Bidding for Reputation *

Jingyi Cui †

Job Market Paper

November 4, 2025

Please click here for the latest version

Abstract

Reputation is often important in markets for experience goods. New sellers commonly invest in reputation by offering introductory pricing or other incentives. By encouraging buyers to try new sellers, these investments generate information externalities for future buyers while diverting business from other sellers. I study reputation investment behavior by workers in the context of a large online labor platform. I show that employers value worker reputation and experience, and that new workers initially bid low wages but raise their bids after obtaining experience and public reviews. I estimate a dynamic equilibrium model where forward-looking workers bid anticipating the impact of reputation and experience on future employment outcomes. Compared to a counterfactual with bidding based only on immediate payoffs, forward-looking bidding increases the equilibrium number of reviewed workers by 52% and quadruples the number of matches on the platform. However, workers' investments remain below the social optimum. The socially optimal platform-funded subsidy for hiring new workers raises total surplus by 22% while increasing platform profit. The subsidy level that maximizes platform profit is lower, but achieves 80% of the total surplus gain.

*I'm deeply grateful to Joe Altonji, Steve Berry, Phil Haile, and Charles Hodgson for their mentorship. I also thank Costas Arkolakis, Gabriel Dias, Rebekah Dix, Johannes Hörner, Jack Liang, Chris Neilson, Katja Seim, Matthew Schwartzman, Sam Solomon, Michael Sullivan, Kevin Williams, and participants at the Yale Industrial Organization, Labor, and Econometrics workshops. I thank Freelancer.com for providing data and expertise. This work was supported by funding from the Cowles Foundation.

†Yale Department of Economics; Contact information: jingyi.cui@yale.edu.

1 Introduction

Reputation can play a central role in markets where quality is hard to assess before purchase. New sellers often invest in building reputation by offering introductory pricing or other incentives (Klein and Leffler, 1981; Bergemann and Välimäki, 1996; Villas-Boas, 2004). Such investments encourage buyers to take a chance on new sellers, generating information about seller quality that is then shared with future buyers through word of mouth or a reputation system. These responses by forward-looking sellers can help mitigate the classic information free-riding problem, whereby buyers do not internalize the information spillovers from their purchase decisions and under-experiment with new sellers (Bolton and Harris, 1999; Keller et al., 2005). However, as with any investment, if the investor cannot fully appropriate the returns, the level of investment will be inefficiently low. While sellers internalize their own gains from reputation, they do not internalize the portion that accrues to buyers. The efficiency of sellers' reputation investments therefore depends on how surplus is shared between buyers and sellers. At the same time, sellers' investments can harm their competitors, both contemporaneously and in the future, by diverting business away. Together, these forces make the welfare implications of reputation investments an empirical question.

In this paper, I study reputation investment behavior in the context of Freelancer.com, a large online labor platform where international workers bid in auctions and build reputation through employer reviews. Two additional features of this setting are noteworthy. First, in addition to reputation, workers may accumulate human capital from experience, creating another investment incentive with spillovers on future employers. Second, workers and employers are not the only players in this market. The platform charges a commission on transactions. While the commission is just a transfer from workers and employers to the platform, it is important because it potentially gives the platform an incentive to adopt policies that enhance total surplus. I develop an empirical framework incorporating the complex forces surrounding workers' reputation and human capital investments. These forces include the externalities on employers, competition among workers, and the incentives of the platform. Using this framework, I quantify the extent to which workers' forward-looking behavior, combined with privately optimal policies for the platform, aligns the market with the social optimum.

Online labor platforms are increasingly important with the rise of digitization and remote work. A recent World Bank report estimates that 4–13% of the global labor force engage in online gig work (Datta et al., 2023). On these platforms, after completing a job, a worker may receive a public rating from the employer. These ratings are informative to future employers. The number of ratings signals experience, while the mix of good and bad ratings signals quality. Because most jobs are one-off transactions, employers do not internalize the information and human capital gains that their hiring decisions generate.[1] This makes worker investment and platform policy choices particularly important for long-run market efficiency.

I use proprietary data covering three million wage bids submitted by over 100,000 international workers to 83,000 jobs, the universe of job auctions under a popular skill category in a seven-month period in 2018. The data allows me to follow workers over time as they build reputation. I first present evidence that employers value worker reputation and experience. For a given worker, controlling for wage bids and other characteristics, the probability of winning a job rises with the number of good reviews, and to a smaller extent, with bad reviews. Using data on both successful and unsuccessful bids, I show that new workers initially bid low wages. They raise their bids after receiving good reviews and do so more modestly after bad reviews.

To evaluate the effects and efficiency of workers' investments in reputation and human capital, I develop a dynamic equilibrium auction model where workers bid for jobs and employers decide whom to hire and leave reviews. The model features a discrete set of public *types* that represent persistent and commonly known worker-level heterogeneity. Allowing for persistent heterogeneity is essential for assessing true state dependence—in this case, the impact of experience and reputation (Heckman, 1981). Conditional on types, workers have latent *quality*, about which workers and employers learn symmetrically from reviews by past employers. Workers also accumulate human capital with job experience. Different types of workers differ along four dimensions: (i) opportunity costs; (ii) baseline productivity; (iii) distribution of latent quality; and (iv) rate of human capital accumulation. After workers submit their bids, employers observe workers' types, bids, past reviews, and match-specific productivity shocks. Employers form expectations of workers' productivity and choose a worker or the outside option. Anticipating the effects of reputation and experience on future outcomes, workers bid to maximize the value of lifetime payoffs. In equilibrium, workers' perceived winning probabilities coincide with their actual chances of winning.

To estimate the model, I extend existing two-step estimation methods from the literature on dynamic games. In the first step, I estimate bid policy functions and winning probabilities, allowing each to vary with the unobserved worker type. I adapt the expectation–maximization algorithm to accommodate several complexities in my model.[2] First, actions are continuous. Second, unobserved heterogeneity is multidimensional. Third, the state transition function—the probability of winning an auction—depends on not just the worker's own type but also on her

[1] High turnover among employers or consumers is a common feature of other labor platforms such as Uber, Grubhub, and TaskRabbit, as well as product marketplaces like Amazon. This paper abstracts from the more complex incentive structures that arise in traditional long-term employment relationships, including incentives to accumulate firm-specific human capital, to retain workers, or to write long-term contracts.

[2] The EM alogirthm was first used by Arcidiacono and Miller (2011) in the estimation of dynamic games with serially correlated unobservables.

competitors' types. In the second step, I apply the Method of Simulated Moments to recover employers' preferences and workers' opportunity costs by matching moments from employers' choices and workers' bids. Match-specific shocks known to workers and employers but unobserved by the econometrician can make bids endogenous. To address this, I use exchange rates to instrument for bids. International workers typically bid in USD, so when their home currencies appreciate, bids in USD rise. I further extend the E-M algorithm to incorporate these exchange rate instruments. I distinguish between learning about worker quality and human capital accumulation under the assumption that human capital accumulates at the same rate across all jobs, regardless of whether the associated reviews are good or bad. I use the first-step estimates to forward simulate continuation values in workers' bid first-order conditions.

The estimates indicate meaningful differences across worker types, highlighting the need to control for persistent unobserved heterogeneity. I find that worker types with higher baseline productivity also have higher opportunity costs. This correlation is intuitive as more capable workers are likely to have better off-platform options. Within each worker type, employers place a high value on reputation and experience. Employers are willing to pay 22-24 percentage point higher wages (different depending on worker type) for a worker with five good reviews compared to a worker with no reviews. Out of this wage premium, about half are attributable to higher expected quality and the rest to to human capital accumulation. Moreover, the estimates indicate that the informational and human capital gains from a job match often exceed the job's flow surplus and are larger for less experienced workers. The estimates also confirm that workers and employers share the value created by a review. As an example, for a model worker, 35% of the value of the first review accrues to the worker herself and 65% to her potential employers in the future.

Equipped with the estimated model, I study the effects of workers' reputation and human capital investments. To this end, I compare the baseline equilibrium to a counterfactual one where workers choose bids to maximize the payoffs from the current auction only. I find that workers' forward-looking bidding increases the equilibrium number of reviewed workers by 52%, improving the choice sets for employers and leading to three times more matches. Workers as a whole benefit from the investments. Even though each worker invests solely to maximize her own surplus, workers' forward-looking behavior creates large positive externalities for employers and the platform. Employer surplus goes up by more than four times compared to the myopic counterfactual, and platform surplus doubles. The results show that while reputation systems create the potential for information flow, workers' (or sellers') costly investments in reputation—such as through introductory pricing—are important for realizing that potential. Absent investments to encourage employers (or buyers) to experiment, markets can stagnate in low information.

Although workers' investments generate large social gains, I find that the equilibrium level of investments remains below the social optimum. Platform-funded subsidies for hiring unreviewed workers can raise total surplus. In deciding the optimal subsidy level, the social planner faces a trade-off. Subsidizing hiring of new workers generates valuable information, but by shifting hiring away from well-reviewed incumbents it reduces employers' ability to exploit existing information. The optimal subsidy balances the dynamic information and human capital gains against the losses from reallocating hiring away from known high-productivity matches. The socially optimal subsidy would raise total surplus by 22%. While the platform sees a revenue gain under the socially optimal subsidy, it prefers a lower subsidy—20% vs. 30%. The wedge arises because the platform treats the subsidy as an expenditure, while the social planner views it as a transfer, and because the platform internalizes only a fraction of the subsidy's benefits through transaction commissions. Despite the misalignment, the platform's preferred subsidy level delivers 80% of the gains from the socially optimal subsidy. Take together, the results show that workers' optimal investment behavior, combined with a profitable platform subsidy, comes close to achieving the social optimum.

Related Literature This paper contributes to a broad literature on incomplete information and reputation across economics. It is related to theoretical work studying information free-riding when experimental outcomes are publicly observed (Bolton and Harris, 1999; Keller et al., 2005), sellers' strategic responses to reputation incentives (Bergemann and Välimäki, 1996; Bergemann and Välimäki, 2000; Villas-Boas, 2004), and mechanism design to improve social learning (Kremer et al., 2014; Che and Hörner, 2018; Vellodi, 2018). I contribute by combining data from online settings with tools from industrial organization to study these questions empirically.

This paper is also related to the empirical literature on reputation mechanisms in online product markets, especially with regards to the cold-start problem and to sellers' responses (Bolton et al., 2004; Cabral and Hortacsu, 2010; Li et al., 2020). Closely related is Dendorfer and Seibel (2024), who estimate the cost of inefficient experimentation with new listings on Airbnb and design platform interventions. My paper differs by focusing on the role that sellers' pricing plays in easing the cold start problem. It also extends these analyses to a labor context, where unobserved heterogeneity and human capital accumulation are essential features.

In labor, this paper bridges two strands of literature that have previously been studied separately. The first is the literature on employer learning and general skills training (e.g. Altonji and Shakotko, 1987; Farber and Gibbons, 1996; Acemoglu and Pischke, 1998; Altonji and Pierret, 2001; Autor, 2001; Kahn and Lange, 2014; Pallais, 2014; Barlevy and Neal, 2019). The second concerns labor market power (e.g. Manning, 2003; Card et al., 2018; Dube et al., 2020; Azar et al., 2022; Lamadon et al., 2022; Berger et al., 2022; Rubens, 2023; Roussille and Scuderi, 2025).

In closely related work, Pallais (2014) shows experimentally that employer reviews improve workers' employment outcomes on a similar platform. She then argues using a theoretical model that the hiring of new workers is inefficiently low due to the minimum wage. I innovate on this work theoretically by showing that the inefficiency result does not rely on the minimum wage and follows naturally once we relax the assumption of perfect competition. I also provide an equilibrium empirical model that allows for evaluation of alternative platform design.

Methodologically, the paper contributes to the industrial organization literature on two-step estimation of dynamic games (Rust, 1987; Hotz and Miller, 1993; Hotz et al., 1994; Bajari et al., 2007). I extend the expectation-maximization approach of Arcidiacono and Miller (2011) who first used it to estimate dynamic games with serially correlated unobservable (Chung et al., 2014; Igami and Yang, 2016). My extension allows for the use of an instrumental variable, for multidimensional unobserved heterogeneity (drawing on insights from Arcidiacono and Jones (2003)), and for an outcome to depend on the unobserved heterogeneity of multiple agents. The first-step policy functions are identified following Kasahara and Shimotsu (2009), the large-market equilibrium concept builds on the Oblivious Equilibrium of Weintraub et al. (2008), and the counterfactuals are computed using the policy function iteration method of Sweeting (2012).

This paper is also related to work on online labor platforms (e.g. Horton et al., 2011; Agrawal et al., 2015; Stanton and Thomas, 2016; Filippas et al., 2018; Barach et al., 2020; Barach and Horton, 2021; Holtz et al., 2022; Krasnokutskaya et al., 2020; Galdin and Silbert, 2024). Following Brinatti et al. (2021) and Stanton and Thomas (2025), I exploit exchange-rate fluctuations for identification. The focus on workers' forward-looking incentives in bidding and the estimation of a dynamic model are new to this literature.

Outline Section 2 introduces the setting and data. Section 3 provides preliminary evidence. Section 4 presents a dynamic equilibrium model. Section 5 discusses estimation method. Section 6 presents estimation results. Section 7 studies the equilibrium effects of workers' investments. Section 8 evaluates alternative platform design. Section 9 concludes.

2 Setting and Data

2.1 Setting

The setting of my study is Freelancer.com, one of the world's largest online labor platforms that matches global employers and workers to collaborate on short-term, often remote, jobs. According to its 2024 Annual Report, the platform has over 80 million registered users.[3] The

[3] All of the information in this paragraph is based on Freelancer's 2024 Annual Report.

top countries for employers are the United States, India, Australia, the United Kingdom, and Canada. Workers are similarly international. In 2024, Freelancer facilitated $85 million worth of contracts, with the average job paying $334 in the last quarter. The most popular job categories were design, media, and architecture (33% of completed projects), followed by websites, IT, and software (29%), writing and content (9%), and sales and marketing (7%).

Online labor platforms such as Freelancer, Upwork, and Fiverr meet businesses' demand for skilled, project-based work. These platforms also expand income opportunities for workers, especially those in less developed countries or excluded from traditional labor markets due to gender, disability, or other barriers (WTO, 2019; ILO, 2021). Market research firm Modor Intelligence estimates that the freelance platform sector reached a market size of $7.6 billion in 2025.

The primary mode of matching on Freelancer is through job auctions. Employers post public job descriptions under specific skill categories and specify a minimum budget (most jobs in my data are fixed-price as opposed to hourly). Although employers could, in principle, set minimum budgets strategically, the evidence suggests otherwise. The platform offers a drop-down menu with preset options such as "Task $10–30" and "Micro $30–50." Appendix Figure A2 shows that most minimum budgets cluster at these preset values, suggesting the lack of strategic behavior. Accordingly, throughout the paper, I use the minimum budget to normalize workers' bids and costs.

Workers submit bids that include a proposed pay at or above the minimum budget, an expected completion time, and a short description explaining their fit for the job. Workers' prior reviews—in particular, their average star rating (from 1 to 5) and total number of ratings—are prominently displayed alongside their bids. Employers can click on workers' profiles to view the full review histories. Appendix Figure A1 displays a screenshot of a posted job and selected bids.

The default order in which workers' bids are displayed on the employer's page is based on a proprietary algorithm that accounts for ratings from previous employers, skills, profile, and past experience.[4] Employers can re-sort bids in various ways (e.g. by bid amount). Employers are free to reach out to workers for a private conversation before deciding whom to hire, and to award the job to any or none of the bidders.

Bids become immediately visible to all participants upon submission. Bidding remains open until the employer awards the job or the posting expires. Appendix Figure A3 shows the distribution of auction durations. Although job auctions are by default open for seven days, most jobs in my sample—when awarded—are awarded within two hours of posting. The fact that bids arrive rapidly after a job is posted motivates my later modeling of the auction as following a sealed-bid format. While workers can, in principle, observe previously submitted

[4] https://www.freelancer.com/support/freelancer/profile/bid-ranking-factors

bids, in practice they are unlikely to condition on them given the speed of bidding.

After job completion, both parties rate each other, but can only see the rating they receive after submitting their own.

2.2 Data and Summary Statistics

I use proprietary data from Freelancer obtained through a research collaboration agreement. The primary dataset consists of the universe of job auctions posted under the popular skill category PHP—a programming language—from June 1 to December 31, 2018. I restrict the sample to fixed-payment auctions, which account for 89% of jobs in this sample, and further to those denominated in USD (59% of fixed-payment auctions). The final dataset includes 2,738,310 bids submitted to 83,210 jobs, with an average of 33 bids per job.

For each job, I observe the minimum budget and posting timestamp. For each bid, I observe the worker ID, bid amount, and whether the bid was selected. If the employer awarded the job to a worker, I observe whether the employer left a rating and the content of the rating. I also observe the workers' countries. Appendix Section A contains a list of variables.

In addition, I obtained the full review history for all the workers in my sample. The review data include the timestamp and rating of each review received. Appendix Figure A5 shows that about 80% of the ratings in my sample are five out of five. Given the skew towards five-star ratings, I treat reviews as binary signals: a five-star rating is classified as a good signal, and any lower rating as a bad signal. I merge the review history with the auction panel to reconstruct each worker's numbers of good and bad reviews at the time of bidding.

I supplement the auction panel with daily exchange rates. Theses are daily interbank exchange rates that Freelancer uses to process payments.

Table 1 presents summary statistics of the main data set. The median job in my data sample has a minimum budget of $30, while the mean is considerably higher at $230. A subset of jobs are awarded.[5] Among awarded jobs, the median winning bid is $61 or about 2.5 times the minimum budget. Jobs are typically short-term, with a median completion time of 40 hours. Appendix Figure A4 presents a further breakdown of jobs' duration.

Most bids come from experienced workers. The median bid is submitted by a worker with 23 good reviews and 4 bad reviews. The average worker in the sample accounts for 25 bid observations, while the median worker has significantly fewer, indicating a skewed distribution of bidding activity.

An auxiliary dataset tracks a random sample of workers who registered around the same time in 2018 and records all bids they submitted across skill categories. I restrict to a subsample

[5] According to the research collaboration agreement, I cannot disclose the project award rate.

Table 1: Summary Statistics

(a) Auction Panel

	Mean	p25	p50	p75
Project-Level ($N=83{,}210$)				
Number of bidders	33	10	22	44
Minimum budget (usd)	230	10	30	250
Winning bid (usd)	213	25	61	194
Winning bid / min. budget	3.4	1.5	2.5	3.6
Paid amount (usd)	228	20	50	180
Auction duration (hour)	71	.4	2.0	31
Job duration (hour)	862	1.1	40	281
Bid-Level ($N=2{,}738{,}310$)				
Bid / min. budget	3.3	1.5	2.2	5.0
Bidder's number of good reviews	81	4	23	93
Bidder's number of bad reviews	15	0	4	15
Bidder-Level ($N=109{,}919$)				
Number of bids in sample	25	1	2	4

(b) Worker Panel

	Mean	p25	p50	p75
Bidder-Level ($N=8{,}973$)				
Number of total bids	596	16	53	194
Active months	22	2	15	37

Notes: Tables present statistics from (i) the main dataset, an auction panel of all fixed-budget, USD-denominated job auctions in the popular PHP skill category (a web programming language) from June–December 2018, and (ii) an auxiliary dataset comprising the complete bid histories across skill categories of a random sample of workers who registered in 2018 and won at least one job. Active months are the number of months between a worker's first and last bids. For the summary statistics of bids, bids above $\exp(3)$ times of minimum budget are excluded. The statistics are based on my data samples and do not reflect the platform averages.

of workers who won at least one job on the platform, which contains 8,973 workers. Among these workers, the mean total number of bids submitted is 596 and the median is 53. The amount of time between the worker's first and last bids has a mean of 22 months and median of 15 months. These statistics suggest that successful workers are active for a long period of time on the platform, and it is reasonable to assume that they have dynamic incentives.

The vast majority, or 89%, of jobs done by workers in the worker panel involve employer–worker pairs that collaborated only once. This pattern is consistent with the freelance nature of the platform, where employers typically hire for one-off jobs. Consequently, employers have limited incentives to invest in learning workers' quality over time or training workers, and instead rely on the reviews and experience that workers received from prior employers.

3 Preliminary Evidence

I provide evidence that employers value worker reputation and experience and that workers change their wage bids after obtaining experience and public reviews.

3.1 Employers Value Reputation and Experience

I show that, conditional on the bid, a given worker's probability of winning increases with the number of good reviews and, to a smaller extent, with the number of bad reviews. This suggests that employers are willing to pay for reputation and human capital.

I use the main dataset to examine the relationship between winning and a worker's numbers of good and bad reviews, controlling for worker fixed effects, bids, and other characteristics. Unobserved match-level productivity shocks known to both workers and employers—for example, through pre-award communication—can render bids endogenous. To address this, I use exchange rates as an instrument for bids. Exchange rate fluctuations shift international workers' wage bids in USD. When a worker's home currency appreciates relative to the USD, her bid in USD increases. The exchange rate–induced variation in bids should be orthogonal to changes in match-specific productivity.

Let b_{it} denote worker i's bid for job t divided by the job's minimum budget; $y_{it}=1$ if worker i won job t and 0 otherwise; z_{it} the log of worker i's local currency per USD exchange rate demeaned at the country level; $n_{it}=(n_{it}^{\text{Good}}, n_{it}^{\text{Bad}})$ the numbers of good and bad reviews at the time of bidding; and x_t the job's minimum budget.

