Information Frictions and Observable Experience

Christopher Stanton Catherine Thomas

Harvard LSE

July 2015*

ESTIMATES ARE PRELIMINARY

Abstract

In many markets, interactions can be conditioned on partners' observable experience levels. In the online labor market oDesk, workers submit higher wage bids for projects posted by inexperienced employers. This is due to two potential effects: First, interacting with an inexperienced employer may result in predictable hassle costs to workers. Second, inexperienced employers may have less precise information about the value of a match, reducing the elasticity of hiring probability to wage bids and increasing the optimal bid markup. These two frictions are decomposed using a model of labor demand. Because workers on oDesk come from different countries but all bids are made in US dollars, the model is identified from shifts in workers' outside wages caused by differential exposure to exchange rate movements. Counterfactuals explore how platform fees that vary with employer experience affect platform profits.

*We thank seminar participants at the AEA Meetings, Harvard, LSE, Mannheim, NBER Summer Institute, and Yale along with Nava Ashraf, Ed Lazear, Arnaud Maurel, Luis Rayo, Yona Rubinstein, Scott Schaefer, Kathryn Shaw, and Nathan Seegert for helpful comments.

1 Introduction

In many markets, participants have or may attempt to acquire information about their trading partners. In settings such as online marketplaces, information about market experience is provided through publicly observable feedback. As a consequence, market participants have the ability to condition transactions on their partners' prior experience.[1] Much existing work on internet markets like oDesk, Amazon, eBay and Yelp (among others) shows the importance of seller reputation for buyer demand and market efficiency (for example, Pallais, 2013; Stanton and Thomas, 2012; Luca, 2013).

Less is known about how observable buyer or employer experience affects prices or wages offered by the supply side of the market. On oDesk, the market studied here, employers post projects and then workers bid hourly wages, allowing us to observe the set of offered and contracted wages over the entire history of employers' use of the market. To facilitate comparison of employers posting similar jobs over time, we concentrate on relatively homogeneous administrative support jobs where the modal job posting involves data entry work. When posting administrative support jobs, inexperienced employers receive wage offers that are, on average, more than 10 percent higher than employers with five or more prior hires. This positive gap is present in OLS regressions and in models with fixed effects for workers, employers, and for workers and employers. An even larger gap in mean wages was found in a pilot experiment when posting identical jobs from experimental employers with different experience levels.[2]

This paper asks why inexperienced employers receive higher wage offers. In contrast to the existing literature on the importance of observable measures of worker feedback, we find little evidence that employer reputation contributes to differences in quoted wage offers. Instead, frictions appear to be driven primarily by inexperienced employers' *own* uncertainty about the market. When entering unfamiliar markets like oDesk, each user is unlikely to have uniform needs or characteristics, making the actual value of a transaction stochastic and individual-specific. When workers can observe employer experience, they may infer that inexperienced employers will behave differently in evaluating a wage offer, in interactions leading up to a job offer, or after work begins on the job.

A structural model of labor demand allows us to estimate the extent to which inexperienced employers' uncertainty about workers causes workers to submit higher bids. Uncertainty about the market plays a key potential role in explaining why employer demand may differ with their

[1] In markets without posted prices, sellers often go to great lengths to elicit information about the extent to which buyers have sampled other outlets or products.

[2] This experiment is ongoing, so exact results are not presented here.

experience. Several comparative static results suggest that new employers learn while conducting searches to fill early job openings. This has implications for how workers bid for jobs due to a discontinuity in their payoffs around being hired. A distinct job opening on oDesk involves matching with a small number of workers (typically one), making employers' willingness to hire an individual applicant a function of the highest perceived order statistic of wage-adjusted productivity among available workers. As the precision of employers' beliefs about the distribution of match quality improves, the variance of the highest expected order statistic decreases. As a consequence of their own uncertainty, inexperienced employers may be less sensitive to the wage bid offered. If wage-offers are marked up over the opportunity cost of work, workers' payoffs are discontinuous around the hiring threshold and the increased variance in the distribution of the highest order statistic due to uncertainty may rationally result in workers submitting experience-dependent bids. This intuition is general beyond labor or matching environments and applies to other contexts with unit demand; noise in the environment when evaluating products in goods markets or applicants in labor markets has the potential to change price elasticities, altering equilibrium markups.

Other frictions may be a more direct consequence of employer inexperience, such as the expectation that there will be hassle costs associated with applying to or working for an inexperienced employer. Workers may anticipate that inexperienced employers interview more applicants and hire fewer interviewed workers. In the data, an employer on the fifth job posting conducts about half as many interviews as on the first job posting, and, on average, employers on the fifth job are around twice as likely to hire as those on the first job. Another possible source of hassle is that workers may have to explain to their employer how to use the market, how to initiate a contract, or how to review work. To the extent that any of these potential costs are not compensated, they may be passed through to employers in the form of higher wage offers.

Assuming that observed wage bids are the equilibrium of a differentiated products Bertrand game allows us to recover workers' outside wages and hassle costs. We find that differences in demand are responsible for about 25 percent of the difference in bids, with hassle costs–the residual–responsible for the remainder.

Because wage bids are likely correlated with unobserved worker and job-match quality, estimating demand requires an instrument that shifts workers' costs but not employer demand. We use worker-country exchange rates as the instrument. Workers on oDesk come from different countries, and they are exposed to differential cross-country macroeconomic fluctuations that are plausibly orthogonal to individual worker-level unobserved characteristics and employer demand after accounting for workers' participation on the platform. Because the worker's outside wage is typically paid in their local currency while the oDesk contract is paid in US dollars, we use the standardized log of the local currency to dollar exchange rate to instrument for the log of workers' wage offers. The residuals from this first stage regression serve as control functions, allowing us to estimate demand using logit or mixed logit models.[3] The estimates of the sum of the outside wage and hassle costs are plausible; they decrease with employer experience and are correlated with country-level gdp per capita.

Several additional empirical regularities motivate the conclusion that uncertainty and employers' need to learn about the market drive differences in bids. The most interesting relationship concerns search intensity and search outcomes. One might expect that employers who interview more candidates or search more intensely would have weakly better project outcomes. Additional search does not result in better outcomes in the data. In fact, the relationship between outcomes and search intensity is negative in the most homogeneous job categories.

This result is consistent with the predictions of a search model with learning about an unknown distribution.[4] With an unknown distribution of match quality, as employers become more certain about the distribution, the hiring reservation value declines. This is because the option value of waiting falls as the precision of beliefs increases. That is, the value of the marginal interview declines with cumulative experience. When inexperienced employers hire after only a small number of interviews, they forgo further learning opportunities. These employers are likely revealed-preferred to have found a very good match on an early interview that exceeds a high reservation value. Employers who conduct more interviews are likely to have worse job outcomes because the reservation value declines with cumulative search and they are revealed-preferred to have not found a match that exceeds the high initial reservation value. In the data, an inexperienced employer who hires after having interviewed 5-8 workers is 10 percentage points less likely to report a successful job outcome than an employer that interviews a single employee.

The oDesk revenue model currently imposes a tax of 10 percent on all transactions on the site.

[3] We do not estimate a dynamic model of employer demand. An existing literature estimates empirical models of optimal search with learning, but they are challenging because the parameters of the distribution of beliefs are unobserved, serially correlated state variables. Allowing for experience-dependent prices adds an additional layer of complexity that requires more restrictive modeling assumptions.

[4] For theory, see Lars Ljungqvist and Thomas Sargent's textbook treatment of the job search model with an unknown offer distribution. The labor economics literature has extensively considered learning models over the career. See Gibbons and Farber (1996), Altonji and Pierret (2001), Lange (2007), Kahn (2013), and Kahn and Lange (2014). For empirical work in the context of learning one's own type, see Miller (1984) and Arcidiacono, Aucejo, Maurel, and Ransom (2013). This paper is closer to the literatures in industrial organization and marketing about how demand changes as consumers learn about product quality (see Erdem and Keane, 1996; Ackerberg, 2003).

Using estimated differences in demand according to employer experience and estimated differences in workers' costs of interacting with employers of different experience levels, we examine whether a tax rate that varies with employer experience would increase transactions and generate additional profits for the platform. Although inexperienced employers are less elastic, we find that oDesk should charge lower fees to the inexperienced employer segment to account for spillovers to future demand. Optimal platform pricing should account for how decentralized actors' wage offers depends on the information that the platform provides.

The paper proceeds as follows: Section 2 provides stand-alone details about oDesk and how employers search in this market. Section 3 introduces evidence on differential pricing by observable experience. Section 4 estimates demand by employer experience levels and assesses sources of frictions. Section 5 evaluates how changes in the market affect hiring rates and profitability. Section 6 provides additional evidence that a search model with learning can fit important moments of the data, lending additional support to the hypothesized mechanism driving the result. Section 7 concludes.

2 The Search Process on oDesk.com

oDesk.com is an online platform that allows employers to contract with remote workers.[5] It facilitates both search and matching for workers as well as remote task and project management and payments. The site includes a range of jobs where output can be delivered electronically, and the most frequently observed job categories are Web Development and Administrative Support. Jobs tend to be short-term spot transactions, with the majority lasting less than three months. Around 85 percent of the transactions in the market span international borders and, therefore, constitute international labor services trade. This paper focuses on oDesk's matching role.

An employer who has a job arrive can create an account, free of charge. To post a job opening, the employer is asked to select the job category, give the job a title, and describe the work to be done as well as the skills needed. The expected duration can also be noted. Once the posting is in the system, potential employees can learn about the job through search on the site or through automatic notification. Interested workers can submit applications for the job posting. Employers also have the option of searching worker profiles directly and inviting candidacies from individual workers. Workers' profiles, visible to potential employers, contain information about their skills,

[5] Other prominent platforms included elance and Guru, the first of which merged with oDesk in 2014. The merged company has recently changed its name to Upwork.

education, prior offline work experience, and experience on oDesk. The country where they are located is also displayed prominently on the profile. For those workers that have experience on the site, the profile shows a summary feedback score out of five, and all worker profiles include a requested hourly wage rate that the worker is free to change at any time.

After receiving applications or initiating candidacies, employers can opt to interview any number of workers for the job. If applicants agree to be interviewed, the interview usually takes place via Skype. Whether an interview actually occurs is not recorded in the oDesk database, so the remainder of the paper will refer to the intention to interview as an interview. An employer can choose to hire an applicant with or without having interviewed them first. Upon hiring, the employer can monitor workers via software provided by oDesk, and oDesk manages all payments. When a job is complete, each party is asked to give feedback on the other, and the employer is also asked whether or not the job has been completed successfully.

The employers' search sequence on a single job, as described above, is observed on each of their successive job postings. For each posting, the data contain information about the complete applicant pool; which candidates, if any, are interviewed; and which candidate, if any, is hired. While several recent papers use oDesk data (Agrawal et al., 2013a and 2013b; Horton, 2013; Lyons, 2013; Pallais, 2013; Ghani, Kerr, and Stanton, 2014; and Stanton and Thomas, 2012), this is the first paper to use this empirical setting to study how employers' hiring processes evolve over time and how workers respond to employer experience.

3 Data and Summary Measures

The sample used in this paper consists of those employers who posted administrative support job on oDesk between January 2006 and June 2010. The employers are mostly located in the United States, and the person posting jobs and evaluating workers may either be a private individual or someone hiring on behalf of a firm. Of the $16,744$ employers who post jobs without ever hiring previously, $10,543$ received all job applications before going on to post a second job. For these jobs, the hiring periods can be segmented by time. Figure 1a presents the histogram of the distribution of total number of jobs posted per employer. 99 percent of employers in the data post fewer than 30 jobs throughout the period. Figure 1b presents the same histogram, this time looking at the distribution of the number of job openings resulting in a hire per employer. Figure 1c plots the distribution of the number of postings conditional on having hired once.

The paper analyzes Administrative Support jobs because, among the most frequently observed job categories, these jobs are the least likely to require on-the-job coordination among the workers hired by the same employer. As a consequence, the identity of previous hires is less likely to affect subsequent hiring behavior through a production-related channel. The intent is to narrow in on search-process-related dependencies between past and future hiring behavior. Around one third of the Administrative Support jobs in the data are for Data Entry, around 20 percent are Personal Assistant, and another 20 percent entail Web Research.[6] In an effort to ensure that jobs are similar, later analyses that estimate demand restrict the sample to Data Entry jobs or Data Entry and Web Research jobs.

3.1 Employer Behavior

Tables 1A-1D present summary statistics about job postings, grouped by the number of hires conducted (across all job categories) for the relevant employer. Table 1A includes all job posts in the sample. Table 1B includes only the subset of these jobs for which the job search was complete by the time the prior or next job search was complete. We call this the non-overlapping or sequential sample, and this restriction reduces the total number of postings from $53,803$ to $37,737$. Tables 1C and 1D present summary statistics for the sample restricted to Data Entry jobs.

Those employers without hiring experience in Table 1A receive an average of 33.4 applications prior to closing the job or making a hire. In the sequential sample, inexperienced employers receive an average of 34.7 applications per job. The number of applications per job tends to fall with employer experience, with the largest decline taking place between hiring zero and one prior worker. The average hourly wage bid for inexperienced employers in the overall sample is 4.94 USD, and in the sequential jobs sample it is 5.13 USD. Column 3 shows that the average hourly bid declines with employer experience. By job two, average bids fall to 4.43 USD (4.54 USD in the sequential sample). By job four, or later, employers receive bids that average 3.56 USD (3.56 USD in the sequential sample).

Learning about the market is one potential mechanism explaining why employer behavior changes with experience, and some evidence of learning is present in these summary measures. Employers plan to conduct an average of about 3.8 interviews per job posting when inexperienced, both for all jobs and in the sample of non-overlapping jobs. The data record that an interview takes place whenever an employer sends an interview request to a worker, whether that worker was

[6] The remainder are: Email Response Handling, Transcription, and Administrative Support - Other.

invited to apply for the job or initiated the candidacy himself. Column 4 of Table 1 shows how this statistic evolves with employer experience. Employers tend to interview more applicants on earlier job postings, again with the largest decline taking place after the first hire.

The overall probability that a job opening results in an eventual hire also differs substantially with employer experience. Column 5 shows that inexperienced employers are least likely to hire. The probability of hiring doubles for employers with at least 1 prior hire.

To motivate that differences in demand may play a role in the observed differences in bids documented in Table 1, Figure 2 provides a summary sketch of the relationship between wages bid and the probability of being hired for workers applying to data entry jobs. To attempt to remove unobserved heterogeneity, the sample is restricted to workers from the Philippines (the most frequent country of origin for job applicants). Bids are also residualized based on observable worker characteristics and time fixed effects. The two connected lines come from separate local polynomial regressions evaluated at the plotted nodes. It is important to note that this figure is only an illustration the demand may potentially differ based on experience on oDesk. With this caveat in mind, the relationship between bids and the probability of hire looks dramatically different for employers with and without experience.

3.2 Workers' Hourly Bids and Hourly Wages

Table 2 investigates how wage bids vary with employer experience. The omitted category is employers prior to making their first hire. The coefficients show the percentage change in the average wage bid received when workers observe that the employer has hired one, two, three, four or five or more prior hires. Worker-level controls include the applicants' feedback and an indicator for missing feedback, an indicator for having done a previous job without receiving feedback, and the number of previous hires. Job-level controls include job category fixed effects and expected job duration fixed effects. Monthly time fixed effects are also included, and, to isolate changes in what workers observe when bidding, the sample includes only those jobs that are sequential.

