Beliefs, Reputation, and Barriers to Entry in Online Labor Markets*

Hamna Ahmed Zunia Tirmazee Ruoxuan Wu [JMP]

November 14, 2025

(*Updated regularly - please click here for latest version*)

Abstract

Hundreds of millions of workers in developing countries seek digital jobs online but face entry barriers without an established reputation. Novice workers could offset this by lowering initial wages, yet few do. Our baseline survey points to two explanations: workers believe employers interpret low wages as low quality signals and are uncertain about their own abilities. We conduct two field experiments on a leading global freelancing platform to examine how these beliefs shape worker outcomes. In the demand-side experiment, we randomize wage offers by novice workers to 703 jobs and find employers respond favorably to lower wage offers from novices, contrary to what workers believe. In the supply-side experiment with 481 novice workers, we randomly provide them with accurate information about employer responses and their performance. Correcting workers' beliefs increases their willingness to lower wages. Consistent with reputation models, effects are driven by high-ability novices with high returns to reputation once these frictions are removed. Simulations show that without external intervention, worker learning about employer responses is slow and costly. Our findings highlight that information interventions can help workers in developing countries overcome reputation barriers and accelerate talent discovery in online labor markets.

*Ahmed: Lahore School of Economics, Tirmazee: Lahore School of Economics, Wu (corresponding author): University of Chicago, wur@uchicago.edu. We are very grateful for the continuous guidance and encouragement from Marianne Bertrand, Christina Brown, and Michael Dinerstein. We thank Fiona Burlig, Joshua Dean, Jeanne Sorin, Ed Jee, Erin Kelley, Michael Kremer, Xianglong Kong, Gregory Lane, Camilla Schneier, Haruka Uchida, Emma Zhang, and participants at Dev Tea and Public Labor Lunch at UChicago for their support and helpful feedback. We thank Sarah Asif, Seemab Sajid, Khawaja Dawood Tariq, especially Sania Zeb and Yijing Zhang, for their excellent research assistance, and the outstanding data collection team from the Center for Public Opinion Research for their hard work. We gratefully acknowledge funding from IGC, the Development Economics Research Fund by Becker Friedman Institute and the Development Innovation Lab, the Weiss Fund, and the Roman Family Center for Decision Research. The research protocols are approved by UChicago IRB and the experiments were pre-registered in the AEA RCT registry (AEARCTR-0015840). All errors are our own.

1 Introduction

Inexperienced workers often struggle more than experienced ones to find jobs, particularly in low- and middle-income countries (LMICs). In some regions, they can be three times more likely to be unemployed than the general population (Alfonsi et al., 2020). A key challenge is their inability to credibly signal their skills to employers, particularly when their qualifications are difficult to evaluate or verify (Abebe et al., 2021; Bertrand et al., 2024). This problem is especially pronounced in the rapidly expanding online labor market, where hundreds of millions of workers from these countries seek digital job opportunities on global freelancing platforms (Datta et al., 2023). Although these platforms expand access to opportunities, information frictions about worker ability create substantial entry barriers for inexperienced LMIC workers applying to foreign employers (Agrawal, Lacetera and Lyons, 2016; Hilbert and Lu, 2020). As employers cannot observe worker ability prior to hiring, they favor those with established reputations, creating a "reputation trap" for inexperienced workers: they need jobs to build reputation but need reputation to secure jobs. This trap limits market entry for potentially qualified workers.

Classic reputation models by Tirole (1988) and Stiglitz (1989) suggest that inexperienced workers can overcome this trap by offering low wages initially to attract employers and demonstrate their skills. Yet, in practice, inexperienced workers rarely do so. This pattern may reflect workers' beliefs about demand responses or their own abilities. Workers may hesitate to lower wages if they believe employers interpret low prices as signals of low quality, a pattern documented in product markets where consumers infer lower quality from lower prices when evaluating unfamiliar goods.[1] Alternatively, workers may refrain from offering low wages if they lack accurate information about their ability to perform well in a global market. This uncertainty can distort their perceived returns to this reputation investment.

This paper investigates whether beliefs about demand responses and their own ability deter inexperienced workers from lowering wages to secure jobs in online labor markets. We conduct two complementary field experiments on a leading global freelancing platform, guided by the reputation model of Tirole (1988). On the demand side, we randomize wage offers to test whether employers' responses align with treating low wages as signals of low quality. On the supply side, we randomly provide inexperienced workers with information about either demand responses or their own relative ability, and measure the effect on decisions to offer low wages. Identifying which side of the market holds inaccurate beliefs is crucial for designing interventions to reduce entry barriers in online labor markets.

We study data entry freelancers from LMICs on one of the largest global freelancing plat-

[1] See Rao and Monroe, 1989; Wathieu and Bertini, 2007; Erdem, Keane and Sun, 2008; Ashraf, Jack and Kamenica, 2013.

forms. Data entry is a popular service category among these workers as it requires less advanced technical skills than other digital services (Datta et al., 2023). On the platform, employers post jobs with a stated budget, and freelancers apply by submitting wage bids and paying an application fee. We define *wage undercutting* as bidding below the posted budget. When evaluating applicants, employers observe each freelancer's wage bid and reputation, represented by the number of completed jobs and job success rate on the platform. Prior research shows that reputation from initial jobs has heterogeneous effects: public reviews by employers increase hiring chances and future earnings for high-performing workers but prevent low-performing ones from securing further work (Pallais, 2014; Fazio, Freund and Novella, 2025). Given that high-ability novices can benefit from building reputation, one might expect them to undercut wages to secure initial jobs. Yet, analyzing 2,332 applications for data entry jobs, we find that only 17% of inexperienced applicants do so, despite competing against both experienced workers and fellow novices.

To shed light on whether this pattern stems from homogeneous worker abilities or beliefs about returns to undercutting, we hire 481 inexperienced freelancers from 37 LMICs for a standardized data entry task on the platform. We find substantial variation in performance, yet workers possess limited knowledge of their abilities relative to other novices. Moreover, while workers identify high competition as their primary challenge on the platform, they worry that low wage bids signal poor quality to employers.

We extend the model from Tirole (1988) to formalize frictions in reputation acquisition and derive testable predictions. In the original model, a separating equilibrium can emerge where high-ability novices offer low wages initially to build a good reputation for repeated sales, while low-ability novices do not. Low wage offers by high-ability novices both increase their chances of being hired and serve as credible quality signals to employers, thereby accelerating talent discovery and improving match quality relative to a pooling equilibrium. We show that the separating equilibrium breaks down under two frictions: employers interpret low wages as negative quality signals and novices possess imperfect information about their own abilities. Causal identification using observational data is challenging because wage bids, beliefs, and reputation are jointly determined. We address this through field experiments that generate exogenous variation in wages and beliefs to test the model's predictions.

Our first experiment examines whether employers interpret low wages as poor quality signals for novices in the absence of other credible signals. If so, demand should respond less to wage undercutting by novices than by veterans with established reputations, a prediction we test directly. To create exogenous variation in wage bids and reputation, we create four freelancer profiles located in LMICs specializing in data entry. Half are novices with no job history; the other half are veterans for whom we build platform histories. We compile relevant job postings on the platform and submit four applications to each. This yields a sample of 2,812 applications sent to 703 jobs. Within profile type, we randomly assign one application to bid at 80% of the budget and the other at the full budget, holding all other application characteristics constant.

We find results contrary to workers' pessimistic beliefs about demand responses to wage undercutting. Although novices are much less likely to receive employer responses than veterans, low wage bids significantly improve outcomes for both groups: they increase the likelihood that employers read the application and callback the candidate. Crucially, employers are more responsive to low wage bids by novices than by veterans: novice callback rates increase by 54% from 3.41 percentage points, while veteran callback rates increase by only 16% from 5.26 percentage points. This differential response is consistent with employers interpreting low bids from novices as a willingness to invest in reputation rather than as poor quality signals. Moreover, wage undercutting does not sort workers into exploitative contracts. High-rated employers respond more favorably to low-wage offers from novices than low-rated employers do. Lastly, we quantify novices' returns to undercutting through back-of-the-envelope calculations. If a novice consistently undercuts, they should expect to earn twice as much on the platform as those who bid at the full budget within their first year. These results suggest that undercutting is an effective and profitable strategy.

Our second experiment examines on the supply side whether correcting workers' beliefs about either demand responses or their own ability affects their wage bidding behavior. We conduct this experiment with the 481 novice freelancers hired for the standardized data entry task. After workers complete the initial task, we evaluate their output and cross-randomize two treatments at the individual level, stratified by baseline performance. The first treatment, the *feedback treatment*, provides half of the workers with private information about their task performance relative to other novices in the sample. A few days later, we post a higher-value data entry job from another employer account and refer it to all experimental participants by sharing the job link. The second treatment, the *employer evaluation info treatment* (hereafter, *employer info treatment*), adds one sentence for randomly selected workers, stating that the employer will not judge worker quality based on wage bids. Baseline and endline survey data confirm treatment validity at correcting workers' beliefs about their performance and alleviating concerns with wage undercutting. We measure treatment effects on workers' application decisions and wage bids for the referred job. Finally, to measure downstream outcomes, we randomly select a small set of applicants to receive job offers and track acceptance rates and task performance.

The sample represents a substantial underemployed labor force with limited local job opportunities and significant barriers to online market entry. Nearly half report no offline work, and more than 80% have no completed jobs on this or any other freelancing platform. For these workers, overcoming initial entry barriers and building a reputation to secure digital jobs is especially valuable.

Our experimental results show that providing information about demand responses and workers' ability induces novices to bid low wages but through different mechanisms. In the control group, only 5.5% of workers bid below the job budget, with no difference between high- and low-performers. The employer info treatment increased this proportion to 20.5%, nearly quadrupling the control group mean. This indicates that workers' pessimistic beliefs about employer responses strongly discourage wage undercutting. Moreover, treatment effects are homogeneous across performance types, consistent with the pattern in our baseline survey that most workers believe they are high-performers.

By contrast, the feedback treatment increased the proportion of wage-undercutting workers to 14.4%. The effect is driven solely by high-performers who held optimistic beliefs about returns to undercutting for high-ability workers at baseline. This suggests that high-ability workers understand which strategies work for their type but are uncertain about their own standing. Receiving validation of their ability prompts them to adopt the undercutting strategy. When workers receive both employer info and performance feedback, we observe separation in wage bids: high-performers are five times more likely to undercut than those in the control group, while low-performers are only twice as likely. Taken together, workers' misbeliefs about employer responses and uncertainty about their own abilities both deter them from adopting wage undercutting to acquire reputation. Information provision addresses these frictions through different channels. The employer info treatment works by updating novices' pessimistic beliefs about how employers interpret low wages, while the feedback treatment works by resolving workers' uncertainty about their ability.

The demand- and supply-side experiments show that inaccurate beliefs prevent novice workers from adopting wage undercutting to overcome entry barriers. A natural question is whether workers could eventually correct these beliefs on their own through experience in the market. Using estimates from both experiments, we simulate Bayesian learning to assess how long it would take a perfectly rational worker to learn the true demand curve through trial and error. Even under systematic experimentation and perfect updating, learning is slow and costly: a median worker with moderately pessimistic beliefs would require roughly 200 job applications and incur about $270 in application fees, equivalent to one-quarter of annual income in low-income countries. For novice workers from our settings, these costs make it difficult to adapt bidding strategies purely through experience. This highlights the value of low-cost external interventions, such as guidance on employer hiring behavior, that can correct misperceptions and expand access to digital work.

This paper bridges three fundamental issues that constrain young jobseekers in LMICs: workers' limited ability to signal their skills, their lack of information about the labor market and their own abilities, and the central role of initial reputation in securing jobs - particularly in online labor markets. We study these frictions in the increasingly important context of global freelancing platforms. Employers on these platforms rely heavily on workers' job histories to infer quality (Pallais, 2014; Agrawal, Lacetera and Lyons, 2016; Barach and Horton, 2019), and alternative signals available to novice workers, such as micro-credentials or advertising, have limited effectiveness (Kässi and Lehdonvirta, 2022; Hannane, 2024). Building on classic reputation models (Tirole, 1988; Stiglitz, 1989), we provide causal evidence that novice workers can compensate for missing reputation by initially lowering wages. However, they under-adopt this strategy because of the belief frictions documented above. While prior work emphasizes upskilling freelancers in LMICs and finds limited effects on platform success (Baptista, Freund and Novella, 2023; Das et al., 2024; Fazio, Freund and Novella, 2025), our results show that when workers cannot credibly signal their skills, correcting belief frictions is essential for adopting strategies that overcome entry barriers.

Second, we contribute to growing body of work documenting jobseekers' inaccurate beliefs as a barrier to employment (see Caria et al., 2024 for a review). These beliefs are central because they shape job search strategies and, in our setting, determine whether workers are willing to sacrifice current earnings to build reputation. Prior work emphasizes two types of misbeliefs. The first concerns overly optimistic expectations about wages (Banerjee and Sequeira, 2023; Alfonsi, Namubiru and Spaziani, 2024; Kelley, Ksoll and Magruder, 2024) or hiring probabilities (Chakravorty et al., 2024; Abebe et al., 2025; Bandiera et al., 2025), which can lead jobseekers to reject viable opportunities. We identify a new mechanism: misbeliefs about how employers interpret low wages as signals of worker quality. We experimentally show that even though novices are willing to work for less to prove themselves, pessimistic beliefs about employer responses deter them from adopting effective wage bidding strategies. The second concerns workers' imperfect information about their own abilities (Bassi and Nansamba, 2022; Carranza et al., 2022; Kiss et al., 2023). Correcting these misperceptions helps jobseekers target roles aligned with their strengths. We extend this insight by embedding it in a reputation model and show that performance feedback induces high-ability novices to lower wages to build reputation, generating separation in bidding behavior and potentially accelerating talent discovery on the platform.

Finally, we contribute causal evidence to the theoretical literature on prices as signals of quality, from both sides of a labor market. Theory predicts that market conditions determine whether sellers signal high quality through *high* prices (Wolinsky, 1983; Milgrom and Roberts, 1986; Judd and Riordan, 1994) or *low* prices (Tirole, 1988; Stiglitz, 1989). Empirical work has focused almost entirely on product markets. Empirical tests of these predictions in labor markets are scarce because worker types and beliefs are inherently unobservable. Roussille (2024) shows suggestive evidence that employers interpret higher wage requests as signals of higher ability on a US online labor platform. Closest to our work, Huang, Li and Zuo (2025) examine how novice domestic-service workers in China use prices to signal quality, relying on initial worker characteristics as proxies for unobserved types. We advance this literature by collecting novel, directly measured data on workers' true abilities and their beliefs, and by experimentally varying both wage bids and information. This approach allows us to show that novice workers and employers hold systematically misaligned beliefs about how wages signal quality, and that these belief gaps are costly to correct in a market characterized by high search costs and intense competition.

The remaining sections are organized as follows. Section 2 describes the online freelancing platform setting and presents stylized facts. Section 3 develops a conceptual framework following Tirole (1988) and derives testable predictions for our experiments. Section 4 presents the demand-side experiment design and results, and estimates returns to wage undercutting. Section 5 presents the supply-side experiment, results, and underlying mechanisms. Section 6 simulates worker belief updating, discusses policy implications, and examines the generalizability of our findings. Section 7 concludes.

2 Setting

Online freelancing markets have become an important component of the global workforce. Recent estimates suggest that between 154 and 435 million people engage in online gig work, with the majority based in lower-income countries (Datta et al., 2023). Workers connect with employers on online freelancing platforms to perform and deliver one-time tasks remotely. In this section, we describe the job matching process on the online platform and present stylized facts that motivate the experiments.

2.1 Job Matching in Online Labor Markets

We focus on one of the world's largest online freelancing platforms, which shares key features with other major platforms. On this platform, employers post job listings with a task description, required skill level, expected duration, and budget set as either an hourly rate or fixed payment. Workers can apply directly to postings or receive invitations from employers. For new workers, direct applications are the most common way to find jobs. To apply, workers submit a wage bid and cover letter, and pay an application fee (typically between \$1.20 and \$1.65 per application). Wage bids can be above or below the posted budget as long as they meet minimum requirements (\$5 for fixed-payment projects, \$3 per hour for hourly projects). Bids are visible only to employers, not to other applicants.

Employers browse applicants through a dashboard displaying basic information: profile title, platform job history, self-reported skills, wage bid, and cover letter preview. As shown in Appendix Figure A1, job history and wage bid are the most salient application features. Employers can click individual applications to review complete cover letters and full profiles, which typically include self-reported credentials and detailed work history. They contact workers through the platform's messaging system to initiate contracts or request additional information. While wage negotiation is possible at this stage, it rarely occurs (Barach and Horton, 2019). When a contract is signed, the job automatically appears on the worker's profile. After completion, both parties publicly rate each other on a one-to-five scale.

We collect primary data and conduct two experiments on this platform, focusing on data entry jobs with fixed payment. This choice reflects two considerations. First, data entry allows us to generate accurate and objective skill measurement, unlike more subjective tasks such as software development or graphic design. It is also accessible to workers from LMICs, particularly new workers, as it requires less advanced skills and represents a popular entry point to the platform. Second, fixed-payment contracts dominate initial employment for novices: 74% of data entry workers in our data receive fixed-payment contracts in their first job. These characteristics make fixed-payment data entry jobs a relevant setting for examining how novice workers from LMICs enter online labor markets.

2.2 Stylized Facts from Platform Data

We analyze comprehensive platform records on data entry freelancers to document stylized facts about competition and wage bidding in this market. Our analysis draws on three complementary data sources from the platform.