The first stage is

$$\log b_{it} = \delta_i + \pi_z z_{it} + \pi_G \log n_{it}^{\text{Good}} + \pi_B \log n_{it}^{\text{Bad}} + \pi_x \log x_t + \epsilon_{it} \tag{1}$$

and the second stage is

$$y_{it} = \alpha_i + \gamma_b \log b_{it} + \gamma_G \log n_{it}^{\text{Good}} + \gamma_B \log n_{it}^{\text{Bad}} + \gamma_x \log x_t + \varepsilon_{it}. \tag{2}$$

Table 2 presents the estimation results. Column (1) reports the OLS estimates, with all values normalized by the sample mean of the winning probability. For a given worker, the probability of winning declines with log bid and, somewhat counterintuitively, with the number of good

Table 2: The Effects of Reviews on Winning

	(1) Winning: OLS	(2) Bid (log)	(3) Winning: IV
Bid (log)	-0.76***		-16.4***
	(0.013)		(1.9)
No. of good reviews (log)	-0.43***	0.073***	0.66***
	(0.04)	(0.0019)	(0.14)
No. of bad reviews (log)	0.06	0.015***	0.22***
	(0.05)	(0.0024)	(0.06)
Min. budget (log)	-0.57***	-0.35***	-6.04***
	(0.006)	(0.00024)	(0.68)
Exchange rate (log)		0.17***	
		(0.013)	
Observations	2,684,310	2,684,310	2,684,310
Worker FE	Yes	Yes	Yes

Notes: Table reports the within-worker relationship between winning and reviews received. Bids are divided by jobs' minimum budgets. Column (1) presents OLS estimates. To address potential endogeneity in bids, I instrument for log bids using the log of the worker's local-currency-per-USD exchange rate demeaned at the country level. Column (2) reports the first-stage results, and Column (3) the second stage. The estimates in Columns (1) and (3) are normalized by the sample mean winning probability.

reviews, while the coefficient on the number of bad reviews is imprecisely estimated. These estimates are likely biased due to the presence of match-specific unobservables. Column (2) shows estimates of the first-stage regression of log bid on the exchange rate instrument. I estimate an exchange rate pass-through to workers' wage bids of 0.17, which is close to the 0.20 estimate of Brinatti et al. (2021) and above the 0.08 estimate of Stanton and Thomas (2025), both from similar settings.

Column (3) reports the 2SLS results using exchange rates as an instrument for log bids, where all values are similarly normalized by the mean winning probability. Relative to the OLS results, the coefficient on log bid becomes significantly more negative, consistent with the presence of match-specific unobservables. The results indicate that workers tend to bid higher on jobs for which they have higher match-specific productivity. Consequently, failing to instrument for bids leads to an underestimation of the negative effect of bids on winning probabilities. The coefficient on the log number of good reviews is positive. The first good review increases the likelihood of winning by 46%, while the first ten good reviews raise it by 158%. The effect of the log number of bad reviews is also positive, but smaller than that of good reviews. Two lessons are drawn. First, reviews raise a worker's probability of winning regardless of the information signal, consistent with a theory of general human capital accumulation with experience. Second, good reviews raise winning probability more than bad reviews, consistent with a theory of reputation—good reviews provide more positive information about workers' latent quality.

3.2 Workers Increase Bids after Receiving Reviews

I show that workers increase their wage bids after building experience and receiving reviews.

I begin by regressing normalized wage bids (normalized by jobs' minimum budgets) on quintiles of the number of good reviews, controlling for the number of bad reviews and the job's minimum budget.[6] Figure 1 plots the coefficients on the quintile indicators. Bids rise with the number of good reviews. Moving from the first to the fourth quintile increases bids by over 30% of the job's minimum budget, and moving to the fifth quintile raises bids by more than 50%.

Figure 1: Wage Bids and the Number of Good Reviews

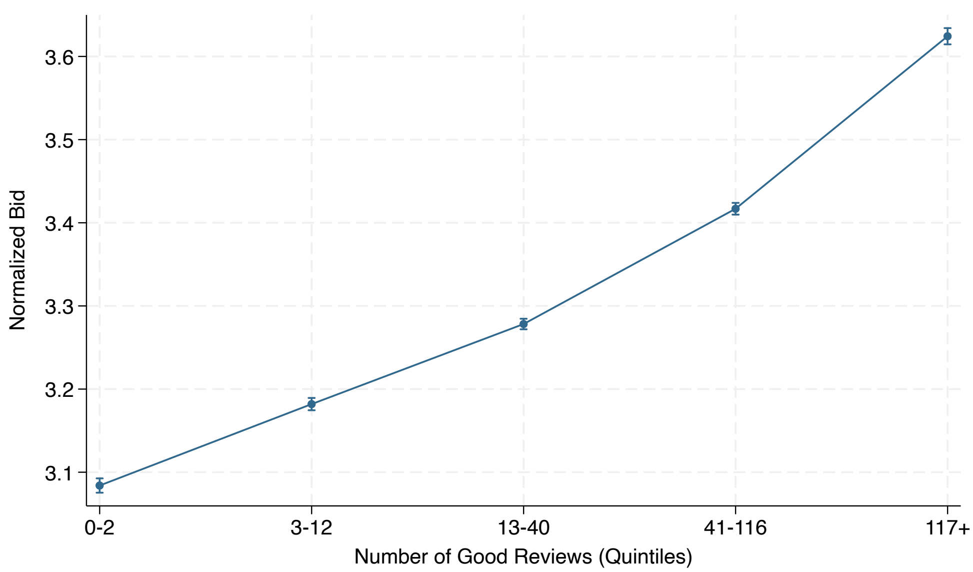

Notes: Figure plots the average wage bid normalized by the job's minimum budget against quintiles of the number of good reviews, holding the number of bad reviews and the minimum budget at their mean values.

To study within-worker bid changes, I regress the worker's log normalized bid on the numbers of good and bad reviews, controlling for the job's minimum budget and worker fixed effects:

$$\log b_{it} = \alpha_i + \beta_G \log n_{it}^{\text{Good}} + \beta_B \log n_{it}^{\text{Bad}} + \beta_x \log x_t + \epsilon_{it}.$$

[6]For all the analyses in this subsection, wage bids above 30 times the minimum budget are dropped from the calculation, as they are unlikely to reflect strategic bidding behavior.

The results are presented in Table 3. Column (1) reports estimates from a regression without worker fixed effects, and Column (2) includes worker fixed effects. Column (1) shows that bids are higher from workers with more good reviews, and higher to a smaller extent from those with more bad reviews. Column (2) shows that for a given worker, she raises her bid following a good review and raises her bid to a smaller extent following a bad review. The first good review is associated with a 5% increase in the wage bid, while the first ten good reviews correspond to a 17% increase. The first bad review is associated with an increase of less than 1%.

Bids could increase following bad reviews for at least two reasons. First, the worker may become more pessimistic about her own quality and decide to invest less for future jobs, because they anticipate receiving more negative reviews. Second, the worker may learn from the experience despite receiving a bad review and become more productive. The model will formalize these channels.

Comparing Column (2) with Column (1), I find that wage bids are more responsive to the number of good reviews when worker fixed effects are included. This pattern is consistent with workers who have accumulated more good reviews generally bidding lower, reflecting persistent worker heterogeneity and the endogeneity of reviews.

Table 3: The Effects of Reviews on Wage Bids

	(1) Bid (log)	(2) Bid (log)
No. of good reviews (log)	0.0429***	0.0699***
	(0.000454)	(0.00185)
No. of bad reviews (log)	0.0124***	0.0101***
	(0.000594)	(0.00234)
Min. budget (log)	-0.256***	-0.351***
	(0.000244)	(0.000243)
Observations	2,738,310	2,684,310
Worker FE	No	Yes

Notes: Table presents the relationship between bids and reviews. Bids are normalized by jobs' minimum budgets.

A worker may adjust her bid after receiving reviews either because of standard static incentives to maximize current auction revenue or due to changes in her dynamic incentives to invest in reputation and experience. In the following sections, I use a structural model to disentangle these static and dynamic components of bidding incentives and to examine how workers' strategic bidding behavior shapes equilibrium outcomes.

4 Model

Motivated by the preliminary evidence, I develop a dynamic equilibrium model to study the effects and efficiency of workers' investments in reputation and human capital. The model includes symmetric Bayesian learning about worker quality, human capital accumulation, bidding for jobs, and hiring. The model builds on the general class of dynamic game models introduced by Ericson and Pakes (1995), with innovations to accommodate a large number of market participants and to account for heterogeneity at both the worker and match levels.

Environment I develop a continuous-time model with infinite horizon (Arcidiacono et al., 2016; De Silva et al., 2023). The high-frequency nature of the setting makes continuous-time models a natural fit.

Workers, each denoted by i, are characterized by their country $o(i)$ and public *type* $k(i) \in \{1,...,K\}$. The public type is common knowledge and determines the worker's productivity and opportunity cost in ways described below. The state variables are the numbers of good and bad reviews the worker has received by time t: $n_{it} = (n_{it}^{\text{Good}}, n_{it}^{\text{Bad}})$. The publicly visible number of job experiences is $(n_{it}^{\text{Good}} + n_{it}^{\text{Bad}})$.[7]

At each point in time, new workers enter the market exogenously at Poisson rate λ, each starting with zero reviews. Job auctions arrive at Poisson rate λ_a. Jobs are posted by different employers, so employers face no dynamic considerations—an assumption consistent with the platform's freelance nature and statistics presented in Section 2.2. In what follows, t refers to time, the job auction posted at that time, and the employer posting the auction.

Worker i participates in an auction with an exogenous probability that can depend on the worker's review state $a(n_{it})$. Afterwards, the job's minimum budget x_t, exchange rates z_{ot}, and match-specific shocks ν_{it} realize, following the distributions described below. Observing these, participating workers submit total-payment bids. Let b_{it} represent the worker's bid divided by the job's minimum budget.

The employer is free to award the job to one or none of the bidders. If the job is awarded, the employer leaves a review with exogenous probability r, updating the hired worker's reputation state. Finally, workers exit the market according to a Poisson rate that may depend on their review state, $\delta(n_{it})$.

[7] Because employers do not always leave a review after a job, the worker's true number of jobs done on the platform may exceed the visible total. I assume that without a posted review, the worker cannot credibly signal that the job occurred. The human capital accumulation defined below should therefore be interpreted as reflecting human capital growth from both reviewed and unreviewed experiences. The model implies that the expected number of unreviewed jobs rises with the number of reviewed ones.

Review Workers have latent *quality* $\tilde{q}_i$. The prior over the latent quality follows a Beta distribution with parameters that depend on the worker's public type $\tilde{q}_i \sim \text{Beta}(a_{k(i)}, b_{k(i)})$. The latent quality is unknown to the worker and employers. They learn about it symmetrically through reviews provided by past employers.

A worker's performance on the job sends a binary signal of their quality to the employer. The probability that the signal is good is equal to the worker's quality. The employer leaves the signal as a public review with exogenous probability r.[8] Based on Bayes's theorem, the expectation that the worker and employers hold of the worker's latent quality, given her reviews and public type, is

$$\mathbb{E}[\tilde{q}_i | n_{it}, k(i)] = \frac{a_{k(i)} + n_{it}^{\text{Good}}}{a_{k(i)} + b_{k(i)} + n_{it}^{\text{Good}} + n_{it}^{\text{Bad}}}.$$

Employer's problem The employer receives workers' bids and chooses one of the bidders or the outside option to maximize her utility. The employer faces a tradeoff between wages and expected productivity. Her utility of hiring worker i for the job at time t is

$$u_{it} = \alpha_b \log b_{it} + q(k(i), o(i), n_{it}, \nu_{it}, x_t) + \epsilon_{it}, \tag{3}$$

where b_{it} is the normalized wage bid. The expected worker productivity $q(k(i), o(i), n_{it}, \nu_{it}, x_t)$ equals

$$\underbrace{q_{o(i)} + q_{k(i)} + q_{k(i),x} \log x_t}_{\text{baseline productivity}} + \underbrace{\alpha_q \mathbb{E}[\tilde{q}_i | n_{it}, k(i)]}_{\text{expected latent quality}} + \underbrace{h_{k(i)} \log(n_{it}^{\text{Good}} + n_{it}^{\text{Bad}} + 1)}_{\text{human capital accumulation}} + \underbrace{\nu_{it}}_{\text{match-specific shock}}. \tag{4}$$

The first terms $q_{o(i)} + q_{k(i)} + q_{k(i),x} \log x_t$ represent the average expected productivity of a new type k worker from country o. The effect of the worker's type can differ depending on the job's minimum budget x_t. The minimum budget is independent over time and drawn from an exogenous distribution. The match-specific shock ν_{it} is common knowledge and is i.i.d. over i,t. The influence of the expected latent quality is $\alpha_q \mathbb{E}[\tilde{q}_i | n_{it}, k(i)]$. And $h_{k(i)} \log(n_{it}^{\text{Good}} + n_{it}^{\text{Good}} + 1)$ captures the human capital accumulation. The speed of human capital accumulation can differ by worker's type. Finally, ϵ_{it} is the employer's private taste shock, distributed i.i.d. over i,t following T1EV distribution. The employer's outside option is normalized to have mean 0: $u_{0t} = \epsilon_{0t}$.

[8] The assumption that the employer reviews a worker's latent quality rather than her total productivity aligns with the theory of expectation (dis)confirmation, which posits that consumers' evaluations reflect how experiences meet or fall short of prior expectations. Here the prior expectations are given by all the components of the worker's expected productivity in Equation 4 except the expected latent quality. Recent empirical evidence in support of the theory includes Li et al. (2025) on restaurant ratings, Bondi et al. (2025) on movie ratings, and Meister and Reinholtz (2025) on Airbnb.

Worker i's winning probability given a set of bidders B_t $(i \in B_t)$ is, therefore,

$$\frac{\exp(\alpha_b \log b_{it} + q(k(i), o(i), n_{it}, \nu_{it}, x_t))}{1 + \sum_{l \in B_t} \exp(\alpha_b \log b_{lt} + q(k(l), o(l), n_{lt}, \nu_{lt}, x_t))}. \tag{5}$$

Worker's problem Workers bid to maximize the sum of present and future payoffs. I assume that worker i's opportunity cost of working on job t, expressed in USD and divided by the job's minimum budget, follows

$$c_{it} = (c_{k(i)} + c_{o(i)} + c_x \log x_t + c_{k(i),\nu} \nu_{it}) \times (z_{o(i)t})^{c_z}, \tag{6}$$

where $c_{k(i)} + c_{o(i)}$ is the baseline opportunity cost for type-k workers from country o; x_t is the minimum budget; ν_{it} is the match-specific shock whose impact on opportunity costs can be type-specific; and $z_{o(i)t}$ is the worker's local currency per USD exchange rate normalized at the country level. I assume the exchange rates z_{ot} follow country-specific distributions F_Z^o, and are i.i.d. over time within the time window of my data. If $c_z = 1$, that means complete pass-through from exchange rate shocks to workers' opportunity costs in USD, consistent with the case where workers' outside options are entirely domestic. If workers' outside options include working for foreign employers in other channels, the pass-through parameter c_z would be less than 1.

I summarize the four dimensions along which workers of different public types differ: (i) opportunity costs; (ii) baseline productivity; (iii) the latent quality distribution; and (iv) the speed of human capital accumulation.

When a worker decides how much to bid, she does not know the identities or the number of her competitors.[9] Workers hold beliefs about their chances of winning against the distribution of competitor sets. I denote a particular bidder set that includes worker i by B. Worker i's belief of her winning probability is given by

$$G_{o(i)}^{k(i)}(b_{it}, n_{it}, x_t, \nu_{it}) = \int \frac{\exp(\alpha_b \log b_{it} + q(k(i), o(i), n_{it}, \nu_{it}, x_t))}{1 + \sum_{l \in B} \exp(\alpha_b \log b_{lt} + q(k(l), o(l), n_{lt}, \nu_{lt}, x_t))} f(B) dB. \tag{7}$$

The probability of a set of competitors $f(B)$ is determined by the auction participation probabilities of all potential competitors. It is is given by $\Pi_{l \neq i, l \in B} a(n_{lt}) \Pi_{l \notin B} (1 - a(n_{lt}))$. Equation 7 mimics the distribution of the highest rival bid, an object often used in the literature on sealed-bid auctions to describe the bidder's problem. The difference is that in my context, employers care about non-wage worker characteristics, so the winning probability is not simply a function of

[9] In practice, workers can observe the identities and bids of earlier bidders. However, 50% of all bids arrive within 10 minutes of a project's posting and 75% within the first two hours, making it unlikely that workers carefully review prior bidders' profiles or tailor their bids accordingly. It is therefore reasonable to model workers as bidding against an average, expected set of competitors.

the distribution of the highest rival bid.

Given these beliefs, each worker selects a bid to maximize the sum of the present and future payoffs. Let $V_o^k(n)$ represent the value of being in the market for a worker of public type k from country o and with reputation state n. Because I assume that the job's minimum budget x_t, exchange rates z_{ot}, and match-specific shocks ν_{it} are distributed independently over time, the model implies that conditional on k and o, the reputation state n is the only state variable.[10]

I assume no discounting of future payoffs.[11] The Hamilton-Jacobi-Bellman equation is

$$\delta(n)V_o^k(n) = \lambda_a a(n)\mathbb{E}_{z,x,\nu}\left[\max_b \overbrace{G_o^k(b,n,x,\nu)}^{\text{expected winning prob}} \left(x(b-c)+r\mathbb{E}[V_o^k(n'|n)-V_o^k(n)]\right)\right],$$

where the left-hand side is the instantaneous loss from exiting the market and the right-hand side the worker's expected gain from participating in an auction—including both the auction profit and the change in the worker's continuation value as a result of a new review—multiplied by the instantaneous auction participation probability.

The bid first-order condition is

$$b = c \underbrace{- \frac{G_o^k(b,n,x,\nu)}{\partial G_o^k/\partial b}}_{\text{static markup}} \underbrace{- \frac{1}{x} r\mathbb{E}[V_o^k(n'|n)-V_o^k(n)]}_{\text{dynamic incentives}}. \tag{8}$$

The static markup arises from employers' private taste shocks and from heterogeneity in expected worker productivity, which reflects worker types, information contained in reviews, human capital accumulated through experience, and match-specific productivity shocks. In addition, workers respond to dynamic incentives. If they expect that winning a job will increase their continuation values on the platform, they will shade their bids down in response.

Equilibrium I consider a large market equilibrium concept similar to an oblivious equilibrium, which is an approximation to the Markov-perfect equilibrium in dynamic games with many agents (Weintraub et al., 2008; Iyer et al., 2014).[12] This equilibrium concept fits my setting

[10]Here I ignore the potential dependence of the value function on the number and state variables of other workers, because as mentioned below, the equilibrium concept features a stationary distribution of workers. In equilibrium, the worker's belief of winning is averaged over this stationary distribution and does not depend on the time t realization.

[11]This is driven by the observation that most workers in the worker panel sample are active in the market for less two years. However, discounting can be easily incorporated into the model and estimation.

[12]The equilibrium concept is similar to OE because workers do not track the time t realization of the number of competitors and their review states; rather, workers bid against the stationary distribution of competitors. It is not exactly OE because of aggregate shocks from exchange rate fluctuations. However, because exchange rate shocks are assumed i.i.d., the distribution of workers' states is still stationary conditional on exchange rate realizations.

because the worker population exceeds 100,000, making it unrealistic for each worker to track every other worker's reputation state. Instead, workers condition on an aggregate statistic—the expected winning-probability function—similar to Backus and Lewis (2025). An equilibrium is a pair of bid strategies $b_o^k(n,z,\nu,x)$ and expected winning probabilities $G_o^k(b,n,x,\nu)$ such that

1. Workers' bids maximize the sum of present and future payoffs, given workers' expected winning probabilities;

2. Workers' expected winning probabilities are equal to the average winning probabilities in a stationary distribution where workers follow the bid strategies and employers choose the option that generates the highest utility.

5 Estimation

To estimate my model, I adapt the two-step method for estimation of dynamic games (Rust, 1987; Hotz and Miller, 1993; Hotz et al., 1994; Bajari et al., 2007). The key estimation challenge is the unobservability of workers' types k, which complicates the first-step estimation. I address this by applying and extending the Expectation-Maximization (EM) algorithm, first used by Arcidiacono and Miller (2011) in the context of dynamic games with serially correlated unobservables. Estimates from the first step are then used in a second step, where I estimate parameters of employers' demand and workers' opportunity costs.

5.1 First-Step Estimation via the EM Algorithm

In the first step, I use maximum likelihood to estimate the population distribution of types (π_k) and the type-specific bid policy function, winning probability function, prior over latent quality, and distribution of initial conditions.[13] These are collectively represented by the parameter vector θ.

Let $d_t^{-i_t^*}$ denote the vector of participation decisions in auction t for all workers other than the winner i_t^*. If the employer selects the outside option, then $d_t^{-i_t^*} = d_t$.

[13]The parameters of the bid policy functions and the distributions of the initial conditions are not structural parameters, because the functional forms assumed in the first step of the estimation are not necessarily consistent with the structural model. This is a common drawback of two-step estimators, as discussed in Bajari et al. (2007).

