Experience Markets:
An Application to Outsourcing and Hiring

Christopher T. Stanton
Harvard, NBER, and CEPR

Catherine Thomas
LSE, CEP, and CEPR

April 2018*

Abstract

Trying out a market often allows new buyers to learn about the distribution of products or services available. We study how hiring varies with employer experience in an online labor market platform where employers contract with remote workers. Experience affects employer selection into posting further jobs, increases the perceived value of using the market, and alters how employers evaluate individual workers. These changes in demand for workers are consistent with experience reducing employer uncertainty about the distribution of applicants. Job applicants respond with higher wage bids to new employers. We evaluate policies that encourage employer learning about individual market fit.

*A previous version of this paper was circulated under the title "Information Frictions and Observable Experience: The New Employer Price Premium in an Online Market." We thank seminar participants at the AEA Meetings, the CEPR Workshop on Incentives, Management and Organisation, Chicago Booth, Columbia, Copenhagen Business School, ESSEC, Harvard, HBS, LMU, LSE, Mannheim, NBER Summer Institute, Queens, Stanford, Sussex, Utah, and Yale, along with Ajay Agrawal, Ricardo Alonso, Nava Ashraf, Heski Bar-Isaac, Zoe Cullen, David de Meza, Ricard Gil, John Horton, Lisa Kahn, Bill Kerr, Ed Lazear, Jin Li, Rocco Macchiavello, Kristof Madarasz, Arnaud Maurel, Ramana Nanda, Amanda Pallais, Luis Rayo, Yona Rubinstein, Scott Schaefer, Kathryn Shaw, Andrei Shleifer, and Nathan Seegert for helpful comments.

1 Introduction

Online labor platforms are new technologies that allow remote employment across country borders, with huge potential allocative efficiency gains (Clemens, 2011; Cairncross, 2001). They have also become popular laboratories for studying how information affects labor market outcomes. Several papers show that having work experience in these markets improves an applicant's job prospects by providing the employer with information about them.[1] Little is known, however, about how employers' market experience affects their use of these markets and how encouraging new employers to gain experience affects market growth.[2] Studying the employer side of this market is instructive about demand for gig economy services (Katz and Krueger, 2016), for online offshoring in particular, and, more generally, for how employers hire in new labor markets.

We analyze administrative data from oDesk.com, the largest online labor market, covering approximately 80,000 employers, more than 300,000 job postings, and over five million job applications. The data reveal two stylized facts that suggest that employer experience in the platform affects hiring processes and outcomes. First, over half of the new employers who begin the recruiting process by posting a job never ultimately hire anyone on the platform. The probability that an employer hires any applicant for a job opening then doubles if the employer has previously hired a worker online. Second, workers on oDesk bid an hourly wage when applying to jobs. Inexperienced employers, who lack visible evidence of past hiring, receive hourly wage bids that are over 7% higher than the bids made to employers who have several past hires displayed in their profiles.

Given these stylized facts, our analysis has two main goals. The first goal is to provide evidence distinguishing different explanations for why hiring and worker applications vary so dramatically with employer experience. Our main finding is that a change in employer demand with experience—shorthand for the probability that a single applicant is hired from the set of workers applying to a job—is the main reason for these patterns. It is potentially plausible that workers have higher expected costs of applying to or working for new employers. However, there is limited evidence that cost differences, or other labor supply channels—like applicants' sorting to job openings or changes in competition for jobs, explain the patterns in the data. The second goal of the paper is to propose and interpret a model of employer demand in which market experience provides an employer with information that affects his hiring processes and decisions in ways that match the patterns observed in the data.

Our evidence that demand differs with employer experience comes from estimating employer hiring

[1] See, for example, the papers based in the oDesk.com labor market platform, now Upwork.com, including Pallais, 2014; Stanton and Thomas, 2016; and Agrawal et al., 2016. Moreno and Terwiesch, 2014 study similar phenomenon on a platform called vWorker.