First, we scraped profile characteristics for all data entry workers based in Bangladesh, India, and Pakistan in May 2024, as these three countries account for over 50% of the global freelancer population (Online Labour Observatory, 2025). For 10,818 freelancers, we obtained the basic information visible to employers: number of completed jobs, total earnings, total hours worked, job success rate, preferred hourly rate, and agency association. Second, we complement the profile data by compiling job-level records for all 4,887 data entry workers in Pakistan, who represent over 45% of freelancers across the three countries. These records include job title, duration, worker rating, earnings, and contract type for all jobs completed through May 2024. Third, as part of the freelancer experiment, we gathered job application data from 2,332 workers outside our experimental sample. The application data captures their wage bids, cover letters, and basic profile characteristics.

We highlight three key patterns. First, novice workers face significant entry barriers. Among data entry workers, 44% have not completed any jobs and 62% have completed two or fewer. This aligns with broader evidence: International Labour Organization (2021) reports that 27-77% of freelancers across five major platforms have never completed a job. Similarly, Hilbert and Lu (2020) estimates only 10% of registered freelancers in Latin America have ever secured platform work, with jobs highly concentrated among established workers.

Second, freelancers' job prospects improve with their number of completed jobs. Appendix Figure A2 plots wait time until the next job and log earnings against the number of cumulative jobs completed. Those with more experience experience shorter gaps between jobs and faster earnings growth. While these patterns likely reflect selection on worker ability, they illustrate why securing initial jobs may be particularly valuable for high-ability novices seeking to build their reputations.

Third, despite intense competition and returns to initial employment, wage bids show little variation. Application data shows that each vacancy receives an average of 64 applications, 20% from experienced workers with at least three completed jobs. Yet the vast majority of workers bid exactly at the posted budget as shown in Appendix Figure A3). Only 15% of freelancers propose wages below the budget and 3% above the budget. Among inexperienced freelancers with fewer than three jobs, only 17% bid below the budget. This is surprising given that high-ability novices compete at the same wage as experienced workers, while they benefit from landing initial jobs and establishing reputation. In the next subsection, we dive into potential explanations of the lack of variations in wage bids.

2.3 Performance and Beliefs of Novice Freelancers

The lack of wage variation could reflect workers having similar abilities or holding similar beliefs about returns to undercutting.[2] Testing these hypotheses with platform data alone is difficult since ability and beliefs are not directly observable. To address this limitation, we collect performance measures and belief data through a standardized task and baseline survey. Below we describe our sample selection, measures of performance and belief, and descriptive findings.

Recruitment and Sample Selection: We recruited 481 novice freelancers by posting a beginner-level data entry and online research job on the platform over eight rounds between May and August 2025.[3] To target the relevant population, we restricted offers to applicants with no more than one completed job, located in LMICs, and specialized in data entry. Within each round, we selected eligible applicants on a first-come-first-served basis until all slots were filled.[4] We describe detailed sample characteristics in Section 5.1.2.

Hired freelancers received detailed instructions for the data entry task, which involved cor-

[2] An alternative explanation is that workers avoid bidding down wages due to threats of social sanction. This is unlikely in our setting. Prior literature documenting such effects focuses on village economies where workers are in close networks. In online marketplaces, freelancers operate independently and their action is not observed by other workers. Moreover, in our pilot, few workers expressed concerns about bidding low wages being unfair to others.

[3] We conduct the hiring in batches because hiring all workers simultaneously would be logistically challenging and delay performance evaluation across the sample.

[4] To minimize attrition and reduce the risk of scams, we excluded freelancers who submitted unprofessional cover letters or proposed unreasonable wage bids.

recting existing entries against the original sources and adding new data from scanned PDF files into an Excel Spreadsheet within 24 hours.[5] To incentivize effort, we announced performance bonuses for the three top-performing freelancers. Upon task completion, workers participated in a baseline survey that collected information such as demographics, beliefs about their task performance relative to other novices in the same round, and perceptions of the effectiveness of various application strategies.

Performance Measure: We evaluate worker performance on the data entry task along three objective dimensions: accuracy, speed, and following instructions.[6] Accuracy was measured as the error rate, determined by comparing worker output against verified entries prepared by professional data entry specialists before the experiment. Speed measured the time between task release and final submission. Instruction compliance was scored from one to six, with each point reflecting adherence to a specific formatting requirement. To create comparable performance measures across workers, we ranked all freelancers within each round on each dimension and assigned quartile scores from 1 (lowest) to 4 (highest). The total performance score equals the sum of these three quartile scores, ranging from 3 to 12.[7] Workers scoring above the median score within their round were classified as high-type, and others as low-type.

Belief Measure: We elicited workers self-assessments of their performance relative to other novices in the same round through the baseline survey after task completion. Workers were told that their performance would be evaluated on the three dimensions described above and asked to predict their performance quartile for each dimension, with financial incentives for correct guesses. The perceived performance score sums the predicted quartiles across the three dimensions, also ranging from 3 to 12. This design allows direct comparison between workers' perceived and actual performance. Additionally, the survey elicited workers' beliefs about the effectiveness of different application strategies for novices to secure jobs, such as cover letter writing, paying the platform to boost their application, and bidding below the job budget.

Descriptive Findings: First, we find substantial variation in task performance among novice workers, yet they have limited knowledge of their own ability. Panel A of Figure 1 shows the distribution of performance scores for all 481 workers. Panel B compares actual and perceived performance, revealing that the vast majority misperceive their performance

[5]Although workers were allowed to use OCR or AI tools, manual work was necessary since the scanned PDFs were of poor quality and the spreadsheet required a specific format.

[6]To ensure these measures capture relevant dimensions of quality, we conducted extensive interviews with freelancers prior to the experiment and confirmed these three aspects were considered the most important performance indicators. In addition, this approach is similar to the performance measurement in Pallais (2014) and Kala and Lyons (2025).

[7]When freelancers had the same score, we broke ties by sequentially comparing their accuracy, ability to follow instructions, and speed.

quartiles.[8] This pattern suggests that limited wage variation is more likely to stem from workers' uncertainty about their own ability rather than from homogeneity in actual performance.

Figure 1: Worker Performance

Panel A. Distribution of Performance Scores **Panel B.** Actual versus Perceived Performance

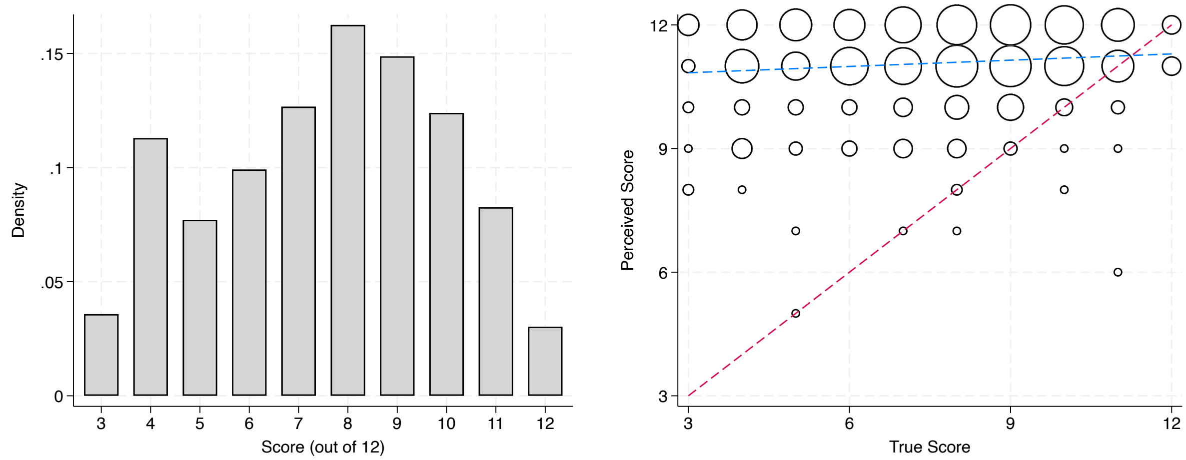

Notes: The figures show the distribution of the performance scores (between 3 and 12) and the correlation between the actual performance scores and workers' perceived scores from 481 novice freelancers. Panel A shows the histogram of the true performance scores. Panel B plots the true scores on the x-axis and the perceived scores on the y-axis. The red dotted line is the 45-degree line and the blue line shows the best linear fit. The marker size is relative to the number of observations.

Second, novice freelancers express concerns that low wage offers signal low service quality to employers. As shown in Figure 2, 44% of freelancers agreed with the statement in the baseline survey: *"If a new freelancer offers a lower price, the client may think that their work is low quality."* This concern is consistent across worker types by baseline performance, suggesting that pessimistic beliefs about employer responses may discourage novices from proposing lower wages to avoid sending negative quality signals.

These findings point to two mechanisms explaining the lack of wage undercutting: workers' uncertainty about their own ability and pessimistic beliefs about how employers interpret low wages. Next, we extend the classic reputation model to incorporate these mechanisms and derive testable predictions that guide our experimental design.

3 Conceptual Framework

We present a simple conceptual framework motivated by Tirole (1988) to formalize how novice workers may overcome entry barriers in online labor markets through wage undercutting and to examine implications for aggregate employment outcomes. We show how the

[8]Appendix Figures A4 and A5 present the distribution and the comparison between actual and perceived scores for each dimension.

Figure 2: Novices' Beliefs of Employers Interpreting Low Bids = Low Quality

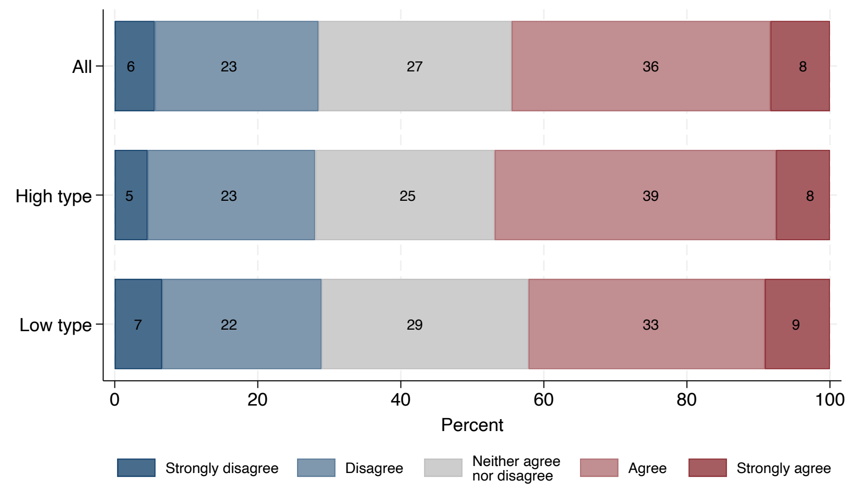

Notes: The figure displays novice workers' opinion on the statement: *"If a new freelancer offers a lower price, the client may think that their work is low quality."* The bars from top to bottom represent the entire sample, "high-type" novices with baseline performance above median, and "low-type" novices with baseline performance below median.

market equilibrium deviates from the standard model when employers interpret low wages as signals of low quality and when novice workers have imperfect information about their ability. The framework generates distinct testable predictions under each friction, which directly guide our experimental design.

3.1 Setup

Consider a labor market with a unit mass of workers and employers. Workers, indexed by i, live for two periods, while employers, indexed by j, live for one period. Workers hired in the first period become *veterans* in the second period, while those in their first period or not hired previously are referred to as *novices*. At the end of each period, all employers and second-period workers exit the market and are replaced by a new generation.

At the beginning of the period, each employer posts an identical job on the online platform with an exogenous budget $\bar{w}_j$. The budget $\bar{w}_j$ represents the maximum amount the employer is willing to pay for an online freelancer, based on the cost of hiring a local worker offline.

Workers differ in their innate ability $\theta_i \in \{H, L\}$ to execute the task. We assume that workers know their own type upon entering the market. Workers decide whether to apply to a job by submitting a wage bid $w_{ij} \leq \bar{w}_j$ or not applying at all. [9] If hired, the worker is paid w_{ij}; otherwise, the worker receives the outside option $c_\theta = \{c_H, c_L\}$.

[9] This is empirically supported: fewer than 3% freelancers in our application data proposed wages higher than the budget.

Employers derive value $v_\theta = \{v_H, v_L\}$ from a worker's output but cannot observe the worker type prior to hiring. Instead, they observe proposed wages for all applicants and public reviews r_i for veteran workers. For veterans, employers infer their type from their review history. For novices, employers form beliefs about their type based on their proposed wages w_{ij}: a novice is high type with probability $\gamma(w_{ij})$ and low type with probability $1-\gamma(w_{ij})$. If all novices propose the same wage, or if wages are uninformative about worker type, employers may screen candidates through interview at costs s_j and obtain additional quality cues λ_{ij}. The employer's objective is to maximize the expected return from hiring:

$$\max_i \{E\left[v_i | r_i, w_{ij}, \lambda_{ij} \times 1[\text{screening}]\right] - w_{ij} - s_j \times 1[\text{screening}], \bar{v}_j\}$$

where $\bar{v}_j$ is the employer's outside option. Employers hire from the platform only if expected returns exceed $\bar{v}_j$.

Once a task is completed, employers observe the output and leave a public rating on the platform that reveals the worker's type. Following a bad review, low-type veterans will not receive repeated offers in the next period. High-type veterans, in contrast, will be hired again at $\bar{w}_j \in \{c_H, v_H\}$. The worker's problem is therefore to set wages in the novice period to maximize expected lifetime earnings on the platform:

$$\max_{w_{ij}} (w_{ij} - c_i) + \beta(\bar{w}_j - c_i) \times 1[\theta_i = H]$$

The timeline of events within each period is as follows:

(i) A new generation of workers and employers joins the platform. Each employer posts a job with budget $\bar{w}_j$.

(ii) Workers observe $\bar{w}_j$ and decide whether to apply by proposing a wage w_{ij}.

(iii) Employers observe the proposed wage w_{ij} and public reviews r_i if they exist, and possibly incur s_j to screen inexperienced candidates. They make at most one offer to the candidate that generates the highest return.

(iv) Hired workers produce the output and their type is revealed.

(v) All employers and workers in their second period exit the market.

The model rests on assumptions that are consistent with empirical evidence from online labor markets. First, the assumption that employers live for only one period rules out incentives on the demand side to invest in worker training or build long-term relationships. This is realistic in our setting: the average contract length among novices in our data is 6.4 hours, and most freelancers work with unique employers (Pallais, 2014). Second, we assume that low-type workers do not improve their productivity through accumulating platform experience because once their types are revealed, they will not be able to acquire more jobs on the platform. Fazio, Freund and Novella (2025) find suggestive evidence that low ratings on initial jobs substantially reduce workers' chances of securing further work. Unlike traditional labor markets where inexperienced workers may accept low initial wages to learn on the job through internships or apprenticeships, online platforms provide limited scope for skill accumulation after a bad initial signal. As a result, only high-type novices have strong incentives to offer low wages in their first period in order to establish a reputation and secure future employment.

3.2 Market Equilibrium

We examine labor market outcomes under the three scenarios: (i) the benchmark case as in Tirole (1988), (ii) employers interpreting low wage offers as signals for low quality for novices, (iii) novices having incomplete information about own ability. The experimental treatments are designed to shift employers and workers across these scenarios.

Proposition 1. *In the benchmark case, a perfect Bayesian equilibrium exists when returns to reputation for high-types exceed the difference in outside option between types:* $\beta(\bar{w}_j - c_H) \geq c_H - c_L$*. High-type novices propose wages below the job budget at* $w_{Hj} = c_L < v_H$ *in period 1, while low-types propose* $\bar{w}$*. Employers interpret low wages as credible quality signals and hire all high-type novices and veterans in equilibrium.*

Proof. See Appendix B.

The benchmark case follows the classic reputation model in Tirole (1988). High-type novices offer low prices initially to build reputation and generate repeated sales, while low-types maximize earnings in period 1 as they won't get repeated sales once hired. To prevent low-types from mimicking their strategy, high-types set wage bids equal to the low-type's outside option c_L, making undercutting unprofitable for low-types. This separating equilibrium can sustain when future reputation returns exceed the current earnings sacrifice. This is likely to hold in our setting given high returns to reputation and small differences in outside options between types. Given these bidding strategies, employers interpret low wage bids as credible quality signals without incurring screening costs, making them more likely to hire undercutting novices. As a result, high-type workers overcome reputation barriers. In equilibrium, all high-type novices are hired and deliver high quality services.

Proposition 2. *When employers interpret low wage offers by novices as signals of low quality, labor demand is more responsive to undercutting by veterans than by novices. If quality concerns dominate cost considerations in hiring decisions, novice workers pool at the job budget* $\bar{w}_j < \lambda_{ij} v_H + (1 - \lambda_{ij}) v_L - s_j$ *and high-type veterans with* $\bar{w}_j \in \{c_H, v_H\}$ *are* *hired. Compared to the benchmark case, hiring novices becomes more costly and fewer are employed.*

Proof. See Appendix B.

When employers interpret low wages as low quality signals, wage bids by novices affect demand through both cost and perceived quality. If cost considerations dominate, high-type novices can still secure employment through low wage offers, which allows employers to update their beliefs over time. However, when quality concerns dominate hiring decisions, novices pool their wage bids at the job budget to avoid being perceived as low quality. This pooling eliminates the credible quality signal that low wage bids provide in the benchmark case. Employers must then incur additional screening costs to distinguish among novices, and with identical bids, cannot update their beliefs about the relationship between wages and quality. In equilibrium, fewer (high-type) novices are hired compared to the benchmark case, slowing talent discovery and reducing match quality.