The log likelihood I maximize is:

$$
\begin{aligned}
LL(\theta) = \sum_{i=1}^{N} \log \Bigg[\sum_{k} \pi_k \Bigg[\Bigg(\prod_{t} \Big(\underbrace{Pr(b_{it} \mid \theta, k, x_t, z_{it}, n_{it}, o_i)}_{\text{Bid}} \times \underbrace{\big(\mathbb{E}_{d_t^{-i_t^*}} \big[Pr(i_t^* \mid \theta, x_t, k, b_{i_t^* t}, z_{i_t^* t}, n_{i_t^* t}, o_{i_t^* t}, d_t^{-i_t^*}) \big]\big)^{\mathbb{I}\{i_t^* = i\}}}_{\text{Winning}} \\
\times \underbrace{Pr(r_{i_t^*} \mid \theta, k, n_{i_t^* t})^{\mathbb{I}\{i_t^* = i\}}}_{\text{Review}} \Big) \Bigg) \times \underbrace{Pr(n_{i1} \mid \theta, k)}_{\text{Initial condition}} \Bigg] \Bigg] + \sum_{t} \log \underbrace{\mathbb{E}_{d_t^{-i_t^*}} \big[Pr(i_t^* \mid \theta, x_t, d_t^{-i_t^*}) \big]^{\mathbb{I}\{i_t^* = 0\}}}_{\text{Employer chooses outside option}}.
\end{aligned} \tag{9}
$$

In Appendix Section B.1, I show that Equation 9 represents the log likelihood of the observed data integrated over the participation decisions of non-winners in each auction. By integrating over the non-winners' participation decisions, I replace the probability that a worker wins the auction—which depends on her actual competitors' unobserved types—with its expected value before competitors are drawn. In other words, this is the log likelihood computed as if the auction identifiers for non-winning bids were unobserved.

Although the likelihood in Equation 9 ignores available information, it offers a key advantage that facilitates the application of expectation-maximization algorithm to my model. Each term depends only on the realization of a single worker's unobserved type and, for a subset of terms, on the population distribution of types—but never on the realizations of multiple workers' unobserved types. Although maximizing this likelihood is less efficient than maximizing the full information likelihood, the estimator remains consistent.

The need for the partial information likelihood arises from a feature of my model that is absent in the Monte Carlo examples in Arcidiacono and Miller (2011) and in other applications (Chung et al., 2014; Igami and Yang, 2016). In my setting, the transition probability of the worker's review state depends not only on her own unobserved type but also on the unobserved types of other workers, even after conditioning on current-period actions. This is because the employer's choice depends on the whole set of bidders' types (see Equation 5). This contrasts with the examples in Arcidiacono and Miller (2011), where, conditional on firms' exit and entry decisions in the current period, the transition probability of a firm's observed state—its incumbency status—depends on neither its own unobserved state (demand shock), nor other firms' unobserved states.

Identification The identification of the type-specific functions relies on Kasahara and Shimotsu (2009). In my context, the argument is that conditional on a worker's current review state, her past bid, win, and review outcomes should not be correlated with current behavior or outcomes. This is effectively an exclusion restriction according to Berry and Compiani (2023).

Parametrization For the winner's probability of winning and the likelihood that the review is good, I use functional forms implied by the model.

I specify the bid policy function to take the following form

$$\log b_{it} = \beta_{o(i)} + \beta_z \log z_{o(i)t} + \beta_k + \beta_{k,x} x_t + \beta_{k,q} \mathbb{E}[\tilde{q}_i | n_{it}, k(i)] + \beta_{k,e} \log(n_{it}^{Good} + n_{it}^{Bad} + 1) + \beta_\nu \nu_{it}, \tag{10}$$

where $\beta_{o(i)}$ is the fixed effect for workers from country o; $z_{o(i)t}$ is the demeaned exchange rate of worker's local currency per USD; $\mathbb{E}[\tilde{q}_i | n_{it}, k(i)]$ is the posterior mean of the worker's latent quality given her review state; $\log(n_{it}^{Good} + n_{it}^{Bad} + 1)$ is the human capital accumulation effect; and ν_{it} is the match-specific shock.[14] I assume that $\{o(i), z_{o(i)t}\}$ are uncorrelated with the rest of the variables, which allows me to first estimate $\beta_{o(i)}$ and β_z using an ordinary least squares regression outside of the EM algorithm.

I represent the distribution of workers' initial states, defined as their numbers of good and bad reviews at the time of their first bids in the sample period, with a simplified three-point discrete approximation that can differ by type:

$$Pr(n_{i1}|k) = \begin{cases} \rho_k^1 & \text{if } \max\{n_{i1}^{\text{Good}}, n_{i1}^{\text{Bad}}\} = 0 \\ \rho_k^2 & \text{if } \max\{n_{i1}^{\text{Good}}, n_{i1}^{\text{Bad}}\} > 0,\ \frac{n_{i1}^{\text{Good}}}{n_{i1}^{\text{Good}} + n_{i1}^{\text{Bad}}} \leq \text{median} \\ 1 - \rho_k^1 - \rho_k^2 & \text{if } \max\{n_{i1}^{\text{Good}}, n_{i1}^{\text{Bad}}\} > 0,\ \frac{n_{i1}^{\text{Good}}}{n_{i1}^{\text{Good}} + n_{i1}^{\text{Bad}}} > \text{median}, \end{cases}$$

where median refers to the median ratio of good reviews among workers.

EM algorithm I adapt the EM algorithm with a sequential maximization step (developed by Arcidiacono and Jones, 2003), an extension of the EM algorithm, to estimate the parameter vector θ. The likelihood comprises four types of outcomes: bids, employers' worker choices, reviews received by workers, and workers' initial reputation states. Some parameters contribute to the likelihood of multiple outcomes. For example, the priors over workers' latent quality affect bids, reviews, and employers' choices. Arcidiacono and Jones (2003) show that these parameters can be estimated sequentially in the maximization step, substantially reducing computational burden.

I conduct the EM algorithm assuming three worker types $(K = 3)$.[15] The algorithm alternates between an expectation step and a maximization step. For iteration $m+1$, the algorithm performs the following steps:

[14] I specify a functional form for how the unobserved match-productivity shock enters the bid function for convenience. For nonparametric identification, it is sufficient to assume that b is strictly monotonic in ν.

[15] The Bayesian Information Criterion is often used to determine the number of types in EM estimation. In applying the EM algorithm to dynamic models, Igami and Yang (2016) follow the approach of Kasahara and Shimotsu (2009) and find that the minimal number of types needed to rationalize their data is three.

1. **Expectation step** Compute $\pi_{ik}^{(m+1)}$, the posterior probability that worker i belongs to type k, given the current parameter estimates $\theta^{(m)}$ and worker i's actions and outcomes:

$$\pi_{ik}^{(m+1)} = \frac{\pi_k^{(m)} L_i(k;\theta^{(m)})}{\sum_{k'} \pi_{k'}^{(m)} L_i(k';\theta^{(m)})}$$

where $\pi_k^{(m)}$ is the estimate from iteration m of the population distribution of type and $L_i(k;\theta^{(m)})$ is the likelihood of worker i's actions and outcomes if she belongs to type k, given estimates from iteration m, and is equal to

$$L_i(k;\theta^{(m)}) = \left(\prod_t \left(Pr(b_{it} \,|\, \theta^{(m)}, k, \ldots) \left[Pr(i_t^* \,|\, \theta^{(m)}, k, \ldots) \right]^{\mathbb{I}\{i_t^*=i\}} Pr(r_{i_t^*} \,|\, \theta^{(m)}, k, \ldots)^{\mathbb{I}\{i_t^*=i\}} \right) \right) Pr(n_{i1} \,|\, \theta^{(m)}, k).$$

2. **Maximization step** Given the vector of workers' posterior type probabilities $\pi_{ik}^{(m+1)}$, compute $\theta^{(m+1)}$ to maximizes an auxiliary function[16]

$$\begin{aligned} Q(\theta^{(m+1)}; \theta^{(m)}) = \sum_i \sum_k \pi_{ik}^{(m+1)} \Bigg[\log \pi_k^{(m=1)} + \sum_t \log Pr(b_{it} | \theta^{(m+1)}, k, \ldots) + \sum_t \log Pr(i_t^* | \theta^{(m+1)}, x_t, k, \ldots)^{\mathbb{I}\{i_t^*=i\}} \\ + \sum_t \log Pr\left(r_{i_t^*} | \theta^{(m+1)}, k, \ldots\right)^{\mathbb{I}\{i_t^*=i\}} + \log Pr(n_{i1} | \theta^{(m+1)}, k) \Bigg] + \sum_t \log Pr\left(i_t^* \middle| \, \theta^{(m+1)}, x_t\right)^{\mathbb{I}\{i_t^*=0\}}. \end{aligned}$$

As mentioned, this is done sequentially:

(a) Update the prior over types

$$\pi_k^{(m+1)} = \frac{\sum_i \pi_{ik}^{(m+1)}}{\sum_i \sum_{k'} \pi_{ik'}^{(m+1)}};$$

(b) Update the priors over workers' latent quality by maximizing with regards to ${a_k}^{(m+1)}$, ${b_k}^{(m+1)}$

$$\sum_i \sum_k \pi_{ik}^{(m+1)} \left[\sum_t \log Pr\left(r_{i_t^*} \middle| {a_k}^{(m+1)}, {b_k}^{(m+1)}, n_{i_t^* t} \right)^{\mathbb{I}\{i_t^*=i\}} \right];$$

(c) Update the bid policy functions by estimating a weighted OLS regression (Equation 10), using $\pi_{ik}^{(m+1)}$ as weights;

[16]An essential insight of the EM algorithm is the observation that any parameter $\theta^{(m+1)}$ that increases $Q(\theta^{(m+1)}; \theta^{(m)})$ beyond $Q(\theta^{(m)}; \theta^{(m)})$ must also increase $LL(\theta^{(m+1)})$ beyond $LL(\theta^{(m)})$ (see for example, Martin Haugh's Machine Learning for OR&FE notes for more details).

(d) Update the demand parameters by maximizing

$$\sum_i \sum_k \pi_{ik}^{(m+1)} \left[\sum_t \log Pr\left(i_t^* \middle| \theta^{(m+1)}, x_t, k, \ldots\right)^{\mathbb{I}\{i_t^*=i\}} \right] + \sum_t \log Pr\left(i_t^* \middle| \theta^{(m+1)}, x_t\right)^{\mathbb{I}\{i_t^*=0\}},$$

where the expected probability of winning is simulated by drawing competitor sets from the data. Appendix Section B.2 provides further details on the simulated estimator;

(e) Update the initial condition distribution

$$\rho_k^{j\,(m+1)} = \frac{\sum_i \left(\pi_{ik}^{(m+1)} \times \mathbb{I}\{j(i)=j\}\right)}{\sum_i \pi_{ik}^{(m+1)}},$$

where $j(i)$ represents the category of i's initial reputation state.

The algorithm continues until the increase in log likelihood falls below a criterion. Appendix Section B.2 provides additional details of the EM implementation.

5.2 Second-Step Estimation: Employers' Demand

I use Method of Simulated Moments (McFadden, 1989; Pakes and Pollard, 1989) to estimate demand parameters to match employers' choices, drawing on results from the first-step estimation.

As part of the E-M algorithm, I already obtained estimates of employers' demand (see step (d) of the maximization step). However, the moments I used there are the winning probabilities of workers against the average, expected set of competitors.[17] Here, I re-estimate employers' demand taking advantage of more data, including the observed competitor characteristics. These finer moments can provide more precise estimates.

The MSM estimator is

$$\underset{\theta}{\operatorname{argmin}} (D - P(\theta))' W' W (D - P(\theta)),$$

where $D - P(\theta)$ is a vector of size $M \times 1$, consisting of individual moments $d_{it} - p_{it}$ stacked together. The data moments are $\{d_{it}\}_{i,t}$, where d_{it} is equal to 1 if the bid by worker i for job t won and 0 otherwise. The model-implied probability that worker i wins auction t is represented by p_{it}. W is a $K \times M$ matrix of instruments. Examples of instruments include the workers' bids and their numbers of good and bad reviews.

[17] As explained in Section 5.1, this ensures the applicability of the EM algorithm.

The model-implied winning probability p_{it} is conditional on the observed characteristics and bids of all bidders in this auction, and is integrated over the unobserved types of these bidders. It is given by

$$p_{it} = \int ... \int \frac{\exp(f(k_i, x_t, b_{it}, o_i, n_{it}, \nu_{it}))}{1 + \sum_{l \in B_t} \exp(f(k_l, x_t, b_{lt}, o_l, n_{lt}, \nu_{lt}))} \prod_l Pr(k_l | \{b_{m\tau}, x_\tau, o_m, n_{m\tau}, z_{m\tau}\}_{m,\tau}) dk_l, \tag{11}$$

where

$$f(k_i, x_t, b_{it}, o_i, n_{it}, \nu_{it}) = \alpha_b \log b_{it} + q_o + q_k + q_{k,x} \log x_t + \alpha_q \mathbb{E}[\tilde{q}_i | n_{it}, k_i] + h_k \log(n_{it}^{\text{Good}} + n_{it}^{\text{Bad}} + 1) + \nu_{it},$$

and

$$\begin{aligned} Pr(k_l | \{b_{m\tau}, x_\tau, o_m, n_{m\tau}, z_{m\tau}\}_{m,\tau}) &= Pr(k_l | \{b_{l\tau}, x_\tau, o_l, n_{l\tau}, z_{l\tau}\}_\tau) \\ &= \frac{\pi_k \times Pr(\{b_{l\tau}, x_\tau, o_l, n_{l\tau}, z_{l\tau}\}_\tau | k)}{\sum_{k'} \pi_k' \times Pr(\{b_{l\tau}, x_\tau, o_l, n_{l\tau}, z_{l\tau}\}_\tau | k')}. \end{aligned}$$

In the last set of equations, the left-hand side is the probability that worker l belongs to type k given the bids, win outcomes, and characteristics of all workers across all auctions in the sample. The first equality is based on the observation that the left-hand side probability should only depend on worker l's bids, win outcomes, and characteristics. And the second equality follows from Bayes's rule. It is computed using the first-step estimates from the EM algorithm.

The wage elasticity of employers' demand is estimated from exchange rate-induced variation in bids. From the first-step estimates, I recover linear transformations of the unobserved match-productivity shocks $\beta_\nu \nu_{it}$ based on Equation 10. These are plugged in for the estimation of Equation 11, serving as control functions.[18] Conditional on $\beta_\nu \nu_{it}$ and the other variables included in $f(k_i, x_t, b_{it}, o_i, n_{it}, \nu_{it})$, the remaining variation in bids comes solely from exchange rate fluctuations. Although I specify a particular functional form for the bid policy function, the functional form is not necessary for identification of the effect of bids in employer demand. A sufficient condition for non-parametric identification is that bids are strictly monotonic in the unobserved match-specific shocks.

Appendix section B.3 contains the full list of instruments and more information on the estimator.

[18] For a given i,t, $\beta_\nu \nu_{it}$ depends on the type drawn. β_ν is estimated in this step as the inverse of the coefficient of f on $\beta_\nu \nu_{it}$.

5.3 Second-Step Estimation: Workers' Costs

I use MSM to estimate cost parameters to match observed bids, using the bid first-order conditions. I follow Hotz et al. (1994) and Bajari et al. (2007) and forward simulate continuation values in the first-order conditions using the previously estimated bid policy functions and employer preference parameters.

The opportunity cost is specified in Equation 6. I begin by estimating the non-linear parameter c_z, which governs the pass-through of exchange rate shocks to dollar-denominated opportunity costs. The bid first-order condition provides a link between the pass-through of exchange rate shocks to costs and the pass-through of exchange rate shocks to bids. I use this relationship evaluated at the median bid to estimate c_z. Appendix Section B.3 provides more details. With the non-linear parameter estimated, the opportunity costs and the workers' payoffs are linear in the unknown cost parameters c_k, c_o, and c_x.

I estimate the remaining opportunity cost parameters using MSM. The MSM estimator is

$$\underset{\theta}{\operatorname{argmin}}(B-\hat{B}(\theta))'W'W(B-\hat{B}(\theta)),$$

where $B-\hat{B}(\theta)$ is a vector of size $M\times 1$, consisting of individual moments $b_{it}-\hat{b}_{it}$ stacked together; b_{it} is the observed bid by worker i in auction t and $\hat{b}_{it}$ is the model-implied bid. W is a $K\times M$ matrix of instruments, which include workers' numbers of good and bad reviews.

I use the bid first-order conditions (Equation 8) to compute the model-implied bids. The first-order conditions include the static markup terms and the dynamic markup terms. I start by drawing the types of all workers (also using the probabilities described in Section 5.2). I simulate the static markup component of each bid using the estimated demand parameters. For the denominator in the static markup, I compute the change in a worker's winning probability, averaged over 2,000 potential competitor sets drawn from the data, when she raises her bid by 1×10^{-4}.

The dynamic markup terms depend on continuation values. I forward simulate value functions $V_o^k(n)$ for each worker type k, country o, and on a grid of review states n. For each $\{k,o,n\}$, I simulate the trajectories of 5,000 workers starting with those characteristics. Each simulated worker participates in up to 10,000 auctions, though in most cases exits occur well before reaching that limit. In each auction, the job's minimum budget, exchange rate, and the worker's match-specific productivity shock are randomly drawn. Using the estimated bid policy functions, I then compute the worker's bid. Competitors' characteristics and bids are sampled from the empirical distribution of auctions, with types assigned according to the probabilities described in Section 5.2. Employers' logit taste shocks are drawn, and the winner is determined based on the estimated employer preference parameters. Conditional on winning, I draw whether the employer leaves a review. When a review is left, I draw whether the rating is good or bad. Finally, I draw an exit shock to determine whether the worker remains active for the next auction. Value functions at points on the reputation grid are given by results from the forward simulation and a set of cost parameters. I apply linear interpolation for values between grid points and assign the value at the nearest grid point for observations outside the grid.

Appendix Section B.3 provides a list of the instruments and more details on the forward simulation.

6 Results

In this section, I first present the model estimates. Using the estimates, I quantify the value of a review and experience and the impact of forward-looking incentives on workers' bids.

6.1 Estimates

Table 4 presents estimates of type-specific bid policy functions, priors over workers' latent quality, and initial condition distributions from the EM algorithm in the first step of the estimation.

The three types differ in their bid levels. Type 1, the most common type, submits lower baseline bids than type 2, who in turn bids less than type 3 (average normalized log bid is 0.75, 1.06, and 1.23, respectively). These large, persistent gaps among different types of workers highlight the need to control for enduring worker heterogeneity. They may reflect persistent differences in opportunity costs (c_k), baseline quality ($q_k + q_{k,x} \log x$), or both—an issue addressed in the second-step estimation. The three types also differ in the distribution of the latent quality, with type 1 having a higher prior mean—and therefore a higher average probability of getting a good review—than type 2 and type 2 in turn than type 3. Across all types, bids rise with the posterior mean of latent quality and with accumulated job experience, and log-normalized bids are higher when a job's minimum budget is lower.

Table 5 reports the demand estimates. Employers dislike higher wage bids. The estimated wage-bid coefficient is -3.94. It is close to the - 4.38 found for the largest employer type in Stanton and Thomas (2025), who study a similar online labor platform. The own-wage elasticity is about 3.91, implying a wage markup of roughly 26 percent. Despite facing many competitors in an auction, workers secure a sizable markup, consistent with imperfect worker substitution.

Across types, employers place value on improvements in workers' expected latent quality and on accumulated experience. Figure 2 illustrates the estimated willingness to pay for additional good reviews for type-2 workers, separating the contribution of information about latent quality from the human capital accumulation. Depending on the worker type, employers are

Table 4: First-Step Estimates from Expectation-Maximization Algorithm

	Type 1	Type 2	Type 3
Population Distribution			
	57%	36%	7%
Bid Policies			
Constant	-0.60	-0.33	-0.22
Expected latent quality	0.15	0.41	1.02
Job experience (log)	0.02	0.02	0.02
Minimum budget (log)	-0.20	-0.34	-0.45
Mean normalized bid (log)	0.75	1.06	1.23
Priors over Latent Quality			
Beta distribution parameters	(12.2, 2.2)	(7.2, 1.7)	(11.6, 4.9)
Mean	0.85	0.81	0.71
Std dev	0.09	0.12	0.11
Initial Condition Distributions			
No review	85%	82%	73%
Reviewed (good-review ratio $<$ median)	7%	9%	15%
Reviewed (good-review ratio $>$ median)	8%	9%	12%

Notes: Table displays type-specific first-step estimates. Bids are first normalized by the auction's minimum budget, taken log, and then demeaned by worker country, with variation due to exchange rates also removed. The remaining variation is specified to be log-linear in the posterior mean of a worker's latent quality, the log number of reviews, and the log minimum budget, with the corresponding type-specific coefficients reported here. The mean normalized bid is the average bid among workers of a type with types drawn according to their posterior probabilities. The initial condition distribution is over the worker's reputation state at the point of her first bid submission in the sample. Demand is estimated as part of the EM algorithm; however, because demand is re-estimated in the second step using more data, which yields more efficient estimates, I omit the demand estimates here.

willing to pay 22–24 percentage point higher wages for a worker with five good reviews compared to one with no reviews. Of this wage premium, 8–12 percentage points are attributable to higher expected quality, and 12–15 percentage points to human capital accumulation.

The demand estimates reveal sizable productivity differences across worker types. In Table 5, I report the estimated average expected productivity of an unreviewed worker by type. Type 1—the group who bid persistently the lowest according to the first-step estimates—is estimated to also have the lowest expected productivity among unreviewed workers. Among unreviewed workers, employers are willing to pay about 35 percent more for a type-2 worker than for a type-1 worker, and about 30 percent more for a type-3 worker than for a type-2 worker. The results indicate that the persistent bid hierarchy reflects underlying productivity differences.