Column 1 shows that, in the cross section, employers with one observable prior hire receive bids that are 7 percent lower than employers who have no observable experience hiring. This falls to about 9 percent and then 11 percent for employers with two and three prior hiring spells, respectively. Employers with five or more hiring spells receive bids that are 17 percent lower than employers with no experience. Column 2 adds controls for changes in qualitative features of the job post, including a third-order polynomial in the number of characters in the description and fixed effects for expected duration. The results change very little.[7]

This pattern of examining changes in the qualitative feature of job posts is repeated throughout the table. Columns 3 and 4 add employer fixed effects. A given employer receives bids that are about 4.5 percent lower with one prior hire and about 10 percent lower with 5 or more prior hires compared to bids when he has no experience. Columns 5 and 6 are estimated with worker fixed effects. The result is robust within worker, as a given worker submits lower bids to experienced employers. Within worker, bids to employers on their fifth or more hiring spell are, on average, about 9 percent lower than bids to employers with no observable experience. The last two columns are estimated with both employer and worker fixed effects, and the parameter estimates show the same pattern of declining wage bids with experience.

Panel B of Table 2 shows that it is employer experience, rather than public employer feedback, that is associated with lower bids. In these columns, the sample is limited to employers' with no observable hiring experience and to employers with five or more prior hires. Indicators for the employer having no observable feedback and for the employer having observable feedback of 4.5 or higher are interacted with the indicator for having five or more prior hires. Including these controls does not change the main result that experienced employers receive lower bids. In fact, the main point estimate captures the effect for experienced employers with bad feedback; these employers still receive bids that are significantly lower than inexperienced employers. The interactions of employer experience and good employer feedback are, in most cases, a fraction of the main effect. In specifications comparing how bids change within employer, the estimates are statistically indistinguishable from zero. With worker fixed effects, the feedback interactions have positive parameter estimates. However, if a lack of feedback was the primary driver of bid differences, one would instead expect that the interaction of good feedback and experience would result in lower bids. When adding both worker and employer fixed effects, these interactions are both small. The main message from these regressions is that the main effect of experience tends to swamp any differences in how employer feedback evolves over time.

Panel C of Table 2 shows that the relationship between employer experience and lower bids is robust to including controls for feedback that the employer has given to workers on prior jobs. As in Panel B, the sample is employers with no prior hiring experience and employers with five or more prior hires. Indicators for the employer having given no observable prior feedback and for having given good observable prior feedback are interacted with an indicator for having five or more prior

[7] The results of a pilot experiment that varies only employer experience also help to alleviate concerns about changing characteristics of job postings with employer experience.

hires. The significant negative coefficients on having experience remain, and increase in magnitude, when including these controls. The estimated coefficients on these controls are positive in sign, which is consistent with the interpretation that workers attribute having left no feedback or giving good feedback as positively correlated with employers' willingness to pay.

To further visualize these bid differences, Figure 3 presents a binned scatterplot of within worker-month mean bids to data entry jobs posted by inexperienced employers against the within-worker month mean bid to jobs posted by experienced employers. The 45-degree line is also plotted in red. Workers' who bid very low wages for experienced employers appear more likely to increase their bid for inexperienced employers; the difference between bids appears to decline for workers who bid higher wages for experienced employers. Although it is not causal, the bid differences appear to be greatest in regions where the demand curves plotted in Figure 2 differ most. We now turn to estimating demand more credibly. The goal from this exercise is to assess how much predictable differences in demand explain differences in bids by experience levels.

4 Demand, Wage Offer Equilibrium, and Markups

4.1 The Employer's Problem

Employer i has tasks requiring a finite amount of quality-adjusted worker time. Workers arrive to job openings stochastically, and the employer observes a signal of worker quality that is commonly valued in all production tasks, q_j, a signal of employer-specific quality, m_{ij}, and a purely idiosyncratic Type-1 extreme value distributed error, ε_{ij}. In the model, workers are assumed to be available when they initiate an application, which is reasonable given aggregate statistics on potential unobserved availability.[8] To complete required production, the employer evaluates workers and chooses that

[8]This requires that the probability that a worker receives two offers over a time interval Δ is sufficiently small. Similar justifications are used when formulating many stochastic processes. For example, if, from the worker's perspective, job offer arrivals follow a Poisson process, then the probability of receiving a job offer in the interval $(t, t+\Delta)$ is $\lambda\Delta + o(\Delta)$, and the probability of two offer arrivals in $(t, t+\Delta)$ is $o(\Delta)$.

Although declined job offers are not observed, this assumption on arrival of offers seems to be reasonable given data from two other measures. First, the observable arrival rate of interview requests fits this assumption. When the worker-day is the interval of analysis, only 1.6% of the worker-days in the sample have more than 1 interview request arriving; 0.15% of the worker days have more than 2 interview requests arriving. Second, a post-candidacy survey asks employers for reasons why particular workers were not hired or workers for reasons why they exited the active candidate set. In 0.02 percent of responses, employers or workers explicitly report a realized scheduling conflict. In some cases, employers invite workers to apply for a job. In 8.7% of the instances when an employer initiates contact with a worker, the worker reports that they are too busy to follow through on the invitation to apply or are not interested in the job. Explicit rejections of application invitations are dropped when estimating demand.

worker j of the J_i applicants that has the highest expected wage-adjusted quality,

$$\frac{q_j m_{ij} \exp(\varepsilon_{ij})}{w_{ij}},$$

subject to the maximal worker being greater than an outside option. With this outside option, the choice set for job opening i is $C_i = \left\{ \{q_j, m_{ij}, \varepsilon_{ij}, w_{ij}\}_{j=1}^{J}, \{0, 0, \varepsilon_{i0}, 0\} \right\}$. The employer's demand is a corner solution and is equivalent to choosing the option that provides the highest perceived $\log(q_j) + \log(m_{ij}) + \varepsilon_{ij} - \log(w_{ij})$. The probability that employer i chooses worker j is then

$$\Pr\left(\log(q_j) + \log(m_{ij}) + \varepsilon_{ij} - \alpha \log(w_{ij}) > \log(q_k) + \log(m_{ik}) + \varepsilon_{ik} - \alpha \log(w_{ik})\right)$$
$$\text{for all } k \in C_i. \tag{1}$$

The parameter α scales the sensitivity of log wages relative to the variance of the idiosyncratic error.

The perceived opening-specific match quality is modeled as

$$\log m_{ij} = \eta_{ije}(\mu_i) + \xi_{ij}, \tag{2}$$

where $\eta_{ije}(\mu_i)$ is a random effect with employer-specific mean μ_i that is orthogonal to common quality by construction and ξ_{ij} is an idiosyncratic error component. Although it is not necessary for the general intuition of the model, the familiar normal-normal learning is a convenient way to parameterize uncertainty and learning. Suppose that employers don't know μ_i with certainty but believe the distribution of μ_i is normal with prior mean and variance $\mu_{i(0)}$and $\sigma^2_{\mu(0)}$. The random effect η_{ije} is also normally distributed with mean μ_i and variance σ^2_η, and ξ_{ij} is iid normal noise in evaluations with mean zero and variance σ^2_ξ. With this setup, when an employer observes $\log m_{ij}, \eta_{ije}(\mu_i)$ is not learned with certainty.

To simplify the intuition, assume any interaction with a worker, whether an interview or a hire, provides the same signal.[9] As employers interview or hire workers, beliefs about the distribution of μ_i change with the observation of additional workers. As beliefs about μ_i change, a Bayesian employers' perception of the signal changes with experience and becomes more precise. To see this, if the employer interacts with a worker and observes $\log m_{ij}$, the distribution of the posterior mean

[9] A detailed analysis of forward-looking strategic search and stopping where an interview provides a less precise signal than a hire is available on request.

is

$$\mu_{(e+1)}|\mu_{(e)},\sigma^2_{\mu(e)} \sim N\left(\left[\frac{\log m_{ij}}{\sigma^2_\eta+\sigma^2_\xi}+\frac{\mu_{(e)}}{\sigma^2_{\mu(e)}}\right]/\left(\left(\sigma_\eta^{-2}+\sigma_\xi^{-2}\right)+\sigma^{-2}_{\mu(e)}\right),\frac{1}{\left(\sigma^2_\eta+\sigma^2_\xi\right)^{-1}+\sigma^{-2}_{\mu(e)}}\right). \quad (3)$$

This updating affects how employers estimate productivity having observed $\log m_{ij}$. The conditional distribution of the individual random effect given the signal is

$$\eta_{ije}|\log m_{ij},\mu_{(e)},\sigma^2_{\mu(e)} \sim N\left(\left[\frac{\log m_{ij}}{\sigma^2_\xi}+\frac{\mu_{(e)}}{\sigma^2_{\mu(e)}+\sigma^2_\eta}\right]/\left(\sigma_\xi^{-2}+\left(\sigma^{-2}_{\mu(e)}+\sigma_\eta^{-2}\right)\right),\frac{1}{\sigma_\mu^{-2}+\left(\sigma^{-2}_{\mu(e)}+\sigma_\eta^{-2}\right)}\right), \quad (4)$$

which has conditional expectation

$$\hat{\eta}_{ij}=E\left(\eta_{ij}|\log m_{ij},\sigma^2_{\mu(e)}\right)=\left[\frac{\log m_{ij}}{\sigma^2_\xi}+\frac{\mu_{i(e)}}{\sigma^2_{\mu(e)}+\sigma^2_\eta}\right]/\left(\sigma_\xi^{-2}+\sigma^{-2}_{\mu(e)}+\sigma_\eta^{-2}\right). \quad (5)$$

As employers gain experience, the prior belief about $\mu_{(e)}$ receives more weight and the noisy signal receives less weight. This is because $\sigma^2_{\mu(e)}+\sigma^2_\eta$, the total variance around $\mu_{i(e)}$, declines. The total variance around $\mu_{i(e)}$ itself comes from two sources, the actual variance around the true mean, σ^2_η, and uncertainty about the mean through the variance of the prior at experience level e, $\sigma^2_{\mu(e)}$. Because $\partial\sigma^2_{\mu(e)}/\partial e<0$, the variance of the conditional expectation, $\sigma^2_{\hat{\eta}e}$, also declines with experience

$$\frac{\partial\sigma^2_{\hat{\eta}e}}{\partial e}<0, \quad (6)$$

As $\sigma^2_{\hat{\eta}e}$ becomes smaller, the expected highest order statistic in the employer's choice set,

$$\left\{\log\left(q_j\right)+E\left(\eta_{ijc}|\log m_{ij},\sigma^2_{\mu(e)}\right)+\varepsilon_{ij}-\alpha\log\left(w_{ij}\right)\right\}_{j=1}^{J},$$

also falls.

Experienced employers are also likely to have a better understanding of their actual expected benefit from hiring. The presence of a heterogeneous (across employers) but persistent μ_i means that $\log m_{ij}$ contains a common component for all workers in the choice set. This reflects that, for the same choice set, different employers will have heterogeneous probabilities of choosing the outside option.

The main focus is on workers' responsiveness to observed inexperience and not on the speed of learning. Therefore we decide to estimate demand separately for experienced and inexperienced employers. An advantage of this approach is its robustness to the form of learning, or to misspecification. This allows us to recover $\overline{\sigma^2_{\hat{\eta}Experienced}}$ for the experienced and $\overline{\sigma^2_{\hat{\eta}Experienced}}$ for the inexperienced, but we do not separately identify σ^2_{ξ} and the path of $\sigma^2_{\mu(e)}$. With our approach all that is required for identification is that the conditional expectation in equation (5) becomes more precise and centered around the true μ_i over time. When this is the case, experienced employers are less likely to hire based on noise, ξ_{ij}, and are less likely to search and experiment because of uncertainty about μ_i.

To complete the model, we introduce a time subscript, t, that is intended to make clear which worker characteristics (or wage bids) may change with calendar time. We assume that the commonly valued component of productivity has functional form

$$\log q_j = X_{jt}\beta + \upsilon_{jt}, \tag{7}$$

where observable common factors X_{jt} are "priced" in the market, including feedback, the workers' past experience, and observable resume characteristics. Market participants likely observe more information about common quality than the econometrician, and this υ_{jt} is likely to be correlated with the wage offered.

4.2 Estimation

4.2.1 First Stage

To address concerns about the endogeneity of wage offers due to correlation with unobserved common (υ_{jt}) or worker-job match (that portion of $\log m_{ij}$ that is observed to workers when bidding) quality, we use an instrumental variables strategy that exploits changes in the local currency to dollar exchange rates for individual workers and their competitors from other countries. Workers' local labor market opportunities are indexed in the local currency; with frictions in the offline labor market due to sticky offline wages or imperfect financial markets, local offline opportunities are likely to adjust more slowly to exchange rate changes than online transactions. With appreciation of the local currency relative to the dollar (one dollar provides more local currency units), workers' wage offers when bidding for jobs reflect differences in their alternative labor market opportunities and hence their opportunity costs of online work. This change in the attractiveness of outside labor markets has a substantive and statistically significant effect on workers' wage bids.[10]

The instrument induces substantial variation in wages. Figure 4 illustrates the variation in the instrument; the vertical axis is the ratio of the monthly mean of Indian workers' log bids and the monthly mean of Filipino worker's log bids; the horizontal axis is the log of the Filipino to Indian exchange rate. The ratio of bids tracks the exchange rate ratio closely. The Philippines and India have the largest shares of oDesk workers, so this source of variation is most important in the data. Similar patterns are found across different countries. To make the instrument comparable across countries, we take z-scores using the aggregate time series for each country over the sample period.

The ideal instrument holds fixed the distribution of unobserved quality for applicant workers and shifts relative prices among workers. This enables estimation using control function methods.[11] However, with any cost-shifting instrument, decisions to participate on oDesk are likely to vary with the instruments as well as the bid decision conditional on bidding. To see how endogenous participation decisions affect the ability to estimate demand, suppose the instrument using local exchange rates is denoted Z_{jt} and the first stage regression is

$$\log\left(w_{ij}\right) = a_1 + a_2 e_{it} + Z_{jt}\gamma_1 + X_{jt}\gamma_2 + u_{jt}. \tag{8}$$

The residual contains worker and match quality, both of which enter the optimal bid with unknown coefficients λ_1 and λ_2, resulting in

$$u_{jt} = \upsilon_j \lambda_1 + \left(\eta_{ijt} + \xi_{ij}\right)\lambda_2 + \nu_{jt}.$$

For exposition purposes, we ignore correlation between Z_{jt} and $\left(\eta_{ijt} + \xi_{ij}\right)$, but illustrate the consequences of correlation between Z_{jt} and common unobserved quality υ_{jt}. If $Cov\left(Z_{jt}, \upsilon_{jt}\right) \neq 0$, then the exclusion restriction is violated. Consistent estimation in the control function setting requires that u_{jt} contains any part of unobserved quality that is correlated with price, as fitted values of $\hat{\gamma}$ are uncorrelated with unobserved υ_{jt} by construction. When the exclusion restriction fails,

$$E\left(\hat{\gamma}_1\right) = \gamma_1 + E\left\{\left(Z'_{jt}Z_{jt}\right)^{-1} Z'_{jt}\upsilon_j\lambda_1\right\},$$

[10] Calendar time fixed effects remove aggregate market changes.