Proposition 3. *When workers have limited information on own type, there is no complete separation in proposed wages by unobserved type. Employers rely on costly screening to distinguish among novice workers. In equilibrium, novices proposing $w_{ij} < \lambda_{ij} v_H + (1 - \lambda_{ij}) v_L - s_j$ and high-type veterans are hired. Compared to the benchmark case, fewer novices are employed and average output quality is lower.*

Proof. See Appendix B.

When workers have limited information about their own type, wage bids reflect workers' expected ability rather than actual ability. This breaks the credible signaling mechanism in the benchmark case: employers cannot infer quality from wages alone and must incur costly screening to distinguish among novices. While workers may still compensate employers by lowering wage offers, fewer novices are hired compared to the benchmark case and average output quality is lower.

Across these scenarios, the second case yields the worst outcomes: when employers interpret low wages as quality signals and quality concerns dominate, no novices undercut and employers must incur screening costs, resulting in the fewest novice hires. In both non-benchmark cases, output quality falls and high-type novices fail to signal their quality through wage undercutting, which slows talent discovery and prevents them from overcoming entry barriers caused by lack of reputation.

3.3 Comparative Statics in Wage Bidding and Hiring

To identify the source of frictions, we derive the following testable predictions from the framework above and directly map them to our experimental design.

Employer Side:

- Prediction 1: If employers interpret low wage bids from novices as signals of low quality, labor demand will be more responsive to undercutting by veterans than by novices, holding other worker characteristics constant. Otherwise, wage undercutting improves hiring prospects for both groups equally, or more so for novices.

Worker Side:

- Prediction 2: If novice workers have limited information about their own ability and hold pessimistic beliefs about employer responses to undercutting, there will be pooling in wage bids and wages do not reflect worker ability.
- Prediction 3: If novice workers do not hold pessimistic beliefs about employer responses to undercutting, providing them with information about their own ability induces separation in proposed wages: high-type novices engage in wage undercutting.
- Prediction 4a: Alleviating workers' pessimistic beliefs that employers interpret low wages as quality signals encourages wage undercutting by novices.
- Prediction 4b: In addition, if novices have complete information about their types, there will be separation in proposed wages: high-type novices engage in wage undercutting.

For the remainder of the paper, we test Prediction 1 using the employer experiment, which elicits employers' demand responses to wage undercutting by novice and veteran workers, and we test Predictions 2-4 using the freelancer experiment, which examines how information about employer responses and workers' own ability affects their bidding decisions.

4 Employer Experiment

4.1 Experimental Design

Building on the conceptual framework, we conduct an audit-style experiment on the same freelancing platform to test whether employers interpret low wages as signals of low quality for novice workers and whether wage undercutting improves their hiring prospects. A key empirical challenge is that hiring decisions depend on wage bids and other application characteristics, making it difficult to isolate the causal effect of wages on employer responses.

Moreover, employers' perceptions of worker quality are not directly observable. Our experimental design addresses these challenges. We submit fictitious job applications with randomly assigned wage bids while holding all other application attributes constant, and compare employer responses to wage bids by novice and veteran applicants. Our design exploits a key theoretical insight: for veteran workers with established reputations, employers do not need to infer quality from their wage bids. Therefore, any differential response to low wage bids between novice and veteran applicants reflects whether employers use wage bids as negative signals of quality when worker ability is uncertain.

4.1.1 Implementation

Freelancer profile setup: We created four freelancer profiles located in Pakistan, specializing in data entry and virtual assistance.[10] Profile characteristics, such as preferred hourly wages, self-reported skills, and language proficiency, were selected to match typical profiles in our descriptive data. We randomly designated two profiles as novices with no job history, and two as veterans for which we built job histories between February and April 2025. All veteran profiles had the same number of completed jobs, job-specific ratings, and overall success rates to hold observable quality signals constant.

Job listing and application: We compiled daily job listings from the platform between May and August 2025. Eligible jobs were data entry or virtual assistant positions with a fixed budget of at least \$7 USD,[11] requiring entry- or intermediate-level skills. We submitted four applications per eligible job, each with a wage bid and cover letter.[12] Within each profile type, one application was randomly assigned to undercut by bidding 80% of the posted job budget, while the other bid the full budget. We chose the 20% threshold because it corresponds to the most common range of wage undercutting observed in our job application data, as shown in Appendix Figure A6. This design generated four application types per job: novice bidding 80% of the budget, novice bidding the full budget, veteran bidding 80% of the budget, and veteran bidding the full budget. Appendix Figure A1 illustrates how these applications appeared on the employer dashboard.

We track two measures of employer responses using unique platform features. First, *employer view* captures when employers click on an application to review the cover letter and freelancer profile, which automatically generates a platform notification. Second, when employers contact applicants to request information or initiate an offer, the platform automatically

[10] Pakistan has the world's third largest online freelancer population and the highest proportion of data entry freelancers in our platform data.

[11] The platform sets a minimum price of \$5 USD for fixed-budget projects. We applied a \$7 USD threshold to ensure all submitted bids complied with this rule.

[12] To ensure the authenticity of the application and minimize suspicion by prospective employers, we created a unique cover letter for every application that tailored to job-specific requirements and information from the freelancer profile.

sends an email notification to the applicant, which we record as *callback*.

4.1.2 Data

Application characteristics and outcomes: We submitted 2,972 applications to 743 job vacancies on the freelancing platform. For each application, we recorded the profile type (novice or veteran), wage bids, cover letter, submission time, employer view, and callback. In some cases, an application deviated from its pre-assigned wage or could not be submitted because the vacancy had already been closed or removed. We restrict the analysis sample to vacancies where all four applications were submitted with correct wages, yielding 2,812 applications to 703 vacancies.

Job and employer characteristics: For each vacancy, we record basic job information, such as job title and budget, at the application time. Two weeks later, we track hiring outcomes and employer characteristics from the original job post, including the number of freelancers hired, employer location, tenure on the platform, number of past jobs posted, and employer ratings from previously hired freelancers. We use these characteristics to examine heterogeneity in the main results. Appendix Table A1 summarize these characteristics.

4.2 Empirical Strategy

We estimate the causal effect of low wage bids on the likelihood of receiving employer responses by worker experience. The main regression specification is

$$Y_{ij} = \alpha + \beta_1 \text{Novice}_i + \beta_2 \text{LowBid}_{ij} + \beta_3 \text{Novice}_i \times \text{LowBid}_{ij} + \mu_j + \varepsilon_{ij} \tag{1}$$

where Y_{ij} is the response to the application from profile i to job j, $Novice_i$ is an indicator equal to one if profile i has zero completed jobs, $LowBid_i$ is an indicator equal to one if the wage bid is 20% below the job budget, and μ_j is the job fixed effects to improve precision of our estimates. We cluster standard errors at the job level.

By comparing employer responses to low wage bids across worker experience, we can map Prediction 1 in Section 3.3 to the signs of coefficients. Specifically,

- If employers interpret lower wages as a negative signal of quality in absence of other credible signals, labor demand will be more responsive to low wage bids by veteran workers than novice ones: $\beta_2 > 0$, $\beta_3 < 0$.
- If not, low wage bids will improve the chance of getting hired for both novices and veterans, equally or more for novices: $\beta_2 > 0$, $\beta_3 \geq 0$.

Moreover, if the first statement holds, low wage bids will have competing effects on demand due to quality concerns and hiring costs. When quality concerns dominate, low bids will hurt novices' chance of getting hired: $\beta_2 + \beta_3 < 0$. Alternatively, if hiring costs dominate, novices can improve their hiring chances by bidding low wages: $\beta_2 + \beta_3 > 0, \beta_3 < 0$.

4.3 Impact of Lowering Bids by Worker Experience

4.3.1 Employer View and Callback

We find that low wage bids increase the likelihood of both employer views and callbacks for novices and veterans, consistent with employers not treating low bids as quality signals. Figure 3 presents the effects by applicant type from specification 1. The findings are robust to alternative specifications and sample restrictions as shown in Appendix Table A2.

Figure 3: Employer Responses to Low Bids

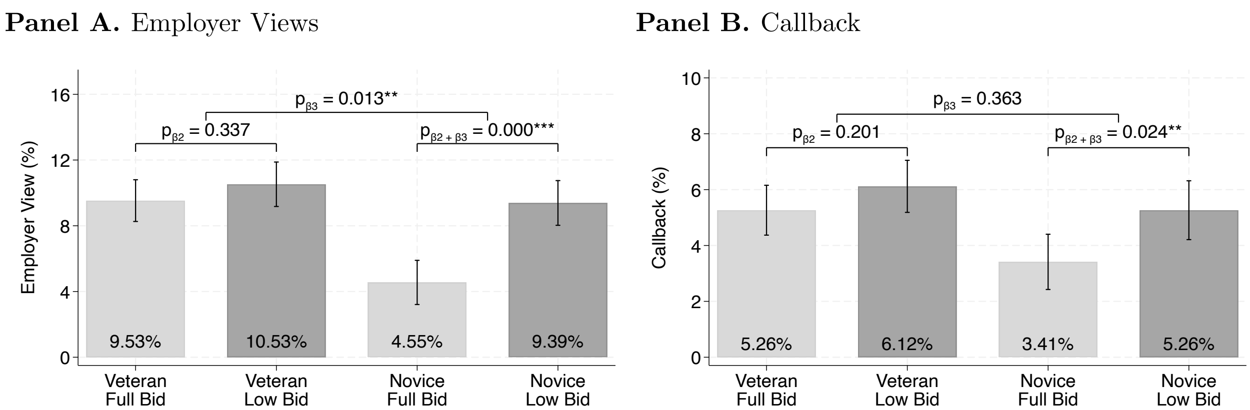

Notes: The figures show results from regression 1 for employer views and callback. Each bar represents the sum of the control mean (veteran bidding full budget) and the relevant regression coefficients. We report the 95% confidence intervals based on the estimated standard errors of the linear combinations of the regression coefficients and p-values from Wald tests.

Panel A reports treatment effects on employer views. When bidding at the job budget, veterans are more than twice as likely as novices to have their applications read. Bidding at 80% of the budget more than doubles the likelihood that employers read novices' applications (from 4.55 to 9.39 percentage points), while veterans see only a 10% increase (from 9.53 to 10.53 percentage points) Statistical tests on β_3 and $\beta_2 + \beta_3$ reject the hypothesis that employers interpret low bids from novices as negative quality signals ($p < 0.001$ and $p = 0.013$).

Panel B shows similar patterns for callbacks. Low wage bids raise callback by 1.85 percentage points for novices and 0.86 percentage points for veterans, corresponding to increases of 54 percent and 16 percent from the baseline values, respectively. Although the difference in treatment effects by worker experience is not statistically significant, low wage bids improve novices' likelihood of callback as $\beta_2 + \beta_3$ is significantly different from zero ($p = 0.024$).

These results confirm that employers do not interpret low wage bids as quality signals.

Low bids improve employer views and callback rates for both novices and veterans, with particularly strong effects for novices. This suggests employers view undercutting by novices as willingness to invest in reputation rather than as poor quality signals, contrary to novices' pessimistic beliefs.

4.3.2 Heterogeneity by Employer Characteristics

Although low wage bids improve hiring chances for novice workers, the mechanisms through which this strategy operates and whether it leads to better job outcomes remain unclear. To address this, we examine heterogeneity in employer responses by job and employer characteristics in Appendix Table A3.

First, we examine whether employers with higher screening costs rely more on wage bids as quality signals. Our conceptual framework implies that employers with limited ability to assess applicants should respond more strongly to low wage bids as a screening mechanism for novice workers. Consistent with this, we find that less experienced employers, measured by their platform tenure and number of jobs posted, are much less likely to consider novices bidding at the job budget but respond much more favorably to those offering low wage bids, compared to more experienced employers. This pattern suggests that employers with weaker screening capacity use low wage bids by novice workers as straightforward quality cues.

Second, we address concerns that exploitative employers may take advantage of inexperienced workers who offer low wages. Using ratings as a proxy for employer quality, we find that higher-rated employers are twice as likely to respond positively to low bids by novices than lower-rated counterparts, ruling out concerns about adverse selection of employers.

Taken together, our results demonstrate that lowering wage bids is an effective strategy for novice workers seeking to improve hiring chances without worsening job quality.

4.4 Returns to Lowering Wage Bids

Given the experimental results, we now quantify novices' returns to offering low wages. Using a back-of-the-envelope calculation, we compare novice workers' job offers and net earnings under two bidding strategies: always bidding at the full budget versus undercutting by 20%.

Suppose a novice submits a applications, paying a fee of c_a per application and receiving an offer with probability p. Because job durations are short, the novice is never capacity constrained and does not need to decline job offers. The expected earning from the first job is e, which grows at rate $\rho(n)$ as workers accumulate n completed jobs. Workers' expected number of jobs are $n(a) = a \times p$ and net earnings are $\Pi(a) = a \times (p \times e \times \rho(n) - c_a)$.

We use callback rates from the demand experiment to approximate p, median earnings from jobs where novices received callbacks to approximate e, and earnings growth rates from scraped job history data for data entry workers to estimate $\rho(n)$. Application costs come from in-sample job posts. Appendix Table A4 reports all parameter values and data sources. Figure 4 plots $n(a)$ and $\Pi(a)$ under the two bidding strategies, with discontinuities in $\Pi(a)$ reflecting earning growth as workers complete jobs on the platform.

Figure 4: Expected Number of Jobs and Net Earnings

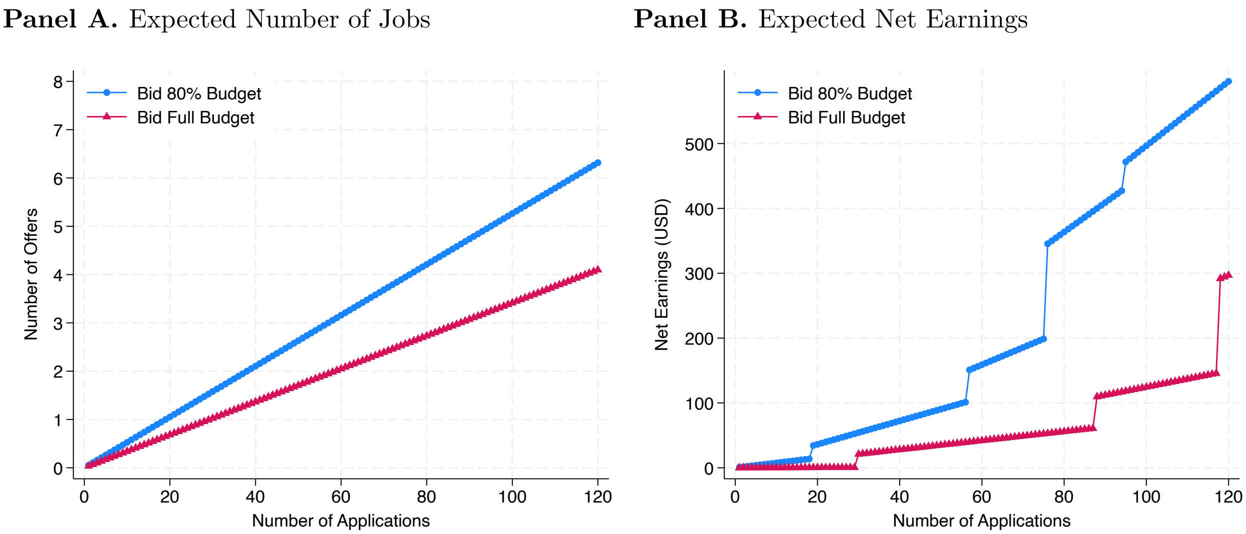

Notes: The figures show the relationship between number of applications and the expected number of jobs (in Panel A) and net earnings (in Panel B) under the two bidding strategies.

We find that undercutting dominates bidding at the full budget: novices secure jobs faster and earn more from experiencing faster earnings growth and saving application costs. To reach \$100 USD in net earnings, novices need 56 applications with consistent undercutting versus 88 applications at full budget. Given that baseline survey freelancers apply to 10 jobs monthly, this translates to 3 additional months of job search. Over one year (120 applications), workers who never undercut earn only \$297 USD on the platform, while those who always undercut earn \$596 USD, nearly double. This \$299 USD difference is equivalent to 26% of annual income in low-income countries. Incorporating application effort costs further widens this gap (Appendix Figure A7).

These estimates are illustrative and may not fully capture actual returns for several reasons. First, we only measure effects of undercutting by 20% on callback, while actual undercutting amounts vary widely (Appendix Figure A6). Second, we assume one-to-one conversion from callbacks to job offers, which likely provides an upper bound on platform earnings since we cannot observe actual conversion rates. Third, we compare pure strategies rather than the mixed strategies workers might employ in practice. Workers may optimize by undercutting selectively based on job characteristics or market conditions. Lastly, as workers gain experience, job offers arrive faster, so those who consistently undercut accumulate jobs more rapidly, further widening the earnings gap over time.

Despite these limitations, the results demonstrate that wage undercutting offers substantial economic gains for workers in developing countries seeking to overcome entry barriers. This raises a key question: given these returns, why don't inexperienced workers offer lower wages to break into the market?

5 Freelancer Experiment

5.1 Experimental Design

Motivated by the demand-side results, the freelancer experiment examines why few novice workers offer low wages to overcome entry barriers. The baseline survey identifies two potential frictions: workers' limited information about their own abilities and the belief that employers interpret low wages as low quality signals. The experiment tests how correcting novice workers' beliefs about their abilities or about employer responses affect their bidding decisions.