[2] Two papers by John Horton study employers' decisions in these platforms, but not how their behavior changes as they gain market experience. Horton (2017a), focuses on employers' reactions to the introduction of price floors in an online labor market and Horton (2017b) studies how candidate recommendations change employers' propensity to fill jobs.

probabilities that vary depending on whether the employer is inexperienced or experienced. We model demand as the probability that an applicant is hired as a function of her resume characteristics and the wage bid that she submits to the job opening, accounting for the set of applicants who apply. The model used is a modified conditional logit estimator. Because the hourly wage bids in job applications may be correlated with unobserved attributes of both workers and job openings, we use an instrumental variables strategy and Petrin and Train's (2000) control function estimation approach. The main instrument exploits the fact that all contracts are denominated in US dollars but applicants' local currencies fluctuate relative to the dollar, leading to exogenous variation in the relative attractiveness of workers' supplying online and offline work. Other instrumentsbased on variation in competition for similar jobs posted by other employersshift bids for workers in countries with dollar-linked currencies.

Our estimates separate three channels through which employer experience affects employer hiring in this market. First, holding fixed the set of applicants to any given job, hiring probabilities are more elastic with respect to individual workers' bids among experienced employers. That is, the percentage change in the probability an individual applicant is hired in response to a percentage change in the wage bid submitted is larger for experienced employers than for inexperienced employers. Second, there is a within-employer effect that resembles learning-by-doing (or learning-by-hiring), shifting out experienced employers' demand. Third is an employer composition effect—the employers with the lowest market valuations select out of the market and do not appear as experienced employers.[3]

The change in the demand curve due to these last two channels is substantial—wage bids to inexperienced employers would have to be 30% lower than they actually are (or about $2.80 per hour lower on average) for hiring probabilities for the inexperienced to be as high as those for experienced employers. Although experienced employers are more likely to hire, their relative demand elasticity is consistent with the fact that they receive lower wage bids. When workers are differentiated from one another, the difference in estimated elasticities across employer experience segments attributes almost all of the wage bid reduction observed for experienced employers to lower equilibrium markups.[4]

The data also permit reduced-form analysis of the supply side of the market and an evaluation of alternative explanations for observed bid differences. We again find little evidence that differences in workers' bids reflect employer-experience-varying costs. Although workers' costs are unobserved, we look for variation in how workers bid to experienced employers who have, for example, good versus bad reputations from their past hiring.[5] We compare the wage bids made to experienced employers with bad reputations

[3] Using a finite mixture model, we allow the distribution of employer heterogeneity to differ between inexperienced and experienced employers, accounting for selection of low value employers out of the market.

[4] Inverting the demand estimates yields estimates of worker-job level marginal costs, which are only slightly higher, on average, for jobs posted by inexperienced employers.

[5] The importance of reputation mechanisms in online platforms is widely documented, in online labor markets and elsewhere. See, for example, Tadelis (2002) and Cabral and Hortacsu, (2010). Our data allow us to separate the effect of worker and employer reputations from other consequences of market experience.

to the bids made to experienced employers with good reputations and to inexperienced employers with no reputation. We find that experienced employers who themselves have received poor individual feedback actually receive lower wage bids than experienced employers who have good feedback. Given that experienced employers with poor reputations are expected to be costly to work with, this pattern is inconsistent with higher anticipated costs driving differences in workers' bids. Wage bid regressions show significant declines in bids with employer experience even after accounting for applicant sorting to jobs and employer composition changes through the use of worker and employer fixed effects.

It remains possible that job postings or the nature of work varies in unobserved ways as employers gain experience. While opportunities for large-scale experiments within the platform are limited, a small field experiment where we posted jobs with different levels of employer experience helps to rule out unobserved changes in how employers advertise jobs as an explanation for observed bid differences. Other possibilities—like differing expectations about the possibility of long-term work, additional hours, or future raises—do not appear consistent with patterns in the data.

