, generative AI may substitute for some tasks previously performed by freelancers with AI-related skills. As AI tools become widely accessible, clients may rely more on automated solutions or require fewer freelancer inputs for standardized tasks. Second, generative AI may lower entry barriers and intensify competition in AI-related task markets. AI tools enable a larger pool of freelancers to perform tasks that previously required specialized skills. As a result, competition increases, leading to lower earnings and fewer contracts for workers who were previously specialized in AI-related tasks.

Table 2: Difference-in-Differences Estimates of AI Shock

DV	(1) Total earnings	(2) Average earnings	(3) Total hours	(4) Number of contracts
Post × High AI Exposure	-0.2411***	-0.2091***	-0.1039***	-0.0458***
	(0.0293)	(0.0266)	(0.0136)	(0.0056)
Constant	2.2895***	2.0950***	0.5125***	0.4786***
	(0.0802)	(0.0743)	(0.0389)	(0.0149)
Observations	465,048	465,048	465,048	465,048
R-squared	0.0093	0.0092	0.0025	0.0117
Workers	29,467	29,467	29,467	29,467

Notes: Post is an indicator for months from November 2022 onward. Robust standard errors are in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 1 provides a model-free comparison of labor outcomes for freelancers with high and low AI exposure before and after the launch of ChatGPT. The figure shows that workers with higher AI exposure experience a more pronounced decline in earnings, hours worked, and job opportunities following the AI shock, suggesting potential substitution effects of generative AI.

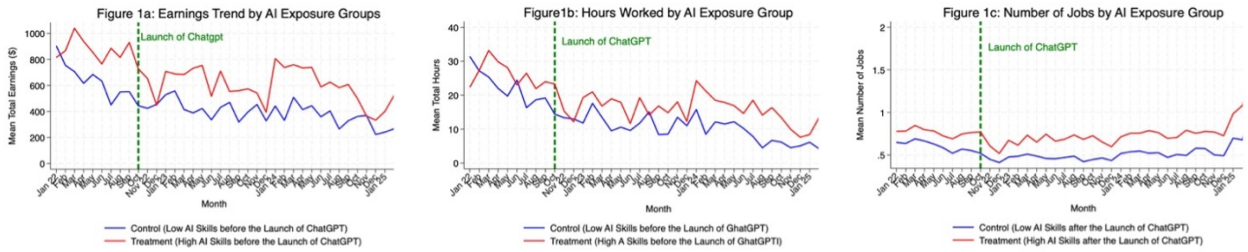


Figure 1: Model-Free Comparison After the Launch of ChatGPT

Heterogeneous Effects by Educational Attainment

Table 3 reports heterogeneous effects of the AI shock across freelancers with different levels of educational attainment. The key coefficients of interest are the interaction terms estimated separately for workers with high and low education levels.

For workers with higher educational attainment, the interaction coefficients are negative and statistically significant across all outcome variables, including total earnings ($\beta = -0.1047$, $p < 0.05$), total hours ($\beta = -0.0885$, $p < 0.01$), and the number of contracts ($\beta = -0.0174$, $p < 0.1$). These results indicate that freelancers with high AI exposure experienced relative declines in labor market outcomes following the introduction of ChatGPT, even among highly educated workers.

However, the magnitude of the negative effects is substantially larger among workers with

lower educational attainment. For the low-education group, the estimated interaction coefficients are -0.2080 ($p < 0.01$) for total earnings, -0.0853 ($p < 0.01$) for total hours, and -0.0399 ($p < 0.01$) for the number of contracts. In particular, the earnings effect for low-education workers is approximately twice as large as that observed among highly educated workers.

These findings suggest that although freelancers with pre-existing AI-related skills experienced relative declines after the AI shock, higher education mitigates the magnitude of these negative effects. One possible explanation is that highly educated workers possess stronger complementary capabilities—such as analytical skills, domain knowledge, and the ability to integrate AI tools into more complex workflows—which allow them to adapt more effectively to technological change. In contrast, freelancers with lower educational attainment may be more concentrated in tasks that are more easily substituted by generative AI.

Table 3: Heterogeneous Effects by Educational Attainment

DV	(1)	(2)	(3)	(4)	(5)	(6)
	Total earnings		Total hours		Number of contracts	
	High edu	Low edu	High edu	Low edu	High edu	Low edu
Post × High AI Exposure	-0.1047**	-0.2080***	-0.0885***	-0.0853***	-0.0174*	-0.0399***
	(0.0517)	(0.0394)	(0.0242)	(0.0186)	(0.0098)	(0.0076)
Constant	2.8477***	2.8489***	0.6736***	0.6674***	0.5517***	0.5568***
	(0.0712)	(0.0731)	(0.0338)	(0.0346)	(0.0140)	(0.0141)
Observations	465,048	465,048	465,048	465,048	465,048	465,048
R-squared	0.0079	0.0080	0.0019	0.0019	0.0100	0.0101
Workers	29,467	29,467	29,467	29,467	29,467	29,467
Clusters	29,467	29,467	29,467	29,467	29,467	29,467

Notes: Columns split the sample by educational attainment (high vs. low education). Robust standard errors are in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Heterogeneous Effects by University Rank Group

Table 4 presents heterogeneous effects of the AI shock across freelancers from universities with different ranking levels, focusing on total earnings. The key coefficients of interest are the interaction terms *Post* × *High AI Exposure*, which capture the differential post-shock effect for AI-exposed workers within each university rank group.