Table 5: Estimates of Employer Demand

	All	Type 1	Type 2	Type 3
Bid (log)	-3.94 (0.72)			
Expected latent quality		7.47 (0.32)	5.65 (0.19)	7.03 (0.86)
Minimum budget (log)		-1.45 (0.24)	-1.80 (0.15)	-2.11 (0.32)
Job experience (log)		0.33 (0.02)	0.27 (0.02)	0.27 (0.02)
Constant		-9.74 (0.46)	-6.61 (0.30)	-5.82 (0.57)
China	0.22 (0.05)			
India	0.05 (0.07)			
Pakistan	0.21 (0.05)			
United States	0.94 (0.11)			
Other countries	0.75 (0.11)			
Avg. new worker productivity		-3.39	-2.03	-0.83

Notes: Table reports employer demand estimates obtained from the method of simulated moments, matching observed employer choices. Types refer to persistent worker types recovered from the first-step EM algorithm. The omitted country is Bangladesh. Bootstrapped standard errors are in parentheses.

Figure 2: Decomposing Employers' Willingness to Pay for Reviews

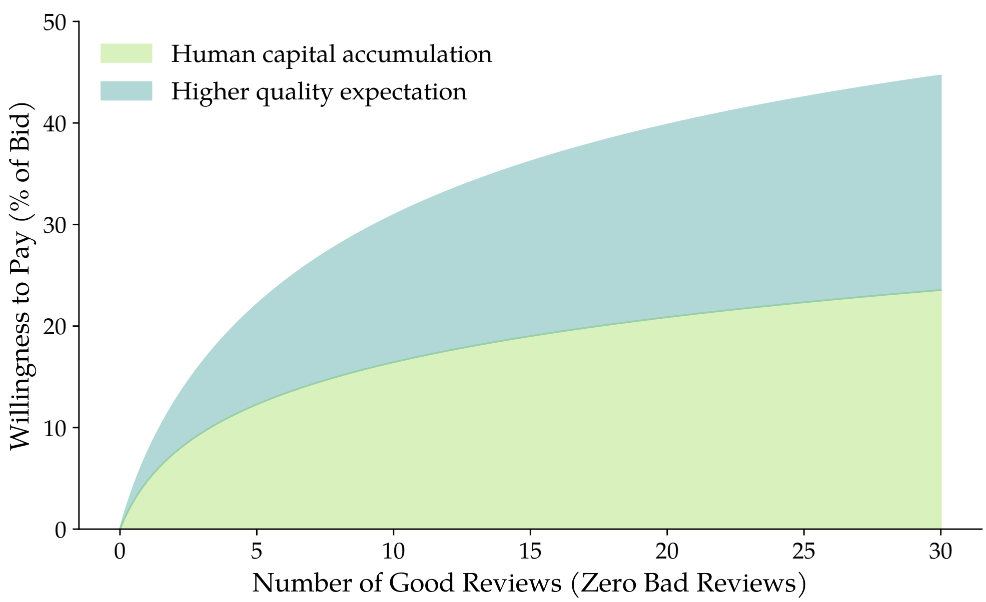

Notes: Figure shows estimated changes in employers' willingness to pay as a percentage of the wage bid for type-2 workers as their number of good reviews increases, and decomposes these changes into components reflecting higher expected quality and human capital accumulation.

Worker types who bid higher have higher expected productivity among unreviewed workers. Finally, workers from different countries have different expected productivity. Workers from the U.S. are estimated to have higher productivity than workers from other major countries.

Table 6: Estimates of Normalized Opportunity Costs

	All	Type 1	Type 2	Type 3
Constant		2.00 (.03)	3.18 (.06)	4.49 (.16)
Match productivity shock		0.57 (.15)	0.62 (.16)	0.47 (.12)
Minimum budget (log)	-1.16 (.01)			
China	-0.06 (.01)			
India	0.11 (.03)			
Pakistan	0.01 (.02)			
United States	0.52 (.01)			
Other countries	0.48 (.03)			

Notes: Table presents parameter estimates of normalized worker opportunity costs. Opportunity costs divided by minimum budget are specified to be linear in a type constant, log minimum budget, match-specific productivity shock (which enters into the employer's utility), and a worker country effect. The omitted country is Bangladesh. Bootstrapped standard errors are in parentheses.

The pass-through of exchange rate shocks to workers' opportunity costs is estimated to be 0.31. The fact that it is less than 1 suggests that workers' outside options are not entirely domestic but might include working for other foreign employers. Table 6 presents estimates of the other parameters governing workers' costs. Workers' costs are higher for jobs for which they have a higher match productivity shock, potentially due to the higher effort they put in. Workers from different countries have different levels of opportunity costs, with workers from the U.S. having higher costs than workers from other major countries.

The cost estimates indicate that worker types who bid persistently higher also face higher opportunity costs of work. Without imposing any restrictions on the correlation between quality and opportunity cost across types, I find a positive relationship between the two. This is intuitive, as higher-quality workers may have better outside options, leading to higher opportunity costs.

For the other parameters, I estimate auctions' arrival rate and workers' auction participation rates directly from the data, and workers' entry and exit rates are estimated to match the observed number and distribution of bidders in an average auction. Appendix Section B.4 contains more details.

6.2 The Value of a Review

I use the estimated model to quantify the value of a review and experience. This is the information and human capital value of a job match today, which accrues to workers and employers in future auctions.

I estimate the value of a reviewed experience by simulating the payoff differences between two scenarios: one in which a worker is hired today and receives a review, and one in which she is not hired. After today, the worker freely participates in auctions and may win additional jobs; the only difference between the scenarios is the worker's initial reputation state. I simulate auction participation, bids, employers' choices, and review dynamics for 5,000 draws of workers starting from a given review state. I record the payoffs of the worker, her competitors, and the employer in each auction.[19] I then compare these payoffs to a counterfactual scenario in which the worker starts with one extra experience and one more review, which may be good or bad with probabilities given by the estimates. The difference in the payoffs between the two scenarios is the value of a single review and experience.

Figure 3 reports the value of a review and experience for a type-2 worker from India. This worker can be considered the modal worker, because type 2 is the median type and India is the most common worker country. I show the value separately for when the worker has 0, 5, 10, or 50 good reviews and 0 bad review.

The value of a single review and experience is sizable. For a worker with no prior reviews it is \$305; with five prior good and zero bad reviews it is \$169. These magnitudes are comparable to the average job payment (\$210) and the estimated average flow surplus from a job (\$148). A substantial share accrues to the worker herself (for a worker with no prior reviews, 35%), because an additional review and experience improve her prospects in future auctions. Future employers also take a significant share (for a worker with no prior reviews, 65%), because the review helps employers make more informed hiring decisions and employers also benefit from the worker's higher human capital. The total change in future competitors' surplus as a result of the review is at most \$2 spread across thousands of workers, which translates into a per-competitor effect that is effectively zero. This finding supports the model's large-market assumption, under which workers do not track the reputation updates of individual competitors.

A review and experience carry substantially more value for a new worker than for a well-reviewed worker: \$305 for a new worker, \$111 for a worker with 10 good reviews, and \$44 for a worker with 50 good reviews. This pattern holds for the total value and for the components

[19] This approach abstracts from second-order effects. For example, a competitor may lose in an auction as a result of the worker's extra review and experience. This will affect this competitor's competitive status in other auctions, an example of second-order effects that I do not consider. Since the estimated first-order externalities on competitors are already small, any second-order effects are likely even smaller and unlikely to alter the conclusions.

Figure 3: The Value of a Review and Experience

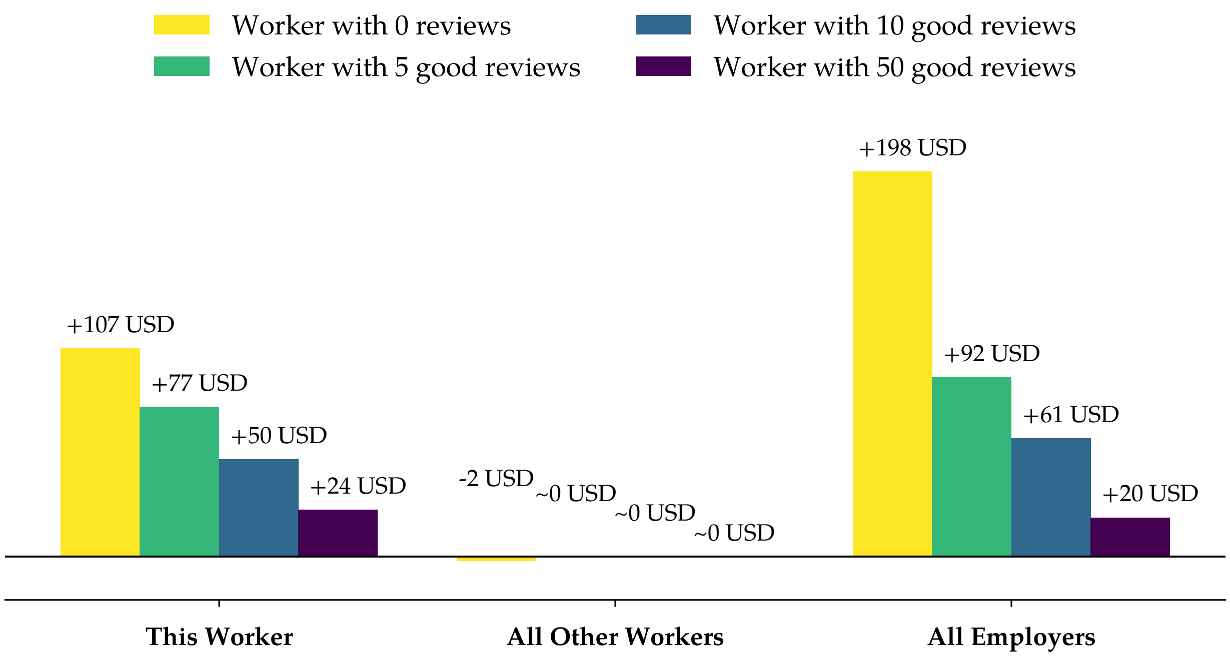

Notes: Figure presents the estimated value of a review and experience for the reviewed worker, the other workers, and for employers. These are calculated for a type-2 worker from India (the median worker type and most represented country). For reference, the average job payment is $210 and the estimated job's average contemporaneous surplus is $148.

accruing to the worker and to future employers. This pattern is due to both larger uncertainty around less experienced workers' quality and larger scope for human capital accumulation. Consequently, hiring a new worker generates larger future benefits than hiring an experienced worker.

Beyond these examples, I compute the value for all potential matches observed in the data and plot the distribution in Appendix A6. The results show that oftentimes the information and human capital value of a job match exceeds its flow surplus. For the median bidder, the value of a review and experience is $55.

Although the employer making today's hiring decisions does not internalize these future benefits, a forward-looking worker internalizes some of these benefits and responds with bid adjustments. I turn my attention to workers' bid responses to the value of a review.

6.3 The Effect of Forward-Looking Incentives on Wage Bids

I use the estimated model to analyze how forward-looking incentives shape workers' bids. I do so by comparing dynamically and statically optimal bidding strategies.

The statically optimal bid maximizes the expected payoff from the current auction alone.

The dynamically optimal bid also accounts for how winning the job and receiving a review influences the worker's future payoffs. I compute the optimal bids by simulating the distribution of competitors using first-step estimates and forward simulating continuation values. To simulate competitors, I draw 2,000 auctions from the data, randomly remove one bidder from each, and assign types to the remaining bidders by sampling from the posterior type distribution implied by their observed actions and outcomes (an output of the EM algorithm). The statically and dynamically optimal bids are defined as the points on the bid grid that maximize the respective objective functions.

Figure 4: The Effect of Forward-Looking Incentives on Wage Bids

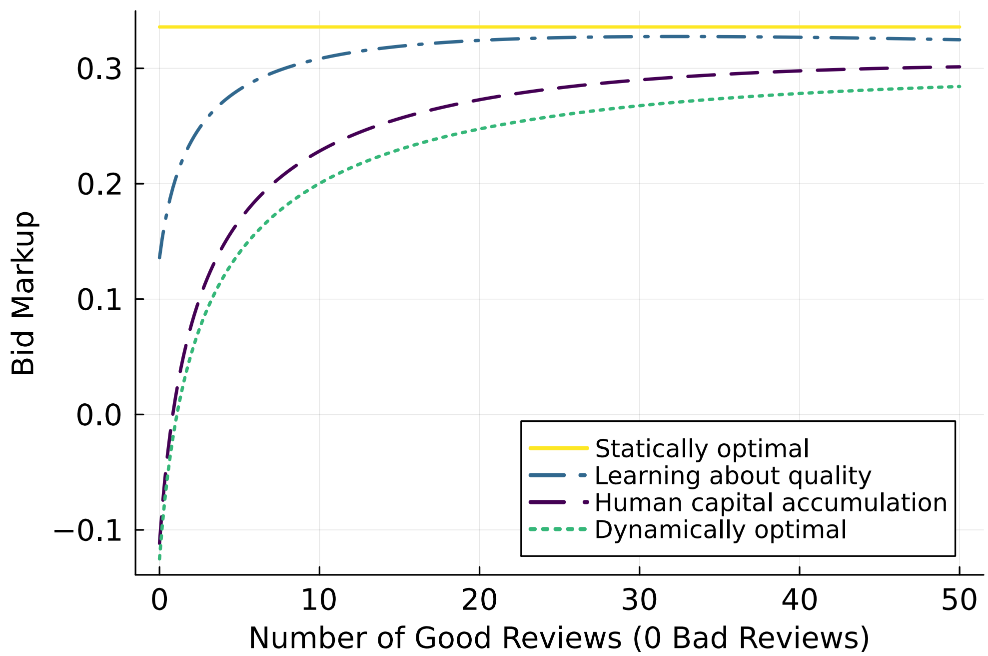

Notes: For a modal worker (type 2, India), the figure plots simulated statically and dynamically optimal bids in an average auction against the worker's number of good reviews. Forward-looking incentives from learning about quality and from human capital accumulation are separated by computing counterfactual continuation values when only one channel is active. Simulated bid values are smoothed using a second-order polynomial in the log number of good reviews.

Figure 4 plots the estimated statically optimal and dynamically optimal bid markups—defined as (bid-cost)/cost—at different reputation levels for a modal worker (a type-2 worker from India) bidding on a job with average characteristics: average minimum budget, average match productivity, and average exchange rate. I find that forward-looking incentives, rather than static ones, are the main drivers of workers' wage bid responses to changes in reviews. While a worker's probability of winning rises substantially with the number of good reviews—at the dynamically optimal bid, the estimated winning probability for a worker with ten good reviews is three times that of a worker with no reviews—the statically optimal bid increases only min-

Figure 5: Optimal Bids and Jobs' Minimum Budgets

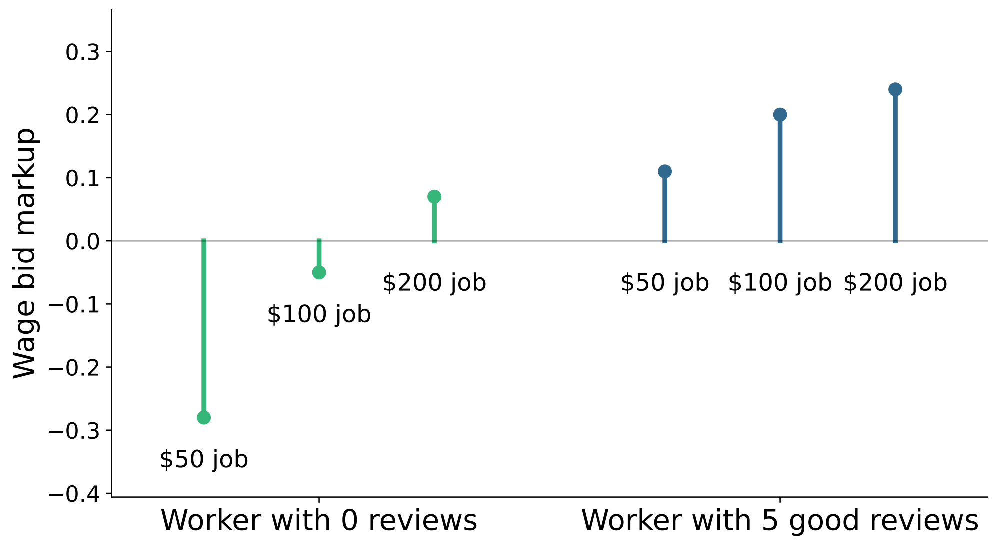

Notes: Figure plots the estimated optimal bid markups for a modal worker (type 2, India) for jobs with different minimum budget. The statically optimal bid markups for all combinations are around 0.36.

imally with the number of reviews.[20] By contrast, dynamic considerations lead workers with fewer reviews to shade their bids more aggressively, producing a steep profile of wage bids.

To illustrate the roles played by the two sources of forward-looking incentives—learning about latent quality and human capital accumulation—I simulate the dynamically optimal bids when only one source is present. For example, to isolate learning about quality, I forward simulate counterfactual continuation values assuming that employers do not factor in workers' productivity gains from experience when making hiring decisions. As Figure 4 shows, each source on its own pushes workers to bid below what they would if they were myopic.

I also examine how a job's minimum budget affects the relative weights of static and forward-looking considerations. In Figure 5, I plot the dynamically optimal bid markups across jobs with different minimum budgets. Although the statically optimal bids are comparable, workers bid higher markups for jobs with larger minimum budgets. The intuition is as follows: while the dynamic incentives from receiving a review are constant across jobs, the static payoffs are larger for higher-budget jobs. As a result, static considerations weigh more heavily when jobs are larger, leading workers to bid higher to capture the greater immediate earnings.

[20]The inelasticity of statically optimal bids with respect to the number of good reviews stems from the cross second-order derivative of the utility function being close to zero: $\frac{\partial^2 G}{\partial b \partial n^{\text{Good}}} = \frac{1}{\alpha_b} \times \frac{du}{dn^{\text{Good}}} \times \frac{G}{1-G}$, which is close to zero because the winning probability G is small. Under a more flexible utility specification—for example, one that directly interacts bid with the number of good reviews—statically optimal bids could be more elastic in the number of reviews.

7 Workers' Investments and Market Efficiency

I use the model to quantify the effects of workers' forward-looking bidding on market efficiency, first in an illustrative example and then in full equilibrium. I follow Sweeting (2012) and use policy function iteration to compute model equilibria, employing approximation techniques to handle the large state space. Appendix C provides details on the computational method and model fit.

Figure 6: Hiring a Worker with 0 Reviews vs. a Worker with 10 Good Reviews

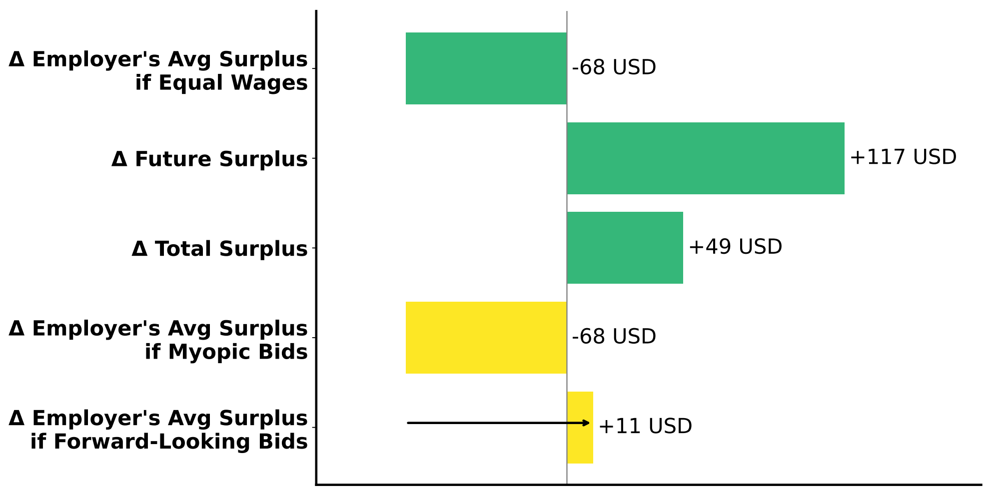

Notes: Figure illustrates that workers' forward-looking bidding can nudge an employer toward the social optimum. It does so in an example where the employer chooses between hiring a modal worker (type 2, India) with 0 reviews or 10 good reviews for a representative job. The difference in the employer's average surplus is averaged over logit taste shocks. The third row is the sum of the first two rows, because difference in the employer's average surplus if equal wages is equivalent to difference in the average flow surplus (the two workers have the same opportunity costs by construction).

An Example I illustrate through an example that workers' forward-looking bidding can move employers' choices toward the social optimum. Consider an employer choosing between a worker with no reviews and one with ten good reviews, holding all else equal. As shown in Figure 6, if the two workers bid the same amount, the employer would, on average, prefer the worker with ten good reviews because of higher expected latent quality and human capital accumulation. With logit taste shocks, the choice is probabilistic. The expected utility advantage of the worker with ten good reviews is \$68, which is sizable relative to the average job payment of \$210. Because wages are transfers and the two workers have the same opportunity costs of work, this \$68 equals the difference in total flow surplus. A social planner, however, would also account for the effect of today's match on future surplus—the informational and human capital gains—which are \$117 higher for hiring the new worker. Absent workers' wage responses, this \$117 wedge separates the employer's private incentives from the social planner's.

How do workers' wage responses change the employer's decision? I use the model to simulate the statically and dynamically optimal bids for these two workers. The statically optimal bids are very similar across the two workers and thus do not change the employer's decision. Under forward-looking bidding, however, the new worker bids significantly lower than the worker with ten good reviews, by \$79, which shifts the employer's choice toward hiring the new worker. In this example, forward-looking bids push the employer's decision closer to the social optimum, though they do not fully eliminate the wedge.