5.1.1 Timeline and Randomization

Our experimental design consists of two phases. In the first phase, we hired novice freelancers for a standardized data entry task to measure their baseline performance and beliefs. In the second phase, we cross-randomized two treatments that provided workers with information about employer responses and their performance. We measured workers' application behavior for a job opportunity on the platform as main outcomes. Figure 5 presents the experimental design.

Section 2.3 describes sample recruitment, the standardized task and performance evaluation, and the baseline survey from the first phase in detail. Following the performance evaluation, we cross-randomized two treatments at the individual level, stratified by whether workers performed above the median (high-type) or below the median (low-type).

Feedback treatment: To address workers' limited knowledge about their performance, we randomly assigned half of the sample to receive private feedback about their performance. Workers were told that we evaluated their work based on accuracy, speed, and following instructions relative to other novice freelancers hired for this job. We confirm these three dimensions as the most important performance indicators through extensive interviews with data entry freelancers at different experience levels. For each dimension, they received a score from 1 to 4 stars, representing their quartile rank. Below is an example of the feedback.

Hi! We reviewed your data entry work and compared it with other new freelancers

Figure 5: Freelancer Experiment Timeline

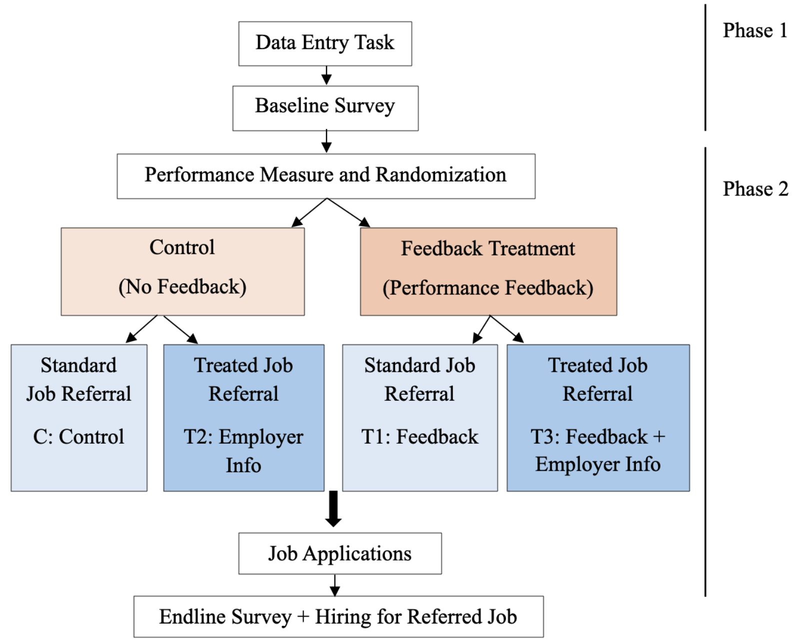

Notes: This figure shows the freelancer experimental design. This design was repeated for 8 rounds and each round took place over 2 weeks.

> *we hired. We looked at three things: accuracy of the entry, the time taken to complete the task, and how well you followed instructions. Your results:*
> *Accuracy:* ★★ *- below average, but not in the bottom 25%*
> *Speed:* ★ *- bottom 25%*
> *Following instructions:* ★★★ *- above average, but not in the top 25%*
> *Please note that this will not affect our review for you. We ONLY share this feedback with you and hope this helps you grow and succeed on this platform.*

A few days after the feedback treatment, we posted a data entry job from another employer accounts on the platform. Appendix Figure A8 shows that the job budgets were comparable to typical data entry jobs as measured in the demand experiment and well above the $5 minimum for fixed-payment projects. We then shared the job links with workers in the experimental sample from our original account.[13] The standard message was as follows:

> *Hi! A former colleague of ours is hiring for a data entry job and asked us to help spread the word widely. They mentioned they are open to working with new freelancers. If you're interested, feel free to apply directly! Since we are not involved in the hiring process, we won't be able to make individual recommendations.*

This referral design allows us to provide credible information about the employer while

[13] Within each round, we posted one job from two different employer accounts to avoid over-crowding one application pool. Workers were randomly assigned to one job opportunity.

preserving novices' incentives to attract employers unfamiliar with their abilities.

Employer info treatment: To address workers' belief that employers treat low wage bids as negative signals, we cross-randomized half the sample to receive additional information about the employer in the message. The treated message was as follows:

> *Hi! A former colleague of us is hiring for a data entry job and asked us to help spread the word widely. They mentioned they are open to working with new freelancers and* ***won't judge quality of new freelancers based on their proposed prices.*** *If you're interested, feel free to apply directly! Since we are not involved in the hiring process, we won't be able to make individual recommendations.*

One concern is that the treated message might eliminate workers' incentives to signal high quality by bidding low wages. However, baseline survey data show that less than 4% of workers believed that employers would interpret low wage bids from novices as high-quality signals. The neutral message also avoids priming effects or making workers suspicious about the job.

After the treatment, we tracked applications received within two days, as most freelancers responded within 24 hours. Application data were scraped from the employer dashboard and included both in-sample and out-of-sample applicants.

We then administered the endline survey to collect workers' opinions about the referred job and re-elicit their beliefs about performance relative to other novices. For those who did not apply, we also elicited reasons for non-application. After completing the survey, we closed the initial contract and left a public but uninformative review.

Finally, we made hiring decisions for the referred job and recorded acceptance. Among in-sample applicants, workers were classified by baseline performance type and undercutting status, and one hire was randomly selected from each group. For out-of-sample applicants, we randomly selected one novice and one veteran. The referred job involved extracting data from online reports and compiling them into an Excel spreadsheet. We measured offer acceptance rates and worker performance using the same approach described earlier.

5.1.2 Sample Characteristics and Experiment Validity

Freelancer Sample: The experimental sample consists of 481 novice freelancers from 37 LMICs. Freelancer characteristics were collected through the baseline survey, including demographics, work status, platform experience, and performance measures. We complement these with public profile characteristics scraped from the platform, including number of jobs completed, preferred hourly wages, location, English proficiency, professional certificates, and other background details. Column 1 of Appendix Table A5 reports summary statistics.

Our sample is close to evenly split by gender, has an average age of 29, and 83% hold at least a bachelor's degree. Three features are particularly notable. First, the sample represents a large idle workforce: nearly half (48%) are not currently employed offline. Second, workers struggle to break into the online market. Despite being registered for an average of 19 months, the vast majority (87%) have completed no jobs. Third, the entry barrier extends across platforms. While 44% have tried other freelancing platforms, only 18% have earned any money from them. These patterns suggest that for these workers, building reputation to access digital jobs is especially valuable given their limited offline employment opportunities and persistent difficulty breaking into online markets.

These characteristics are balanced across treatment and control groups. Appendix Table A5 compares freelancers along 20 baseline and profile characteristics. Only three show significant differences across groups at the 5% or 10% level, as expected under random chance.

Treatment Validity: The validity of our treatments rests on whether workers correctly interpreted the information provided. Figure 6 shows the treatments worked as intended.

Figure 6: Treatment Effects on Worker Beliefs

Panel A. Beliefs about Own Abilities

Panel B. Beliefs about Employers Not Inferring Quality from Wage Bids

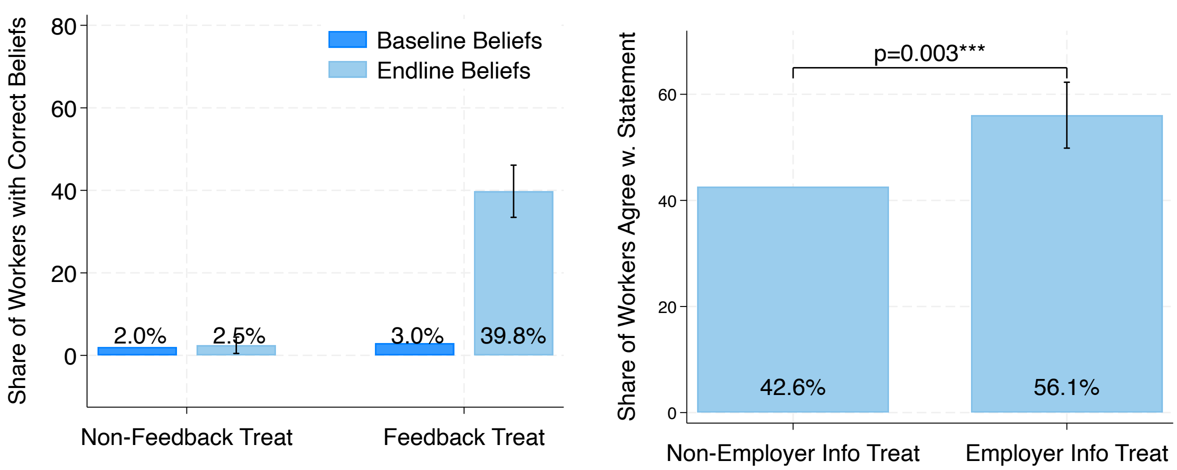

Notes: This figure reports the treatment effects on worker beliefs about own abilities and employer responses. Panel A reports the share of workers with correct beliefs of their quartiles across all three dimensions of performance (accuracy, speed, and following instructions) at baseline and endline. Panel B reports the share of workers who "agree" or "strongly agree" with the statement: *"The employer (of the referred job) will not judge my quality based on my wage bids."* We report the 95% confidence intervals in both panels and p-values from the Wald test in Panel B.

Panel A displays the proportion of workers who correctly identified their performance quartile across all three dimensions at baseline and endline. While fewer than 3% of workers accurately gauged their performance quartiles at baseline, 39.8% of those who received private feedback correctly identified their quartile at endline. In contrast, workers without feedback showed no improvement in their assessments. We show that the treatment effect holds across worker performance in Appendix Figure A9 and across alternative measures of worker beliefs in Appendix Figures A10 and A11. This confirms that the feedback treatment enables workers to correct their beliefs about their performance relative to other novices.

To examine whether the employer info treatment diminished concerns that employers would interpret low wage offers as indicators of poor quality, the endline survey asked workers whether they agreed with the statement: "The employer (of the referred job) will not judge my quality based on my wage bids." Panel B reveals that in the treated groups, the proportion of workers agreeing with the statement rose by 13.5 percentage points, representing a 32% increase from the non-treated mean. This indicates that the employer info treatment significantly alleviated workers' pessimistic beliefs about employer responses compared to the non-treated groups ($p = 0.003$). In addition, we show in Appendix Figure A12 that the treatment did not affect workers' perceptions of the job's credibility or competitiveness.

5.2 Empirical Strategy

We use the following empirical specification to measure treatment effects on the likelihood of wage undercutting.

$$\begin{aligned}(2) \qquad \text{Undercut}_i =& \alpha + \beta_1 \text{Feedback}_i + \beta_2 \text{EmployerInfo}_i + \beta_3 (\text{Feedback+EmployerInfo})_i \\ &+ \eta_r + \varepsilon_i\end{aligned}$$

where $Undercut_i$ is an indicator equal to one if novice i proposes a wage bid below the posted budget, Feedback_i, EmployerInfo_i and $(\text{Feedback+EmployerInfo})_i$ are treatment indicators, η_r denotes the experiment-round fixed effects, and ε_i are the Huber-White robust standard errors. In this specification, β's measure the average treatment effects on worker i's decision to undercut. In the presence of workers' pessimistic beliefs about employer responses and limited information about their abilities, we expect to detect non-zero estimates of the β coefficients. In particular, we can directly test Prediction 4a from Section 3 by signs of β_2 and β_3.

- Prediction 4a: Alleviating workers pessimistic beliefs encourages wage undercutting by novices That is, $\beta_2 > 0$ and $\beta_3 > 0$.

Note that the conceptual framework does not give a clear prediction on the sign of β_1 because the direction of the feedback effects depends on workers' prior beliefs of own skills.

To test wage separation by worker ability, we use the following specification for heterogeneous treatment effects by workers' baseline performance:

$$\begin{aligned}(3)\qquad \text{Undercut}_i =& \alpha + \theta_1\text{Feedback}_i + \theta_2\text{EmployerInfo}_i + \theta_3(\text{Feedback+EmployerInfo})_i \\ &+ \delta\text{High}_i + \gamma_1\text{Feedback}_i \times \text{High}_i + \gamma_2\text{EmployerInfo}_i \times \text{High}_i \\ &+ \gamma_3(\text{Feedback+EmployerInfo})_i \times \text{High}_i + \eta_r + \varepsilon_i\end{aligned}$$

where $High_i$ indicates worker i's baseline performance was above the median. In this specification, θ's measure the average treatment effects among low-type workers and γ's measure the differential treatment effects among high-type workers. Again, we map the signs of the coefficients to Predictions 2, 3, and 4b.

- Prediction 2: When both belief frictions exist, there will be pooling in wage bids and wages do not reflect worker ability. That is, $\delta = 0$.
- Prediction 3 and 4b: Providing novice workers with information about their ability induces separation in proposed wages where high-type engage in wage undercutting. That is, $\gamma_1 > 0$, $\gamma_3 > 0$.

We now take these predictions to experimental results.

5.3 Effects on Job Application and Wage Bids

Before turning to the main results on wage bids, we first analyze treatment effects on workers' application decisions. Figure 7 Panel A demonstrates that over 80% of freelancers applied to the referred job. We find no evidence of differential attrition across treatment groups or by treatment-worker type. Endline survey results support this pattern: workers' most common reasons for not applying are the perception of being unlikely to be selected due to high competition (28%) or lacking sufficient funds to submit an application (25%).

We analyze treatment effects on workers' wage bids. Panel B of Figure 7 compares the proportion of freelancers who proposed wages below the job budget across treatment groups. Two key findings emerge. First, consistent with our stylized fact that undercutting is rare among novices, only 5.5% of freelancers in the control group bid below the budget.[14] Second, all three treatments significantly increased the proportion of workers who undercut: treated workers were 8.9-15.0 percentage points more likely to propose wages below the job budget than the control group, representing a 162-273 percent increase relative to the control mean. This demonstrates that worker beliefs substantially influence their bidding decisions. Moreover, the result supports Prediction 4a that alleviating workers' pessimistic

[14] This rate is lower than what we observed in the descriptive data, likely because experimental participants received a job referral, which they may have perceived as providing a competitive advantage over other applicants.

Figure 7: Treatment Effects on Job Application and Wage Bids

Panel A. Application Decisions **Panel B.** Bids Below Job Budget

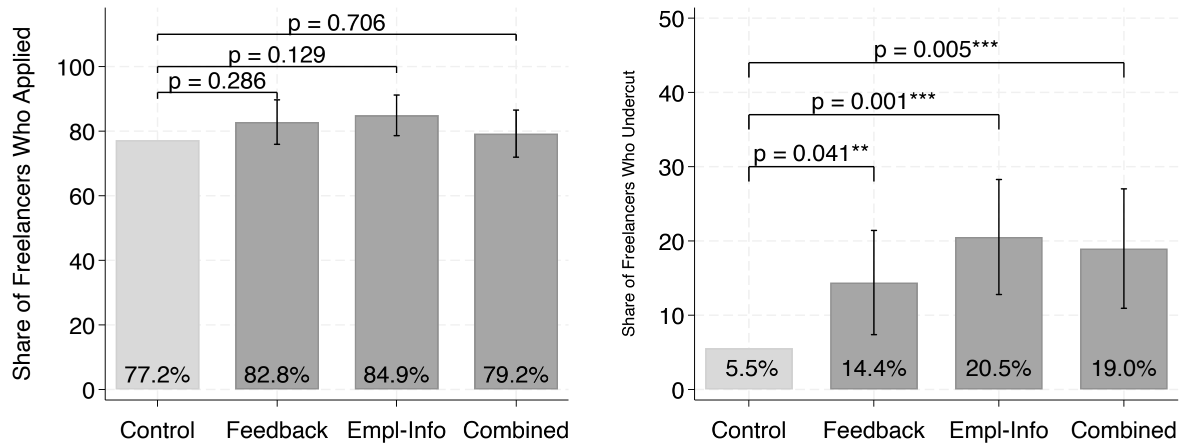

Notes: The figures show treatment effects on job application and wage bids from regression 2. Panel A shows the treatment effects on whether workers apply to the referred job and Panel B shows effects on whether workers bid below the job budget conditional on applying. Each bar represents the sum of the control group mean and the relevant regression coefficients. We report the 95% confidence intervals based on the estimated standard errors of the linear combination of regression coefficients and p-values from Wald tests. $p < 0.10^{*}, p < 0.05^{**}, p < 0.01^{***}$.

beliefs about employer responses to low wage bids increases their willing to bid lower to secure employment.

We next examine whether treatments induce separation in wage bids by worker ability. Figure 8 disaggregates each treatment group by worker type and compares bidding decisions. First, when workers don't have performance feedback, no separation emerges. In the control group, high- and low-type workers are equally likely to bid at the job budget[15] ($p = 0.742$), as predicted when both belief frictions exist in Prediction 2. In the employer info treatment, both types exhibit similar propensity to bid below the budget ($p = 0.614$), consistent with the baseline pattern that most workers believe they are high performers.

Second, when workers learn about their abilities from the performance feedback, undercutting is driven by the high types. In the feedback only group, high-type workers are four times more likely to bid below the budget than low-type ones ($p = 0.014$), consistent with Prediction 3. In the combined treatment, where both belief frictions are addressed, high-types remain twice as likely to undercut, though the difference is only marginally significant ($p = 0.135$), supporting Prediction 4b.

The experimental results raise two questions. First, given that the vast majority of workers believe they are high performers at baseline, why does feedback induce high-types to undercut when they presumably already know their quality? Second, why does the combined

[15] Only one worker in our sample proposed a wage above the job budget.