[11] Results from control function estimation may be sensitive to estimation error in the control function. The precision of the first-stage estimates suggests that this problem is unlikely to be severe. Petrin and Train (2009) discuss the properties of the estimator. We derive a new bias correction method in the appendix and find that results with and without the bias correction are very similar.

which results in residuals for a control function that contain a part of unobserved quality

$$\begin{aligned} CF &= \log\left(w_{ijt}\right) - Z_{jt}\left\{\gamma_1 + \left(Z_{jt}'Z_{jt}\right)^{-1} Z_{jt}'\left(\xi_j\right)\lambda_1\right\} \\ &= \log\left(w_{ijt}\right) - Z_{jt}\gamma_1 - P_Z \upsilon_{jt}\lambda_1 \\ &= \nu_{jt} + \left(I - P_Z\right)\upsilon_j\lambda_1 + \left(\eta_{ijt} + \xi_{ij}\right)\lambda_2. \end{aligned}$$

This means that, due to sorting, the instrument does not eliminate endogeneity concerns. With sorting, only the part of unobserved quality orthogonal to the projection matrix of υ_{jt} is absorbed in the residual.

To account for how differences in participation change the distribution of unobserved worker quality, we use a second instrument that captures competition among workers. Because workers can see statistics about aggregate bids for each job opening, changes in observed bids by other workers are likely to affect participation decisions; to capture this, we use other workers' average bids (aggregate) excluding own country. We then remove a monthly time fixed effect is removed to ensure that only cross-sectional variation is used for identification. This instrument picks up cross-country differences in the intensity of competition from other workers. The instrument is correlated with participation behavior, but other workers' bids are plausibly uncorrelated with unobserved quality from the perspective of employers after controlling for aggregate sources of time effects. With this additional instrument, we estimate a participation equation at the monthly level using a probit model to correct for any potential sorting based on the instrument. We include the inverse mills ratio estimated from this probit model into the first stage when estimating the control function.

Table 3 provides details about the participation probit and the first stage regression. The results make clear that the instruments are strong. However, the parameter estimate on the local exchange rate is sensitive to having the inverse mills ratio included, suggesting that there is sorting on the instrument.[12]

4.2.2 Choice Model

With the employer's problem defined and a strategy for accounting for endogeneity, the estimated model must still account for heterogeneity across employers through μ_i and changing perceptions about $\hat{\eta}_{ij}$. A modified random-parameters logit model accomplishes both of these objectives. Choice

[12] We have also estimated models that do not include the inverse mills ratio in the first-stage regression used to obtain the control function. Both experienced and inexperienced employers are estimated to be less own-bid elastic than in richer specifications with the inverse mills ratio included.

probabilities conditional on $\hat{\eta}_{ij}$ and μ_i are

$$\begin{aligned} p_{j|i,\hat{\eta},\mu} &= \exp\left(\hat{\eta}\left(\mu_{i(e)}, \sigma^2_{\mu(e)}\right) + X_{jt}\beta - \alpha \log\left(w_{ij}\right) + CF_{jt}\psi\right) / \\ & \left(1 + \Sigma_j^{J_i} \exp\left(\hat{\eta}\left(\mu_{i(e)}, \sigma^2_{\mu(e)}\right) + X_{jt}\beta - \alpha \log(w_{ij}) + CF_{jt}\psi\right)\right). \end{aligned} \tag{9}$$

Marginal choice probabilities for employer i and each potential choice j come from integrating over the conditional distribution of $\eta_{ije}\left(\mu_i\right)$ and the distribution of μ_i

$$p_{ij} = \int\int p_{j|i,\hat{\eta},\mu} f\left(\hat{\eta}_{ij}|\mu_i, \sigma^2_{\hat{\eta}e}\right) f\left(\mu_i|\bar{\mu}_e, \sigma^2_{\mu(e)}\right) d\hat{\eta}_{ij} d\mu_i, \tag{10}$$

where $\hat{\eta}_{ij}\tilde{\ }N\left(\mu_i, \sigma^2_{\hat{\eta}e}\right)$ and $\mu_i\tilde{\ }N\left(\bar{\mu}_e, \sigma^2_{\mu(e)}\right)$.

$$\hat{\eta}_{ij} = E\left(\eta_{ij}|\log m_{ij}, \mu_{i(e)}, \sigma^2_{\mu(e)}\right) = \left[\frac{\log m_{ij}}{\sigma^2_\xi} + \frac{\mu_{i(e)}}{\sigma^2_{\mu(e)} + \sigma^2_\eta}\right] / \left(\sigma_\xi^{-2} + \sigma^{-2}_{\mu(e)} + \sigma_\eta^{-2}\right).$$

The distribution of random parameters is identified, as the outside option does not contain $\eta_{ij}\left(\mu_i\right)$ or μ_i. The distribution of μ_i is identified in the data by variation in the attractiveness of employers' candidate pools. When an employer with a very favorable candidate pool chooses not to hire, part of this decision is due to the idiosyncratic Type-1 extreme value error and part may be due to μ_i; similarly, when an employer with an unfavorable candidate pool chooses to hire, he may have a high μ_i or may have received a good draw of ε_{ij}. The extent to which employers with unfavorable pools hire and the extent to which employers with favorable pools choose the outside option identifies σ^2_μ and $\bar{\mu}$. Identification of $\sigma^2_{\hat{\eta}e}$ comes from the extent to which employers make choices that differ from the distribution of ranks implied by the extreme value error and the parameters α, β, and ψ given $\log\left(w_{ij}\right)$, X_{jt}, and CF.[13]

The parameters in equation (10) are estimated using just-identified simulated method of moments with moment conditions that exploit orthogonality between the control functions, characteristics, wage bids, and the difference between the actual choice and the simulated choice probabilities under $f\left(\hat{\eta}_{ij}|\mu_i, \sigma^2_{\hat{\eta}e}\right) f\left(\mu_i|\bar{\mu}_e, \sigma^2_{\mu(e)}\right)$. In the simulation procedure, MLHS draws for each employer are generated for μ_i. These sequences are designed to provide good coverage. For each of these initial draws, a larger sequence of standard normal random variables are drawn for each applicant j, allowing for the estimation of the distribution of $\hat{\eta}$ conditional on each draw of μ_i.

[13]Note that $\sigma^2_{\hat{\eta}e}$ is identified relative to the variance of the extreme value error, which is normalized for all alternatives. Because the outside option does not have $\hat{\eta}_{ij}$, $\sigma^2_{\hat{\eta}e}$ is identified based on variation in inside-alternative choices relative to the choice of the outside option. The identification argument is similar to the identification of branch-wise variances in nested logit models.

The model is estimated separately for inexperienced employers and for experienced employers on their first job posting in Data Entry after having hired on two prior jobs in any job category.[14] This sample split allows us to compare the conditional distribution of $f(u_i|Experience > 0)$ and the unconditional distribution taken from estimates on inexperienced employers.

With estimates of demand parameters $\hat{\theta} = \left\{\beta, \alpha, \bar{\mu}_e, \sigma^2_{\mu}, \sigma^2_{\hat{\eta}e}, \psi\right\}$, in hand, workers' outside wages (opportunity costs) plus other application costs are recovered from the equilibrium of a Bertrand Nash game in prices.

4.3 The Workers' Problem and Markups

Each worker's objective when deciding what bid, if any, to make takes into account employers' unit demand, p_{ij}, which is a function of the wage bid, w_j, and (left implicit) all other wage bids w_{-j}. Because oDesk has an ad-valorem fee τ, the wage upon hiring is $w_j (1+\tau)^{-1}$ and the wage the employer pays is w_j. If the worker is not hired, the flow outside opportunity cost is c_j. The worker then trades off the hiring probability and higher wages when bidding on job opening i to maximize

$$\max_{\log w_j} p_{ij}(\log w_j) \times \exp\left(\log w_j - \log(1+\tau)\right) + \left[1 - p_{ij}(\log w_j)\right] \times c_j.$$

Note that c_j captures opportunity costs plus any hassle costs of work.[15]

This objective is static, but the first order conditions for the workers' problem are the same if one assumes that no other jobs offers arise over the duration of work on job i. The system of equations determining equilibrium bids, holding constant the number of job applicants observed on each job, comes from each worker's first order condition.

$$\frac{\partial p_{ij}}{\partial \log w_j}\left(w(1+\tau)^{-1} - c_j\right) + p_{ij} w (1+\tau)^{-1} = 0$$

[14] We are also in the process of estimating models that allow for truncation of $f\left(u_i|\bar{u}, \sigma^2_{\mu}, Experience > 0\right)$, as employers selectively use the platform with experience. The equation for p_{ij} is modified in this case such that the integrand is $\Pi_e p_{j|i}$. We don't report results for this case, as it is more restrictive, but truncation is found to be important.

[15] The objective could alternatively be written as

$$\max_{\log w_j} p_{ij}(\log w_j) \times \exp\left(\log w_j - \log(1+\tau) - c_H\right) + \left[1 - p_{ij}(\log w_j)\right] \times c_{jO}$$

where c_H is a hassle cost paid upon finding employment and c_{jO} is the outside wage. The first order condition in this case makes clear that only $c_j := c_{jO} + c_H$ is identified.

which implies

$$c_j = w\left(1+\tau\right)^{-1} + p_{ij} w_j \left(1+\tau\right)^{-1} / \frac{\partial p_{ij}}{\partial \log w_j}. \tag{11}$$

In the mixed logit model, the term $\frac{\partial p_{ij}}{\partial \log w_j}$ in equation (11) is

$$\frac{\partial p_{ij}}{\partial \log w_j} = \int\int -\alpha p_{j|i,\hat{\eta},\mu}\left(1 - p_{j|i,\eta,\mu}\right) f\left(\hat{\eta}_{ij}|\mu_i, \sigma^2_{\hat{\eta}e}\right) f\left(\mu_i|\bar{\mu}_e, \sigma^2_{\mu(e)}\right) d\hat{\eta}_{ij} d\mu_i.$$

4.4 Demand Results

Table 4 shows estimates from a conditional logit model that makes the distribution of random parameters degenerate. The results are presented for two different samples of employers. The first two columns present estimates for employers on their first job posting; the model is estimated using data from 1781 postings. Columns 3 and 4 present results for the subset of employers who post at least one job in Data Entry after their second hire on the site. This subset comprises 1672 employers. The important result to note deals with differences in estimated elasticities for the two groups; our preferred estimates are in Columns 1 and 3 and include the inverse mills ratio in the first stage. The mean own-bid elasticity for inexperienced employers is -6.97. The mean own-bid elasticity is greater for experienced employers, at -11.57. However, with limited sample sizes, the 95% confidence intervals for the elasticities overlap.

With elasticities in hand, workers' costs are recovered by inverting the first order conditions that determine bids. In the data entry sample, the mean cost of working for an inexperienced employer is 3.08 USD per hours (before the oDesk fee) compared to 2.60 USD per hour for in the experienced sample. At the mean, this represents a difference in estimated costs of about 18.5%. In the sample of data entry jobs considered here, the mean bid gap between experienced and inexperienced employers is 26%. Thus, about 75% of the gap is due to cost differences. Although it seems that inexperienced employers are less elastic with respect to hiring, the differences in elasticities are estimated to be too small to account for all of the difference in bids.

The results in Table 5 are deprecated. We had estimated these results before instrumenting for participation with other countries' bids; they were similar to the conditional logit results without this additional instrument. Our goal with using these results was to assesses whether differences in uncertainty about workers, through changes in $\sigma^2_{\hat{\eta}e}$, changes in the distribution of μ_i with experience, or both explain some of these differences in elasticity. Panel A first presents the most relevant estimated values of $\hat{\theta} = \left\{\beta, \alpha, \bar{\mu}, \sigma^2_{\mu}, \sigma^2_{\eta e}, \psi\right\}$ from the demand equation. The first row of the Panel reports the estimates of α, the coefficient on wage bids, from the mixed logit model, including the control function after instrumenting for wages using the exchange rate instrument described in the previous subsection. Column 1 shows that the estimated coefficient for inexperienced employers is XX. The experienced sample of employers has an estimated coefficient of XX. The parameter estimates on the control function, ψ, account for unobserved worker heterogeneity.[16] The value of σ_μ is estimated imprecisely for experienced employers, but the distribution of latent μ for the experienced has mean of XX compared to a mean of XX for the experienced. This change in the distribution reflects that inexperienced employers include those who learn the market is of limited value to them and exit after the first job.

Turning to the estimates of the parameters associated with employer learning in the model, the random effects estimates decline with employer experience. The estimated standard deviation for inexperienced employers is XX and is XX for experienced employers. These parameters, however, are too imprecisely estimated to reject that the distributions are the same.

Panel B of the table shows the implied average wage bid elasticities, mark ups and cost estimates from the model. Inexperienced employers have a wage bid elasticity of XX. The experienced employers have more elastic demand, with an average mean elasticity of XX. Using the Betrand equilibrium assumptions as described in the previous subsection, these elasticity estimates, together with observed wage bids, allow us to decompose wages into a mark up and a marginal cost term. The greater wage elasticity for experienced employers leads workers to charge them lower markups. Despite inexperienced employers' lower probability of hiring, the estimated average mark up over cost for inexperienced employers is XX percent, compared to XX percent for experienced employers. This reflects differing elasticities across these segments of employers.

The estimated marginal costs of applying for job openings also differs by employer experience. The mean cost across all workers applying to jobs posted by inexperienced employers is estimated to be XX USD. This falls to XX USD for workers applying to jobs posted by experienced employers. This difference suggests that the additional expected hassle cost of applying to an inexperienced employer is XX cents. The estimated marginal costs also vary by country in plausible ways. Workers in the United States have estimated marginal costs of applying to inexperienced employers of XX USD, which falls to XX USD for experienced employers. The equivalent figures are XX USD and XX USD for workers in India, and XX and XX USD for workers in the Philippines.

[16] The β parameters measuring the mean value attributed to worker characteristics other than wage bid are not reported.

5 Counterfactual Estimates: Hiring Rates and Platform Fees

The first part of Table 6 presents a decomposition of how much of the difference in wage bids can be attributed to the supply-side friction of differences in hassle costs and how much is an implication of workers' responses to inexperienced employers demand differences. Because the sample of employers used in estimation is in data entry, the bid differences by experience are greater in this subsample than those that were previously reported in the overall summary statistics for all Administrative Support jobs. Simulating the costs that applicants to inexperienced employers would face if those employers were experienced suggests that hassle costs make up a substantial fraction of the differences in costs, as bids fall dramatically when assuming that costs are held fixed at the experienced level. Average bids (post oDesk fee) to inexperienced employers fall to \$3.34 from\$3.95 when workers are assumed to have the implied opportunity and hassle costs that are estimated from the experienced employer sample. This change due to differences in costs represents about 75% of the bid gap between the experienced and inexperienced.

We now turn to the issue of differential platform fees. We are interested in understanding whether a different oDesk fee structure would improve platform profits. The current fees are ad-valorem, which, in this setting, makes analysing counterfactuals computationally difficult. Instead we focus on specific fees and illustrate the tradeoff that the platform faces in the simple situation where only employers who hire on the first job can hire on the future jobs.