Full Equilibrium Motivated by the example, I use the model to evaluate the market-wide, equilibrium effects of workers' forward-looking behavior. I also examine the extent to which these investments are necessary for realizing the gains from the public reputation system. An important distinction between the full-equilibrium analysis and the preceding example is that, in the previous example, I fix the composition of bidders and therefore abstract from changes in the distribution of workers' review states.

I conduct three counterfactual simulations and compare them with the forward-looking bidding baseline.

1. No reputation system: A counterfactual equilibrium without a public reputation and feedback mechanism. Employers do not observe reviews left by previous employers or the number of jobs a worker has completed on the platform, and workers cannot credibly signal such information.[21] Workers have no incentives to build reputation and experience because they have no way of communicating these to the employers.

2. Myopic bidding: A counterfactual equilibrium with the same public reputation and feedback system as in the status quo but myopic bidding, i.e. bidding that maximizes immediate auction payoffs only. The difference from the previous counterfactual is that now the employers observe the workers' reputation and experience.

3. Myopic bidding with baseline reputation distribution: The same as above but the distribution of workers' review states is held fixed at the forward-looking bidding baseline.

The focus of these exercises is on changes in workers' bids. Therefore, for these simulations, I hold workers' market entry, market exit, and auction participation behavior the same as under status quo, as well as employers' job postings. Table 7 and Figure 7 summarize the results of these simulations.

[21] Employers still benefit from the increased productivity from workers' human capital accumulation from experience, but they do not observe the amount of human capital in advance.

Table 7: The Effects of Workers' Forward-Looking Bidding on Market Outcomes

	No Reviews	Public Reviews		
		Myopic	Myopic w/ Baseline Dist.	Baseline
Job Fill Rate / Baseline	12%	22%	38%	1
No. Unreviewed Workers	1,079	1,099	414	414
No. Reviewed Workers	1,347	1,327	2,012	2,012
Daily Hiring of Unreviewed / Baseline	58%	56%	12%	1
Wage Bill ($m)	2.21	4.06	7.73	7.39

Notes: Table reports simulation results under four scenarios: (i) no reputation system, (ii) myopic bidding, (iii) myopic bidding with the baseline reputation distribution, and (iv) baseline. The job fill rate is defined as the share of job posts for which the employer selects a bidder rather than the outside option. No. unreviewed workers and No. reviewed workers denote, respectively, the average numbers of active workers who have never been reviewed and those who have been reviewed at least once. Daily hiring of unreviewed refers to the daily average number of new workers hired. The wage bill is calculated on an annual basis. All numbers pertain to the PHP skill category.

Figure 7: The Effects of Workers' Forward-Looking Bidding on Surplus

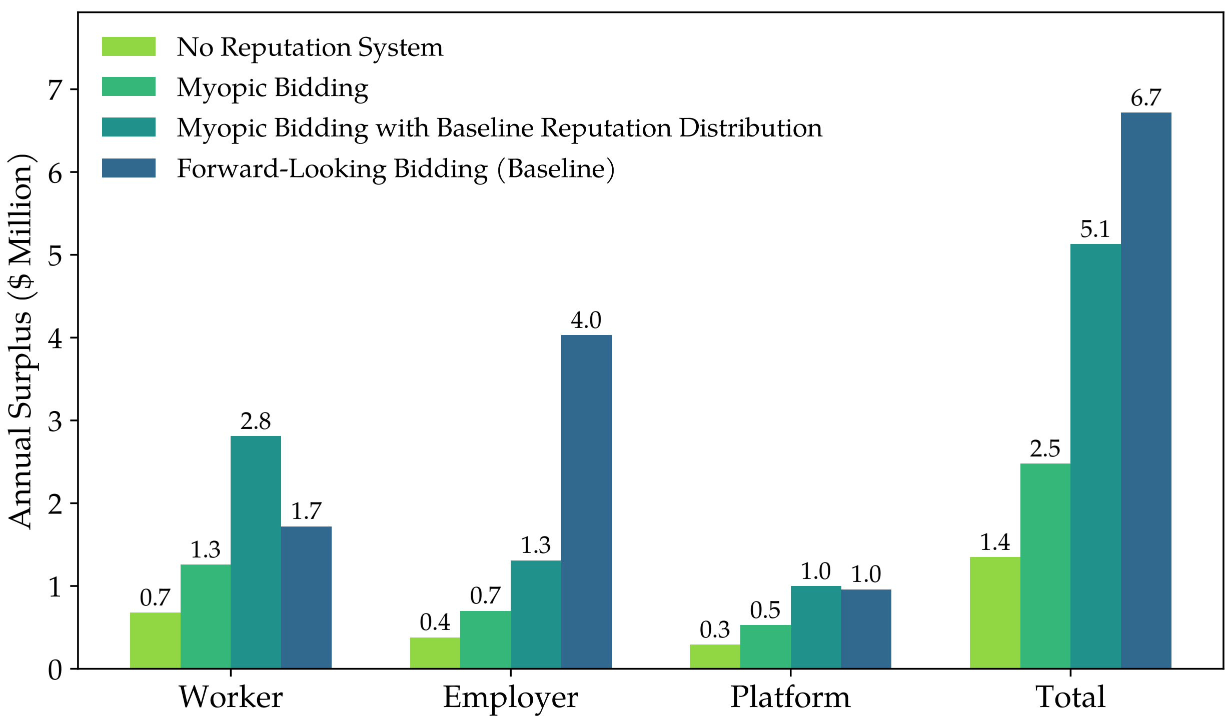

Notes: Figure displays simulation results under four scenarios: (i) no reputation system, (ii) myopic bidding, (iii) myopic bidding with the baseline reputation distribution, and (iv) baseline.

Figure 8: Workers' Investments Increase Information about Worker Quality

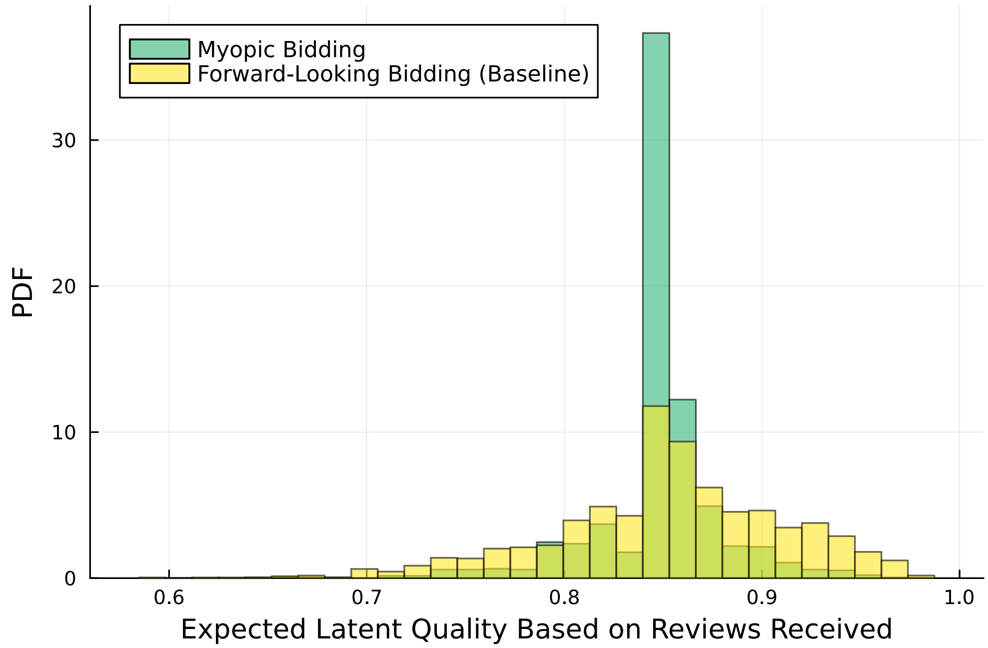

Notes: Figure displays the distribution of expected latent quality for active type 1 (most common type) workers in the equilibrium where workers bid to maximize current payoffs vs. in the equilibrium where workers bid accounting for the effect of reviews and experience on their future payoffs.

I begin by comparing the myopic bidding equilibrium with the baseline equilibrium. If workers were myopic, far fewer workers could accumulate reviews and experience. At a given point in time, the number of workers with at least one review would fall by 34%, limiting employers' options. Figure 8 illustrates the negative impact of myopic bidding on employers' choice sets. It plots the distribution of expected latent quality for a given type of workers (type 1, the most common type) in these two equilibria. Because fewer workers are reviewed in the myopic bidding equilibrium, the distribution of expected worker quality is narrower compared to that in the baseline, despite having the same distribution of actual worker quality. Fewer reviews mean less information that helps employers differentiate the higher-quality workers from the lower-quality ones.

Although there are more unreviewed workers in the myopic bidding equilibrium, these unreviewed workers bid similarly to the reviewed ones (as shown in Figure 9), giving employers little incentive to take a risk on unreviewed workers. The flat wage bids weaken the market's ability to screen and promote new talents.

Absent workers' investments, the platform's job fill rate—the probability that employers choose one of the bidders instead of the outside option—would fall to 22% of the status quo level, and total annual surplus for the workers, employers, and the platform combined would

Figure 9: Bid Policy in Myopic Bidding vs. Baseline Equilibrium

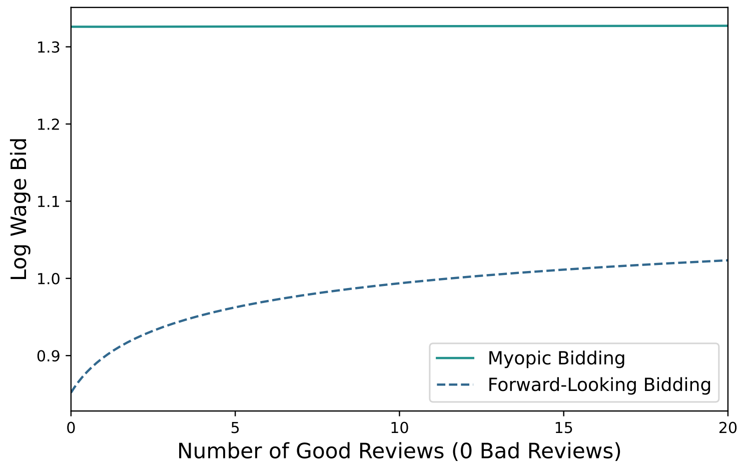

Notes: Figure displays the equilibrium bid policy of type 2 Indian workers in an equilibrium where workers bid to maximize current payoffs vs. in an equilibrium where workers bid accounting for the effect of reputation and experience on future payoffs.

decline by 63% (from \$6.7 to \$2.5 million). While workers' investments create surplus gains for workers as a whole, the largest gain notably accrues to the employers and a significant share to the platform.

To decompose the equilibrium impact of workers' investments, I consider the counterfactual case of myopic bidding with the baseline reputation distribution. This intermediate case helps illustrate the two channels through which workers' forward-looking incentives take effects: (a) workers' forward-looking bidding changes the equilibrium distribution of worker reputation and experience; (b) investment incentives lower wages, which can result in increased hiring. Figure 7 shows that both channels matter. Going from myopic bidding to myopic bidding with baseline reputation distribution achieves 62% of total surplus gains from forward-looking bidding. Holding the worker distribution at the baseline level, going from myopic to forward-looking bidding accounts for the rest of the total surplus gain.

Despite the large, positive effects of workers' forward-looking bidding, I show that it also has a drawback. Moving from myopic to forward-looking bidding changes not only how many jobs are filled but also which workers are hired. Figure 10 plots the distribution of hired workers' productivity in USD, measured as the non-bid component of employers' utility minus that of the outside option. Among auctions in which the employer hires under both myopic and forward-looking bidding, the productivity distribution of the hired workers shifts to the left under forward-looking bidding. On average, the productivity of the hired worker under forward-looking bidding is \$20 lower than that of the hired worker under myopic bidding, if a worker is hired in both cases. This illustrates one channel through which workers' forward-looking bidding can lower total surplus: it reallocates hiring from higher- to lower-quality matches—a manifestation of business stealing. Forward-looking bidding increases the number of matches, as lower wage bids induce employers to choose the inside option more often. But the productivity distribution of hired workers in these new matches enabled by workers' investments is even lower.

Figure 10: Forward-Looking Bidding and Match Productivity

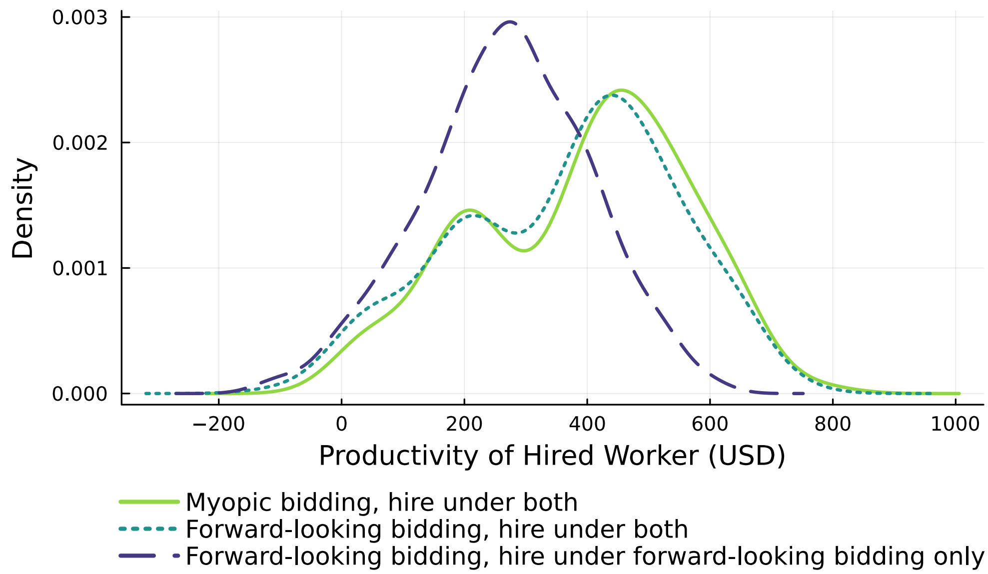

Notes: Figure displays the distribution of hired workers' match productivity—measured as the non-bid component of employers' utility normalized by the outside option—in simulations where the distribution of worker reputation and experience is held fixed at the status quo level. Workers are either myopic, maximizing payoffs from current auctions only, or forward-looking, accounting for the effects of reputation and experience on future payoffs. The figure distinguishes between (i) auctions in which employers hire a worker under both bidding schemes, and (ii) auctions in which employers hire only under forward-looking bidding.

While reputation and feedback systems have been commonly viewed as instrumental for fostering trust in online marketplaces such as eBay (see Tadelis (2016) for a review and Reimers and Waldfogel (2021) for quantification), I show that much of their effectiveness hinges on workers' (sellers') investments in reputation. I compare the myopic bidding equilibrium and the baseline equilibrium with the no reputation system equilibrium. I show that 79% of the surplus gains attributable to the reputation system vanish in the absence of workers' investments. Although the system enables employers to learn from each other, without workers' investments employers have weak incentives to take a chance on new workers, so little information or human capital is produced. In other words, a reputation system is necessary but not sufficient. Much of its benefits relies on workers' or sellers' costly investments in reputation and, potentially, platform-provided incentives that I study next.

8 Counterfactual Platform Policies

In this section, I study how the equilibrium level of worker investments compares to the socially optimal level and whether the platform has incentives to adopt policies that enhance total surplus.

In addition to workers and employers, the platform is an active participant in the market. The platform charges a 13% commission on each transaction. It may internalize part of the information and human capital gains through higher transaction volume and commission revenue, and therefore has incentives to adopt policies that increase total surplus. However, the platform's objective—maximizing commission revenue net of policy costs—generally differs from that of the social planner, who seeks to maximize the sum of worker, employer, and platform surplus. As a result, the platform's preferred policy need not be socially optimal.

For most of this section, I study platform-funded subsidies for hiring workers with no reviews. These subsidies are discounts off workers' posted bids: under a 20% subsidy, the employer pays 80% of the bid and the platform covers the remainder.[22] These subsidies compensate employers for the risk of hiring unreviewed workers while generating information and human capital that benefit future matches. At the same time, they may displace experienced, well-reviewed workers by diverting business away from them.

I simulate counterfactual equilibria across subsidy rates and report the effects on outcomes including match rates; employer, worker, and platform surplus (commission revenue net of subsidy outlays); and total surplus. The results are displayed in Table 8 and Figure 11.

Workers adjust their wage bids in response to the subsidy. Figure 11a plots the equilibrium bid of a modal unreviewed worker (type 2, India) in an average auction across alternative subsidy levels. Higher subsidies induce higher bids, partially offsetting the intended discounts on unreviewed workers' wage bids. Ignoring these bid adjustments would bias estimated subsidy effects on wages, match rates, and surplus.

I find that the total surplus is higher under a range of subsidies (including all of the ones studied here), implying that workers' equilibrium levels of investments in reputation and human capital are sub-optimally low. This is consistent with the large positive externalities created by the workers' investments estimated in Section 7.

[22]The platform subsidy mirrors government hiring and training subsidies for young workers. In the U.S., the Workforce Innovation and Opportunity Act reimburses employers up to 50% of the trainee's wage "for the extraordinary costs of providing the training and supervision related to the training." France provides up to €6,000 to employers for taking on apprentices and young workers.

Table 8: Summary of Counterfactual Policy Results

	Δ **Surplus** (1,000$)			
Counterfactual	Worker	Employer	Platform	Total
Subsidy: 10%	199	400	55	654 (+9.7%)
Subsidy: 20%	409	695	79	1,183 (+17.6%)
Subsidy: 30%	562	847	71	1,479 (+22.0%)
Subsidy: 40%	553	744	-14	1,283 (+19.1%)
Subsidy: 50%	754	536	-144	1,146 (+17.1%)
Subsidy: 60%	1,076	-61	-374	641 (+9.5%)
Verification	1,222	514	309	2,044 (30.4%)

Notes: Table reports estimated annual changes in surplus for employers, workers, the platform, and all three parties combined under alternative policies, for the PHP skill category. "Subsidy: 10%" denotes a 10 percent discount on unreviewed workers' wage bids. Verification refers to granting every new entrant a first review and experience without requiring an employer hire. The platform's surplus change equals the sum of (i) subsidy expenditures—zero under the verification policy—and (ii) the change in platform 13% commissions from the wage bill.

In deciding the optimal level of subsidy, the social planner faces a trade-off. Subsidizing the hiring of new workers generates valuable information and human capital, but by diverting hiring away from experienced, well-reviewed workers, it reduces employers' ability to exploit existing information. As Figure 11b illustrates, the subsidy causes the employers to hire unreviewed workers more. As the subsidy rate rises, conditional on hiring a worker, employers are more likely to hire one who has received no reviews. At moderate levels of subsidies, the increased experimentation with new workers raises market efficiency. The number of matches and total surplus go up.

At higher subsidy levels, however, excessive hiring of unreviewed workers can reduce overall efficiency as the market over-allocates hiring toward workers with uncertain quality and low levels of human capital allocation. Figure 11c shows that the number of reviewed workers hired per day, as well as the total stock of reviewed workers active in the market, both exhibit a hump-shaped pattern. The initial increase reflects greater inflows—more new workers obtain their first reviews. The subsequent decline arises from reduced exploitation of existing information and human capital. This reduction operates along two margins. At the intensive margin, conditional on hiring, employers hire reviewed workers less often. At the extensive margin, reviewed workers exit more frequently, as they are less likely to win jobs.