Figure 8: Treatment Effects on Wage Bids by Worker Type

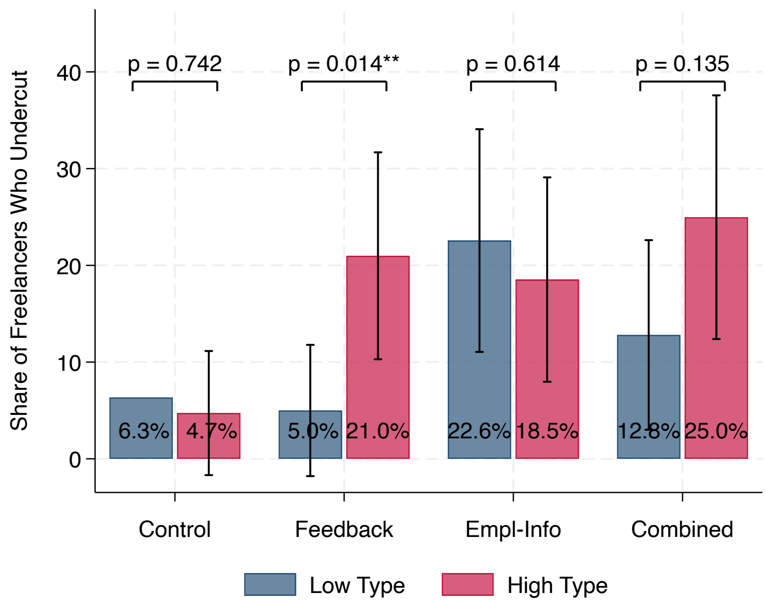

Notes: The figures show heterogeneous treatment effects on wage undercutting by worker type from regression 3. Each bar represents the sum of the control group mean and the relevant regression coefficients. We report the 95% confidence intervals based on the estimated standard errors of the linear combination of regression coefficients and p-values from Wald tests. $p < 0.10^{*}, p < 0.05^{**}, p < 0.01^{***}$.

treatment show little complementarity between the two interventions?

To address these questions, Appendix Table A6 examines heterogeneity in workers' baseline beliefs about bidding strategies. We find that differential treatment effects in the feedback-only group are driven by high-ability workers who were more likely to believe that undercutting is effective for high-types and less likely to hold pessimistic beliefs about employer responses. This reveals two insights. First, high-ability workers understand which strategies work but remain uncertain about their own abilities. The feedback treatment validates their self-assessment and prompts them to adopt the undercutting strategy they already believed was effective. Second, because these high-ability workers already held accurate beliefs about employer responses, they would be less responsive to the employer info treatment. This explains why the two treatments show limited complementarity: workers who respond to feedback already have correct beliefs about employers, so the combined effect is less than a simple sum of the individual effects.

Taken together, the experimental results reject the standard reputation model in Tirole (1988) and demonstrate that belief frictions deter novice workers from proposing lower wages to overcome reputation barriers in the online market. Our information interventions address belief frictions through different channels. Correcting workers' misperception about employer responses to low wage bids encourage those with pessimistic beliefs about employer responses to offer low wages to secure jobs. Resolving workers' uncertainty about own abilities strengthens high-ability workers' incentives to invest in reputation through low wage offers, thereby accelerating talent discovery.

6 Discussion

6.1 How Fast Can Novices Learn about Demand?

Our experimental results reveal that novice freelancers refrain from lowering initial wages to gain market entry because they hold inaccurate beliefs that employers interpret low wage bids as signals of low quality. A natural question is whether novice workers can learn the true demand curve over time. If workers were fully Bayesian and experimented with different pricing strategies, they could infer that lowering prices increases their chances of getting hired. To evaluate the feasibility of such learning, we simulate how long it would take a perfectly rational novice to learn the true demand through trial and error and what this implies for policy design.[16]

We model a novice freelancer choosing between two strategies when submitting job applications: bidding at the full budget or undercutting by 20%. Each application yields a binary outcome on whether the worker receives a callback. The worker starts with a 50/50 prior over two possible states of the world:

- "Low wage bids $\neq$ low quality" world (true model): Employers do not penalize low bids. Undercutting by 20% increases callback rates from 3.41% to 5.26% for novices.
- "Low wage bids = low quality" world: Employers respond less to novices with lower bids. The callback rate at the full budget is 3.41%. We consider two types of pessimistic beliefs about undercutting. Under moderate pessimism, workers believe the callback rate declines by 20 percent, from 3.41% to 2.73%. Under strong pessimism, workers believe the rate falls by 50 percent, to 1.71%.

The novice follows an A/B testing strategy: alternating between the two bidding strategies and updating beliefs based on the application outcome using Bayes' rule. Learning is complete when the posterior probability of the true model reaches 80%. We simulate 3,000 novice workers that each submits 2,000 applications and record the number of applications needed to learn the true model.

[16] This simulation focuses on learning about employer responses to wage bids. Novices face a similar learning challenge in assessing their own abilities, as they need to secure jobs and receive employer feedback to update beliefs about their performance. Therefore, learning about one's ability would impose similar (if not more) time and financial costs as learning about demand.

The simulation reveals that even under systematic experimentation and perfect Bayesian updating, learning is slow. Figure 9 plots the distribution of applications required to learn under the two pessimistic priors. Over 99% novices reach 80% certainty of the true model within 2,000 applications. The median worker with moderately pessimistic beliefs needs 201 applications to reach 80% confidence, while with strongly pessimistic beliefs needs 89 applications. Learning is faster under stronger pessimism because each successful application under a highly pessimistic prior is more informative about the true model and leads to belief update. Nevertheless, learning remains sluggish overall because the contact rate is low and most applications provide little information. Consequently, even perfectly rational workers cannot quickly infer how employers respond to price changes.

Figure 9: Distribution of Applications Required to Learn True Demand

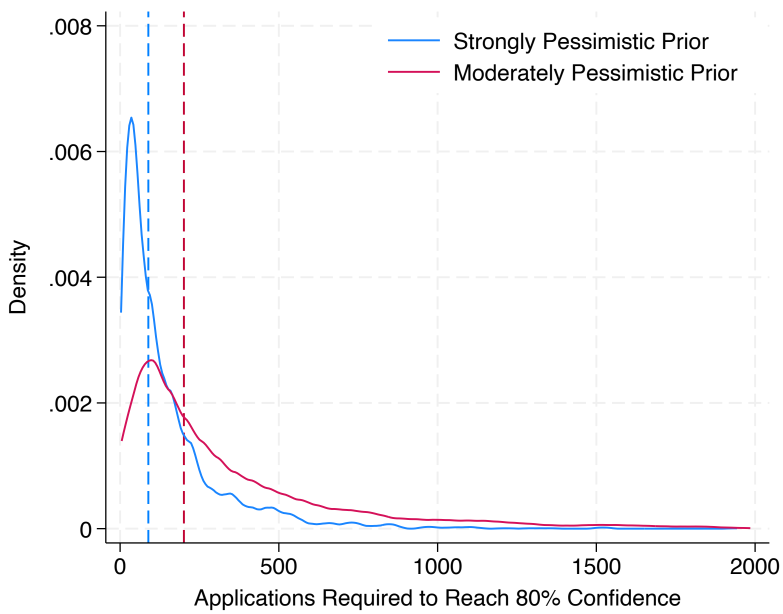

Notes: The figure plots the simulated distribution of applications required for novice freelancers with pessimistic prior beliefs to learn the true demand model. Results are based on 3,000 simulated workers who alternate between high and low bids and update their beliefs using Bayes rule. The dotted line indicates the median number of applications needed to reach 80% confidence in the true model.

The learning process translates into substantial real-world costs. Novice freelancers in our baseline survey submit on average 10 applications per month at an average cost of $1.35 USD per application (excluding workers' effort costs). A novice with moderately pessimistic beliefs would need 20 months of experimenting at the cost of over $271 USD, which is equivalent to 24% of the average annual income in low-income countries according to the World Bank.

These findings demonstrate that the learning process poses considerable time and financial burdens on workers, especially for those in LMICs, that likely discourage them from experimenting sufficiently to update their beliefs about demand. As a result, many novices may not adopt effective bidding strategies to overcome entry barriers in online labor markets. These challenges point to the potential value of light-touch interventions that lower the cost of learning rather than requiring workers to discover these patterns on their own. Existing initiatives, such as the Ajira Digital Program in Kenya and the Mastering the World of Online Freelancing projects in Jordan and Lebanon, incorporate mentorship and business strategy guidance to help new workers navigate bidding in online freelancing platforms (Datta et al., 2023). Similar components could be integrated into digital training programs to help workers form accurate expectations about job search in global markets.

6.2 Generalization to Other Settings

Although this paper focuses on online marketplaces, our findings have broader implications for traditional labor markets in developing countries as they share important characteristics. First, both operate as spot markets characterized by high unemployment rates, high search costs, and short job spells (Breza and Kaur, 2025). Recent evidence also documents lack of wage variation in labor markets in developing countries (Breza, Kaur and Shamdasani, 2018; Breza, Kaur and Krishnaswamy, 2019; Cefala et al., 2024; K, 2025), where wages vary less than the underlying marginal product of labor. Second, inexperienced workers often struggle more to find jobs than experienced ones, particularly in low- and middle-income countries. In some regions, they can be three times more likely to be unemployed than the general population (Alfonsi et al., 2020). These workers face challenges credibly showcasing their skills to employers, as education quality is often poor and provides little information about actual abilities (Bertrand et al., 2024). Lastly, young workers often possess limited knowledge about their own abilities and employer preferences even within the local market (Carranza et al., 2022; Kiss et al., 2023; Caria et al., 2024). Taken together, our results suggest that supply-side information frictions could potentially explain the lack of wage variation in developing countries and targeted treatments could improve job matching in their entry-level labor markets.

However, three important features of offline labor markets may limit the applicability of our results. First, returns to experience are often low in many developing country labor markets (Donovan, Lu and Schoellman, 2023), which may reduce workers incentives to forgo initial earnings as investments in reputation building. However, this pattern primarily reflects that when workers switch jobs, their performance cannot be credibly conveyed to new employers. Indeed, conditional on staying in the same job, returns to experience are substantially higher (Breza and Kaur, 2025) because worker performance information becomes available over time. This mirrors the reputation-building process on online platforms, where performance history is available to all potential employers. Therefore, high-ability workers in offline labor markets can still benefit from offering low initial wages to reveal their abilities to employers.

Second, workers in LMICs may face subsistence constraints that appear to prevent wage undercutting. However, empirical evidence shows that wage undercutting does occur in these settings. Studies of low-wage manual laborers in rural labor markets in India find that workers accept employment below prevailing wages when their decisions are unobserved by others (Breza, Kaur and Krishnaswamy, 2019) or when workers have dispersed reservation wages (K, 2025). These results demonstrate that workers do engage in strategic wage-setting despite financial constraints.

Third, social norms against wage cuts are prevalent across many traditional labor markets, especially in rural areas (Breza, Kaur and Krishnaswamy, 2019). In such contexts, workers with perfect information may still under-adopt wage-signaling strategies to avoid social sanctions. This suggests that our findings are most relevant for sectors in urban labor markets with higher returns to experience, such as manufacturing, rather than manual labor jobs.

7 Conclusion

In this paper, we conduct two field experiments to examine whether novice workers from LMICs can overcome entry barriers in online labor markets through offering low wages initially. In the demand-side experiment, we randomly assign wage offers by novice and veteran workers to online jobs and track employer responses. We find that employers respond positively to low wage offers from novice workers. Back-of-the-envelope calculations reveal substantial returns to wage undercutting: workers who consistently undercut earn twice as much in their first year compared to those who bid at full budget.

Despite these returns, few novices adopt this strategy. Our survey data points to two potential explanations for why they under-adopt this profitable strategy: pessimistic beliefs that employers interpret low wages as negative quality signals and limited information about their abilities. In the supply-side experiment, we test whether providing novice workers with information about employer responses and their performance affect wage bidding decisions. Both interventions significantly increase undercutting, with effects concentrated among high-ability workers informed about their performance, consistent with the standard reputation model. This separation benefits both workers and employers: it increases high-ability workers' hiring prospects while accelerating talent discovery on the platform.

To assess whether novice workers can correct their beliefs through trial and error, we simulate worker learning under Bayesian updating. Even with systematic experimentation, learning on one's own is costly and time-consuming in this market, creating significant burdens for workers in developing countries with financial constraints. These findings highlight the value of external information interventions, such as mentorship or business strategy programs, that lower novices' cost of acquiring accurate information and help them overcome entry barriers to a new market.

Our study has two limitations. First, our analysis focuses on data entry services at the lower end of the skill spectrum among digital services. For services that require more advanced skills, such as graphic design and website development, inexperienced workers may be better able to showcase their abilities to foreign employers by building portfolios or samples of work offline. Nevertheless, using job history data on graphic designers from the same platform, we still find evidence of high entry barriers, as jobs are concentrated among top-rated freelancers. Second, since we are not able to directly engage with employers on freelancing platforms, we do not have direct evidence on how employers infer worker quality from their wage bids and how they use other potential quality signals from worker profiles. Better information on how employers make hiring decisions would help new workers optimize their application strategies to enter the market.

Nevertheless, the implications of these findings can be extended beyond online labor markets to broader questions of information provision and market efficiency. Our results suggest that policymakers can improve talent discovery and worker outcomes by providing skill assessments for new entrants and sharing information about employer decision-making processes. More broadly, the findings indicate that policies addressing information frictions on the supply side, rather than the demand side, may be more effective for improving job matching in markets characterized by high entry barriers and low-information environments.

References

Abebe, Girum, A Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, and Simon Quinn. 2021. "Anonymity or distance? Job search and labour market exclusion in a growing African city." *The Review of Economic Studies*, 88(3): 1279–1310.

Abebe, Girum, A Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, Simon Quinn, and Forhad Shilpi. 2025. "Matching frictions and distorted beliefs: Evidence from a job fair experiment." *The Economic Journal*, 135(671): 2089–2121.

Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons. 2016. "Does Standardized Information in Online Markets Disproportionately Benefit Job Applicants from Less Developed Countries?" *Journal of Labor Economics*, 34(S1): S225–S268.

Alfonsi, Livia, Mary Namubiru, and Sara Spaziani. 2024. "Meet Your Future: Experimental Evidence on the Labor Market Effects of Mentors." GLM—LIC, IZA Working Paper No. 87, Bonn.

Alfonsi, Livia, Oriana Bandiera, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2020. "Tackling youth unemployment: Evidence from a labor market experiment in Uganda." *Econometrica*, 88(6): 2369–2414.

Ashraf, Nava, B. Kelsey Jack, and Emir Kamenica. 2013. "Information and Subsidies: Complements or Substitutes?" *Journal of Economic Behavior & Organization*, 88: 133–139.

Bandiera, Oriana, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2025. "The search for good jobs: evidence from a six-year field experiment in Uganda." *Journal of labor economics*, 43(3): 885–935.

Banerjee, Abhijit, and Sandra Sequeira. 2023. "Learning by searching: Spatial mismatches and imperfect information in Southern labor markets." *Journal of Development Economics*, 164: 103111.

Baptista, Dulce, Richard Freund, and Rafael Novella. 2023. "Entrepreneurial skills training for online freelancing: Experimental evidence from Haiti." *Economics Letters*, 232.

Barach, Moshe, and John J. Horton. 2019. "How Do Employers Use Compensation History? Evidence from a Field Experiment." *Journal of Labor Economics*, 37(1): S25–S46.

Bassi, Vittorio, and Aisha Nansamba. 2022. "Screening and signalling non-cognitive skills: experimental evidence from Uganda." *The Economic Journal*, 132(642): 471–511.

Bertrand, Marianne, Bruno Crpon, Stefano Caria, and Natalie Valent. 2024. "Improving Job Seekers Employment and Earnings Through Credible Skills Signals." *Policy Insight, The Abdul Latif Jameel Poverty Action Lab (J-PAL)*, Last updated November 2024.

Breza, Emily, and Supreet Kaur. 2025. "Labor Markets in Developing Countries." *Annual Review of Economics*, 17.

Breza, Emily, Supreet Kaur, and Nandita Krishnaswamy. 2019. "Social norms as a determinant of aggregate labor supply." National Bureau of Economic Research.

Breza, Emily, Supreet Kaur, and Yogita Shamdasani. 2018. "The morale effects of pay inequality." *The Quarterly Journal of Economics*, 133(2): 611–663.

Caria, Stefano, Kate Orkin, Alison Andrew, Robert Garlick, Rachel Heath, and Niharika Singh. 2024. "Barriers to search and hiring in urban labour markets." *VoxDevLit*, 10(1): 1–46.

Carranza, Eliana, Robert Garlick, Kate Orkin, and Neil Rankin. 2022. "Job Search and Hiring with Limited Information about Workseekers' Skills." *American Economic Review*, 112(11): 3547–3583.

Cefala, Luisa, Pedro Naso, Michel Ndayikeza, and Nicholas Swanson. 2024. "Under-training by Employers in Informal Labor Markets: Evidence from Burundi." Working paper.

Chakravorty, Bhaskar, Wiji Arulampalam, Apurav Yash Bhatiya, Clément Imbert, and Roland Rathelot. 2024. "Can information about jobs improve the effectiveness of vocational training? Experimental evidence from India." *Journal of Development Economics*, 169: 103273.

Das, Narayan, Md.R̃aied Arman, Anindita Bhattacharjee, Mohima Gomes, and Sakib Mahmood. 2024. "Tackling Women's Unemployment: Evidence from an Online Freelancing Experiment in Bangladesh." SSRN Electronic Journal Working Paper SSRN 4843957. Posted 27 May 2024.