The platform's objective is to maximize total profits, and it can do so with different fees on inexperienced and experienced employers, respectively denoted t_I and t_E. The platform's problem is

$$\max\left[t_I + t_E \times \Pr\left(H_E | H_I\left(w_I\left[t_I\right]\right), w_E\left[t_E\right]\right)\right] \times \Pr\left(H_I | w_I\left[t_I\right]\right).$$

where $\Pr\left(H_I | w_I\left[t_I\right]\right)$ is the probability that an inexperienced employer hires given wages $w_I\left[t_I\right]$ and $\Pr\left(H_E | H_I\left(w_I\left[t_I\right]\right), w_E\left[t_E\right]\right)$ is the probability an experienced employer hires as a function of wages $w_E\left[t_E\right]$ conditional on the first hire, $H_I\left(w_I\left[t_I\right]\right)$.[17] Notice that the platform does not set wages, only fees, but wages that employers face vary with platform fees unless demand is perfectly

[17]With multiple future jobs, the objective function becomes

$$\max\left[t_I + t_E \times F\left(NH_E | H_I\left(w_I\left[t_I\right]\right), w_E\left[t_E\right]\right)\right] \times \Pr\left(H_I | w_I\left[t_I\right]\right).$$

elastic.[18]

When users are uncertain about platform valuation, $\Pr\left(H_E|H_I\left(w_I\left[t_I\right]\right), w_E\left[t_E\right]\right)$ is not simply the truncated part of the demand curve above wages w_I but instead may change depending on how employers' beliefs about the platform evolve with experience. That $\Pr\left(H_E|H_I\left(w_I\left[t_I\right]\right), w_E\left[t_E\right]\right)$ specifically conditions on $H_I\left(w_I\left[t_I\right]\right)$ captures the possibility that experienced hiring may be affected by the identity of the marginal inexperienced employer. Variation in wages, induced by different platform fees, induces variation in the identity of the marginal employer.

This formulation so far says nothing about how beliefs evolve with employer experience. This leaves the learning process completely free, nesting models with myopic or anticipated learning. If employers anticipate that they will learn about their individual valuations for the platform, initial hiring may reflect employers' recognition of an option to no longer use the platform if the initial experience is unsuccessful. On the other hand, myopic employers who have low expectations of platform valuation may require inducement if the option value of gaining information is not recognized.

Using H_I as shorthand for $\Pr\left(H_I|w_I\left[t_I\right]\right)$ and H_E as shorthand for $\Pr\left(H_E|H_I\left(w_I\left[t_I\right]\right), w_E\left[t_E\right]\right)$, the first order conditions are

$$
\begin{aligned}
H_I + t_I \frac{\partial H_I}{dw_I}\frac{\partial w_I}{\partial t_I} + t_E H_E \frac{\partial H_I}{dw_I}\frac{\partial w_I}{\partial t_I} + t_E H_I \frac{\partial H_E}{\partial H_I}\frac{\partial H_I}{dw_I}\frac{\partial w_I}{\partial t_I} &= 0 \\
H_E \times H_I + t_E \frac{\partial H_E}{\partial w_E}\frac{\partial w_E}{\partial t_E} H_I &= 0
\end{aligned}
$$

The solution to the system of equations sets the fee for experienced employers equal to the monetary value of the optimal markup for a monopolist with zero marginal cost

$$
t_E^* = -\frac{H_E}{\frac{\partial H_E}{\partial w_E}\frac{\partial w_E}{\partial t_E}}. \tag{12}
$$

The fee for the inexperienced is

$$
t_I^* = -\frac{H_I}{\frac{\partial H_I}{dw_I}\frac{\partial w_I}{\partial t_I}} - t_E^* H_E - t_E^* H_I \frac{\partial H_E}{\partial H_I}. \tag{13}
$$

The first term in t_I^* is the standard markup for the segment of inexperienced employers in the

[18] For this analysis, wages that inexperienced employers face (after the tax) have an I subscript and wages that experienced employers face (after the tax) have an E subscript. Wages bid by the worker have a j subscript.

absence of dynamic considerations. This markup is reduced by the latter two terms, which adjust t_I to account for the spillover to future demand. The final term is of particular interest, which accounts for composition effects.

Estimates of most terms of interest come directly from the logit models, but some additional statistics are required to compute optimal taxes: the average of $\frac{\partial H_E}{\partial H_I}$ and $\frac{\partial w_I}{\partial t_I}$ and $\frac{\partial w_E}{\partial t_E}$.

Our estimates of the composition effects to get $\frac{\partial H_E}{\partial H_I}$ come from a logit model where a dummy for repeat hiring is regressed on an indicator that the employer hired on the first job. This model includes time fixed effects. The marginal effect of hiring on the first job on the probability of subsequent hiring is 0.26. With a baseline probability of hiring of 0.194, the latter composition term reduces the tax on the inexperienced by 5 percent of the experienced tax.

We use the workers' first order conditions to derive the equilibrium pass-through rate of a tax. This is

$$\frac{\partial w_j}{\partial t} = -c_j \left(1 + \frac{1}{(1-p_{ij})\,\alpha}\right)^{-2} \left(\frac{\frac{\partial p_{ij}}{\partial t}}{\left[(1-p_{ij})\,\alpha\right]^2}\right)\alpha, \tag{14}$$

where α is the coefficient on log price from the conditional logit model and p_{ij} is unit demand by employer i for worker j. Estimates of the average change in the probability of hiring with respect to a tax on all workers suggest that the amount workers receive is little changed. Simulating this change, $\frac{\partial p_{ij}}{\partial t} \approx -.008$; plugging in other parameters and estimates of p_{ij} suggest that workers' bids are reduced by about about .4 percent for a unit tax. Because $\frac{\partial w_j}{\partial t}$ is the wage to the worker, the pass through rate to employers, $\frac{\partial w_E}{\partial t_E}$ or $\frac{\partial w_I}{\partial t_I}$, is approximately 1.

We also simulate how changes in fees change hiring. For the experienced segment, $\frac{\partial H_E}{\partial t_E} \approx \frac{\partial H_E}{\partial w_E}\frac{\partial w_E}{\partial t_E}$ is locally outside of the unit interval (using a numerical derivative), so we also simulate profits under different taxes. The baseline hiring probabilities are $H_E = 0.585$ and $H_I = .302$ with no tax. We find that the optimal tax is $t_E^* = 0.33$ both using analytical results and simulations. This is close to the existing ad-valorem fee, so profits are little changed under between the existing and optimal fee structure. For the inexperienced segment, $\frac{\partial H_I}{\partial t_I} \approx -.517$. When one future job is assumed,

$$t_I^* = -\frac{.302}{-.517} - 0.33 * .585 - .585 * .05 = .36.$$

When two future jobs are assumed (the sample median conditional on hiring), optimal taxes on the inexperienced fall to 0.26. With seven future jobs (the sample mean conditional on hiring), inexperienced employers should receive a subsidy because of the future billing potential that comes from experience.

However, all of this is back-of-the envelope. When the full model is estimated, we will be able to simulate how the distribution of valuations in the experienced segment changes with the fee on inexperienced users directly through the model.

6 Evidence on the Mechanism: Learning

6.1 Employers' Number of Interviews Falls with Experience

Turning to investigate how behavior differs between inexperienced and experienced employers, several findings are consistent with the notion that learning about the quality of matches may result in frictions. One of the immediate observations from Table 1 is that the number of interviews falls, on average, on successive job postings. This is consistent with employers forming more precise beliefs about the distribution of worker matches from their experience interacting with workers on the site. Table 7 presents the results of a regression of the number of interviews conducted on the job posting number and employer fixed effects. The first column does not include employer fixed effects, while subsequent columns include combinations of employer fixed effects, controls for qualitative opening features and fixed effects for expected job duration, and controls for the mean log bid on the opening. Even with different levels of controls, in all specifications with employer fixed effects, the number of interviews decreases at a decreasing rate on successive jobs. Figure 5 expands on this regression by plotting the coefficients in Column 2. The figure shows that the predicted number of interviews falls 50% by the ninth job posting.

A feature of the oDesk marketplace is that interviews take place with a time lag after an interview request is sent. In order to schedule efficiently, it is likely optimal for employers to send out several interview requests at once. Employers tend to send out batches of interview requests. Using information on when each interview request is sent, the interviews on every job are grouped into batches as follows: The first batch of interviews starts with the first interview request sent. Then, taking each successive interview request, when more than two hours elapses between two successive requests, the next of these interviews is categorised into a new batch. Figure 8 shows how the number and size of batches varies with job opening number for employers who hire on the first job.[19]

[19] Figures that do not condition on hiring on the first job are similar but they are more difficult to interpret because some employers do not interview anyone and do not hire; these individuals appear not to begin the search process

Figure 8 shows the discrete number of interview batches for employers on their first, second and fifth plus postings. On the first job postings that lead to a hire, around 20 percent of employers conduct no interviews before hiring, around 50 percent of employers send out only one batch of interviews. The remaining 30 percent send two or more. On the second posting, this distribution is shifted to the left. A larger share of employers conduct no interviews or send out fewer interview batches. By the fifth or later job, very few employers send out more than one batch of interview requests. Figure 8b shows that the number of interviews in the first interview batch is also lower for later job openings. By the fifth and later job opening, the vast majority of employers are doing no interviews or sending only one interview request in the first batch of interviews.[20]

Because the order of interview requests is observed in the data, it is possible to examine whether inexperienced employers' early actions on the site are correlated with later choices and outcomes. If this were the case, many results up to this point (and later) may be driven by unobserved, predetermined, employer characteristics that shape employers' search processes rather than their experience in the market. Figure 9a examines the distribution of the hourly wage bids of workers selected for the first interview for employers who go on to do fewer than five interviews on the first job and employers that go on to do five or more interviews in total on the first job. The figure plots the residual of a regression of the log of the hourly wage on job category and year-month fixed effects. A comparison of the two distributions shows very little difference in the choice of first interviewee for employers who go on to interview few or many applicants. Figure 9b repeats this exercise by whether the employer hires on the first job. These comparisons cast doubt on the hypothesis that inexperienced employers eventual differences in interviewing or hiring behavior is driven by unobservable differences in information or preferences at the time of the initial job posting. Even more important, the overlap in wages on the initial interview request sent by employers suggests that workers are not able to segment inexperienced employers based on the eventual number of interviews that they conduct or the probability that they will hire.

after receiving applications.

[20] It is possible that the number of interview batches and the size of each batch varies because the arrival rate of applicants for a job varies with the job opening number. This is examined in Appendix Table 1. The omitted group is the first job opening. Columns 1 and 2 examine whether inexperienced and experienced employers receive the same number of applications in the first hour. Columns 3 and 4 examine the same for the first two hours; Columns 5 and 6 for the first day. In the cross section, experienced employers tend to receive more applications than employers on the first job. These differences are not present once employer fixed effects are included.

6.2 Employer Outcomes

This section presents additional evidence on how employer learning about the market maps into hiring outcomes.[21] To fix ideas, consider a sequential search model where an employer is trying to find a good employee match. Imagine, first, that the employer has fixed beliefs about the distribution of match quality, denoted by $F(m|\theta)$. The optimal policy rule in this case is a reservation threshold rule. With a mean-preserving spread of $F(m|\theta)$, the reservation threshold increases. The problem is slightly more complicated when $F(m|\theta_t)$ evolves based on the prior $\theta_t \tilde{} \tilde{G}_t(\cdot)$. By Kohn and Shavell (1974), the reservation match quality is greater with uncertainty in $G_t(\cdot)$ compared to $F(m|E(\theta_t))$. This implies that the expected reservation quality falls with Bayesian updates about θ due to employer learning. These implications are similar in a model of fixed sample size search, akin to the batches of interviews in the oDesk setting. While the predicted number of batches isn't clear ex-ante, the optimal batch size is larger with greater uncertainty.

We evaluate whether project results decline with search intensity. There is some support for this hypothesis using the subcategories featuring the simplest tasks, Data Entry and Web Research. There appears to be no relationship for other categories, and this seems to be driven by employers that are searching for a personal assistant. We suspect that an interview for a personal assistant may convey additional important information about horizontal match rather than the distribution of workers available on the platform.

The data reveal that inexperienced employers in data entry or web research who interview more candidates on their first job posting tend to have less successful outcomes for the posting. Table 8 shows that employers who interview fewer candidates are more likely to hire, more likely to report having had a successful hire, and more likely to give good feedback to the employed worker. The sample consists of inexperienced employers who do at least one interview on their first job posting. All specifications include job category and expected job duration fixed effects, along with year-month indicators.

Column 1 shows that inexperienced employers who interview between 5 and 8 workers are 9 percentage points less likely to hire than employers who interview only one worker. Columns 2 and 3 show that they are 3 percentage points less likely to hire and report success and 5 percentage points less likely to hire and give good feedback to the worker.

Columns 4 to 7 look at inexperienced employers who conduct at least one interview on their

[21] There are a set of other predictions that relate to assumptions made about local versus global learning, for example, when employers update their beliefs about subsets of workers based on experiences with workers that share similar characteristics. Evaluating these predictions is work in progress.

first job posting and who also hire on that opening. Employers who do two to four, or five to seven, interviews are significantly less likely to report success on the job than employers who do one interview, excluding and also including the hourly wage paid on the job as a control. This result holds for the sample of all such employers and for the subset of inexperienced employers whose first and second job posting can be segmented in the data because the second job posting (if any) occurs after the last application for the first job.

Table 9 examines a similar result for employers who go on to post a second job. Using the variation in interviews on the first job, these results show that employers who conduct more interviews on the first job are less likely to report success or give workers good feedback on the second job posting.

The number of early interviews is negatively related to job success. This is consistent with the hypothesis that, on average, employers are using early interviews to learn about the distribution of unobserved match quality. In order to forgo the opportunity to learn from additional early interviews, and hire after only one or few interviews, employers must have drawn a very high-quality employee on an early interview, leading to a good job outcome on their employment. A model where employers are not using interviews to learn about the distribution would not predict the negative relationship between search effort and outcomes.

6.3 Country Switching

The data also reveal some evidence that worker country is a salient characteristic that plays a role in an employer's search process. Employers tend to switch country-locations upon unsuccessful searches, and they are then more likely to interview workers from the new country, suggesting there is learning about the distribution rather than simply about the efficacy of interviewing.

7 Conclusion

This paper documents that potential employers on oDesk.com who have no prior experience hiring in this labor market receive higher wage bids from workers compared to similar employers with observable hiring experience. Employers who have made at least five prior hires receive bids that are, on average, around 10% lower than employers who have not previously hired. This finding is robust to controlling for employer, worker and for employer and worker fixed effects. The wages of the workers hired are also lower for experienced employers.

We investigate whether workers submit higher bids to inexperienced employers because these employers have more uncertainty about the distribution of the employer-specific quality of the workers on the site. When employers face greater uncertainty about worker quality, workers' optimal bids can be higher because of two channels: one on the demand side and one on the supply side.

On the demand side, inexperienced employers' hiring probabilities are less elastic to the bids offered because, at greater levels of uncertainty about match quality, the variance of the highest expected order statistic of the quality distribution is larger. A marginal bid increase has a smaller effect on the probability of being hired, resulting in higher optimal mark ups. Estimating market demand for both inexperienced employers show that workers bid as if inexperienced employers have more inelastic demand. The estimated wage elasticity for inexperienced employers is -7, compared to -11.5 for experienced employers.

On the supply side, inexperienced employers are likely to create higher costs for workers–both when applying for the job and when hired. For example, inexperienced employers conduct more interviews and are less likely to hire, creating hassle costs for the workers with less expected benefit. Inexperienced employers are also more likely to need help understanding the mechanics of the site. These greater costs may be passed through to inexperienced employers in higher wage bids.

Making the assumption that observed wage bids are the equilibrium outcome of a differentiated products Betrand game allows us to decompose the wage bid difference into these demand and supply side effects. Differences in equilibrium markups due to differences in the wage elasticity of demand account for an estimated 25% of the wage bid gap between inexperienced and experienced employers. The residual–or 75% of the gap–can be attributed to differences in workers' costs.