Together, these forces generate an internal optimum in welfare, where the dynamic gains from information and human capital accumulation are balanced against the lost employment opportunities for experienced workers. The socially optimal platform-subsidized discount

Figure 11: Equilibrium Effects of Platform-Funded Subsidy for Hiring Unreviewed Workers

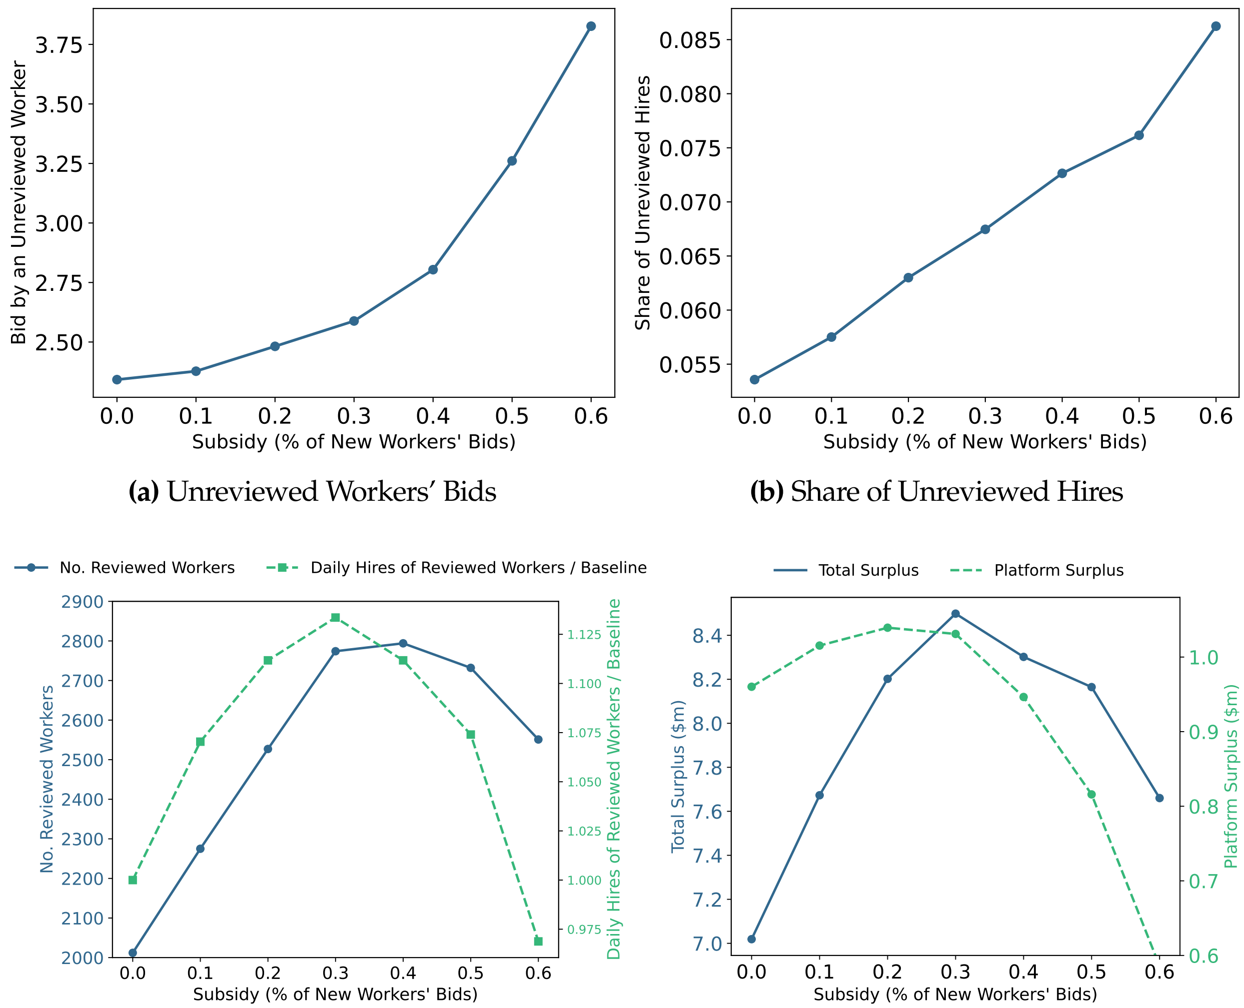

(a) Unreviewed Workers' Bids

(b) Share of Unreviewed Hires

(c) Number and Hires of Reviewed Workers

(d) Surplus: Platform vs. Total

Notes: Figures plot counterfactual results under different levels of platform-funded subsidies for hiring unreviewed workers. The subsidy is a percentage discount of unreviewed workers' bids. Bid by an unreviewed worker is the estimated optimal bid by a type-2 Indian worker (modal worker) with no reviews in an average outcome. Share of unreviewed hires is the percentage of hiring an unreviewed worker conditional on hiring a worker.

is 30% off unreviewed workers' bids. At this rate, annual total surplus in the PHP category would rise by roughly \$1,479,000, a 22.0% increase relative to the status quo. Approximately 57% of the gain accrues to employers, 38% to workers, and 5% to the platform.

As Table 8 shows, the platform benefits across a range of subsidy levels. Subsidies between 10% and 30% yield a net positive return. Under these subsidy levels, the induced increase in transaction volume raises commission revenue by more than the platform's out-of-pocket subsidy expenditure. Appendix Table A2 reports, for each subsidy level, the subsidy outlay, the resulting change in commission revenue, and the net effect on platform profit.

From the platform's perspective, the optimal new-worker discount is 20% off new workers' bids, which raises net platform surplus by about $79,000, an increase of 8%. The platform's preferred subsidy level is lower than the socially optimal level (Figure 11d shows that platform surplus peaks at a lower subsidy). The misalignment arises because the platform considers the subsidy spending as a cost, whereas the social planner treats it as a transfer, and because the platform internalizes only a portion of the subsidy benefits through commissions. Even so, the platform-preferred subsidy raises total surplus by 17.6%, capturing 80% of the welfare gains under the socially optimal subsidy.

I also study the effects of a skill certification program. The certification applies to all new workers, and reveals information that is equivalent to one review from an employer. Effectively, the skill certification program grants all new workers a first review without having to win a job. I report the estimated effects from the certification program in the last row of Table 8. Certification brings large benefits to the platform, increasing total surplus by $2,044,000, or 30.4% of the status quo surplus. The platform experiences a $309,000 increase in revenue. Given the estimated entry rate of 9 new workers per day, this is a revenue increase of $94 per new worker. Although the costs of such certification are difficult to quantify directly, a simple calculation suggests that the benefit of $94 is likely to exceed the implied cost. If we assume that a software engineer spends half an hour reviewing a worker, the associated cost—based on the average half-hourly wage of software engineers in Australia ($23, according to Indeed.com)—is well below the estimated benefit, implying that the certification would be profitable for the platform.

The counterfactual analyses in this and the preceding sections demonstrate that while employers hiring today fail to internalize the future benefits generated by their decisions, the combination of workers' privately optimal investments in reputation and human capital and the platform's privately optimal policies brings the market close to an efficient outcome.

9 Conclusion

In markets where reputation matters, buyers have incentives to free-ride on information generated by others' experiences. Such free-riding can prevent high-quality new sellers from being recognized and reduce long-run market efficiency. This paper uses an empirical case study of a large online labor platform to show that privately optimal behavior by sellers and the platform can substantially mitigate this issue and improve total surplus. It does so by formalizing the incentives of employers, workers, and the platform in a dynamic equilibrium model, estimating the model using proprietary data, and using the model to conduct counterfactuals.

Results on counterfactual platform policies reveal a trade-off between the dynamic gains from encouraging hiring of new workers and the losses from diverting hiring away from experienced ones. As the subsidy for hiring unreviewed workers increases, employers become less able to exploit the information and human capital already accumulated. Moreover, experienced workers may exit the market at faster rates. This tension in platform design—between promoting new sellers and leveraging existing information—parallels the classic exploration–exploitation trade-off and is relevant to a wide range of settings, including other online marketplaces and government procurement auctions.

The qualitative insight—that sellers' reputation investments enhance market efficiency by internalizing dynamic gains from matches and that the platform has a useful role to play—likely applies broadly. The quantitative implications, however, depend on market context. Sellers' incentives to invest in reputation and skills vary with how surplus is divided between buyers and sellers in a market, which in turn depends on the degree of competition and differentiation. In many markets, sellers also face investment frictions: for example, not all students can afford low- or unpaid internships that generate valuable references. Finally, when buyers engage in repeated relationships, as in long-term employment, their incentives complicate. The empirical framework developed here provides a basis for examining how differences in institutional features influence the efficiency of screening and reputation investments in other settings.

References

Acemoglu, Daron and Jörn-Steffen Pischke, "Why do firms train? Theory and evidence," *The Quarterly journal of economics*, 1998, *113* (1), 79–119.

Agrawal, Ajay, John Horton, Nicola Lacetera, and Elizabeth Lyons, "Digitization and the contract labor market: A research agenda," *Economic analysis of the digital economy*, 2015, pp. 219–250.

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

_ **and Robert A Shakotko**, "Do wages rise with job seniority?," *The Review of Economic Studies*, 1987, *54* (3), 437–459.

Arcidiacono, Peter and John Bailey Jones, "Finite mixture distributions, sequential likelihood and the EM algorithm," *Econometrica*, 2003, *71* (3), 933–946.

_ **and Robert A Miller**, "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity," *Econometrica*, 2011, *79* (6), 1823–1867.

— , **Patrick Bayer, Jason R Blevins, and Paul B Ellickson**, "Estimation of dynamic discrete choice models in continuous time with an application to retail competition," *The Review of Economic Studies*, 2016, *83* (3), 889–931.

Autor, David H, "Why do temporary help firms provide free general skills training?," *The Quarterly Journal of Economics*, 2001, *116* (4), 1409–1448.

Azar, José A, Steven T Berry, and Ioana Marinescu, "Estimating labor market power," Technical Report, National Bureau of Economic Research 2022.

Backus, Matthew and Greg Lewis, "Dynamic demand estimation in auction markets," *Review of Economic Studies*, 2025, *92* (2), 837–872.

Bajari, Patrick, C Lanier Benkard, and Jonathan Levin, "Estimating dynamic models of imperfect competition," *Econometrica*, 2007, *75* (5), 1331–1370.

Barach, Moshe A and John J Horton, "How do employers use compensation history? Evidence from a field experiment," *Journal of Labor Economics*, 2021, *39* (1), 193–218.

— , **Joseph M Golden, and John J Horton**, "Steering in online markets: the role of platform incentives and credibility," *Management Science*, 2020, *66* (9), 4047–4070.

Barlevy, Gadi and Derek Neal, "Allocating effort and talent in professional labor markets," *Journal of Labor Economics*, 2019, *37* (1), 187–246.

Bergemann, Dirk and Juuso Välimäki, "Learning and strategic pricing," *Econometrica: Journal of the Econometric Society*, 1996, pp. 1125–1149.

— **and** — , "Experimentation in markets," *The Review of Economic Studies*, 2000, *67* (2), 213–234.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor market power," *American Economic Review*, 2022, *112* (4), 1147–1193.

Berry, Steven T and Giovanni Compiani, "An instrumental variable approach to dynamic models," *The Review of Economic Studies*, 2023, *90* (4), 1724–1758.

Bolton, Gary E, Elena Katok, and Axel Ockenfels, "How effective are electronic reputation mechanisms? An experimental investigation," *Management science*, 2004, *50* (11), 1587–1602.

Bolton, Patrick and Christopher Harris, "Strategic experimentation," *Econometrica*, 1999, *67* (2), 349–374.

Bondi, Tommaso, Michelangelo Rossi, and Ryan L Stevens, "The good, the bad and the picky: Consumer heterogeneity and the reversal of product ratings," *Management Science*, 2025, *71* (8), 7200–7222.

Brinatti, Agostina, Alberto Cavallo, Javier Cravino, and Andres Drenik, "The international price of remote work," Technical Report, National Bureau of Economic Research 2021.

Cabral, Luis and Ali Hortacsu, "The dynamics of seller reputation: Evidence from eBay," *The Journal of Industrial Economics*, 2010, *58* (1), 54–78.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

Che, Yeon-Koo and Johannes Hörner, "Recommender systems as mechanisms for social learning," *The Quarterly Journal of Economics*, 2018, *133* (2), 871–925.

Chung, Doug J, Thomas Steenburgh, and K Sudhir, "Do bonuses enhance sales productivity? A dynamic structural analysis of bonus-based compensation plans," *Marketing Science*, 2014, *33* (2), 165–187.

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

Dendorfer, Florian and Regina Seibel, "The Cost of the Cold-Start Problem on Airbnb," *Available at SSRN 4910881*, 2024.

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri, "Monopsony in online labor markets," *American Economic Review: Insights*, 2020, *2* (1), 33–46.

Ericson, Richard and Ariel Pakes, "Markov-perfect industry dynamics: A framework for empirical work," *The Review of economic studies*, 1995, *62* (1), 53–82.

Farber, Henry S and Robert Gibbons, "Learning and wage dynamics," *The Quarterly Journal of Economics*, 1996, *111* (4), 1007–1047.

Filippas, Apostolos, John Joseph Horton, and Joseph Golden, "Reputation inflation," in "Proceedings of the 2018 ACM Conference on Economics and Computation" 2018, pp. 483–484.

Galdin, Anaïs Soledad and Jesse Silbert, "Equilibrium Provision of Information in Two-Sided Markets: Evidence from an Online Job Platform." PhD dissertation, Princeton University 2024.

Heckman, James J, "Heterogeneity and state dependence," in "Studies in labor markets," University of Chicago Press, 1981, pp. 91–140.

Holtz, David M, Liane Scult, and Siddharth Suri, "How much do platform workers value reviews? An experimental method," in "Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems" 2022, pp. 1–11.

Horton, John J, David G Rand, and Richard J Zeckhauser, "The online laboratory: Conducting experiments in a real labor market," *Experimental economics*, 2011, *14* (3), 399–425.

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

_ , _ , Seth Sanders, and Jeffrey Smith, "A simulation estimator for dynamic models of discrete choice," *The Review of Economic Studies*, 1994, *61* (2), 265–289.

Igami, Mitsuru and Nathan Yang, "Unobserved heterogeneity in dynamic games: Cannibalization and preemptive entry of hamburger chains in Canada," *Quantitative Economics*, 2016, *7* (2), 483–521.

ILO, "The role of digital labour platforms in transforming the world of work," *World Employment and Social Outlook*, 2021.

Iyer, Krishnamurthy, Ramesh Johari, and Mukund Sundararajan, "Mean field equilibria of dynamic auctions with learning," *Management Science*, 2014, *60* (12), 2949–2970.

Judd, Kenneth L, *Numerical methods in economics*, MIT press, 1998.

Kahn, Lisa B and Fabian Lange, "Employer learning, productivity, and the earnings distribution: Evidence from performance measures," *The Review of Economic Studies*, 2014, *81* (4), 1575–1613.

Kasahara, Hiroyuki and Katsumi Shimotsu, "Nonparametric identification of finite mixture models of dynamic discrete choices," *Econometrica*, 2009, *77* (1), 135–175.

Keller, Godfrey, Sven Rady, and Martin Cripps, "Strategic experimentation with exponential bandits," *Econometrica*, 2005, *73* (1), 39–68.

Klein, Benjamin and Keith B Leffler, "The role of market forces in assuring contractual performance," *Journal of political Economy*, 1981, *89* (4), 615–641.

Krasnokutskaya, Elena, Kyungchul Song, and Xun Tang, "The role of quality in internet service markets," *Journal of Political Economy*, 2020, *128* (1), 75–117.

Kremer, Ilan, Yishay Mansour, and Motty Perry, "Implementing the "wisdom of the crowd"," *Journal of Political Economy*, 2014, *122* (5), 988–1012.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler, "Imperfect competition, compensating differentials, and rent sharing in the US labor market," *American Economic Review*, 2022, *112* (1), 169–212.

Li, Lingfang, Steven Tadelis, and Xiaolan Zhou, "Buying reputation as a signal of quality: Evidence from an online marketplace," *The RAND Journal of Economics*, 2020, *51* (4), 965–988.

Li, Xingyi, Yiting Deng, Puneet Manchanda, and Bert De Reyck, "Can lower (ed) expert opinions lead to better consumer ratings?: The case of michelin stars," *Management Science*, 2025.

Manning, Alan, "The real thin theory: monopsony in modern labour markets," *Labour economics*, 2003, *10* (2), 105–131.

McFadden, Daniel, "A method of simulated moments for estimation of discrete response models without numerical integration," *Econometrica: Journal of the Econometric Society*, 1989, pp. 995–1026.

Meister, Matt and Nicholas Reinholtz, "Quality Certifications Influence User-Generated Ratings," *Journal of Consumer Research*, 2025, p. ucaf008.

Pakes, Ariel and David Pollard, "Simulation and the asymptotics of optimization estimators," *Econometrica: Journal of the Econometric Society*, 1989, pp. 1027–1057.

Pallais, Amanda, "Inefficient hiring in entry-level labor markets," *American Economic Review*, 2014, *104* (11), 3565–3599.

Reimers, Imke and Joel Waldfogel, "Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings," *American Economic Review*, 2021, *111* (6), 1944–1971.

Roussille, Nina and Benjamin Scuderi, "Bidding for talent: A test of conduct in a high-wage labor market," Technical Report, National Bureau of Economic Research 2025.

Rubens, Michael, "Market structure, oligopsony power, and productivity," *American Economic Review*, 2023, *113* (9), 2382–2410.

Rust, John, "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica: Journal of the Econometric Society*, 1987, pp. 999–1033.

_ , "Parametric Policy Iteration: An Efficient Algorithm for Solving Multidimensional DP Problems?," Technical Report, Mimeo, Yale University 2000.

Silva, Dakshina De, Matthew Gentry, and Pasquale Schiraldi, "Dynamic Oligopoly in Procurement Auction Markets," *Available at SSRN 4607350*, 2023.

Stanton, Christopher T and Catherine Thomas, "Landing the first job: The value of intermediaries in online hiring," *The Review of Economic Studies*, 2016, *83* (2), 810–854.

_ **and** _ , "Who benefits from online gig economy platforms?," *American Economic Review*, 2025, *115* (6), 1857–1895.

Sweeting, Andrew, "Dynamic pricing behavior in perishable goods markets: Evidence from secondary markets for major league baseball tickets," *Journal of Political Economy*, 2012, *120* (6), 1133–1172.

Tadelis, Steven, "Reputation and feedback systems in online platform markets," *Annual review of economics*, 2016, *8* (1), 321–340.

Vellodi, Nikhil, "Ratings design and barriers to entry," *Available at SSRN 3267061*, 2018.

Villas-Boas, J Miguel, "Consumer learning, brand loyalty, and competition," *Marketing Science*, 2004, *23* (1), 134–145.

Weintraub, Gabriel Y, C Lanier Benkard, and Benjamin Van Roy, "Markov perfect industry dynamics with many firms," *Econometrica*, 2008, *76* (6), 1375–1411.

WTO, *World Trade Report 2019: The future of services trade*, World Trade Organization, 2019.

A Appendix for Setting and Data

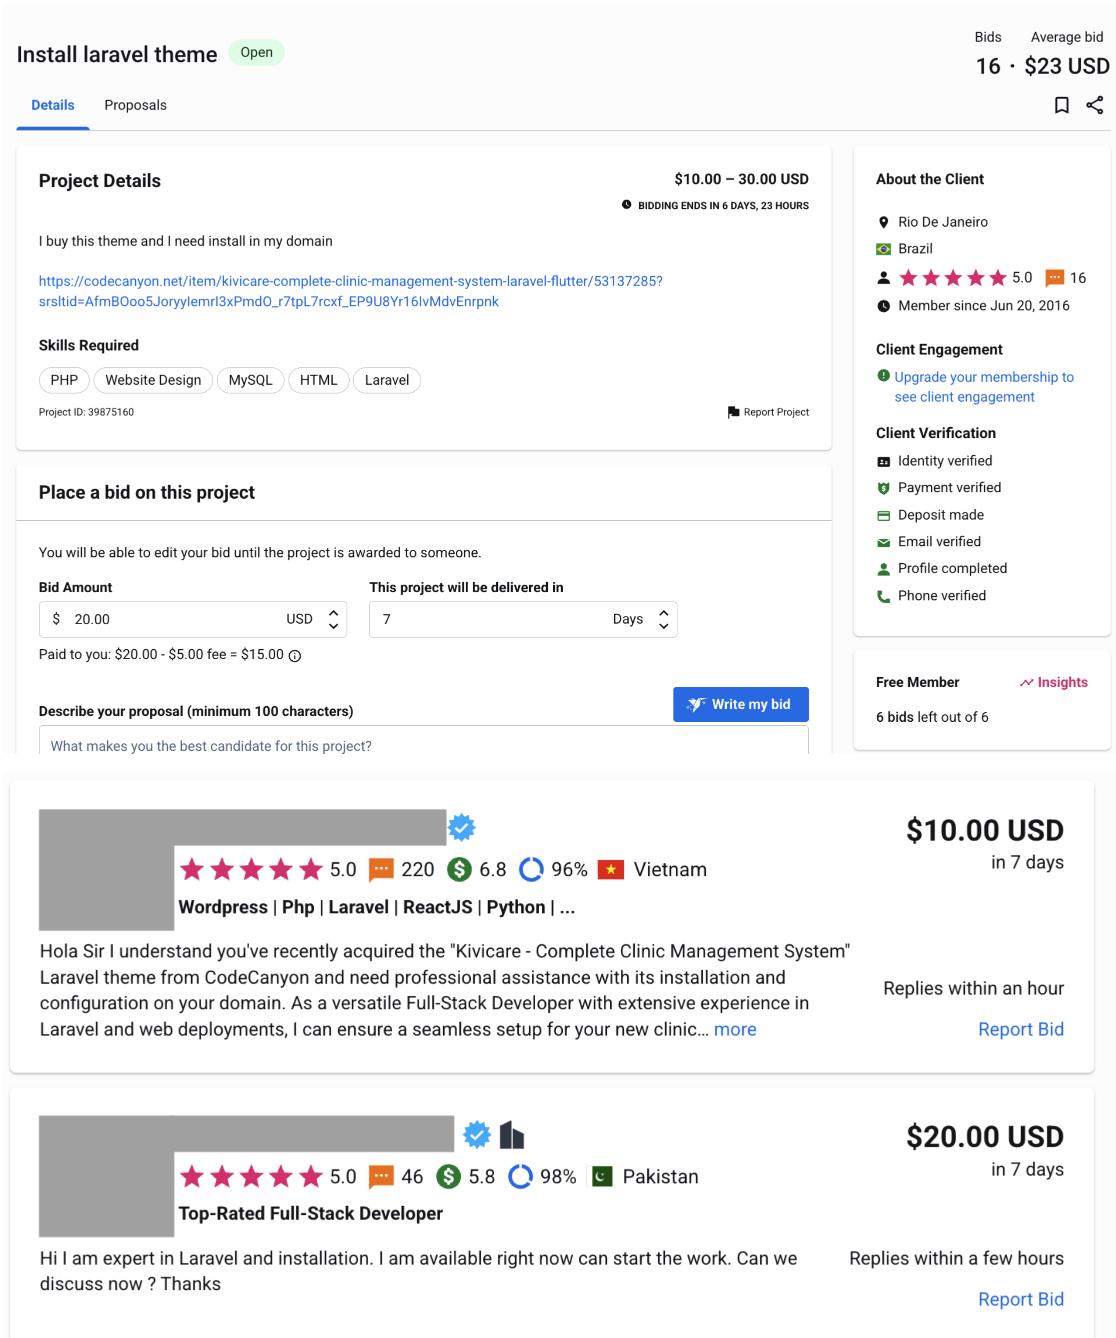

Figure A1: Screenshot of a Job Auction and Selected Bids on Freelancer.com

Figure A2: Cumulative Density Function of Minimum Budget

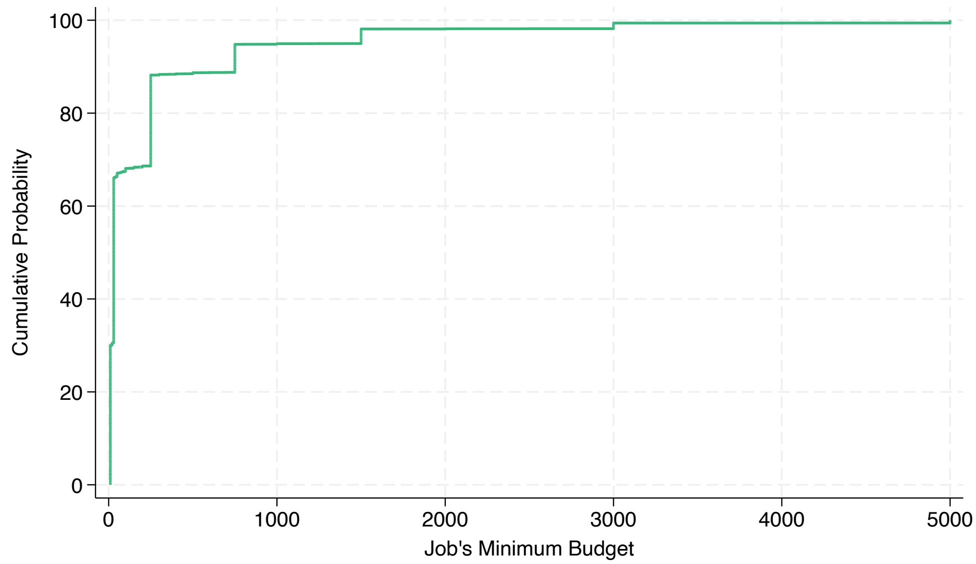

Notes: Employers are required to set a minimum budget, which is the lowest amount that workers may bid. Minimum budgets are denominated in U.S. dollars. Bunching occurs at values corresponding to the platform's drop-down menu options, although employers are free to customize the minimum budget so long as it is at least $10.