Datta, Namita, Chen Rong, Sunamika Singh, Clara Stinshoff, Nadina Iacob, Natnael Simachew Nigatu, Mpumelelo Nxumalo, and Luka Klimaviciute. 2023. "Working Without Borders: The Promise and Peril of Online Gig Work." World Bank Report, Washington, DC.

Donovan, Kevin, Will Jianyu Lu, and Todd Schoellman. 2023. "Labor market dynamics and development." *The Quarterly Journal of Economics*, 138(4): 2287–2325.

Erdem, Tülin, Michael P. Keane, and Baohong Sun. 2008. "A Dynamic Model of Brand Choice When Price and Advertising Signal Product Quality." *Marketing Science*, 27(6): 1111–1125.

Fazio, Maria Victoria, Richard Freund, and Rafael Novella. 2025. "Do Entrepreneurial Skills Unlock Opportunities for Online Freelancing? Experimental Evidence from El Salvador." *Journal of Development Economics*, 172: 103363.

Hannane, Jonas. 2024. "Advertising in online labor markets: A signal of freelancer quality?"

Hilbert, Martin R., and Kangbo Lu. 2020. "The Online Job Market Trace in Latin America and the Caribbean." United Nations Economic Commission for Latin America and the Caribbean (ECLAC) Project Document LC/TS.2020/83, Santiago.

Huang, Yangguang, Chenyang Li, and Si Zuo. 2025. "Price signaling and reputation building: evidence from a service platform." *HKUST Business School Research Paper*, , (2022-062).

International Labour Organization. 2021. "World Employment and Social Outlook 2021: The Role of Digital Labour Platforms in Transforming the World of Work." International Labour Organization Flagship Report, Geneva.

Judd, Kenneth L, and Michael H Riordan. 1994. "Price and quality in a new product monopoly." *The Review of Economic Studies*, 61(4): 773–789.

Kala, Namrata, and Elizabeth Lyons. 2025. "The Effects of Digital Surveillance and Managerial Clarity on Performance." National Bureau of Economic Research.

Kässi, Otto, and Vili Lehdonvirta. 2022. "Do Microcredentials Help New Workers Enter the Market? Evidence from an Online Labor Platform." *Journal of Human Resources.*

Kelley, Erin M, Christopher Ksoll, and Jeremy Magruder. 2024. "How do digital platforms affect employment and job search? Evidence from India." *Journal of Development Economics*, 166: 103176.

Kiss, Andrea, Robert Garlick, Kate Orkin, and Lukas Hensel. 2023. "Jobseekers' Beliefs about Comparative Advantage and (Mis)Directed Search." Centre for the Study of African Economies (CSAE), University of Oxford Working Paper CSAE-WPS-2023-11.

K, Varun. 2025. "Negotiations and Undercutting in Low Wage Labor Markets." Working paper.

Milgrom, Paul, and John Roberts. 1986. "Price and advertising signals of product quality." *Journal of political economy*, 94(4): 796–821.

Online Labour Observatory. 2025. "Online Labour Observatory." http://onlinelabourobservatory.org/, Accessed: 2025-09-18.

Pallais, Amanda. 2014. "Inefficient Hiring in Entry-Level Labor Markets." *American Economic Review*, 104(11): 3565–3599.

Rao, Akshay R., and Kent B. Monroe. 1989. "The Effect of Price, Brand Name, and Store Name on Buyers Perceptions of Product Quality: An Integrative Review." *Journal of Marketing Research*, 26(3): 351–358.

Roussille, Nina. 2024. "The Role of the Ask Gap in Gender Pay Inequality." *Quarterly Journal of Economics*, 139(3): 1557–1610.

Stiglitz, Joseph E. 1989. "Imperfect Information in the Product Market." In *Handbook of Industrial Organization*. Vol. 1, , ed. Richard Schmalensee and Robert D. Willig, Chapter 13, 769–847. Elsevier.

Tirole, Jean. 1988. *The Theory of Industrial Organization.* Cambridge, MA:MIT Press.

Wathieu, Luc, and Marco Bertini. 2007. "Price as a Stimulus to Think: The Case for Willful Overpricing." *Marketing Science*, 26(1): 118–129.

Wolinsky, Asher. 1983. "Prices as signals of product quality." *The review of economic studies*, 50(4): 647–658.

Appendix A Figures and Tables

Appendix Figure A1: Freelancer Profiles for Audit Study

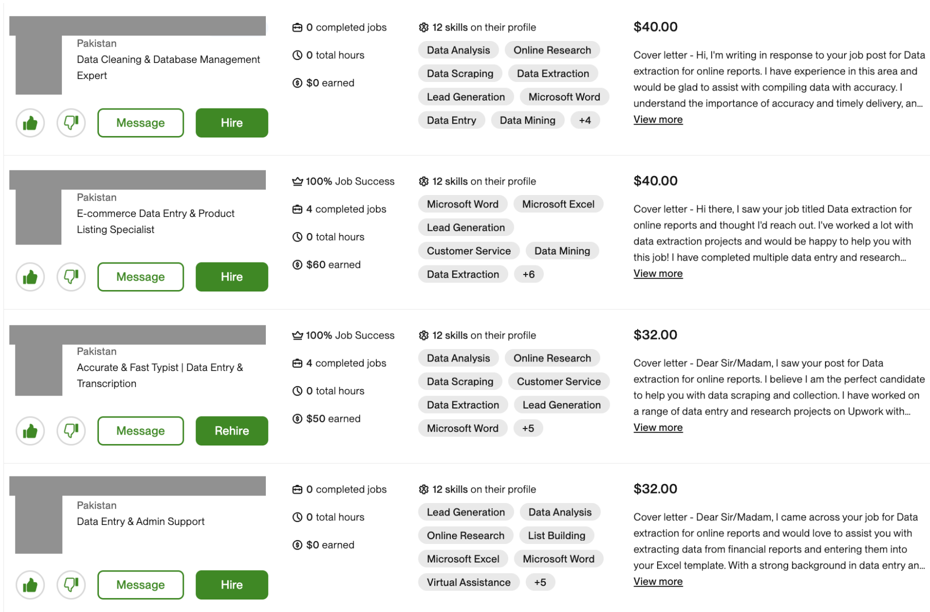

Notes: This is a screenshot taken from the employer's dashboard on the online freelancing platform. The original job post has a fixed budget of 40 USD. The applications are sent by four freelancer profiles in the audit study. The top and the bottom applications are sent by inexperienced freelancers and the middle ones by experienced freelancers.

Appendix Figure A2: Hiring Prospects and Completed Jobs

Panel A. Gap between Jobs (Days)

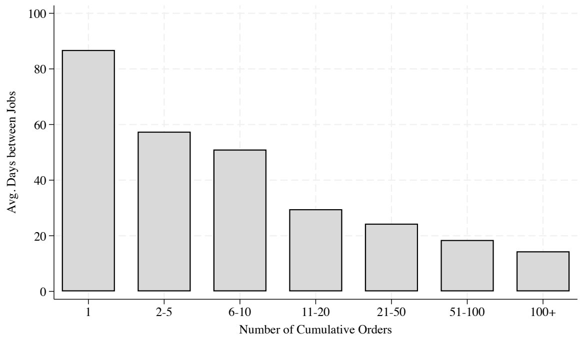

Panel B. Log Earnings of Future Jobs

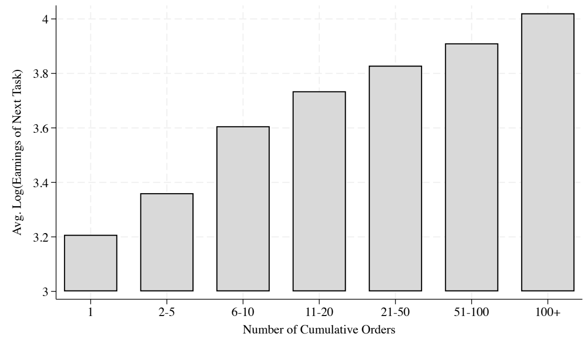

Notes: The figures plot the average number of days between jobs (Panel A) and the log of earnings from the subsequent job (Panel B) against the number of completed jobs. The job-level data was collected for all active Pakistani workers specialized in data entry on the studied platform by June 2024.

Appendix Figure A3: Distribution of Deviation in Proposed Wages

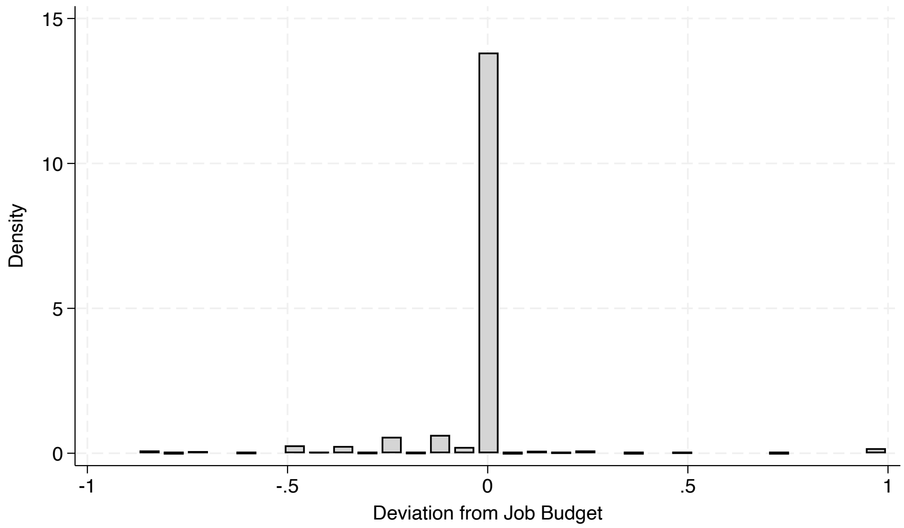

Notes: This figure shows the distribution of deviation in proposed wages from the job budget, measured in the application data.

Appendix Figure A4: Distribution in Performance by Each Dimension

Panel A. Accuracy

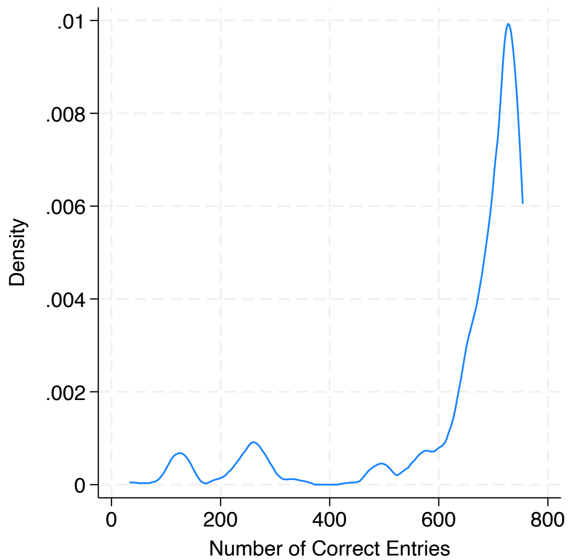

Panel B. Speed

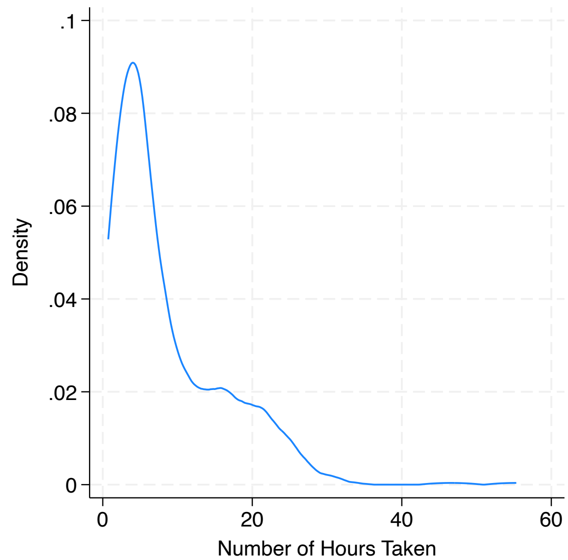

Panel C. Follow Instruction

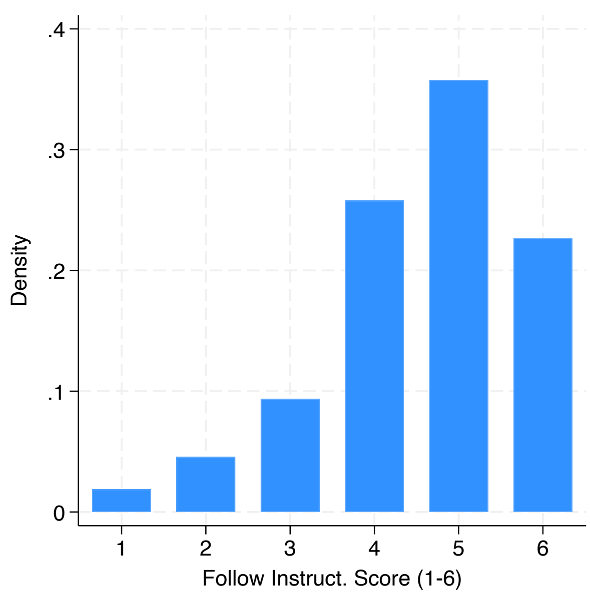

Notes: This figure shows the distribution of the number of correct entries (accuracy), hours taken to complete the task (speed), and the following instruction scores for the standardized data entry task.

Appendix Figure A5: Actual versus Perceived Performance by Each Dimension

Panel A. Accuracy

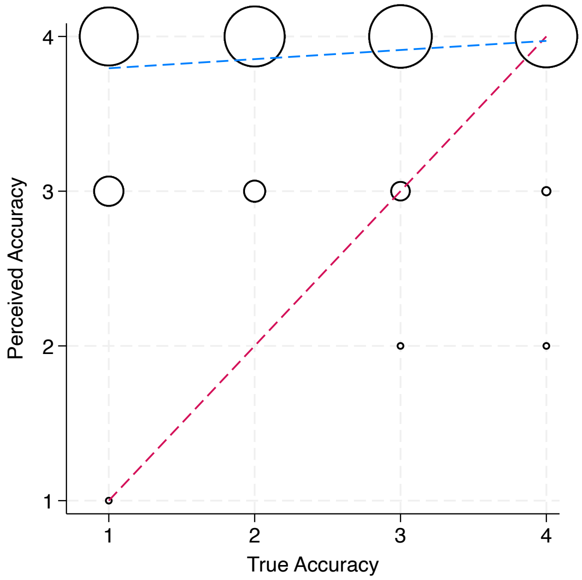

Panel B. Speed

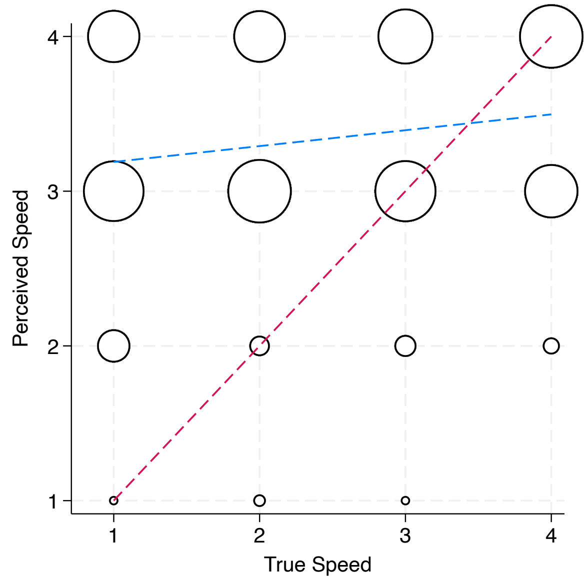

Panel C. Follow Instruction

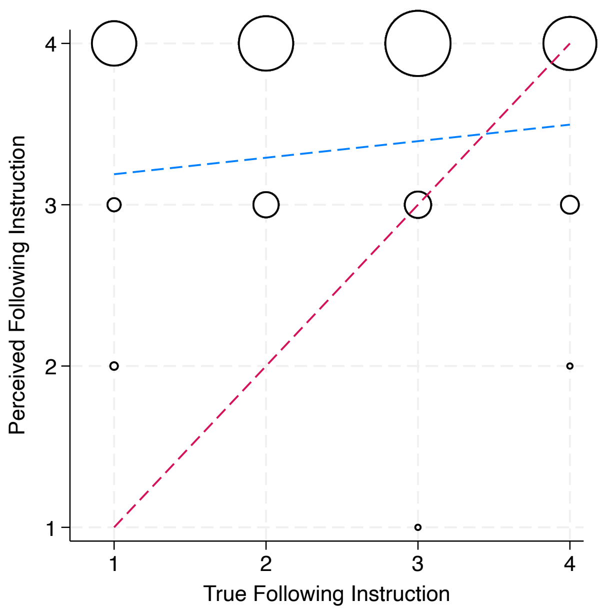

Notes: This figure shows the actual and perceived scores for accuracy, speed, and following instructions, with the true scores on the x-axis and the perceived scores on the y-axis. The red dotted line is the 45-degree line and the blue line shows the best linear fit. The marker size is relative to the number of observations.