Further empirical findings suggest that experience on the site enables employer learning, which corroborates the hypothesis that inexperienced employers' relative uncertainty about match quality explains wage bid differences. Inexperienced employers are shown to interview more candidates but job outcomes are negatively associated with the number of interviews done prior to hiring. A possible explanation for this somewhat counter-intuitive finding is that an interview is valuable for both the information it provides about the quality of the interviewed worker, and also about the distribution of match quality among all workers. When employers have greater uncertainty about this distribution, an interview is more valuable. Opting to stop interviewing and hire an interviewed worker, rather than conduct more interviews, suggests that the revealed match quality of that worker is very high. That is, those that hire early have found good quality workers, which explains their good job outcomes.

Showing that both demand and supply differ in the market faced by inexperienced and experienced employers allows us to find the platform profit-maximising fee that the platform should charge each segment of employers. On the one hand, inexperienced employers' more inelastic demand would lead oDesk to impose a higher fee on these employers but, on the other hand, a lower tax on inexperienced employers will increase the total number of transactions on the site in the future, as these employers are more likely to remain in the market. The net effect is that oDesk could increase profits by setting a fee of 1.05 dollars per hour to experienced employers and a much lower fee to inexperienced employers.

Whether or not potential trading partners have experience transacting in a market is observable in many kinds of markets, including many online product markets. This paper finds evidence that employer (buyer) experience is a salient characteristic for workers (suppliers) in the oDesk labor market because experience affects both the nature of perceived demand and the cost of supply.

8 Appendix on Estimation

This section derives bias correction of the control function estimator when the model is estimated via maximum likelihood.[22] Let δ_j denote the control function for applicant j.

Because the control function is estimated with error and the logit probability involves a nonlinear transformation, estimation error in the control function may affect parameter estimates. We take the expected value of a second-order expansion around the vector δ for the log choice probability for opening i. The bias correction term when considering the log choice probability

$$\log(p_{ij}) = \delta_j\beta - \log\left(1 + \Sigma_{j=1}^{J} \exp(\delta_j\beta)\right)$$

is

$$\frac{1}{2}\Sigma_{m=1}^{J}\Sigma_{k=1}^{j}\frac{\partial^2 \log(p_{ij})}{\partial\delta_m\partial\delta_k}\left(\hat{\delta}_m - \delta_m\right)\left(\hat{\delta}_k - \delta_k\right)$$

$$\frac{\exp(\delta_j\beta)}{1 + \Sigma_{j=1}^{J}\exp(\delta_j\beta)}$$

is

$$\frac{1}{2}\Sigma_{m=1}^{J}\Sigma_{k=1}^{j}\frac{\partial^2 p_{ij}}{\partial\delta_m\partial\delta_k}\left(\hat{\delta}_m - \delta_m\right)\left(\hat{\delta}_k - \delta_k\right).$$

When taking expectations, the only terms that are non-zero are those for which $m = k$. That

[22] The correction is slightly different when the model is estimated via GMM and the level of the choice probability is used rather than the logarithm.

means the bias correction is easy to code up and only requires knowledge of the sampling variances of each $\hat{\delta}_m$.

When $m = j = k$, $\frac{\partial \log(p_{ij})}{\partial \delta_j} = \beta - \frac{\exp(\delta_j \beta)\beta}{1+\Sigma_{j=1}^{J} \exp(\delta_j \beta)} = \beta \left(1 - p_{ij}\right).$ Taking a second derivative when $m = j = k$ yields[23]

$$\frac{\partial^2 \log\left(p_{ij}\right)}{\partial \delta_j^2} = -\beta^2 p_{ij}\left(1 - p_{ij}\right).$$

When $m \neq j = k$, the relevant derivatives are

$$\frac{\partial \log\left(p_{ij}\right)}{\partial \delta_k} = -\frac{\exp\left(\delta_k \beta\right)\beta}{1 + \Sigma_{j=1}^{J} \exp\left(\delta_j \beta\right)} = -\beta p_{ik}.$$

and

$$\begin{aligned} \frac{\partial^2 p_{ij}}{\partial \delta_k^2} &= -\beta \frac{\partial p_{ij}}{\partial \delta_k} p_{ik} - \beta p_{ij} \frac{\partial p_{ik}}{\partial \delta_k} \\ &= \beta^2 p_{ij} p_{ik}^2 - \beta^2 p_{ij} p_{ik}\left(1 - p_{ik}\right) \end{aligned}$$

For the outside option, the derivatives are

$$\begin{aligned} \frac{\partial\left(1/1 + \Sigma \exp\left(\delta_j \beta\right)\right)}{\partial \delta_k} &= -\frac{\beta \exp\left(\delta_k \beta\right)}{\left[1 + \Sigma_{j=1}^{J} \exp\left(\delta_j \beta\right)\right]^2} \\ \frac{\partial^2}{\partial \delta_k^2} &= \frac{-\beta^2 \exp\left(\delta_k \beta\right)\left[1 + \Sigma_{j=1}^{J} \exp\left(\delta_j \beta\right)\right] + 2\beta^2 \exp\left(\delta_k \beta\right)^2}{\left[1 + \Sigma_{j=1}^{J} \exp\left(\delta_j \beta\right)\right]^3} \\ &= \frac{\beta^2 p_{ik}\left(2p_{ik} - 1\right)}{\left[1 + \Sigma_{j=1}^{J} \exp\left(\delta_j \beta\right)\right]}. \end{aligned}$$

To get the sampling variance of $\hat{\delta}_j$, notice that it is just a residual from the first-stage regression.

[23] This comes from

$$\frac{\partial p_{ij}}{\partial \delta_j} = \frac{\beta \exp\left(\delta_j \beta\right)}{\left[1 + \Sigma_{j=1}^{J} \exp\left(\delta_j \beta\right)\right]} - \frac{\beta \exp\left(\delta_j \beta\right) \exp\left(\delta_j \beta\right)}{\left[1 + \Sigma_{j=1}^{J} \exp\left(\delta_j \beta\right)\right]^2} = \beta p_{ij}\left(1 - p_{ij}\right)$$

The fitted values of the residual, $\hat{\delta}_{ij}$ themselves have sampling variance-covariance matrix

$$
\begin{aligned}
& E\left[\left(\hat{\delta}-\delta\right)\left(\hat{\delta}-\delta\right)'|x\right] \\
= \; & E\left[\left(y - x\left(x'x\right)^{-1}x'y-\delta\right)\left(y - x\left(x'x\right)^{-1}x'y-\delta\right)'|x\right] \\
= \; & E\left[\left(xb+\delta-xb - x\left(x'x\right)^{-1}x\delta-\delta\right)\left(xb+\delta-xb - x\left(x'x\right)^{-1}x\delta-\delta\right)'|x\right] \\
= \; & E\left[x\left(x'x\right)^{-1}x'\delta\delta'x\left(x'x\right)^{-1}x'|x\right] \\
= \; & \left[\sigma_{\delta}^{2}|x\right]x\left(x'x\right)^{-1}x'.
\end{aligned}
$$

The diagonals from this matrix are then used in the bias correction procedure.

Estimates of conditional logit models are not sensitive to this bias correction, so these results are omitted.

References

[1] Ackerberg, Daniel. 2003. "Advertising, learning, and consumer choice in experience good markets: an empirical examination." *International Economic Review*, 44(3), 1007-1040.

[2] Agrawal, Ajay, Christian Catalini and Avi Goldfarb. 2011. "Friends, Family, and the Flat World: The Geography of Crowdfunding." NBER Working Paper No.16820.

[3] Agrawal, Ajay and Avi Goldfarb. 2008. "Restructuring Research: Communication Costs and the Democratization of University Innovation." *American Economic Review*, 98(4), 1578-1590.

[4] Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons. 2012. "How Do Online Platforms Flatten Markets for Contract Labor?" Working Paper.

[5] Agrawal, Ajay, John Horton, Nicola Lacetera, and Elizabeth Lyons. 2013. "Digitization and the Contract Labor Market: A Research Agenda." NBER Working Paper w19525.

[6] Altonji, Joseph, and Charles Pierret. 2001. "Employer Learning and Statistical Discrimination." *The Quarterly Journal of Economics*, 116(1), 313-350.

[7] Arcidiacono, Peter, Esteban Aucejo, Arnaud Maurel, and Tyler Ransom. 2013. "College Attrition and the dynamics of information revelation." Working Paper

[8] Autor, David. 2001. "Wiring the Labor Market." *Journal of Economic Perspectives*, 15(1), 25-40.

[9] Bajari, Patrick and Ali Hortaçsu. 2004. "Economic Insights from Internet Auctions." *Journal of Economic Literature*, 42(2), 457-486.

[10] Bergemann, Dirk and Juuso Valimaki. 2006. "Dynamic Pricing of New Experience Goods." *Journal of Political Economy,* 114(4), 713-743.

[11] Cabral, Luis, and Ali Hortacsu. 2010. "The dynamics of seller reputation: Evidence from ebay*." *The Journal of Industrial Economics*, 58.1, 54-78.

[12] Erdem, Tulin, and Michael Keane. 1996. "Decision-making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets." *Marketing Science*, 15(1), 1-20.

[13] Farber, Henry and Robert Gibbons. 1996. "Learning and Wage Dynamics." *The Quarterly Journal of Economics,* 111(4), 1007-1047.

[14] Ghani, Ejaz, William Kerr, and Christopher Stanton. 2014. "Diasporas and outsourcing: Evidence from oDesk and India." *Management Science,* 60(7),1677-1697.

[15] Greif, Avner. 1993. "Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders' Coalition," *American Economic Review*, 83(3), 525-548.

[16] Horton, John. 2010. "Online Labor Markets." Working Paper.

[17] Horton, John. 2013. "The Effects of Subsidizing Employer Search." Working Paper.

[18] Jin, Ginger and Andrew Kato. 2006. "Price, Quality and Reputation: Evidence from An Online Field Experiment." *The RAND Journal of Economics*, 37(4), 983-1005.

[19] Kahn, Lisa. 2013. "Asymmetric Information between Employers." *American Economics Journal: Applied.*

[20] Kahn, Lisa and Fabian Lange. 2014. "Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures." *Review of Economic Studies* (Forthcoming).

[21] Lange, Fabian. 2007. "The Speed of Employer Learning." *Journal of Labor Economics,* 25(1), 1-35.

[22] Lewis, Gregory. 2011. "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors." *American Economic Review*, 101(4), 1535-1546.

[23] Luca, Michael. 2010. "Reviews, Reputation, and Revenues: The Case of Yelp.com." Harvard Business School Working Paper No.13-042.

[24] Lyons, Elizabeth. 2014. "Team Production in International Labor Markets: Experimental Evidence from the Field." Working paper.

[25] Miller, Robert. 1984. "Job matching and occupational choice." *Journal of Political Economy,* 92(6), 1086-1120.

[26] Pallais, Amanda. 2013. "Inefficient Hiring in Entry-Level Labor Markets." Working paper.

[27] Stanton, Christopher and Catherine Thomas. 2012. "Landing the First Job: The Value of Intermediaries in Online Hiring." Working Paper.

Figure 1a: Number of job postings per employer, truncated at 30.

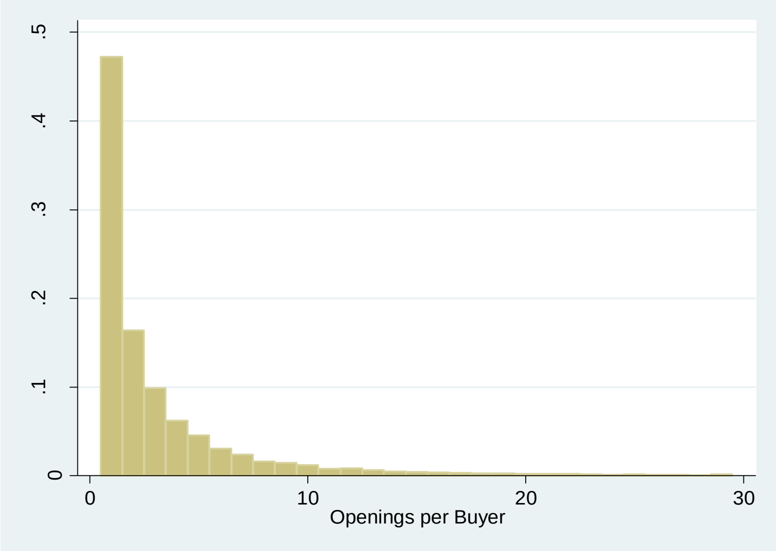

Figure 1b: Number of hires per employer in the truncated sample

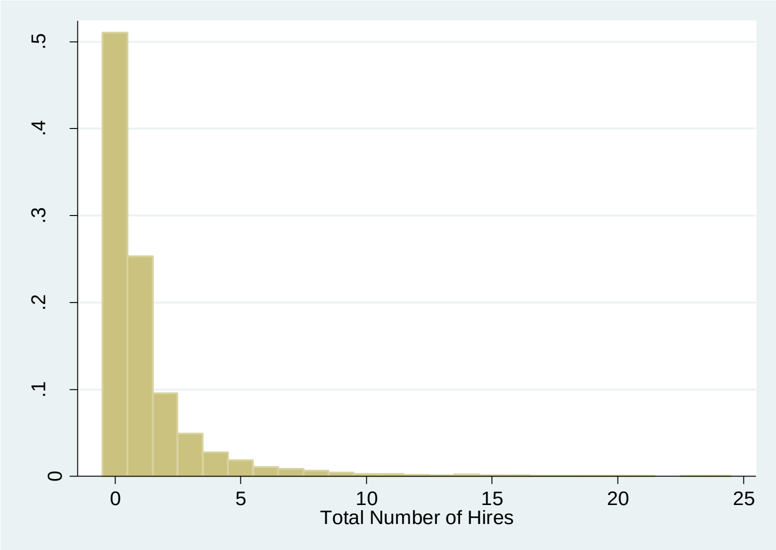

Figure 1c: Number of hires per employer in the truncated sample, conditional on making at least one hire