Figure A3: Distribution of Auction's Duration

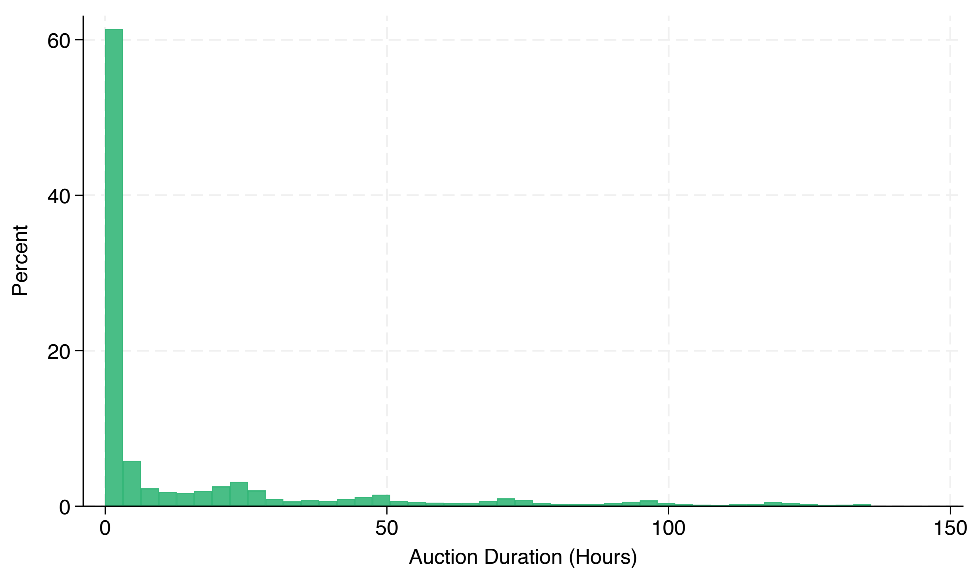

Notes: Auction duration is defined as the time between when the job is posted and when the job is awarded, and is therefore only defined for job auctions that are awarded. Auctions that lasted for over 136 hours (90th percentile) or less than 0 hour (3 observations) are excluded.

Figure A4: Distribution of Jobs' Duration

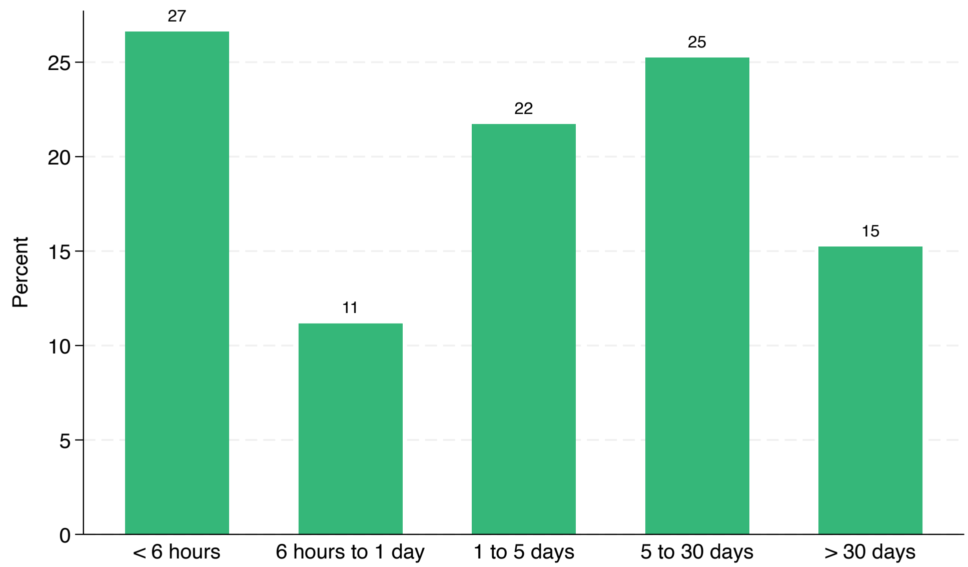

Notes: A job's duration is computed as the time between when the worker accepts the job and when the employer marks the project as complete.

Figure A5: Distribution of Workers' Ratings

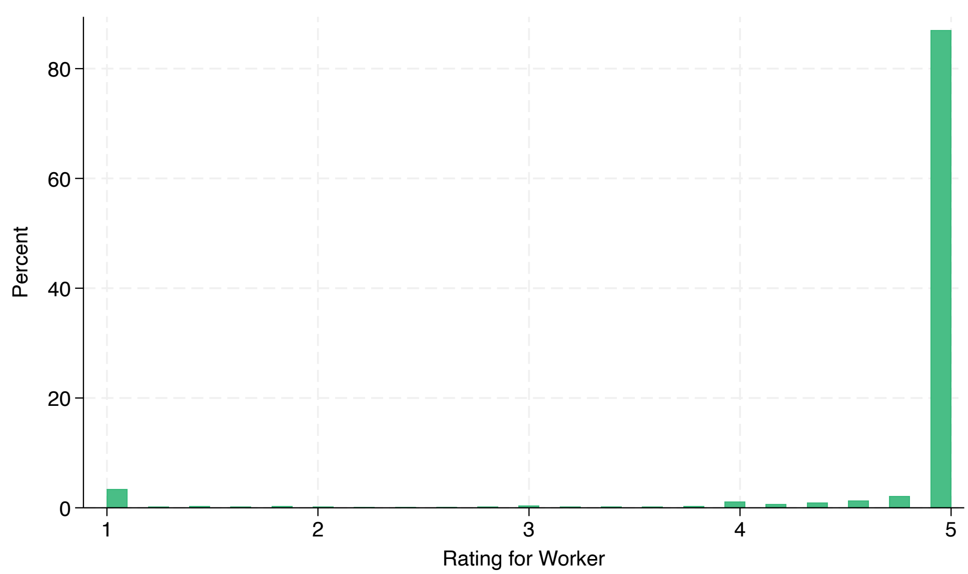

Notes: Employers are asked to give workers one to five stars on five dimensions: work quality, communication, expertise, professionalism, and whether they would hire the worker again. The rating presented here, and displayed on the worker's profile, is the average across these five dimensions.

List of variables

- x_t: the job's minimum budget. Employers are required to specify a minimum budget when posting a job. Workers have to bid at or above the minimum budget.
- b_{it}: normalized bid by worker i for job auction t. Workers' bids divided by the minimum budget, so they fall in $[1,\infty)$.
- $n_{it}=(n_{it}^{\text{Good}},n_{it}^{\text{Bad}})$: worker's reputation state, which includes the number of good (five-star) reviews and the number of bad (less than five-star) reviews.
- $o(i)$: worker country.
- z_{ot}: The exchange rate of country o's currency per USD, divided by its country-specific mean.

B Appendix for Estimation

B.1 Likelihood

Each worker belongs to one of K types: $k_i\in\{1,...,K\}$. Let π_k denote the probability that a worker is of type k. Let $\vec{k}$ denote the vector of type assignments for all N workers. The parameter vector θ comprises the parameters of type distribution and the parameters governing the endogenous outcomes—workers' bids, employers' choices, and employers' reviews of workers.

The observed data is

$$\left\{\left(d_t,\, x_t,\, \{b_{lt},\, z_{lt},\, n_{lt},\, o_l\}_{l\in B_t},\, i_t^*,\, r_{i_t^*}\right)\right\}_t.$$

Here, t denotes an auction. d_t is the vector of participation indicators $d_{lt}\in\{0,1\}$ for all N workers in auction t. B_t is the set of IDs of workers who participated in auction t; b_{lt} is the bid, z_{lt} the exchange rate, n_{lt} the reputation state, and o_l the country of worker l in the bidder set B_t. Finally, the outcomes of the auction are the winner $i_t^*\in B_t\cup\{0\}$, where $\{0\}$ denotes the outside option, and $r_{i_t^*}$, indicating whether the winner received a good or bad review from this auction.

The observed data likelihood integrates over $\vec{k}$:

$$\begin{aligned}
L&=\sum_{\vec{k}}\left[f(\vec{k})P\Big(\left\{\left(d_t,\, x_t,\, \{b_{lt},\, z_{lt},\, n_{lt},\, o_l\}_{l\in B_t},\, i_t^*,\, r_{i_t^*}\right)\right\}_t\Big|\,\theta,\,\vec{k}\Big)\right]\\
&=\sum_{\vec{k}}\left[f(\vec{k})\prod_t\left[P\Big(d_t,\, x_t,\, \{b_{lt},\, z_{lt},\, n_{lt},\, o_l\}_{l\in B_t},\, i_t^*,\, r_{i_t^*}\Big|\,\theta,\,\vec{k},\,\left\{\left(d_\tau,\, x_\tau,\, \{b_{l\tau},\, z_{l\tau},\, n_{l\tau},\, o_l\}_{l\in B_\tau},\, i_\tau^*,\, r_{i_\tau^*}\right)\right\}_{\tau<t}\Big)\right]\right]
\end{aligned}$$

The second equality follows from the chain rule. I now focus on the contribution to the likelihood from period t, given the type assignment $\vec{k}$

$$
\begin{aligned}
L_1 := & P\Big(d_t, x_t, \{b_{lt}, z_{lt}, n_{lt}, o_l\}_{l\in B_t}, i_t^*, r_{i_t^*}\Big|\, \theta, \vec{k}, \big\{\big(d_\tau, x_\tau, \{b_{l\tau}, z_{l\tau}, n_{l\tau}, o_l\}_{l\in B_\tau}, i_\tau^*, r_{i_\tau^*}\big)\big\}_{\tau<t}\Big) \\
= & P\Big(\{b_{lt}\}_{l\in B_t}, i_t^*, r_{i_t^*}\Big|\, \theta, \vec{k}, \big\{\big(d_\tau, x_\tau, \{b_{l\tau}, z_{l\tau}, n_{l\tau}, o_l\}_{l\in B_\tau}, i_\tau^*, r_{i_\tau^*}\big)\big\}_{\tau<t}, d_t, x_t, \{z_{lt}, n_{lt}, o_l\}_{l\in B_t}\Big) \\
& \times P\Big(d_t, x_t, \{z_{lt}, n_{lt}, o_l\}_{l\in B_t}\Big|\, \theta, \vec{k}, \big\{\big(d_\tau, x_\tau, \{b_{l\tau}, z_{l\tau}, n_{l\tau}, o_l\}_{l\in B_\tau}, i_\tau^*, r_{i_\tau^*}\big)\big\}_{\tau<t}\Big) \\
= & P\Big(\{b_{lt}\}_{l\in B_t}, i_t^*, r_{i_t^*}\Big|\, \theta, \vec{k}, d_t, x_t, \{z_{lt}, n_{lt}, o_l\}_{l\in B_t}\Big) \times P\Big(\{n_{lt}\}_{l\in B_t}\Big|\, \theta, \vec{k}, \big\{\big(\{n_{l\tau}\}_{l\in B_\tau}, i_\tau^*, r_{i_\tau^*}\big)\big\}_{\tau<t}\Big) \times C \\
= & P\Big(\{b_{lt}\}_{l\in B_t}\Big|\, \theta, \vec{k}, d_t, x_t, \{z_{lt}, n_{lt}, o_l\}_{l\in B_t}\Big) \times P\Big(i_t^*\Big|\, \theta, \vec{k}, d_t, x_t, \{b_{lt}, z_{lt}, n_{lt}, o_l\}_{l\in B_t}\Big) \\
& \times P\Big(r_{i_t^*}\Big|\, \theta, \vec{k}, d_t, x_t, \{b_{lt}, z_{lt}, n_{lt}, o_l\}_{l\in B_t}, i_t^*\Big) \times \prod_i P\Big(n_{it}\Big|\, \theta, k_i, \big\{\big(n_{i\tau}, i_\tau^*, r_{i_\tau^*}\big)\big\}_{\tau<t}\Big) \times C \\
= & \left[\prod_{l\in B_t} P\Big(b_{lt}\Big|\, \theta, k_l, x_t, z_{lt}, n_{lt}, o_l\Big)\right] \times P\Big(i_t^*\Big|\, \theta, x_t, d_t, \{k_l, b_{lt}, z_{lt}, n_{lt}, o_l\}_{l\in B_t}\Big) \\
& \times P\Big(r_{i_t^*}\Big|\, \theta, k_{i_t^*}, n_{i_t^* t}\Big) \times \prod_i P\Big(n_{it}\Big|\, \theta, k_i, \big\{\big(n_{i\tau}, i_\tau^*, r_{i_\tau^*}\big)\big\}_{\tau<t}\Big) \times C
\end{aligned}
$$

The first equality follows from the chain rule. In the second equality, the first component follows from the model's Markov assumption: conditional on type assignments, workers' reputation states are sufficient statistics for how the present depends on the past. The second component follows from the model's assumption that minimum budget x_t, exchange rates z_{lt}, workers' auction participation decisions d_t, and workers' countries o_l are exogenously determined and thus absorbed into the constant C at the end.

In the third equality, the first three components follow from the chain rule, while the fourth component follows from the assumption that workers' states are conditionally independent across workers and depend only on (i) their own unobserved type, (ii) their own past reputation states, and (iii) wins and reviews from past auctions in which they participated. The current reputation state n_{lt} is not a deterministic function of $\big\{\big(n_{i\tau}, i_\tau^*, r_{i_\tau^*}\big)\big\}_{\tau<t}$ because the data covers auctions for PHP jobs, and workers may have gained reputation from auctions in other skill categories. The key assumption is that the probability of obtaining reputation from other skill areas is independent of workers' unobserved types. Therefore, the only part of this likelihood that depends on workers' unobserved types is workers' initial reputation states during their first bid submission. Formally,

$$
P\Big(n_{it}\Big|\, \theta, k_i, \big\{\big(n_{i\tau}, i_\tau^*, r_{i_\tau^*}\big)\big\}_{\tau<t}\Big) = \begin{cases} C & \text{if } \exists \tau < t \text{ s.t. } i \in B_\tau \\ P(n_{it}|\theta, k_i) & \text{otherwise.} \end{cases} \tag{12}
$$

To simplify notation, in what follows I define a worker's reputation state during the first auction of the sample as her reputation during the first bid I observe her submit $n_{i1} := n_{it}, \nexists\, \tau < t$ s.t. $i \in B_\tau$. Therefore the contribution of the terms in Equation 12 to the total likelihood is $C \prod_i P(n_{i1}|\theta,k_i)$.

In the fourth equality, the first term follows from the model's assumption that a worker's bid is independent of other workers' types or states and that bids are uncorrelated conditional on the auction's minimum budget and the worker's country exchange rate. The second term reflects that the probability of the winner winning depends only on the characteristics of other auction participants, not on non-participants. The third term follows from the model's specification that the probability a review is good equals the posterior mean of the worker's latent quality, which depends only on the worker's own type and past reviews. The other terms remain unchanged.

Let $d_t^{-i_t^*}$ denote the vector of participation decisions in auction t for all workers other than the winner i_t^*. If the employer selects the outside option, then $d_t^{-i_t^*} = d_t$. I take expectation of L_1 with respect to $d_t^{-i_t^*}$. This expectation affects only one component of L_1, the probability that the winner wins.

$$
\begin{aligned}
\mathbb{E}_{d_t^{-i_t^*}}[L_1|d_t^{-i_t^*}] = & \left[\prod_{l \in B_t} P\left(b_{lt} \middle| \theta, k_l, x_t, z_{lt}, n_{lt}, o_l\right)\right] \times P\left(r_{i_t^*} \middle| \theta, k_{i_t^*}\right) \times \prod_i P\left(n_{i1} \middle| \theta, k_i, \left\{\left(n_{i\tau}, i_\tau^*, r_{i_\tau^*}\right)\right\}_{\tau<t}\right) \\
& \times \mathbb{E}_{d_t^{-i_t^*}}\left[P\left(i_t^* \middle| \theta, x_t, \{k_l, b_{lt}, z_{lt}, n_{lt}, o_l\}_{l \in B_t}\right)\right] \times C \\
= & \left[\prod_{l \in B_t} P\left(b_{lt} \middle| \theta, k_l, x_t, z_{lt}, n_{lt}, o_l\right)\right] \times P\left(r_{i_t^*} \middle| \theta, k_{i_t^*}\right) \times \prod_i P\left(n_{it} \middle| \theta, k_i, \left\{\left(n_{i\tau}, i_\tau^*, r_{i_\tau^*}\right)\right\}_{\tau<t}\right) \\
& \times \mathbb{E}_{d_t^{-i_t^*}}\left[P\left(i_t^* \middle| \theta, x_t, k_{i_t^*}, b_{i_t^* t}, z_{i_t^* t}, n_{i_t^* t}, o_{i_t^* t}, \{k_l, b_{lt}, z_{lt}, n_{lt}, o_l\}_{d_{lt}=1,\, l \neq i_t^*}\right)\right] \times C
\end{aligned} \tag{13}
$$

To simplify notation, let

$$
\mathbb{E}_{d_t^{-i_t^*}}\left[P\left(i_t^* \middle| \theta, x_t, k_{i_t^*}, b_{i_t^* t}, z_{i_t^* t}, n_{i_t^* t}, o_{i_t^* t}, d_t^{-i_t^*}\right)\right] := \mathbb{E}_{d_t^{-i_t^*}}\left[P\left(i_t^* \middle| \theta, x_t, k_{i_t^*}, b_{i_t^* t}, z_{i_t^* t}, n_{i_t^* t}, o_{i_t^* t}, \{k_l, b_{lt}, z_{lt}, n_{lt}, o_l\}_{d_{lt}=1,\, l \neq i_t^*}\right)\right]
$$

Given equations 12 and 13, and because $f(\vec{k}) = \prod_{i=1}^N \pi_{k_i}$, I take expectation of the total observed data likelihood with respect to $\{d_t^{-i_t^*}\}_t$

$$
\begin{aligned}
\mathbb{E}_{\{d_t^{-i_t^*}\}_t}\left[L|\{d_t^{-i_t^*}\}_t\right] = & \prod_{i=1}^N \left[\sum_k \pi_k \left[\left(\prod_t P(b_{it}|\theta, k_i, x_t, z_{it}, n_{it}, o_i) \times \mathbb{E}_{d_t^{-i_t^*}}\left[P\left(i_t^* \middle| \theta, x_t, k_{i_t^*}, b_{i_t^* t}, z_{i_t^* t}, n_{i_t^* t}, o_{i_t^* t}, d_t^{-i_t^*}\right)\right]^{\mathbb{I}\{i_t^*=i\}}\right.\right.\right. \\
& \left.\left.\left. \times P\left(r_{i_t^*} \middle| \theta, k_{i_t^*}\right)^{\mathbb{I}\{i_t^*=i\}}\right) \times P(n_{i1}|\theta, k_i)\right]\right] \times \prod_t \mathbb{E}_{d_t^{-i_t^*}}\left[P\left(i_t^* \middle| \theta, x_t, d_t^{-i_t^*}\right)\right]^{\mathbb{I}\{i_t^*=0\}}
\end{aligned}
$$

Taking log of the above, I have

$$\log\mathbb{E}_{\{d_t^{-i_t^*}\}_t}\left[L|\{d_t^{-i_t^*}\}_t\right]=\sum_{i=1}^{N}\log\left[\sum_k\pi_k\left[\left(\prod_t P(b_{it}\big|\theta,k_i,x_t,z_{it},n_{it},o_i)\times\mathbb{E}_{d_t^{-i_t^*}}\left[P\left(i_t^*\middle|\theta,\ldots,d_t^{-i_t^*}\right)\right]^{\mathbb{I}\{i_t^*=i\}}\right.\right.\right.$$
$$\left.\left.\left.\times P\left(r_{i_t^*}\middle|\theta,k_{i_t^*}\right)^{\mathbb{I}\{i_t^*=i\}}\right)\times P(n_{i1}|\theta,k_i)\right]\right]+\sum_t\log\mathbb{E}_{d_t^{-i_t^*}}\left[P\left(i_t^*\middle|\theta,x_t,d_t^{-i_t^*}\right)\right]^{\mathbb{I}\{i_t^*=0\}}. \tag{14}$$

Equation 14 is the likelihood I maximize in Section 5.1 using the expectation-maximization algorithm.