The second goal of the paper is to offer an explanation of why employer experience affects demand. The basis for our preferred explanation is the fact that the demand estimation reveals large variation in new employers' values for the market relative to their alternative sources of labor. Because many low valuation employers post jobs but then fail to hire, we propose that a new employer is uncertain about his initial position in the distribution of values for hiring in this market. The demand differences that we document are consistent with experience reducing the employer's uncertainty about how well he fits with the distribution of workers in the market.[6] Our explanation for the mechanism behind the significant effect of employer experience on labor demand thus centers on employers' resolution of uncertainty about their own fit or type as they gain experience.[7]

Given the novelty of hiring a worker in India, for example, through an online platform for a would-be employer located in the United States, an employer entering the market is unlikely to know exactly what workers offer and how well the available workers map to his particular needs.[8] In our characterization of labor demand in this market, employers have to try out the market to learn their own value for the products available in it. We describe settings like this as "Experience Markets," with three defining features: (i) differentiated products or services are on sale, each of which is an experience good, as in Nelson (1970); (ii) buyers do not know the average of the distribution of their values for the goods or services available; and

[6] A prior literature examines employment decisions when employers and/or employees are uncertain about the quality of the match (Barron, Bishop and Dunkelberg, 1985), and an even more extensive literature considers learning over time (Miller, 1984; Lange, 2007; Kahn, 2013; Kahn and Lange, 2014; Arcidiacono et al., 2016). A separate literature pioneered by Erdem and Keane (1996) considers learning about brand attributes through experience. Israel (2005) examines learning about service experience for auto insurers. Relative to these papers, we consider uncertainty about the set of offers available in a market rather than the attributes of a particular brand or company.

[7] Supporting this interpretation, we find employers' interviews decline with experience in the market. A declining number of interviews suggests exploratory motives to learn about the distribution of workers.

[8] An employer in the United States hiring a worker in India is the modal hire in the data.

(iii) buyers update their beliefs about their own willingness to pay for the average good or service through sampling the products or services. Here the buyers are employers and the sellers are potential workers, but the Experience Market definition is likely to apply to other settings, particularly to new markets for unfamiliar types of goods or services. In these markets, differentiated sellers who can price discriminate will charge a new buyer price premium in equilibrium.[9]

We use a simple theoretical framework to illustrate how the resolution of uncertainty about an employer's mean value for hiring applicants in the market changes his demand elasticities. An employer with a more uncertain value for the market perceives applicants to be more differentiated from each other because he uses the information contained in job applications to assess both his own fit with the market as well as his match quality with each individual applicant. This results in the employer putting more weight on the noise contained in job applications. Because an uncertain employer is less able to filter out noise from the actual productivity of an applicant, he has a higher relative willingness to pay for the applicant who sends the most favorable signal relative to the next-most favourable. His demand is, therefore, less elastic with respect to the wage bid received in the application. Under this interpretation, a comparison of the elasticities of demand between employer segments is informative about the extent of inexperienced employer uncertainty.

A further goal of the paper is to evaluate platform policies to encourage new employers to gain experience. Sellers in spot markets, like the one studied here, fail to internalize the value of market experience for a buyer's future transactions. This may result in prices to new buyers that are too high or, in situations with less price flexibility, in a reluctance to contract with inexperienced users. Understanding how early platform trials translate into subsequent long-term adoption is therefore a key input for platform design. Several possible policies may be relevant, including inducements to encourage better workers to apply to inexperienced employers or fee changes that lower the price of new employer hiring. We concentrate on evaluating pricing changes to encourage early platform use.[10] We calculate optimal ad-valorem fees for employers with different experience and label lower fees for new employers a relative subsidies. Our demand estimates imply that the fee that maximizes platform profits is higher than the 10% fee oDesk charged during our sample period.[11] Although relative subsidies would increase transitions into the experienced segment of employers, the new employers who would be induced to gain experience would not hire enough in the future to recoup the overall reduction in fee revenue. That is, those induced to hire in the future