Appendix Figure A6: Distribution of Wage Bids Below Job Budget

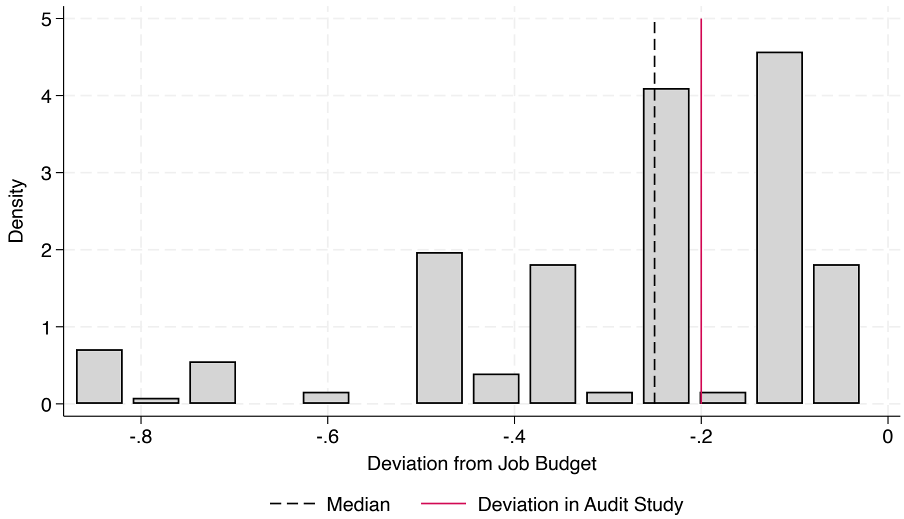

Notes: This figure shows the distribution of wage bids below the job budget, measured in the application data. The dotted black line shows the median value of the deviation and the solid red line represents wage undercutting by 20%, which is what we proposed in the audit study.

Appendix Figure A7: Expected Net Earnings with Effort Costs

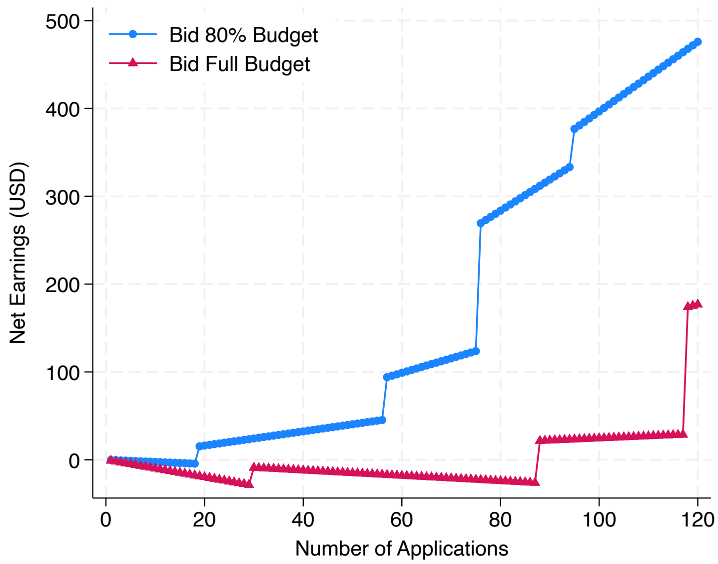

Notes: This figure shows the relationship between number of applications and net earnings under the two bidding strategies. We account for workers' effort costs in the calculation of net earnings.

Appendix Figure A8: Distribution of Budget for Data Entry Jobs

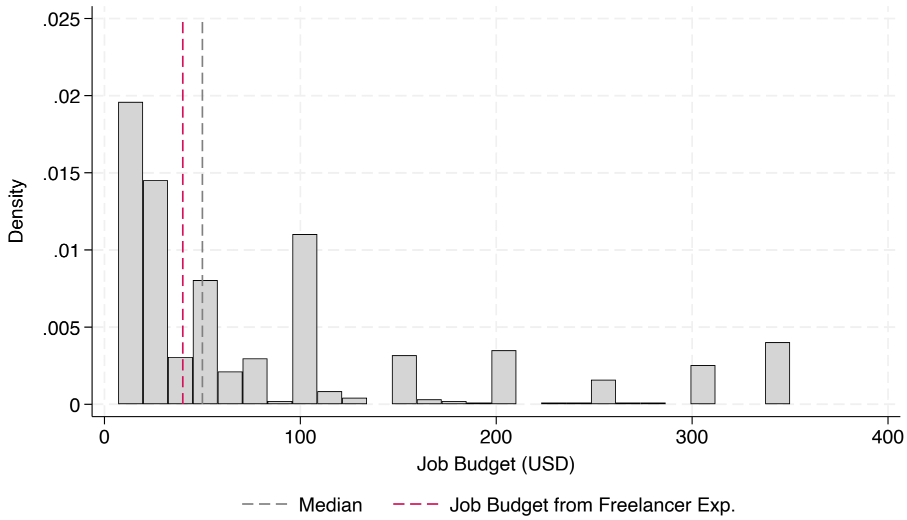

Notes: This figure shows the distribution of job budget from 703 data entry jobs in the demand experiment. The red dotted line shows where the job budget for the referred job falls in the freelancer experiment.

Appendix Figure A9: Treatment Effects on Perceived Performance

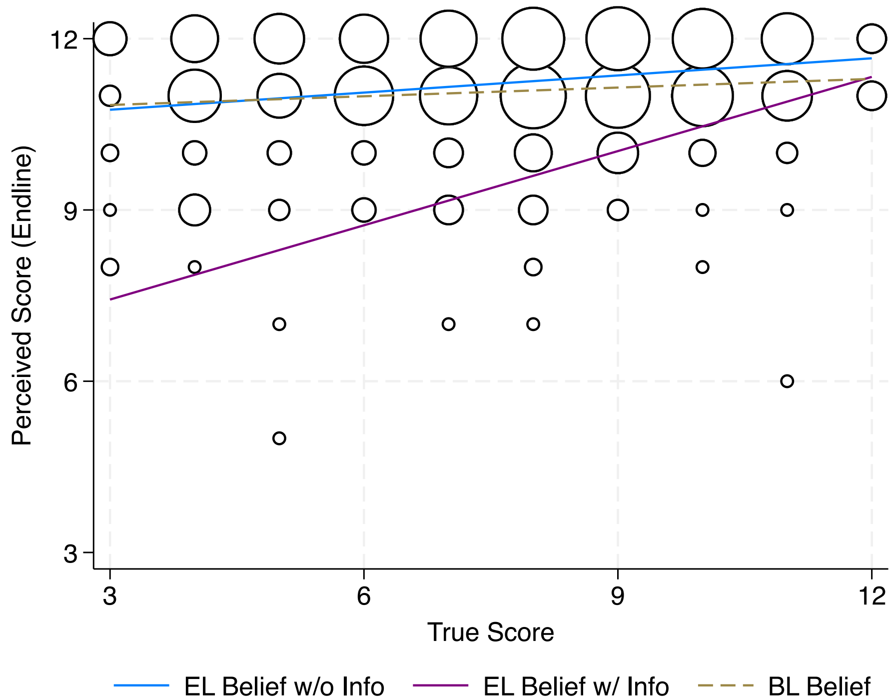

Notes: This figure plots workers' actual performance scores against their perceived scores measured at the endline. The blue line represents the best linear fit for workers without performance feedback and the red line for workers with performance feedback. The dotted line shows the best linear fit of baseline perceived performance against true performance.

Appendix Figure A10: Treatment Effects on Worker Beliefs by Performance Dimension

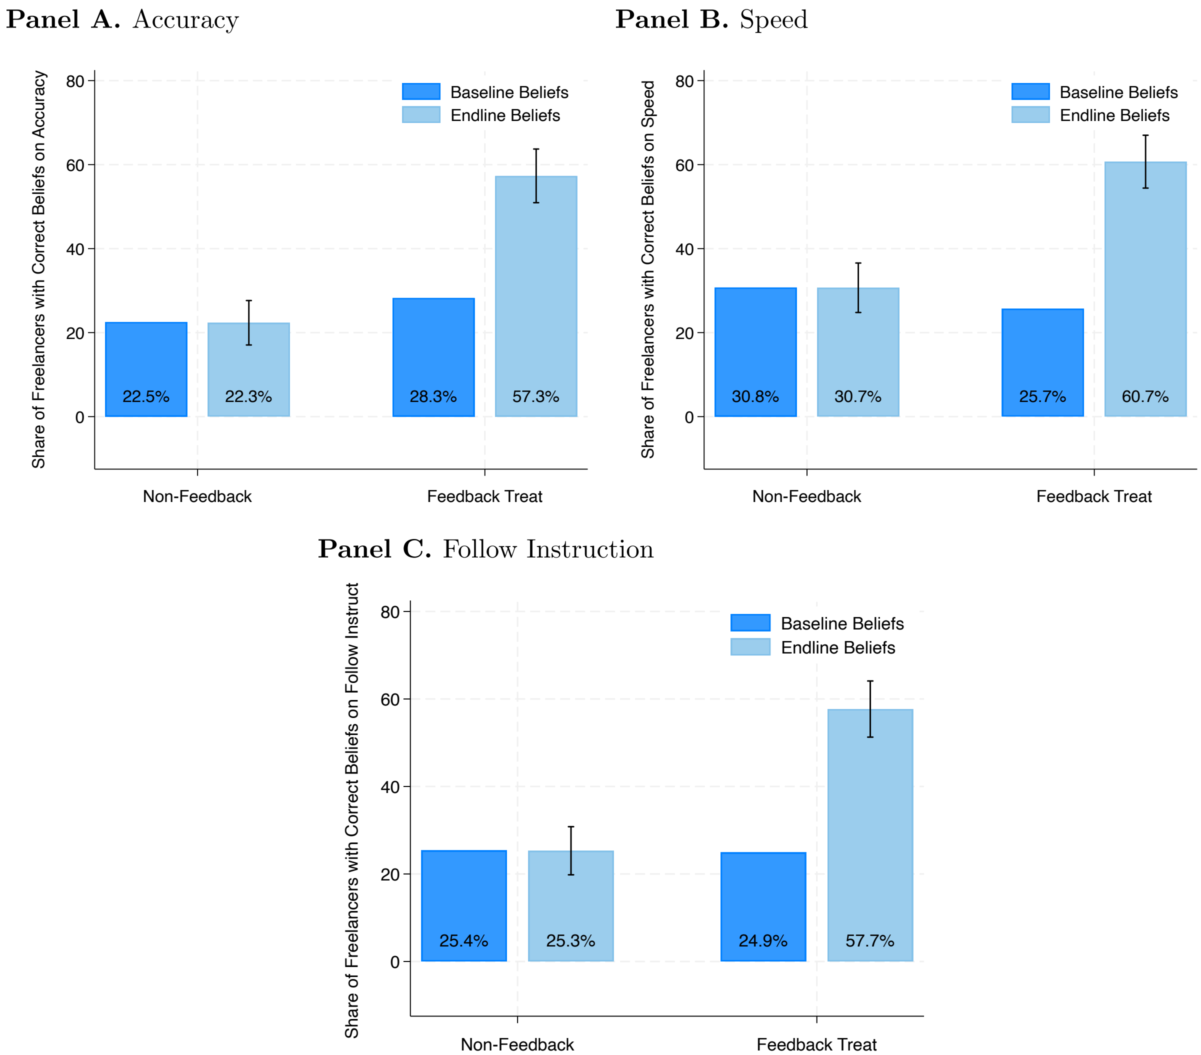

Notes: This figure reports the treatment effects on worker beliefs about own abilities by each dimension. Panels A-C show the share of workers with correct beliefs of their performance quartiles for accuracy, speed, and following instructions, respectively, at baseline and endline.

Appendix Figure A11: Treatment Effects on Number of Correct Beliefs

Panel A. Correct Beliefs on ≥ 1 Dimensions

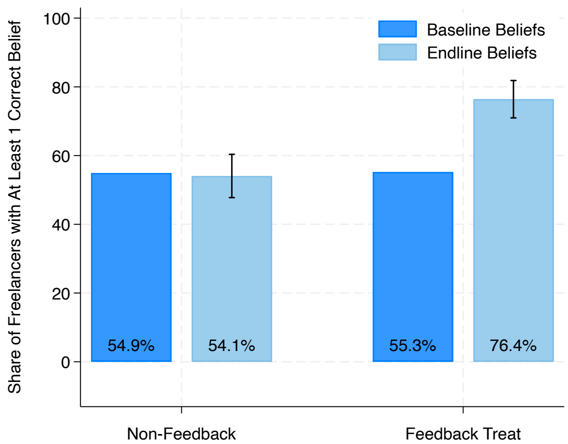

Panel B. Correct Beliefs on ≥ 2 Dimensions

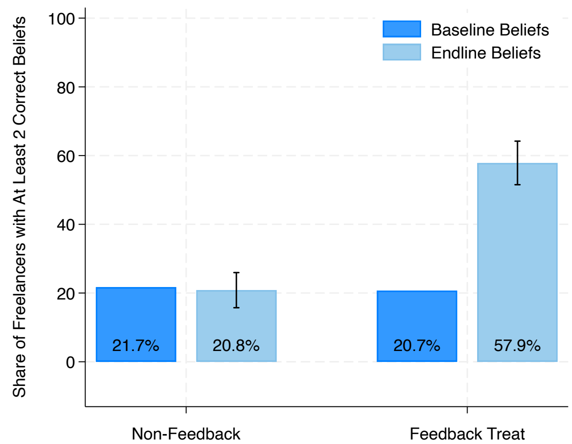

Notes: This figure reports the treatment effects on worker beliefs about own abilities by each dimension. Panels A and B show the share of workers with at least one/two correct beliefs of their performance quartiles across all three dimensions of performance (accuracy, speed, and following instructions) at baseline and endline.

Appendix Figure A12: Treatment Effect on Worker Perceptions of Referred Job

Panel A. This job opportunity seems trustworthy. **Panel B.** Competition may be strong.

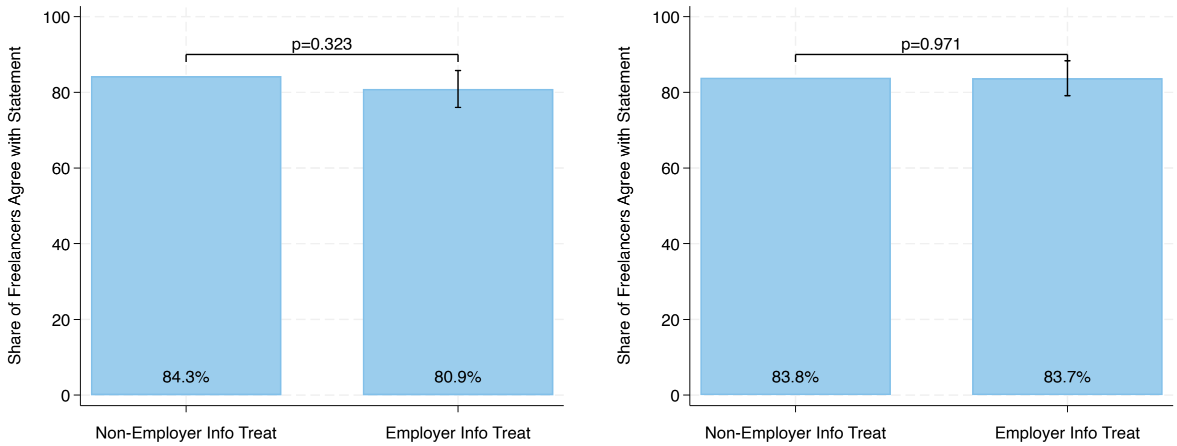

Notes: This figure reports the treatment effects on worker beliefs about the trustworthiness and competition for the referred job. We report the 95% confidence intervals in both panels and p-values from the Wald test.

Appendix Table A1: Job and Employer Characteristics

	Median	Mean	Min	Max	N
Job Characteristics					
Job Budget (USD)	50	112	7	8,000	743
Any Freelancer Hired	1.00	0.54	0.00	1.00	743
No. Freelancer Hired	1.00	0.73	0.00	14.00	743
Employer Characteristics					
Employer Tenure (Month)	24	40	0	281	720
Employer from HICs	1.00	0.85	0.00	1.00	720
No. Jobs Posted	21	139	1	3,468	720
Employer Rating (Out of 5)	4.99	4.86	1.75	5.00	565

Notes: This table reports characteristics of jobs and employers in the analysis sample of the employer experiment. The first panel shows job characteristics and hiring outcomes indicated on the platform, at least two weeks after the initial job post. The second panel shows employer characteristics. *Employer Tenure* is measured by the gap between when the employer first joined the platform and when the job was posted. *Employer from HICs* indicates whether the employer is located in a high-income country, using World Bank country classifications by income level for 2024-2025. *No. Jobs Posted* refers to the total number of jobs the employer has posted since joining the platform. *Employer Rating* is a measure of reputation (from one to five) generated by the platform based on rating by freelancers previously hired by the employer. New employers without any hiring history do not have a rating.

Appendix Table A2: Treatment Effects on Employer Responses

	Client Viewed			Contacted		
	Analysis Sample		Full Sample	Analysis Sample		Full Sample
	(1)	(2)	(3)	(4)	(5)	(6)
Novice (β_1)	-0.050	-0.043	-0.041	-0.018	-0.025	-0.025
	(0.011***)	(0.013***)	(0.012***)	(0.008**)	(0.009***)	(0.009***)
Undercut (β_2)	0.010	0.009	0.010	0.009	0.008	0.008
	(0.010)	(0.010)	(0.010)	(0.007)	(0.007)	(0.006)
Novice × Undercut (β_3)	0.038	0.039	0.035	0.010	0.010	0.010
	(0.015**)	(0.015**)	(0.015**)	(0.011)	(0.011)	(0.010)
Constant	0.095	0.082	0.080	0.053	0.050	0.049
Test: $\beta_2 + \beta_3 = 0$	0.000	0.000	0.000	0.024	0.026	0.019
Job FE	×	×	×	×	×	×
Profile FE		×	×		×	×
No. Applications	2812	2812	2972	2812	2812	2972

Notes: This table reports the main results on employer views and contacts by worker experience and wage undercutting. Column (1) and (4) present the baseline specification and correspond to Figure 3. The rest of columns add freelancer profile fixed effects. In addition, column (3) and (6) extend the analysis to the full sample, including vacancies that received less than four applications or applications with wage bids that deviated from the pre-assigned values. All regressions include job fixed effects. The omitted group is veteran applicants without wage undercutting. We report p-values from the statistical test for $\beta_2 + \beta_3 = 0$ from all specifications. Standard errors are clustered at the job level. $p < 0.10^{*}, p < 0.05^{**}, p < 0.01^{***}$.