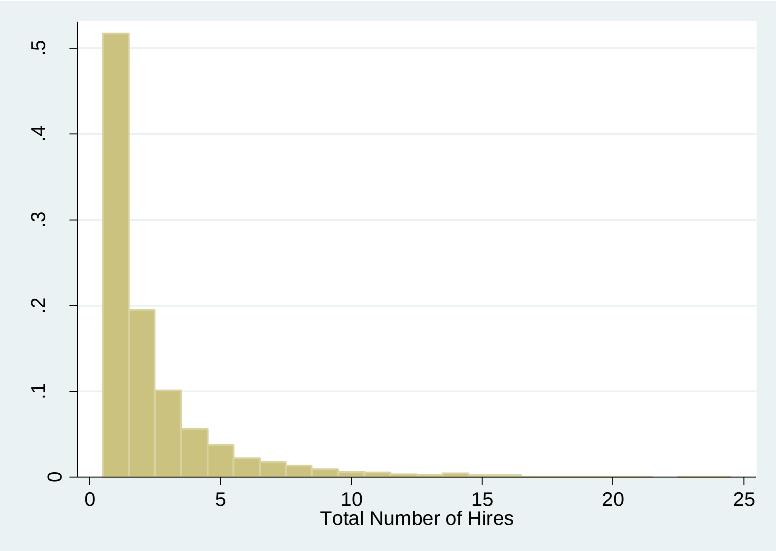

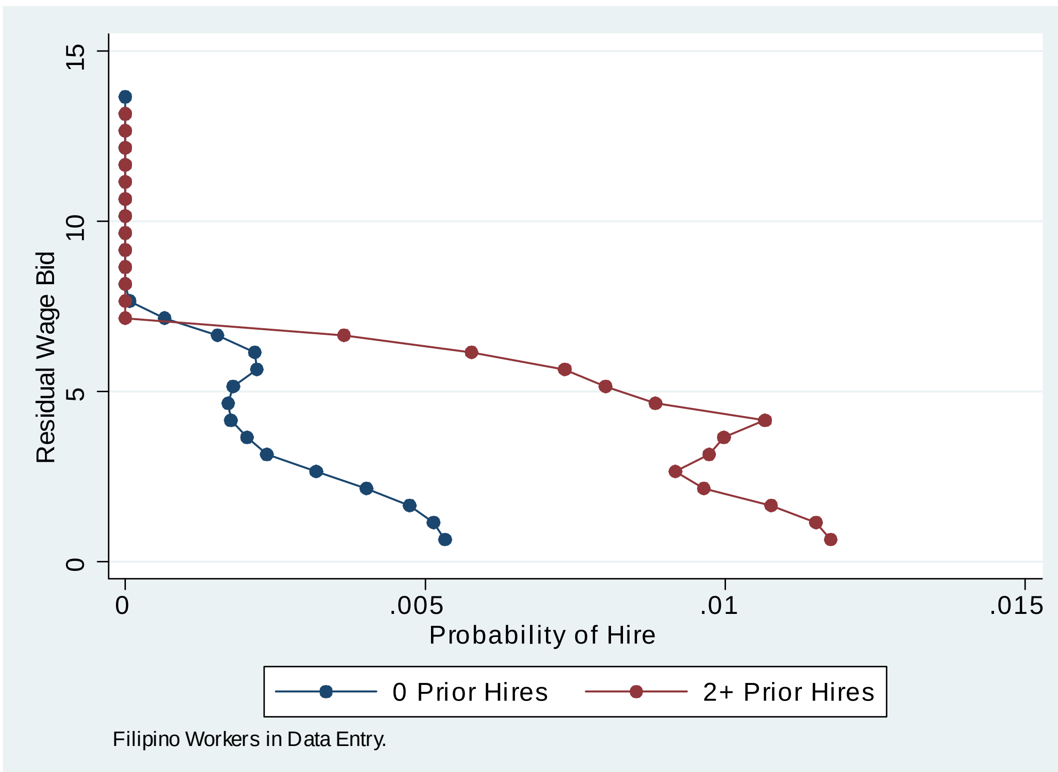

Figure 2: Wage Bids and Hiring Probabilities by Employer Experience. Points are taken from a polynomial smooth of an indicator that the worker was hired on their residual bid. The residuals remove time fixed effects and priced observable characteristics.

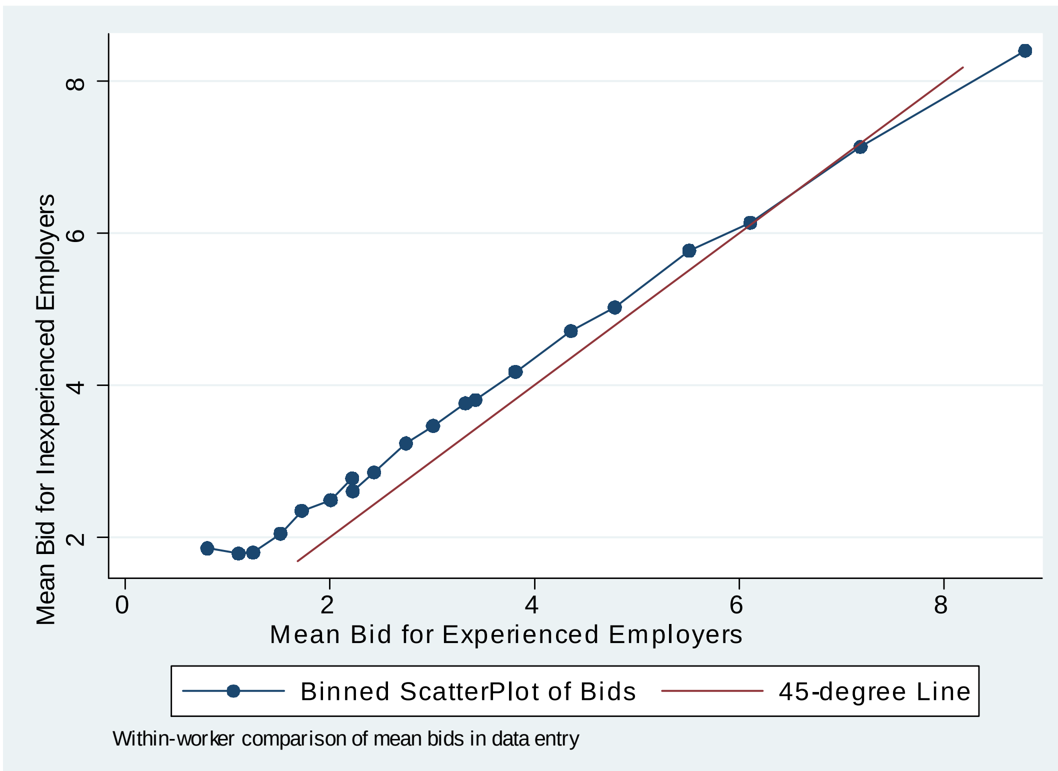

Figure 3: Wage bids by observed employer hiring experience, calculated within worker-month for workers who apply to both experienced and inexperienced employers.

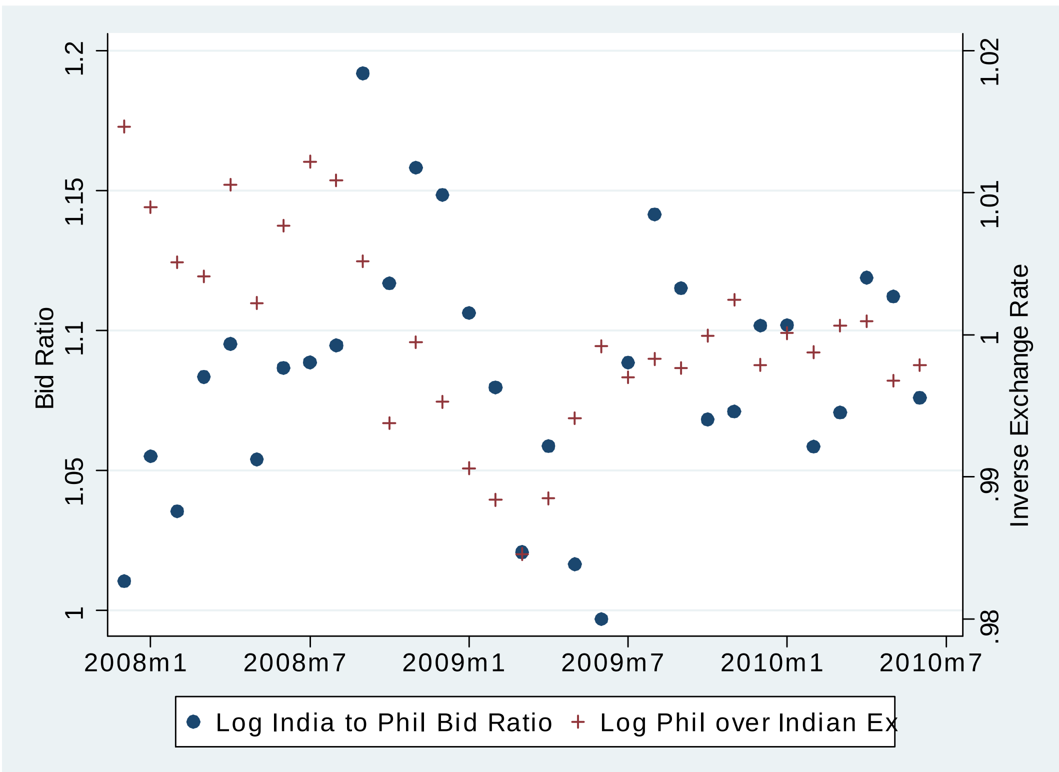

Figure 4: Ratio of Country-Level Mean Log Wage Bids and Log Exchange Rates

Figure 5: Number of interviews on successive job openings (DV is the log of the number of interviews +1)

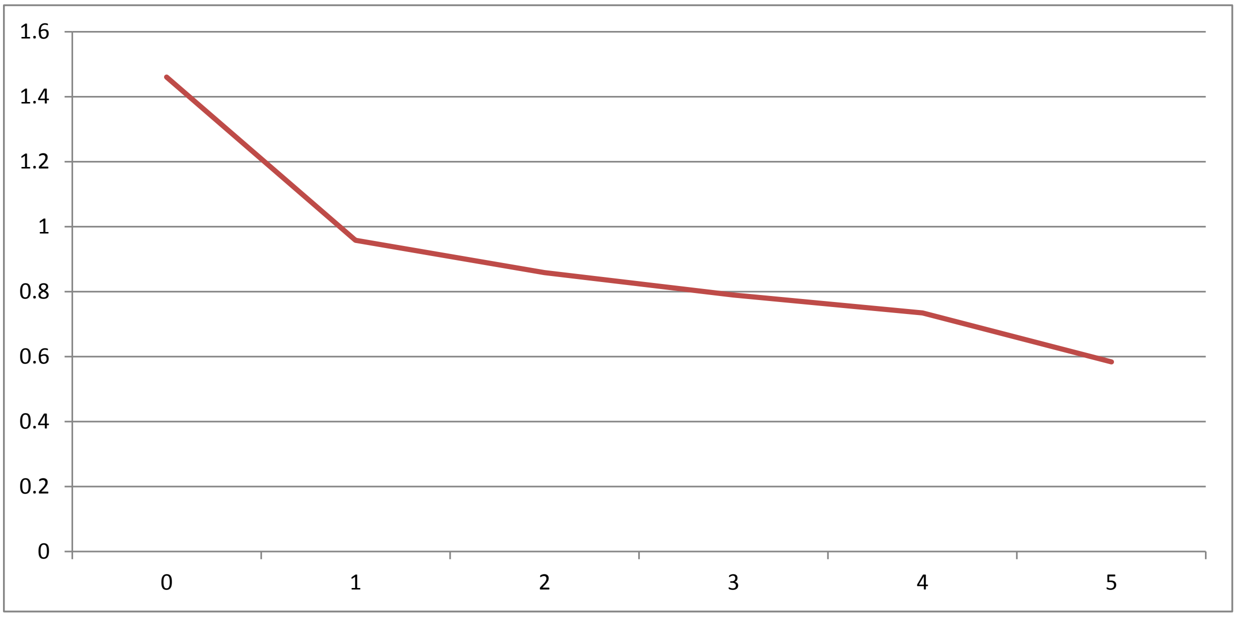

Notes: The series is the predicted number of interviews per posting on successive jobs calculated from the estimated constant term and coefficient estimates from a regression of job posting fixed effects on the total number of interviews. The estimation contains employer fixed effects. All coefficients are significantly different from zero. The sample is job postings where the prior job is complete by the time of posting.

Figure 6a: Number of Batches by Job Opening Number

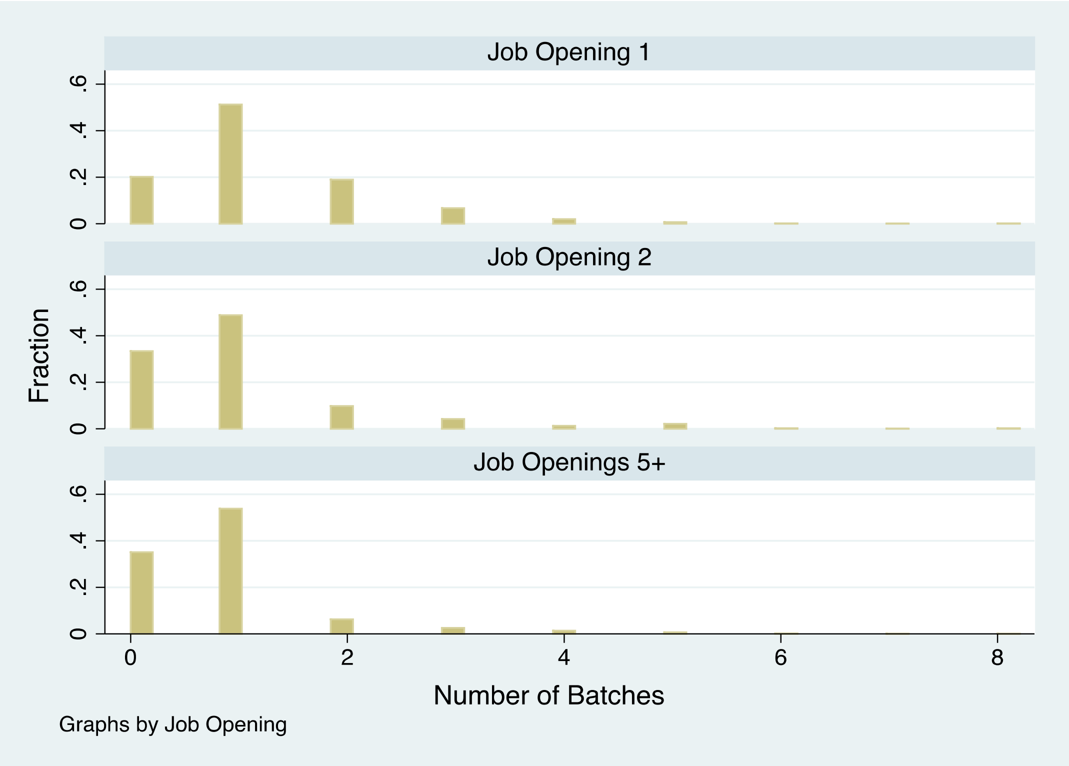

Figure 6b: Number of interviews in batch 1, by Job Opening Number

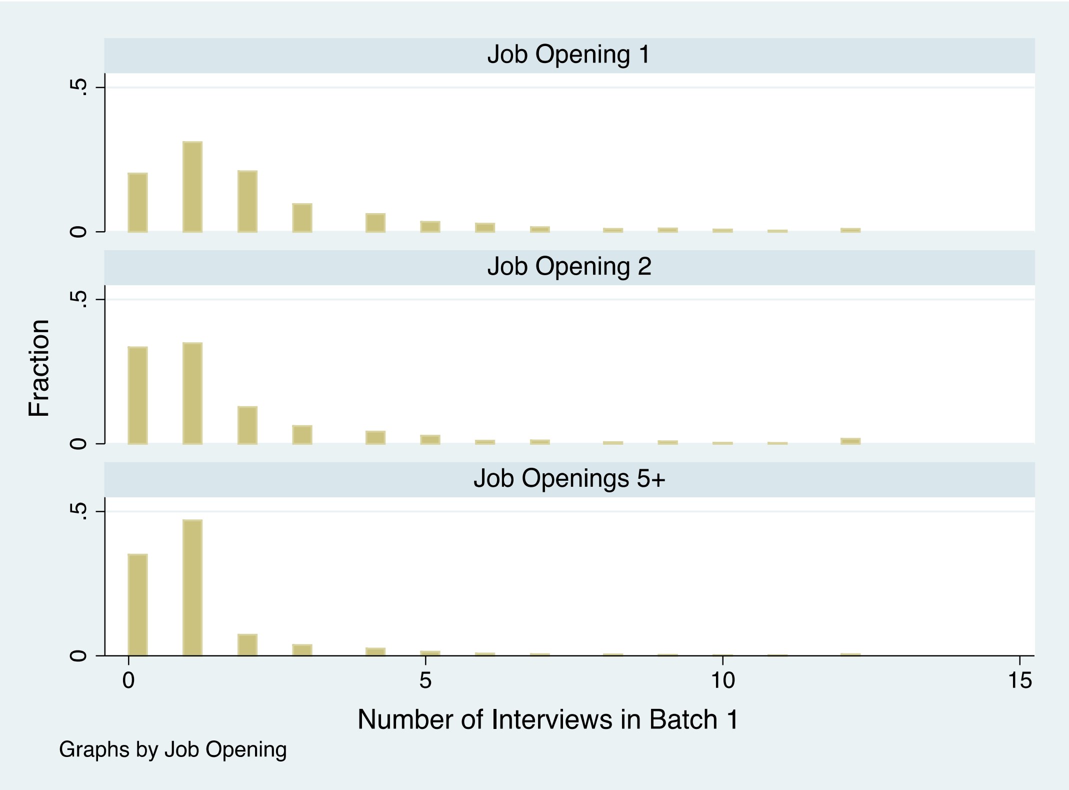

Notes: Both figures include those employers who hire on their first posting, for sequential jobs only. A batch of interviews is defined as all interviews that take place prior to a two-hour gap in sending an interview request.