For workers from Top 100 universities, the interaction coefficient is small and statistically insignificant ($\beta = 0.0036$, $p > 0.1$), suggesting that highly ranked university graduates with pre-existing AI skills did not experience meaningful relative earnings declines following the release of ChatGPT. This finding indicates that elite educational pedigree may provide strong protection against technological disruption, potentially due to superior foundational training, stronger analytical capabilities, or access to more complex and complementary tasks that

remain less susceptible to automation.

In contrast, workers from lower-ranked or unranked universities experienced significant negative effects. The coefficients are -0.6125 ($p < 0.01$) for the 101–500 group, -0.2434 ($p < 0.01$) for the 501+ group, and -0.2274 ($p < 0.01$) for workers from unranked institutions. These results indicate that AI-exposed workers from non-elite universities experienced substantial declines in earnings following the AI shock. Institutional prestige appears to provide a degree of protection against technological disruption, reinforcing the importance of educational signals in platform-based labor markets.

Table 4: Heterogeneous Effects by University Rank Group on Total Earnings

	(1) Top 100	(2) Rank 101–500	(3) Rank 501+	(4) No Rank
Post × High AI Exposure	0.0036	-0.6125***	-0.2434***	-0.2274***
	(0.2281)	(0.2120)	(0.0703)	(0.0331)
Constant	-51.2444	614.4536	146.4112	0.0483
	(431.8722)	(443.1458)	(114.2764)	(.)
Observations	9,968	18,033	82,020	354,805
R-squared	0.0108	0.0102	0.0101	0.0083
Workers	688	1,326	5,325	22,607
Clusters	22,607	22,607	22,607	22,607

Notes: Robust standard errors are in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Conclusion

This study provides early evidence on how generative AI reshapes labor outcomes in online freelance markets. Our findings suggest that while generative AI may increase overall productivity, it also redistributes opportunities across workers depending on their exposure and educational background. Using an Upwork worker–month panel from 2022–2025 and a difference-in-differences design, we find that workers with pre-shock AI exposure suffered substantial declines in earnings and labor supply after the AI shock. The decline is meaningfully smaller for higher-education workers or elite university credentials, indicating that traditional human-capital and signaling mechanisms remain important even when AI can produce many outputs at low marginal cost.

The results are consistent with a task-substitution mechanism in AI-exposed segments of platform work: the democratization of LLM capabilities lowers clients' willingness to pay for standardized outputs and reduces the volume of contracts offered to workers whose services overlap with LLM functionality. At the same time, education and credible credentials appear to matter more—not less—after the shock. This is consistent with education functioning as (i) an adaptability resource that enables workers to reconfigure workflows, adopt complementary tools, and shift toward higher-value tasks, and (ii) a quality signal that remains

valuable when output generation becomes cheaper and quality becomes harder to verify.

These patterns refine the connection between generative AI and inequality. Rather than uniformly eroding returns to education, the AI shock can widen gaps within affected work domains by disproportionately harming workers with weaker signals and lower adaptability. For platforms, the findings imply that interventions such as targeted upskilling, AI-complementary training, and improved credential verification may help mitigate distributional harms. For policymakers, the results highlight that investments in transferable skills and digital literacy may be particularly valuable for workers in modular, globally competitive markets.

Future work can strengthen identification and mechanism tests by estimating event-study dynamics, separating substitution from complementarity at the task level, and linking worker transitions across job categories to changes in skill portfolios.

References

1. Agrawal, A., Gans, J., & Goldfarb, A. (2019). Artificial intelligence: The ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2), 31–50.
2. Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4), 1279–1333.
3. Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942.
4. Dhamani, N., & Engler, M. (2026). *Introduction to generative AI*. Simon and Schuster.
5. Epstein, Z., Hertzmann, A., Investigators of Human Creativity, Akten, M., Farid, H., Fjeld, J., et al. (2023). Art and the science of generative AI. *Science*, 380(6650), 1110–1111.
6. Felten, E., Raj, M., & Seamans, R. (2023). How will language modelers like ChatGPT affect occupations and industries? *arXiv preprint arXiv:2303.01157*.
7. Goldin, C., & Katz, L. F. (2008). *The race between education and technology*. Harvard University Press.
8. Kässi, O., & Lehdonvirta, V. (2018). Online labour index: Measuring the online gig economy for policy and research. *Technological Forecasting and Social Change*, 137, 241–248.
9. Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410.
10. Katz, L. F., & Murphy, K. M. (1992). Changes in relative wages, 1963–1987: Supply and

demand factors. *The Quarterly Journal of Economics*, 107(1), 35–78.

11. Liang, C., Hong, Y., & Gu, B. (2025). Monitoring and the cold start problem in digital platforms: Theory and evidence from online labor markets. *Information Systems Research*, 36(1), 600–620.
12. Liang, C., Peng, J., Hong, Y., & Gu, B. (2023). The hidden costs and benefits of monitoring in the gig economy. *Information Systems Research*, 34(1), 297–318.
13. Liu, Y., Lou, B., Zhao, X., & Li, X. (2024). Unintended consequences of advances in matching technologies: Information revelation and strategic participation on gig-economy platforms. *Management Science*, 70(3), 1729–1754.
14. Möhlmann, M., Zalmanson, L., Henfridsson, O., & Gregory, R. W. (2021). Algorithmic management of work on online labor platforms. *MIS Quarterly*, 45(4), 1999–2022.
15. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.
16. Rivera, L. A. (2012). Hiring as cultural matching: The case of elite professional service firms. *American Sociological Review*, 77(6), 999–1022.
17. Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics*, 87(3), 355–374.
18. Weeden, K. A. (2002). Why do some occupations pay more than others? Social closure and earnings inequality. *American Journal of Sociology*, 108(1), 55–101.
19. Wood, A. J., Graham, M., Lehdonvirta, V., & Hjorth, I. (2019). Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society*, 33(1), 56–75.