Appendix Table A3: Heterogeneous Treatment Effects on Employer Responses

	Job Budget		Employer Tenure		HIC Employer		No. Jobs Posted		Employer Rating	
	< Median (1)	≥ Median (2)	< Median (3)	≥ Median (4)	= 0 (5)	= 1 (6)	< Median (7)	≥ Median (8)	< Median (9)	≥ Median (10)
Panel A: Employer View										
Novice (β_1)	-0.057	-0.043	-0.074	-0.023	-0.029	-0.052	-0.083	-0.014	-0.030	-0.073
	(0.016***)	(0.016***)	(0.018***)	(0.013*)	(0.036)	(0.012***)	(0.018***)	(0.014)	(0.018)	(0.017***)
Undercut (β_2)	0.006	0.013	0.003	0.014	0.000	0.010	0.006	0.011	0.015	-0.015
	(0.014)	(0.016)	(0.016)	(0.014)	(0.024)	(0.012)	(0.016)	(0.013)	(0.017)	(0.015)
Novice × Undercut (β_3)	0.039	0.038	0.057	0.023	0.049	0.038	0.074	0.006	0.023	0.055
	(0.020*)	(0.023)	(0.024**)	(0.020)	(0.043)	(0.017**)	(0.024***)	(0.020)	(0.026)	(0.024**)
Omitted Mean	0.099	0.092	0.116	0.075	0.137	0.088	0.122	0.069	0.098	0.099
No. Applications	1328	1484	1344	1392	408	2328	1344	1392	1064	1092
Panel B: Employer Contact										
Novice (β_1)	-0.021	-0.016	-0.030	-0.003	0.000	-0.019	-0.030	-0.003	-0.011	-0.029
	(0.012*)	(0.011)	(0.013**)	(0.010)	(0.024)	(0.009**)	(0.012**)	(0.011)	(0.016)	(0.011**)
Undercut (β_2)	0.018	0.000	0.006	0.014	0.000	0.012	0.009	0.011	0.011	0.015
	(0.011)	(0.008)	(0.011)	(0.008*)	(0.014)	(0.007)	(0.009)	(0.010)	(0.012)	(0.010)
Novice × Undercut (β_3)	-0.003	0.022	0.018	-0.003	0.029	0.003	0.021	-0.006	-0.004	0.004
	(0.014)	(0.016)	(0.017)	(0.014)	(0.030)	(0.012)	(0.015)	(0.016)	(0.022)	(0.012)
Omitted Mean	0.057	0.049	0.065	0.037	0.088	0.045	0.057	0.046	0.071	0.044
No. Applications	1328	1484	1344	1392	408	2328	1344	1392	1064	1092

Notes: This table examines the heterogeneity in employer's response to wage undercutting by novice and veteran freelancers. For *Job Budget, Employer Tenure, No. Jobs Posted, and Employer Rating*, we divide the analysis sample by the median and for *HIC Employer*, the analysis sample is split by whether the employer is located in a high-income country, using World Bank country classifications by income level for 2024-2025. Panel A reports the heterogeneous treatment effects on whether employers viewed the application and Panel B on whether employers reached out to the applicant. All regressions control for job fixed effects. The omitted group is applications by veteran workers without wage undercutting. Standard errors are clustered at the job level. $p < 0.10^{*}, p < 0.05^{**}, p < 0.01^{***}$.

Appendix Table A4: Data Sources for Back-of-the-Envelope Calculation

	Parameter Values		Source
	Bid 80%	Bid 100%	
Callback Probability (p)	5.26%	3.41%	Demand Experiment
Initial Job Earnings (e)	50	40	Demand Experiment
Application Fee (c_a)	1.35	1.35	Demand Experiment
Application Effort Cost (c_e)	1	1	Assumption
Earning Growth by Completed Jobs ($\rho(n)$)	-	-	Job History Data
Application Sent per Month	10	10	Freelancer Survey

Notes: This table reports the key parameters used in the back-of-the-envelope calculation of the returns to wage undercutting, along with their data sources. The initial job earnings reflect the original job budgets posted by employers before any wage discounting. These values differ across bidding strategies because, in the demand experiment, jobs that responded to lower bids had a higher median budget than those that responded to full-price bids. For the application effort cost, we assume that workers spend 10 minutes per application and have a reservation wage of $6 per hour based on the baseline survey, which implies an effort cost of approximately $1 per application. Earnings growth is computed as a function of the number of completed jobs using the job history data, so no single value is reported here.

Appendix Table A5: Balance Test for Freelancer Experiment

	All	Control	Info		Signal		Info+Signal		p-val
N	481	118	117		126		120		
	Mean (SD)	Mean (SD)	Mean (SD)	β_1 (p-val)	Mean (SD)	β_2 (p-val)	Mean (SD)	β_3 (p-val)	
Panel A: Personal Characteristics									
Female	0.52	0.56	0.55	-0.01	0.47	-0.09	0.50	-0.06	0.45
	(0.50)	(0.50)	(0.50)	(0.86)	(0.50)	(0.16)	(0.50)	(0.36)	
Age	28.76	29.16	28.55	-0.62	28.96	-0.19	28.35	-0.81	0.72
	(5.98)	(6.67)	(5.87)	(0.45)	(5.65)	(0.81)	(5.74)	(0.32)	
Bachelor Degree or Above	0.83	0.79	0.85	0.06	0.83	0.04	0.85	0.06	0.59
	(0.38)	(0.41)	(0.36)	(0.24)	(0.38)	(0.46)	(0.36)	(0.21)	
Employed Offline	0.52	0.52	0.58	0.06	0.50	-0.02	0.48	-0.03	0.44
	(0.50)	(0.50)	(0.50)	(0.32)	(0.50)	(0.80)	(0.50)	(0.61)	
Full-Time Freelancer	0.36	0.35	0.35	0.00	0.33	-0.02	0.40	0.05	0.67
	(0.48)	(0.48)	(0.48)	(0.95)	(0.47)	(0.71)	(0.49)	(0.40)	
Platform Tenure (Month)	19.65	22.25	17.50	-4.75	19.00	-3.28	19.87	-2.36	0.59
	(25.70)	(34.46)	(19.90)	(0.20)	(24.62)	(0.39)	(21.51)	(0.53)	
Present on Other Platforms	0.44	0.47	0.40	-0.07	0.48	0.00	0.42	-0.06	0.52
	(0.50)	(0.50)	(0.49)	(0.24)	(0.50)	(0.98)	(0.50)	(0.38)	
Has Earned from Other Platforms	0.18	0.20	0.20	-0.01	0.19	-0.01	0.14	-0.06	0.56
	(0.39)	(0.40)	(0.40)	(0.89)	(0.39)	(0.80)	(0.35)	(0.22)	
High Type	0.50	0.47	0.53	0.06	0.49	0.02	0.49	0.02	0.86
	(0.50)	(0.50)	(0.50)	(0.40)	(0.50)	(0.79)	(0.50)	(0.79)	
Panel B: Profile Characteristics									
Number of Completed Jobs	0.13	0.13	0.14	0.01	0.09	-0.04	0.18	0.06	0.18
	(0.35)	(0.33)	(0.35)	(0.80)	(0.28)	(0.32)	(0.41)	(0.24)	
Hourly Rate (USD)	7.91	7.71	7.89	0.19	7.93	0.22	8.10	0.39	0.95
	(5.06)	(4.59)	(5.15)	(0.77)	(4.99)	(0.72)	(5.53)	(0.56)	
Low Income Country	0.09	0.09	0.09	-0.00	0.10	0.00	0.07	-0.03	0.80
	(0.28)	(0.29)	(0.29)	(0.99)	(0.29)	(0.97)	(0.25)	(0.44)	
Lower-Middle Income Country	0.78	0.74	0.75	0.02	0.78	0.04	0.84	0.10	0.17
	(0.42)	(0.44)	(0.43)	(0.79)	(0.42)	(0.45)	(0.37)	(0.04**)	
Upper-Middle Income Country	0.13	0.16	0.15	-0.01	0.13	-0.03	0.09	-0.07	0.32
	(0.34)	(0.37)	(0.36)	(0.89)	(0.33)	(0.45)	(0.29)	(0.11)	
Fluent in English	0.97	0.95	0.95	0.00	0.99	0.04	0.97	0.03	0.07*
	(0.18)	(0.22)	(0.22)	(0.99)	(0.09)	(0.05*)	(0.16)	(0.30)	
Portfolio	0.65	0.61	0.69	0.08	0.70	0.09	0.61	-0.00	0.28
	(0.48)	(0.49)	(0.46)	(0.20)	(0.46)	(0.16)	(0.49)	(0.97)	
Certificate	0.43	0.41	0.41	-0.00	0.50	0.08	0.37	-0.04	0.27
	(0.49)	(0.49)	(0.49)	(0.96)	(0.50)	(0.20)	(0.49)	(0.52)	
Verified Certificate	0.08	0.04	0.06	0.02	0.14	0.09	0.09	0.05	0.06*
	(0.28)	(0.20)	(0.24)	(0.55)	(0.35)	(0.01**)	(0.29)	(0.13)	
Number of Skills	11.12	11.48	10.74	-0.73	11.28	-0.20	10.97	-0.51	0.43
	(3.72)	(3.65)	(3.74)	(0.13)	(3.66)	(0.66)	(3.83)	(0.29)	
Employment History	0.95	0.97	0.97	0.01	0.95	-0.01	0.92	-0.05	0.23
	(0.21)	(0.18)	(0.16)	(0.69)	(0.22)	(0.58)	(0.28)	(0.11)	

Notes: This table summarizes personal (panel A) and profile (panel B) characteristics of all participants in the freelancer experiment. Personal characteristics are measured from the baseline survey and the data entry task. Profile characteristics are public information collected from freelancer profiles on the platform. We show means and standard deviations within treatment arms as well as coefficients and p-values on the treatment indicators from the regression: $Y_i = \alpha + \beta_1 \text{Info}_i + \beta_2 \text{Signal}_i + \beta_3 (\text{Info+Signal})_i + \eta_r + \varepsilon_i$. The last column reports a joint test of mean values across treatment groups. $p < 0.10^{*}, p < 0.05^{**}, p < 0.01^{***}$.

Appendix Table A6: Heterogeneity in Treatment Effects by Workers' Baseline Beliefs

Dependent Var: Wage Undercutting	Employer Think Low Price=Low Quality		High-Type Should Undercut		Undercut Works Better for High-Type		Certainty in Own Assessment	
	Disagree (1)	Agree (2)	Disagree (3)	Agree (4)	Disagree (5)	Agree (6)	Decreased (7)	Increased (8)
Panel A: Feedback Only versus Control								
Treated	-0.040	0.131	0.062	-0.069	0.007	-0.049	-0.012	-0.056
	(0.081)	(0.105)	(0.073)	(0.094)	(0.076)	(0.111)	(0.076)	(0.141)
High-Type	0.051	-0.091	0.050	-0.069	0.018	-0.071	-0.010	-0.056
	(0.088)	(0.080)	(0.068)	(0.094)	(0.081)	(0.100)	(0.077)	(0.110)
Treated × High-Type	0.262	-0.051	-0.041	0.373	0.102	0.309	0.158	0.286
	(0.116**)	(0.131)	(0.096)	(0.131***)	(0.108)	(0.148**)	(0.103)	(0.183)
Omitted Mean	0.040	0.091	0.000	0.111	0.034	0.111	0.069	0.056
No. Obs	110	78	84	104	109	79	142	46
Panel B: Employer Info Only versus Control								
Treated	0.240	0.088	0.217	0.122	0.201	0.099	0.175	0.069
	(0.099**)	(0.094)	(0.088**)	(0.098)	(0.086**)	(0.111)	(0.085**)	(0.120)
High-Type	0.051	-0.091	0.050	-0.069	0.018	-0.071	-0.010	-0.056
	(0.102)	(0.099)	(0.091)	(0.104)	(0.101)	(0.105)	(0.090)	(0.112)
Treated × High-Type	-0.170	0.130	-0.172	0.079	-0.078	0.075	-0.048	0.112
	(0.138)	(0.136)	(0.126)	(0.140)	(0.129)	(0.159)	(0.118)	(0.173)
Omitted Mean	0.040	0.091	0.000	0.111	0.034	0.111	0.069	0.056
No. Obs	103	95	84	114	122	76	151	47

Notes: This table examines the heterogeneity in feedback and employer info effects on wage undercutting by workers' baseline beliefs. Panel A reports the heterogeneous treatment effects for feedback-only and control groups and Panel B on employer info-only and control groups. All regressions control for round fixed effects. The omitted group is low-type workers in the control group. Standard errors are clustered at the job level. $p < 0.10^{*}, p < 0.05^{**}, p < 0.01^{***}$.

Appendix B Proposition Proofs

B1 Proof of Proposition 1

Even though employers cannot observe the type of novice workers prior to hiring, high-type novices may distinguish themselves from low-type peers by proposing lower wages in period 1 as an investment in their reputation to generate repeated hiring on the platform. That is, high-type novices offer wages w_{Hj} to maximize their lifetime earnings:

$$\max_{w_{Hj}} \Pi_H(w_{Hj}) = (w_{Hj} - c_H) + \beta(\bar{w}_j - c_H)$$

The optimization is subject to the following constraints:

$$\begin{cases} \Pi_L(w^*_{Hj}) = (w^*_{Hj} - c_L) \leq 0 & \text{Incentive Compatibility (IC)} \\ \Pi_H(w^*_{Hj}) = (w^*_{Hj} - c_H) + \beta(\bar{w}_j - c_H) \geq 0 & \text{Participation Constraint (IR)} \end{cases}$$

High-type novices set wages equal to the low-type's outside option to ensure that at w^*_{Hj}, low-type novices will incur a loss in lifetime earnings, but not themselves. Together, these two constraints imply that wage separation by types can be sustained if for high-type novice workers, returns to reputation for high-type novices exceeds the cost differential between types: $\beta(\bar{w}_j - c_H) \geq c_H - c_L$.

Consistent with this wage-setting strategy, employers can infer high quality from low wages proposed by novices such that $\gamma'(w_{ij}) < 0$. No costly screening is required. Hence, in equilibrium, employers hire all high-type novices at $w^*_{Hj} = c_L$ and veterans.

B2 Proof of Proposition 2

Now suppose employers interpret higher proposed wages as signals of higher quality, even in the absence of other credible signals: $\gamma(w_{ij})$ is increasing in w_{ij}.

Under this condition, demand for novices is less responsive to proposed wages than demand for veterans. If workers offer lower wages, this reduces hiring costs for both novices and veterans, but also reduces perceived returns for novices only. Whether proposing lower wages improves novices' hiring chances depends on the relative strength of the cost and perceived quality channels.

When the cost channel dominates, high-type novice workers can still increase their hiring probability by proposing lower wages as in the standard case. Once employers hire novice workers who undercut and discover their type, they can update their beliefs about $\gamma(w_{ij})$ and thus, employers' perception will be corrected over time.

When quality concerns dominate, novice workers avoid undercutting due to concerns about negative signals and propose wages equal to the job budget, $w_{ij} = \bar{w}_j$. Without variation in proposed wages, employers need to incur screening cost s_j to obtain additional quality cues λ_{ij}. Since novice workers with unknown type charge the same as veteran workers in this case, employers will always prefer the experienced ones. Nevertheless, novices may be hired if the expected returns can compensate for employers' screening costs. That is

$$\lambda_{ij} v_H + (1 - \lambda_{ij}) v_L - \bar{w}_j - s_j \geq 0$$

As employers do not observe proposed wages below $\bar{w}_j$, they cannot update their beliefs about $\gamma(w_{ij})$ over time. In equilibrium, employers hire fewer novices and average output quality is lower than the benchmark case.

B3 Proof of Proposition 3

Alternatively, suppose neither workers nor employers know worker type upon entering the market. Workers believe that they belong to the high type with probability μ_i and the low type with probability $1 - \mu_i$. Once hired, they learn their true type by observing public reviews.

Under imperfect information about their own types, novice workers propose wages w_{ij} to maximize their expected lifetime earnings given their beliefs μ_i:

$$\max_{w_{ij}} \Pi(w_{ij}, \mu_i) = w_{ij} - (\mu_i c_H + (1 - \mu_i)c_L) + \beta \cdot \mu_i(\bar{w}_j - c_H)$$

In this case, proposing lower wages no longer signals true high ability but rather reflects novice workers' belief about being the high-type. Consequently, employers no longer view proposed wages as credible signals of novices' ability and therefore incur costly screening to discern types through the additional signal λ_{ij}. Novice workers may still be hired if they can compensate for the employer's screening cost by offering lower wages. That is

$$\lambda_{ij} v_H + (1 - \lambda_{ij})v_L - w_{ij} - s_j \geq 0$$

However, due to the additional screening costs, employers hire fewer novices in equilibrium and average output quality is lower than the benchmark case.