Figure 7: Residual Hourly Bid made by first applicant selected for interview, after JC2 and YM dummies as controls.

7a: Log Bids for employers who conduct fewer than five and five or more interviews on the first posting.

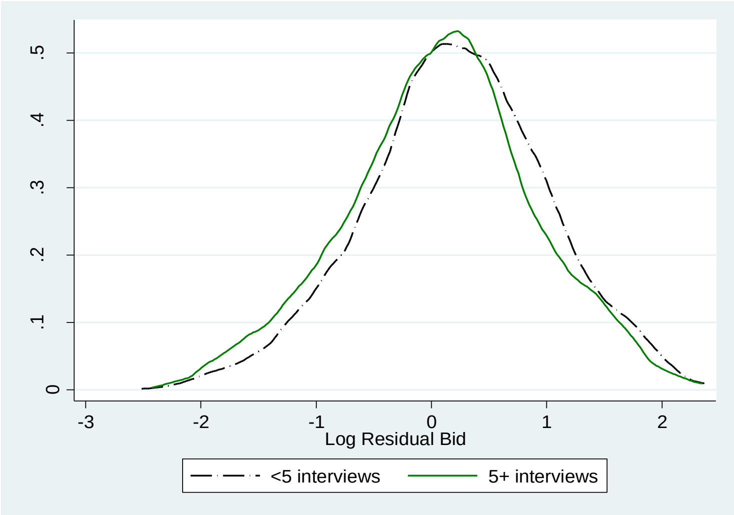

7b: Log bids ror employers who don't hire and who hire on the first posting.

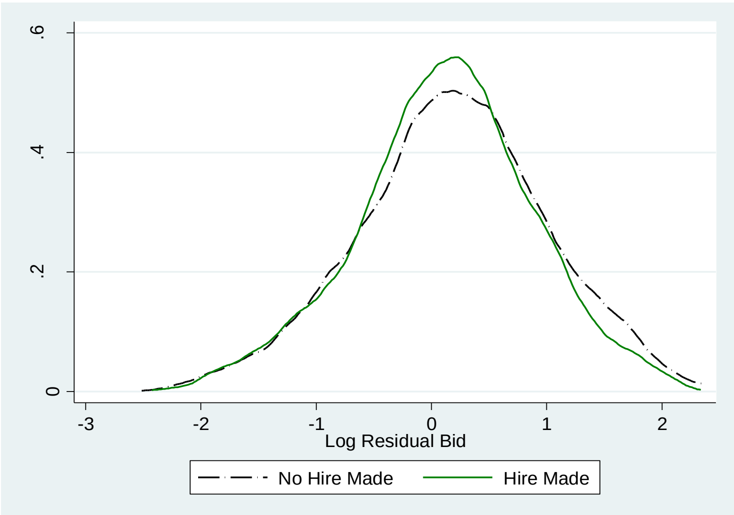

Table 1A: Summary Statistics

Previous Hires	Number of Job Openings	Number of Candidates	Share of Employer-Initiated Candidates	Mean Wage Bid	Number of Interviews	Probability a Hire is Made	Mean Wage Bid of Hired Worker	Median Wage Bid	Share with Good Worker Feedback	Share with Missing Worker FB
							On Postings where a Hire was made			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
0	16,744	33.4	0.17	4.94	3.78	0.25	4.17	3.33	43%	35%
		(55.1)	(0.35)	(4.27)	(8.09)	(0.43)	(3.96)			
1	4,786	26.4	0.24	4.43	2.56	0.54	3.75	2.78	45%	32%
		(56.6)	(0.41)	(3.81)	(5.53)	(0.50)	(3.29)			
2	3,512	25.8	0.25	4.16	2.35	0.56	3.54	2.78	46%	30%
		(63.1)	(0.42)	(3.68)	(5.39)	(0.50)	(3.30)			
3	2,663	25.2	0.23	3.93	2.15	0.59	3.42	2.22	45%	31%
		(55.2)	(0.41)	(3.49)	(4.57)	(0.49)	(3.24)			
4+	26,098	22.7	0.24	3.56	1.98	0.59	3.06	2.22	50%	28%
		(63.5)	(0.42)	(3.56)	(5.10)	(0.49)	(3.46)			

Table 1B: Summary Statistics for Sequential Openings

Previous Hires	Number of Job Openings	Number of Candidates	Share of Employer-Initiated Candidates	Mean Wage Bid	Number of Interviews	Probability a Hire is Made	Mean Wage Bid of Hired Worker	Median Wage Bid	Share with Good Worker Feedback	Share with Missing Worker FB
							On Postings where a Hire was made			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
0	10,543	34.7	0.21	5.13	3.92	0.24	4.38	3.33	44%	34%
		(54.8)	(0.38)	(4.38)	(8.03)	(0.43)	(4.11)			
1	3,273	27.4	0.27	4.54	2.61	0.54	3.87	3.00	45%	32%
		(57.8)	(0.43)	(3.88)	(5.52)	(0.50)	(3.45)			
2	2,502	26.2	0.28	4.27	2.41	0.55	3.64	2.78	47%	28%
		(60.3)	(0.44)	(3.78)	(5.27)	(0.50)	(3.51)			
3	1,916	25.8	0.25	4.00	2.21	0.58	3.53	2.77	46%	30%
		(55.6)	(0.42)	(3.53)	(4.63)	(0.49)	(3.36)			
4+	19,503	24.0	0.27	3.56	2.02	0.59	3.06	2.22	51%	28%
		(65.8)	(0.43)	(3.56)	(5.16)	(0.49)	(3.52)			

Notes: Sample period is from January 2006 to June 2010, with most of the openings posted after 2008. Wage bids are winsorized at the 99th percentile, and applications closed by oDesk as likely spam are excluded. Standard deviations are given in parentheses.

Table 1C: Summary Statistics for Data Entry Job Postings

Previous Hires	Number of Job Openings	Number of Candidates	Share of Employer-Initiated Candidates	Mean Wage Bid	Number of Interviews	Probability a Hire is Made	On Postings where a Hire was made: Mean Wage Bid of Hired Worker	Median Wage Bid	Share with Good Worker Feedback	Share with Missing Worker FB
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
0	7,041	47.9	0.15	3.93	4.1	0.19	3.16	2.22	47%	32%
		(75.3)	(0.34)	(3.73)	(9.7)	(0.39)	(3.00)			
1	1,495	45.4	0.25	3.48	2.5	0.58	2.87	2.22	50%	29%
		(88.4)	(0.42)	(3.12)	(5.7)	(0.49)	(2.20)			
2	1,155	44.9	0.26	3.15	2.3	0.59	2.76	2.22	50%	27%
		(96.8)	(0.43)	(2.63)	(5.5)	(0.49)	(2.25)			
3	913	40.6	0.24	3.02	2.0	0.65	2.50	2.22	49%	29%
		(79.0)	(0.42)	(2.70)	(4.5)	(0.48)	(1.98)			
4+	10,409	32.2	0.26	2.53	1.8	0.65	2.17	1.67	55%	24%
		(89.0)	(0.44)	(2.58)	(4.9)	(0.48)	(2.30)			

Table 1D: Summary Statistics for Sequential Openings for Data Entry Job Postings

Previous Hires	Number of Job Openings	Number of Candidates	Share of Employer-Initiated Candidates	Mean Wage Bid	Number of Interviews	Probability a Hire is Made	On Postings where a Hire was made: Mean Wage Bid of Hired Worker	Median Wage Bid	Share with Good Worker Feedback	Share with Missing Worker FB
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
0	4,195	52.6	0.19	3.95	4.6	0.20	3.31	2.44	49%	32%
		(76.3)	(0.38)	(3.68)	(10.3)	(0.40)	(3.25)			
1	1,031	48.8	0.27	3.56	2.6	0.57	2.96	2.22	50%	29%
		(90.4)	(0.44)	(3.14)	(5.5)	(0.50)	(2.33)			
2	842	44.3	0.30	3.22	2.2	0.59	2.86	2.22	50%	27%
		(90.1)	(0.45)	(2.66)	(4.4)	(0.49)	(2.39)			
3	655	43.1	0.28	3.17	2.1	0.62	2.63	2.22	49%	27%
		(80.8)	(0.44)	(2.77)	(4.7)	(0.49)	(2.08)			
4+	8,050	33.0	0.29	2.51	1.7	0.65	2.17	1.67	55%	23%
		(90.2)	(0.45)	(2.53)	(4.9)	(0.48)	(2.22)			

Notes: Sample period is from January 2006 to June 2010, with most of the openings posted after 2008. Wage bids are winsorized at the 99th percentile, and applications closed by oDesk as likely spam are excluded. Standard deviations are given in parentheses.

Table 2: Wage Bids Decline with Observable Employer Experience

	OLS	OLS	Employer Fixed Effects	Employer Fixed Effects	Worker Fixed Effects	Worker Fixed Effects	Employer and Worker Fixed Effects	Employer and Worker Fixed Effects
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A								
On posts after making 1 hire	-0.0805***	-0.0857***	-0.0482***	-0.0444***	-0.0356***	-0.0376***	-0.0172***	-0.0154***
	(0.00950)	(0.00932)	(0.0149)	(0.0145)	(0.00196)	(0.00195)	(0.00900)	(0.00767)
2 hires	-0.0930***	-0.0985***	-0.0639***	-0.0606***	-0.0463***	-0.0486***	-0.0313***	-0.0299***
	(0.0112)	(0.0111)	(0.0156)	(0.0153)	(0.00220)	(0.00220)	(0.00796)	(0.00713)
3 hires	-0.109***	-0.116***	-0.0704***	-0.0706***	-0.0544***	-0.0572***	-0.0419***	-0.0420***
	(0.0122)	(0.0122)	(0.0170)	(0.0167)	(0.00243)	(0.00243)	(0.010620)	(0.00885)
4 hires	-0.135***	-0.141***	-0.0973***	-0.0952***	-0.0680***	-0.0704***	-0.0560***	-0.0550***
	(0.0179)	(0.0176)	(0.0234)	(0.0232)	(0.00282)	(0.00282)	(0.01606)	(0.01105)
5+ hires	-0.163***	-0.167***	-0.101***	-0.0974***	-0.0900***	-0.0916***	-0.0614***	-0.0595***
	(0.0120)	(0.0119)	(0.0166)	(0.0164)	(0.00153)	(0.00153)	(0.01436)	(0.00500)
Third order polynomial of job description length	No	Yes	No	Yes	No	Yes	No	Yes
Observations	1,006,982	1,006,931	1,006,982	1,006,931	1,006,982	1,006,931	1,006,931	1,006,931
R-Squared	0.187	0.189	0.297	0.297	0.717	0.717	0.758	0.758
Panel B								
On posts after making 5+ hires	-0.160***	-0.244***	-0.133***	-0.129***	-0.143***	-0.143***	-0.0762***	-0.0743***
	(0.0114)	(0.0285)	(0.0261)	(0.0260)	(0.0197)	(0.0196)	(0.0208)	(0.0111)
5+ hires and no observable employer feedback		0.132***	0.0382	0.0398*	0.0904***	0.0895***	0.0247*	0.0258**
		(0.0286)	(0.0235)	(0.0237)	(0.0199)	(0.0198)	(0.0164)	(0.0131)
5+ hires and good observable employer feedback		0.0793***	0.0231	0.0226	0.0547***	0.0531***	0.0116	0.0115
		(0.0280)	(0.0198)	(0.0199)	(0.0195)	(0.0193)	(0.0141)	(0.0126)
R-Squared	0.197	0.201	0.316	0.316	0.734	0.735	0.775	0.775
Panel C								
On posts after making 5+ hires	-0.160***	-0.202***	-0.144***	-0.140***	-0.119***	-0.121***	-0.0948***	-0.0927***
	(0.0114)	(0.0221)	(0.0249)	(0.0247)	(0.0147)	(0.0147)	(0.0171)	(0.0189)
5+ hires and no worker feedback given		0.0534*	0.0417	0.0419	0.0418**	0.0428**	0.0388***	0.0391***
		(0.0279)	(0.0261)	(0.0264)	(0.0173)	(0.0176)	(0.0098)	(0.0122)
5+ hires and good worker feedback given		0.0391**	0.0411***	0.0416***	0.0328***	0.0342***	0.0320*	0.0321
		(0.0192)	(0.0144)	(0.0145)	(0.0126)	(0.0127)	(0.0168)	(0.0199)
Third order polynomial of job description length	No	Yes	No	Yes	No	Yes	No	Yes
Observations	771,957	771,943	771,957	771,943	771,957	771,943	771,943	771,943
R-Squared	0.197	0.199	0.316	0.316	0.734	0.734	0.775	0.775

The sample is limited to worker-initiated applications on sequential job openings. Robust standard errors are clustered by employer. Columns 7 and 8 contain bootstrapped standard errors with resampling over employers. All specifications contain the worker's feedback score, an indicator for zero feedback, number of jobs completed, an indicator for past work and zero feedback, and fixed effects for expected duration of job, expected duration of job interacted with hours per week, time, and detailed job category. The sample in Panels B and C is employers' first jobs and all jobs after their fourth hire.