B.2 Expectation-Maximization Algorithm

I provide further details on the implementation of the expectation–maximization algorithm.

The convergence criterion is that the proportional change in the log likelihood (defined by Equation 9) falls below 10^{-5}. Because the simulated MLE in step (d) of the maximization step can be time-consuming, I first use EM to maximize a similar but misspecified likelihood that does not require simulated MLE, and then use its estimates as starting values for maximizing the correct likelihood.

I start by performing EM to maximize a log likelihood function similar to the one in Equation 9:

$$\tilde{LL}=\sum_{i=1}^{N}\log\left[\sum_k\left[\pi_k\left[\prod_t P\left(b_{it}\middle|\theta,k_i,x_t,z_{it},n_{it},o_i\right)\times\left(P\left(i_t^*|\theta,x_t,k_{i_t^*},b_{i_t^*t},z_{i_t^*t},n_{i_t^*t},o_{i_t^*t}\right)^{\mathbb{I}\{i_t^*=i\}}\right)\right.\right.\right.$$
$$\left.\left.\left.\times\left(\left[1-P\left(i_t^*|\theta,x_t,k_{i_t^*},\ldots\right)\right]^{\mathbb{I}\{i_t^*\neq i\}}\right)\times P\left(r_{i_t^*}|\theta,k_{i_t^*},n_{i_t^*,t}\right)^{\mathbb{I}\{i_t^*=i\}}\right]P\left(n_{i1}\middle|k_i\right)\right]\right]. \tag{15}$$

The differences between $\tilde{LL}$ and the correct log likelihood LL are twofold: the addition of the likelihood of losers losing, and the omission of the likelihood of the outside option being selected. In addition, rather than using the model-implied functional form for the expected winning probability—which requires integration over all possible competitor sets—I approximate the expected probability function with a logistic function:

$$P\left(i_t^*\middle|\theta,x_t,k_{i_t^*},b_{i_t^*t},z_{i_t^*t},n_{i_t^*t},o_{i_t^*t}\right)=\frac{1}{1+\exp(-f(b_{i_t^*t},x_t,k_{i_t^*},o_{i_t^*},z_{i_t^*t},n_{i_t^*t}))},$$

where the function $f(.)$ is the average employer utility from working with i_t^* in a job with minimum budget x_t.

To maximize $\tilde{LL}$, the algorithm is largely the same as that described in Section 5.1, with one key modification. In step (d) of the maximization step, rather than running a simulated MLE to update the demand parameters, I estimate a weighted logistic regression of each bid's win/loss outcome on the bidder's own characteristics, using the current posterior over the worker's type, $\pi_{ik}^{(m+1)}$, as the weights.

I use a range of start values to initialize the EM algorithm to maximize $\tilde{LL}$. The start values consist of a prior over types and parameters of bid policy functions, logit win probability functions, distribution of initial conditions, and the priors over workers' latent quality. I set the initial values for the prior over types to be $(\pi_1^{(0)},\, \pi_2^{(0)},\, \pi_3^{(0)} = 0.3,\, 0.3,\, 0.4)$, the initial Beta parameters to be $(Beta_a^k,\, Beta_b^k = 7.8,\, 1.4,\, \forall k)$, and the initial values for the distribution of the initial condition to be $p_k^I = 0.5,\, \forall k$. To initialize the bid policy parameters, I begin by regressing bid residuals on worker fixed effects and a set of variables:

$$\tilde{b}_{lt} = \beta_l + \beta_x x_t + \beta_G \log(n_{lt}^{Good}) + \beta_B \log(n_{lt}^{Bad}) + \beta_\nu \nu_{lt}, \tag{16}$$

Workers are then ranked by their estimated individual fixed effects, and the observations are divided into three equally sized groups. For each group, I estimate the bid policy function (Equation 10) and the logit winning probability function to obtain type-specific parameter estimates. Using these estimates, I generate five sets of initial values for both functions by adding normally distributed noise with mean 0 and standard deviation 0.3. The parameter σ_ν, the standard deviation of the match-specific shock ν_{it}, is constrained to be positive.

For all five sets of initial values, the algorithm converged in fewer than 50 iterations. The resulting estimates are similar across runs, both in magnitude and in likelihood. The estimates yielding the highest likelihood serve as the base initial values (E1) for maximizing the correct log likelihood LL. Using estimates E1, I draw worker types according to the estimated posterior probabilities and estimate the logistic regression described below. The coefficients from the logistic regression, together with the remaining EM estimates E1, form the initial values for maximizing the correct log likelihood LL.

From these initial values, the algorithm converged in two iterations. This initialization reduced the number of simulated MLE runs required in the maximization step.

I provide more details on the simulated MLE. I set the initial values for the simulated MLE algorithm to be results from a logit regression

$$win_{it} = \Bigg(1 + \exp\Big(\gamma_2 \mathbb{I}\{k_i = 2\} + \gamma_3 \mathbb{I}\{k_i = 3\} + \gamma_b \log(b_{it}) + \gamma_x^k x_t + \sum_o \gamma_o$$
$$+ \sum_k \big(\gamma_q^k \mathbb{E}[\tilde{q} \,|\, n_{it}, Beta^{(m+1)}, k] + \gamma_h^k \log(n_{it}^{Good} + n_{it}^{Bad} + 1) + \gamma_\nu \nu_{it}\big)\Big)\Bigg)^{-1} \quad)$$

after first drawing worker types according to their current type posterior probabilities, $\pi_{ik}^{(m+1)}$. This approach ensures that the simulated MLE starts from a point close to the global maximum, reducing the risk of convergence to a local maximum. To simulate the set of competitors and competitor characteristics in the denominator of Equation **??**, I do the below for each iteration of the simulated MLE:

1. Randomly select 1,000 auctions and the bids submitted to them. In the simulations, I use the observed country, exchange rate, worker reputation, and match shock $o_l, z_{lt}, n_{lt}, \nu_{lt}$, holding these fixed across iterations. For each bid, I draw a type based on the worker's posterior type probability, comparing this probability with a fixed random number so that small changes in posterior probabilities do not produce large changes in the drawn types.

2. Because competitors' bids depend on the auction's minimum budget, I predict competitors' bids using the current bid policy parameters. To reduce computation, I coarsen minimum budgets into five bins and perform bid prediction for each bin.

3. With the simulated competitor set and bids in hand, I compute the expected winning probability for a winner or the outside option, conditional on the auction's minimum budget and the demand parameters.

B.3 Second-Step Estimation

Demand I provide more detail on the second-stage estimation, beginning with the estimation of employer demand. The instruments W I use in the simulated method of moments estimator is a $K \times N$ matrix where K is the number of parameters and N is the number of prediction residuals. The K instruments include a constant, the worker's log normalized bid, the job's log minimum budget, indicators for the worker's type draw and country, the match component of the bid $\beta_k^\nu \nu_{it}$, and the numbers of good and bad reviews—each normalized by the mean of the respective variable.

Cost: exchange rate pass-through I estimate the pass-through of exchange rate shocks to dollar-denominated opportunity costs by rearranging the bid first-order condition to have cost on one side of the equation by itself, taking logs, and then differentiating with respect to the exchange rate:

$$\begin{aligned}\frac{\partial \log c}{\partial \log z} &= \frac{\partial}{\partial z}\log(b+\frac{G(b)}{G'(b)}+X)\\ &= (b+\frac{G(b)}{G'(b)}+X)^{-1}\left(2-\frac{G(b)G''(b)}{(G'(b))^2}\right)\frac{b}{z}\frac{\partial \log b}{\partial \log z},\end{aligned}$$

where $X = -\frac{1}{x_t} r\left[\mathbb{E}[V_o^k(n_{i,t+1})|n_{it}] - V_o^k(n_{it})\right]$.

I evaluate the right-hand side at the median bid. The median bid corresponds to a log normalized wage of 0.8, submitted by a type-2 worker from India whose posterior mean latent quality is 0.83. This worker has 28 reviews, so I assume the continuation value from an additional review is approximately zero. Under these values, the estimated pass-through of the exchange rate to opportunity costs, $\frac{\partial \log c}{\partial \log z}$, is approximately equal to the pass-through of the exchange rate to the bid, 0.31.

Cost: forward simulation The estimation of workers' costs involves forward simulation of value functions. I begin by constructing a grid over reputation states, defined by a 17×5 grid over $\{n^{\text{Good}}, n^{\text{Bad}}\}$, with denser representation at the lower end. This reflects the fact that the marginal impact of an additional review on total payoffs is greater when a worker has few reviews. Specifically, the grid for n^{Good} and n^{Bad} is $[0,1,2,3,4,5,6,7,8,12,18,30,50,120]$.

For each worker type k, country o, and grid point $(n^{\text{Good}}, n^{\text{Bad}})$, I simulate the lifetime payoffs for 5,000 workers starting from that reputation state. Each simulated worker proceeds through the following steps:

1. Draw an auction from the empirical distribution of auctions;

2. Use the auction date and the worker's country to determine the exchange rate. Draw a match-specific term $N(\sigma^k, 0)$. Use the estimated bid policy function to compute the worker's predicted bid;

3. Randomly drop one participant from the auction and compute the simulated worker's win probability. Simulate whether the worker wins by comparing this probability with a draw from Unif$(0,1)$;

4. If the worker wins, record the worker's earning and the cost covariates $x_t \times {z_{it}}^{c_z}$ and $x_t \times \log x_t \times {z_{it}}^{c_z}$;

5. If the worker wins, draw from Unif$(0,1)$ to determine whether the employer leaves a review. If the employer leaves a review, compute the probability that the review is good based on their current review state. Simulate whether the next review is good using a draw from Unif$(0,1)$, and update the worker's reputation state;
6. Draw an exit shock from Unif$(0,1)$ and compare it with the estimated per-auction exit; probability to determine whether the worker continues.
7. If the worker continues, I repeat the above. If the worker continues, repeat the process. Each worker is allowed to participate in up to 10,000 auctions, which far exceeds the average number observed in the data.

Importantly, the random number draws for auctions, wins, reviews, and exits are held constant across $\{k,o,n\}$ to ensure that simulated value functions are comparable.

Cost: instruments For the SMM estimation of workers' cost parameters, the instruments include a constant, the job's log minimum budget, and indicators for the worker's type draw and country.

Match-specific shocks I provide additional details on how the estimation incorporates the match-specific shocks ν_{it}. In the model, ν_{it} enters linearly in the employer's utility function and also influences the worker's opportunity cost of work, with effects that vary by worker type. For estimation, I assume a specific functional form for the bid policy function: after removing country fixed effects and exchange–rate–induced variation, the log bid is linear in ν_{it}. Under this assumption, I obtain $\beta_\nu \nu_{it}$ from the estimated bid policy. In the second-step demand estimation, I include $\beta_\nu \nu_{it}$ and estimate the coefficient of the employer's utility on $\beta_\nu \nu_{it}$, which equals $1/\beta_\nu$. This coefficient allows me to recover ν_{it}. In the second-step cost estimation, I then estimate type-specific cost coefficients on ν_{it}, reported in Table 6.

B.4 Estimation of Other Parameters

In this section, I discuss the estimation of daily auction arrival rates, workers' auction participation probabilities, daily platform entry and exit rates, and the probability of receiving a review conditional on being awarded a job.

Auction arrival rate The daily auction arrival rate is estimated to be 390, equal to the average number of auctions posted per day in the sample.

Auction participation probabilities I use an auxiliary data set to estimate a worker's average daily auction participation, allowing this measure to vary with the worker's total number of reviews. The auxiliary data set contains all bids submitted by a random sample of workers—across all skill areas, not limited to PHP—who registered in the second half of 2018 up till data extraction in 2024. In total, the data set covers over 5 million bids submitted by 8,973 workers. I discretize workers' review counts at the time of bid submission into five bins of approximately equal size (0–2, 3–8, 9–21, 22–62, and 63 or more reviews) and estimate average auction participation rate separately for each experience level. I then estimate parameters of an auction participation function $a(n) = a_1 + a_2 \times \log(n^{\text{Good}} + n^{\text{Bad}} + 1)$ based on the data moments. The estimated auction participation probability function is $a(n) = \min(9.1, 1.27 + 1.84 \times \log(n^{\text{Good}} + n^{\text{Bad}} + 1))/390$.

Market entry and exit rates The market entry rate is defined as the mean of a Poisson distribution governing the number of new entrants each day. Market exit rates are defined as the probability that a worker exits the market on a given day, which I allow to vary with the worker's number of reviews. I parameterize the exit rate as $\delta(n) = (e_1 + e_2 \times \log(n^{\text{Good}} + n^{\text{Bad}} + 1))^{-1}$. Given the auction participation probabilities, I estimate the market entry and exit rates by matching two sets of moments: (i) the average number of bids submitted by workers in each of the five experience levels at a typical auction (estimated from the main data set as 10.10, 2.14, 3.49, 5.71, and 11.52) and (ii) the average total number of lifetime bids submitted by workers (estimated from the auxiliary data set as 596).

Probability of review upon winning Because reviews are not mandatory, workers who are awarded a job are not always reviewed by employers. Using the main data set, I estimate the probability of receiving a review to be 61%.

C Counterfactual Equilibrium Computation and Model Fit

To evaluate model fit, I compute the status quo equilibrium of the estimated model and compare its features to the data. I adapt existing methods to accommodate a large state space.

Solving the model involves an outer loop and an inner loop. In the outer loop, I simulate 5,000 market days, ensuring that the system reaches a stationary distribution of workers' review states. Each day, new workers enter the market; a fixed number of auctions are posted; workers decide whether to participate and submit bids according to the current bid policy functions. Employers' taste shocks then realize, and each employer either hires a bidder or selects the outside option. The chosen workers receive reviews with some probability, and the realized rating is drawn based on the worker's latent quality. At the end of the day, reputation states are updated and some workers exit. I take the last five days of simulated bidder sets—after randomly dropping one bidder from each set—as the stationary distribution of competitors.

Table A1: Model Fit

	Data	Simulation
Project-Level		
Number of bidders	33	32
Median normalized winning bid	2.5	2.9
Bid-Level		
Median normalized bid	2.2	2.5
Median no. good reviews	23	17
Median no. bad reviews	4	3

Notes: Table compares the status quo equilibrium of the estimated model with the data. The simulation statistics are averaged over 2,000 auctions with auction characteristics drawn from the empirical distribution. Bids are normalized by jobs' minimum budgets.

In the inner loop, following Sweeting (2012), I use policy function iteration (Judd, 1998; Rust, 2000) to solve for the worker's dynamically optimal bid policies given the current competitor distribution. The algorithm iterates between (i) a policy evaluation step, in which value functions are computed given the current bid policies using the HJB equation, and (ii) a policy improvement step, in which bid policies are updated given the evaluated value functions. Because the state space—especially the reputation state space (counts of good and bad reviews)—is large, I approximate both the value and policy functions. I draw 1,000 (good, bad) review-count pairs from the data to form a grid and perform policy function iterations on these grid points. To approximate value functions, I use restricted cubic splines (cubic polynomials with four knots). For bid policies, I first regress bids on exchange rates, country dummies, and worker type dummies. I then regress the residuals on $\log n^{\text{Good}}$, $(\log n^{\text{Good}})^2$, $\log n^{\text{Bad}}$, $(\log n^{\text{Bad}})^2$, and log minimum budgets, allowing all coefficients to depend on worker types. I iterate between these two steps until bid policies converge, then advance the outer loop by simulating another 5,000 days using the updated policies. A solution is attained when bid policies change negligibly across successive outer loops—that is, when bid policies are best responses to the stationary distribution of competitors induced by those policies.

Table A1 shows that the model fits the data well in terms of the average number of bidders per auction and the average bid-level numbers of good and bad reviews. The median bid and the median winning bid are somewhat higher in the estimated model than in the data.[23]

D Additional Results

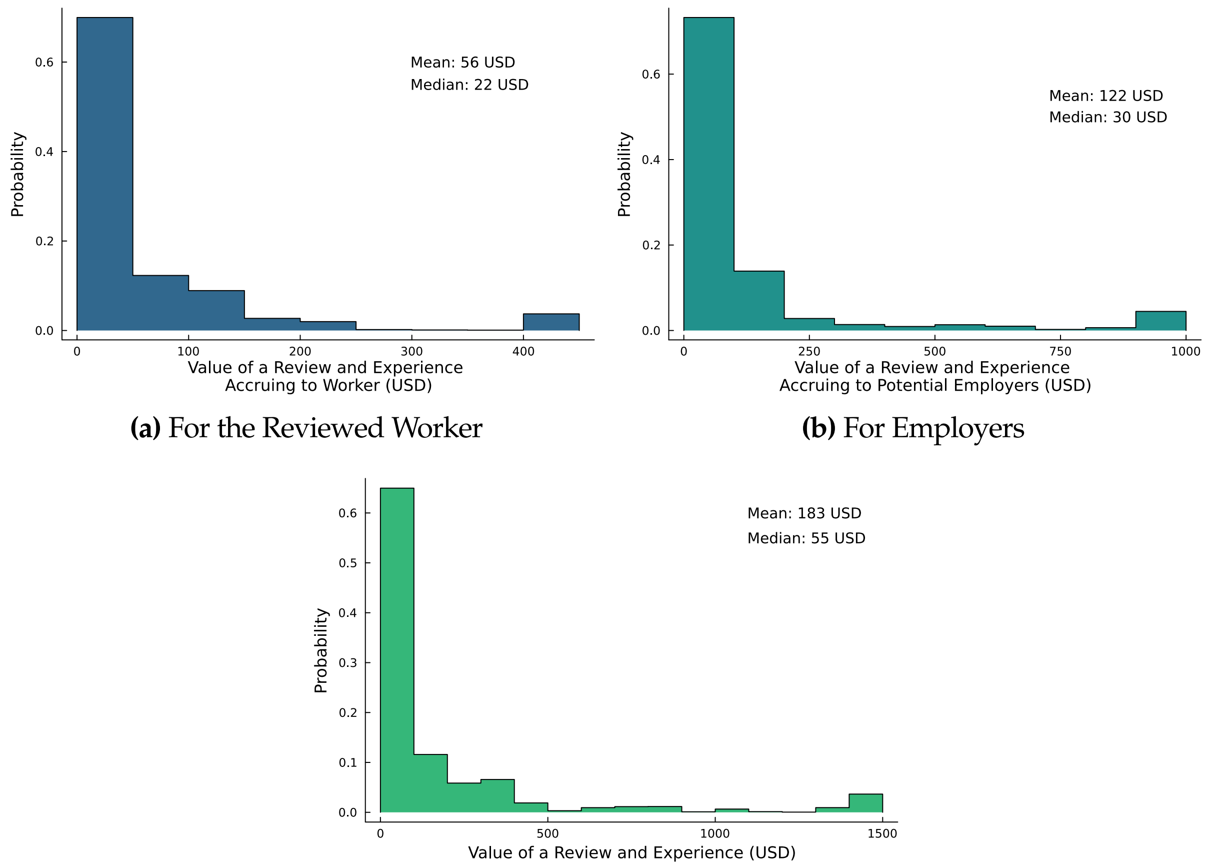

(a) For the Reviewed Worker

(b) For Employers

(c) Total

Figure A6: The Value of a Review and Experience

Notes: Figures plot the distributions, across all bids (potential matches) in the sample, of the estimated value of a review and of experience, reported as (i) the total value, (ii) the component accruing to the reviewed worker, and (iii) the component accruing to employers. Values are computed on a grid of reputation states for all country–type combinations and interpolated between grid points. For reference, the average job payment is \$210 and the estimated average contemporaneous surplus is \$148.

[23] I also evaluate fit by comparing the average probability that a job is filled (as opposed to the employer choosing the outside option). Under my agreement with Freelancer, I cannot disclose the underlying empirical moment, but the model-generated value lies within 20% of that moment.

Table A2: Cost and Benefit Analysis for the Platform

	Change in Platform's Payoffs ($1,000)		
New Worker Discount	Subsidy	Commissions	Total
10%	-31	87	55
20%	-76	155	79
30%	-134	205	71
40%	-215	202	-14
50%	-329	185	-144
60%	-518	143	-374

Notes: Table reports counterfactual equilibrium outcomes under alternative discounts for hiring new workers, focusing on changes in the platform's annual payoffs in the PHP skill category relative to the status quo. Subsidy lists the platform's subsidy expenditure; Commissions reports the change in commission revenue due to changes in total transaction volume (the platform charges a 13% fee); Total is the sum of the subsidy expenditure and the commission change.