Table 3: First Stage Regression of Log Hourly Bids and Participation on Exchange Rate Instruments (Data Entry Openings)

Dependent Variable	Log Hourly Wage Bid (Job Opening Level)				1+ application during the month	
Model	Linear Regression				Probit	
Sample	First Job Post	First Job Post	Experienced Sample	Experienced Sample	First Job Post	Experienced Sample
	(1)	(2)	(3)	(4)	(5)	(6)
Log Local Currency to Dollar (monthly; standardized)	-0.0460***	-0.0472***	-0.0506***	-0.0448***	0.0688***	0.0939***
	(0.00631)	(0.00632)	(0.00644)	(0.00644)	(0.037)	(0.042)
Mean Bid from Workers in Other Countries (Aggregate)					-3.021***	-1.637***
					(0.0395)	(0.0447)
Feedback Score (Out of 5)	0.0788***	0.149***	0.0662***	0.149***	0.5779***	0.560***
	(0.00509)	(0.00839)	(0.00413)	(0.00865)	(0.003)	(0.0027)
Agency Affiliate Indicator	0.102***	0.108***	0.0585***	0.0697***	0.062***	0.100***
	(0.00864)	(0.00866)	(0.00765)	(0.00771)	(0.007)	(0.0072)
Prior Work Experience	-0.0649***	-0.0147	-0.0284***	0.0409***	0.5975***	0.5837***
	(0.00861)	(0.0101)	(0.00753)	(0.0102)	(0.009)	(0.0087)
Weekly Participation Inverse Mills Ratio		0.431***		0.428***		
		(0.0468)		(0.0423)		
Number of Observations	54,122	54,122	59,168	59,168	1,837,515	1,475,701
R-Squared	0.212	0.213	0.186	0.187	0.3007	0.3068
F Statistic on Excluded Instruments	606.2	586.8				

Notes: The inexperienced sample is employers on their first job post. The experienced sample is employers who have hired 2 or more previous workers from any job category and have posted at least 2 previous jobs in adminstrative support. Robust standard errors in parentheses. All models contain a fifth order polynomial in calendar time, fixed effects for 6 country groups, controls for English skills, and an indicator for having zero feedback. The last country group includes many countries with very small application shares. Models in columns 1 - 4 also include a piecewise-linear spline with 4 knots for the application number, an indicator for an employer-initiated application, and an indicator that the worker only applies to this job during the month. The Log Local Currency to Dollar exchange rate is calculated using monthly data and z-scores from the aggregate macro data are used to make the measure comparable across countries. The inverse mills ratio in columns 2 and 4 is taken from columns 5 and 6. Other workers' average bids (aggregate) in the probit models in columns 5 and 6 are first calculated excluding own-country and then a monthly time fixed effect is removed. A separate interaction for workers in the United States and the basket of other currencies (not reported) because these workers do not have any variation in own exchange rates.

Table 4: Conditional Logit Parameter Estimates and Elasticities for Data Entry Jobs

	(1)	(2)	(3)	(4)
Sample:	Own-Exchange Rate and Inverse Mills Ratio in Control Function, Inexperienced Employers	Own-Exchange Rate and no Inverse Mills Ratio in Control Function, Inexperienced Employers	Own-Exchange Rate and Inverse Mills Ratio in Control Function, Employers with 2+ Hires	Own-Exchange Rate and no Inverse Mills Ratio in Control Function, Employers with 2+ Hires
Log Hourly Bid	-7.013	-2.375	-11.68	-2.433
	(2.090)	(2.677)	(1.638)	(1.996)
Control Function	6.546	1.904	10.84	1.583
	(2.089)	(2.678)	(1.637)	(1.995)
Feedback Score Out of 5	1.152	0.593	1.249	0.374
	(0.266)	(0.244)	(0.181)	(0.157)
Agency Affiliate Indicator	0.257	-0.232	0.666	0.113
	(0.282)	(0.326)	(0.157)	(0.166)
Indicator for Prior Experience on oDesk	1.983	2.272	1.532	1.786
	(0.243)	(0.272)	(0.140)	(0.149)
Mean Own-Bid Elasticity	-6.968	-2.360	-11.57	-2.409
	(2.612)	(3.127)	(2.674)	(1.941)
Number of Job Openings	1781	1781	1672	1672
Mean Lerner Markup, Pre-oDesk-Fee	14.30%	42.40%	8.65%	41.53%
Mean Cost (USD, Pre-Fee)	3.08	2.08	2.6	1.66
Median Cost (USD, Pre-Fee)	2.16	1.45	1.84	1.18
Mean Cost in India (USD, Pre-Fee)	2.9	1.95	2.65	1.7
Mean Cost in the Philippines (USD, Pre-Fee)	2.47	1.65	2.18	1.4
Mean Cost in the United States (USD, Pre-Fee)	5.82	3.92	4.95	3.17

Estimates come from a conditional logit model that includes an option not to hire an applicant. Standard errors for the elasticities calculated using 150 cluster bootstrap replications of the procedure. Each job opening forms a cluster. Other controls included for English score, an indicator for no feedback, and an indicator for being experienced without having feedback.

Table 5: Mixed Logit Parameter Estimates, Elasticities, Costs, and Markups

	(1)	(2)
Panel A: Parameter Estimates		
Sample:	Data Entry Job Post 1	Experienced Employers
Log Hourly Bid		
Control Function		
Std Dev of Inexperienced Random Effect ($\sigma_{\eta 0}$)		
Std Dev of Experienced Random Effect ($\sigma_{\eta 2+}$)		
Time-Weighted Average Mean of μ		
Std Dev of Permanent Heterogeneity (σ_{μ})		
Probability of Any Hire (Data)		
Number of Job Openings		
Number of Employers		
Mean Log Likelihood (Per Employer)		
Panel B: Estimated Elasticities, Markups and Costs		
Mean Own-Price Elasticity		
Mean Lerner Markup, Pre-oDesk-Fee		
Mean Cost (USD, Pre-Fee)		
Median Cost (USD, Pre-Fee)		
Mean Cost in India (USD, Pre-Fee)		
Mean Cost in the Philippines (USD, Pre-Fee)		
Mean Cost in the United States (USD, Pre-Fee)		

Notes: The sample is employers on their first job posting (column 1) or on their first job posting after having hired 2 previous workers (column 2). Mixed logit model estimates and costs and markups are described in the text. Standard errors are taken from a variance-covariance matrix calculated as the inverse of the outer-product of the gradient divided by the number of employers (rather than the applications or job applicants). Standard errors are not corrected for first-stage estimation error of the control function.

Table 6: Counterfactuals

Data	
Mean Hourly Post-Fee Bid on Job 1	3.95
Mean Hourly Post-Fee Bid on Experienced	3.12
Job 1 Under Experienced Cost Structure	
Mean Hourly Post-Fee Bid on Job 1 Under Experienced Cost Structure (From Table 4)	3.34
Preliminary Estimate of Optimal Fees ($)	
Estimate of Profit Maximizing Specific Fee on Experienced	0.33
Estimate of Profit Maximizing Specific Fee on Job Post 1 (Static)	0.58
Estimate of Profit Maximizing Specific Fee on Job Post 1 (Dynamic with 1 Future Job)	0.36
Estimate of Profit Maximizing Specific Fee on Job Post 1 (Dynamic with 2 Future Jobs)	0.26

All figures in this table are based on calculations using estimates from Table 4 columns 1 and 3.

Table 7: Log interviews per job opening fall with hiring experience

DV: Log Number of Interviews +1	OLS	Employer Effects	Employer Effects	Employer Effects	Employer Effects
	(1)	(2)	(3)	(4)	(5)
One previous hire	-0.0952***	-0.503***	-0.502***	-0.491***	-0.490***
	(0.0199)	(0.0538)	(0.0539)	(0.0545)	(0.0546)
Two previous hires	-0.133***	-0.602***	-0.603***	-0.586***	-0.587***
	(0.0208)	(0.0622)	(0.0624)	(0.0630)	(0.0632)
Three previous hires	-0.150***	-0.671***	-0.673***	-0.652***	-0.655***
	(0.0223)	(0.0684)	(0.0686)	(0.0694)	(0.0696)
Four previous hires	-0.158***	-0.726***	-0.726***	-0.707***	-0.708***
	(0.0237)	(0.0721)	(0.0723)	(0.0731)	(0.0733)
Five or more previous hires	-0.238***	-0.877***	-0.875***	-0.856***	-0.855***
	(0.0195)	(0.0820)	(0.0823)	(0.0833)	(0.0836)
Mean Log Bid				0.124***	0.121***
				(0.0125)	(0.0124)
Constant	0.833***	1.461***	1.423***	1.355***	1.322***
	(0.0423)	(0.0934)	(0.0933)	(0.0957)	(0.0956)
Observations	37,704	37,704	37,624	37,702	37,622
Includes job duration fixed effects and third order polynomial of job description length	No	No	Yes	No	Yes

Notes: Robust standard errors are clustered by employer. All specifications contain year-month fixed effects as well as job category and job duration fixed effect.

Table 8: Productivity and Search Effort on the First Job, for employers who interview on first job
Data Entry and Web Research Sample

	Hires a Worker	Hires and Reports Success	Hires with Good Feedback	Hires on Openings with 1+ Interview: Success	Success, Sequential Openings	Success, also add Wage Control	Success, also add Controls for Worker Characteristics
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A							
2-4 Interviews	-0.037**	0.005	-0.031**	-0.052*	-0.035	-0.032	-0.033
	(0.017)	(0.022)	(0.015)	(0.030)	(0.032)	(0.031)	(0.031)
5-8 Interviews	-0.061***	-0.028	-0.049**	-0.090**	-0.096**	-0.089**	-0.090**
	(0.023)	(0.024)	(0.021)	(0.037)	(0.041)	(0.040)	(0.041)
9-11 Interviews	-0.091***	-0.009	-0.071***	-0.137**	-0.163**	-0.161**	-0.162**
	(0.032)	(0.037)	(0.026)	(0.059)	(0.067)	(0.067)	(0.067)
12+ Interviews	-0.188***	-0.096***	-0.109***	-0.014	0.009	0.015	0.014
	(0.028)	(0.028)	(0.025)	(0.053)	(0.064)	(0.064)	(0.065)
Mean of DV	0.338	0.197	0.193	0.707	0.746	0.746	0.746
Observations	4,993	2,456	4,502	1,414	1,187	1,187	1,187
Panel B							
2-4 Interviews	-0.072***	-0.053**	-0.069***	-0.059	-0.028	-0.021	-0.023
	(0.023)	(0.021)	(0.019)	(0.039)	(0.042)	(0.043)	(0.042)
5-8 Interviews	-0.118***	-0.100***	-0.090***	-0.094**	-0.052	-0.035	-0.038
	(0.031)	(0.026)	(0.027)	(0.045)	(0.054)	(0.055)	(0.055)
9-11 Interviews	-0.144***	-0.123***	-0.132***	-0.140**	-0.106	-0.100	-0.102
	(0.044)	(0.031)	(0.030)	(0.067)	(0.077)	(0.078)	(0.076)
12+ Interviews	-0.248***	-0.161***	-0.171***	0.010	0.067	0.087	0.082
	(0.034)	(0.027)	(0.028)	(0.078)	(0.082)	(0.078)	(0.080)
Mean of DV	0.353	0.222	0.203	0.707	0.747	0.747	0.747
Observations	4,147	3,897	3,715	1,226	1,038	1,038	1,038

Notes: Panel B includes fixed effects for groups of employers who interview workers with similar characteristics on the first job. Standard errors clustered by detailed job category x time. All specifications contain fixed effects for expected duration of job, time, and detailed job category. Columns 1-3 include all first openings. Columns 4 to 7 include all first openings where a hire is made.

Table 9: Productivity on the Second Job and Search Effort on the First Job, for employers who interview on the first job.
Data Entry and Web Research Sample

	Reports Success	Gives Good Feedback	Reports Success, Sequential Openings	Gives Good Feedback, Sequential Openings	Reports Success, Sequential Openings and Wage Control	Gives Good Feedback, Sequential Openings and Wage Control
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A						
2-4 Interviews	-0.066**	-0.086**	-0.131***	-0.131***	-0.128***	-0.129***
	(0.027)	(0.033)	(0.046)	(0.047)	(0.045)	(0.048)
5-8 Interviews	-0.080*	-0.193***	-0.158**	-0.209***	-0.149**	-0.201***
	(0.046)	(0.040)	(0.064)	(0.067)	(0.066)	(0.068)
9-11 Interviews	-0.100**	-0.080	-0.035	-0.016	-0.028	-0.008
	(0.048)	(0.057)	(0.061)	(0.076)	(0.064)	(0.075)
12+ Interviews	-0.131**	-0.126***	-0.241***	-0.237***	-0.229***	-0.223***
	(0.051)	(0.046)	(0.075)	(0.079)	(0.071)	(0.077)
Mean of DV	0.620	0.690	0.653	0.708	0.653	0.708
Observations	1,463	1,213	614	517	614	517
Includes Buyer Group Fixed Effects	No	No	No	No	No	No
Panel B						
2-4 Interviews	-0.055	-0.078	-0.107	-0.112	-0.105	-0.111
	(0.040)	(0.047)	(0.079)	(0.087)	(0.081)	(0.088)
5-8 Interviews	-0.081	-0.201***	-0.129	-0.191*	-0.126	-0.191*
	(0.060)	(0.055)	(0.102)	(0.096)	(0.103)	(0.097)
9-11 Interviews	-0.086	-0.110	-0.039	0.063	-0.036	0.063
	(0.067)	(0.085)	(0.109)	(0.139)	(0.111)	(0.139)
12+ Interviews	-0.090	-0.120*	-0.137	-0.115	-0.133	-0.112
	(0.063)	(0.062)	(0.110)	(0.127)	(0.113)	(0.127)
Mean of DV	0.614	0.688	0.644	0.700	0.644	0.700
Observations	1,263	1,045	525	440	525	440
Includes Buyer Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Sample is employers who do at least one interview on the first job. Panel B includes fixed effects for groups of employers who interview workers with similar characteristics on the first job. Standard errors clustered by detailed job category x time. All specifications contain fixed effects for expected duration of job, time, and detailed job category.

Appendix Table 1: Arrival Rates of Applicants by Job Opening Number

Posting Number	Number of Applications in First Hour	Number of Applications in First Hour	Number of Applications in First Two Hours	Number of Applications in First Two Hours	Number of Applications in First Day	Number of Applications in First Day
Two	0.974***	0.356	1.311***	0.460	2.422***	0.290
	(0.174)	(0.462)	(0.253)	(0.665)	(0.678)	(1.635)
Three	1.099***	0.384	1.438***	0.417	2.552***	0.870
	(0.201)	(0.577)	(0.287)	(0.825)	(0.776)	(2.212)
Four	0.976***	0.710	1.279***	0.777	1.691**	0.827
	(0.242)	(0.668)	(0.338)	(0.960)	(0.831)	(2.425)
Five	1.400***	0.437	1.893***	0.479	3.176***	-0.638
	(0.302)	(0.708)	(0.436)	(1.052)	(1.174)	(2.942)
Six	1.139***	0.182	1.580***	0.409	2.696**	-0.280
	(0.280)	(0.671)	(0.410)	(0.971)	(1.115)	(2.703)
Seven	1.639***	0.602	2.244***	0.645	4.487***	0.532
	(0.343)	(0.757)	(0.491)	(1.125)	(1.355)	(3.427)
Eight	1.404***	0.102	1.881***	0.106	3.005**	-1.626
	(0.339)	(0.786)	(0.472)	(1.136)	(1.319)	(3.048)
Nine	1.956***	0.432	2.591***	0.467	4.469***	-1.720
	(0.387)	(0.859)	(0.551)	(1.216)	(1.369)	(3.362)
Ten or greater	1.579***	0.277	2.153***	0.325	4.606***	-0.172
	(0.138)	(0.612)	(0.197)	(0.880)	(0.562)	(2.345)
Constant	8.426***	10.01***	12.55***	15.18***	29.89***	38.11***
	(0.735)	(2.407)	(1.118)	(3.632)	(2.459)	(7.513)
Observations	17,977	17,977	17,977	17,977	17,977	17,977
Errors clustered by	Employer	Employer	Employer	Employer	Employer	Employer
Employer fixed effects	No	Yes	No	Yes	No	Yes

Notes: Robust standard errors in parentheses. The sample is sequential jobs where the vacancy duraction exceeds one hour. All specifications contain fixed effects for expected duration of job, time, and detailed job category.