[9] Shapiro (1983) first analyzed the experience good pricing problem. Cremer (1984) shows that a new buyer price premium is an optimal pricing strategy in settings where buyers are initially uncertain of their willingness to pay for a monopolist-supplied good. This is because buyers are willing to pay to experiment and learn their valuation. Even in settings where sellers cannot price discriminate by buyer experience, common prices to all consumers might change to reflect the expected composition of new and experienced buyers, as in Bergemann and Valimaki (2006)'s time-consistent monopoly pricing model where the price depends on the fraction of informed buyers at time t. The aggregate uncertainty about the value of the market, in our characterization, is closest to Benabou and Gertner (1993).

[10] While it may be important in some contexts, we do not find evidence of worker reluctance to apply to inexperienced employers in this market.

[11] After the end of our sample period, the platform did change this fee, and the average fee rate is now higher than 10%.

through lower fees would not hire enough to recover the opportunity cost of subsidizing new employer entry. These results suggest a platform policy targeting high value customers maximizes platform profits. It is these employers who will, in expectation, discover that the platform works for their needs.

Taken together, our findings demonstrate that missing information about an employer does not hamper their market entry. This stands in contrast to how missing information about a worker slows their careers on the platform. Instead, for the majority of employers, a small amount of experience allows them to acquire information that then prompts market exit. Moreover, while this and similar platforms offer significant wage savings relative to hiring workers in developed labor markets, the heterogeneity in employer type that determines platform adoption appears to limit most employers' ability to realize labor savings.

What do our findings tell us about platform policy and labor services offshoring more generally? The fact that a niche positioning is profit maximizing for this platform has implications for the growth of offshoring through online channels. Early commentators forecast that the Internet would match the global workforce to dispersed employers more efficiently. Aggregate data makes it clear, however, that only a small minority of firms currently finds it valuable to offshore labor services online. Temporary staffing was a $375 billion annual business in 2015, but domestic outsourcing or offshoring through third-party companies accounted for around 99% of this total (staffingindustry.com). The low take up rates overall, and the high exit rate for those who experiment with using this particular platform, suggest that hiring individual workers via these channels incurs additional costs that outweigh the potential variable labor cost savings within these markets. Fort (2017) shows that only 2% of U.S. manufacturing plants fragment production by purchasing services abroad. She finds that adopting electronic communication technology is more strongly associated with an increase in domestic contracting than in offshore contracting and suggests this is because coordination costs remain relatively high across country borders. Since exiting employers in the setting studied here had considered entry worthwhile ex ante, we suggest that some of these additional costs were hidden to them prior to market entry.

The paper proceeds as follows: Section 2 introduces the data and the empirical context. Section 3 presents the hiring probability estimation for inexperienced and experienced employer segments. Section 4 provides a framework that relates the estimated differences in employer demand across experience segments to variation in the extent of the employer's knowledge about their own market fit. Section 5 presents reduced form evidence about the supply-side response to employer inexperience. This section also presents some additional evidence consistent with employer learning about fit with experience. Section 6 presents the counterfactual analysis of employer learning, market volumes, and platform profits under varying platform fees. Section 7 concludes.

2 The Setting, Data, and Summary Statistics

2.1 oDesk.com: How It Works

oDesk.com (re-branded as Upwork after our sample ends) is an online platform that allows employers to contract with remote workers who sell labor services.[12] The platform facilitates search and matching, remote task and project management, and payments. Work 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 of postings requiring less than three months of work. Around 85 percent of the transactions in the market span international borders.

An employer who wants to purchase online labor services creates an account on the platform, with no up-front charge. To post a job opening, the employer must select the job's work category and its expected duration, give the job a title, and describe the work to be done and the necessary skills. Once the posting is in the system, potential applicants learn about the job by searching on the site or through automatic notification. Like the example in Figure 1, the postings contain information about the employer and the job. The employer's experience in the market (in the bottom right-corner: "About the Client") is prominently displayed.

Figure 1: A Job Posting.

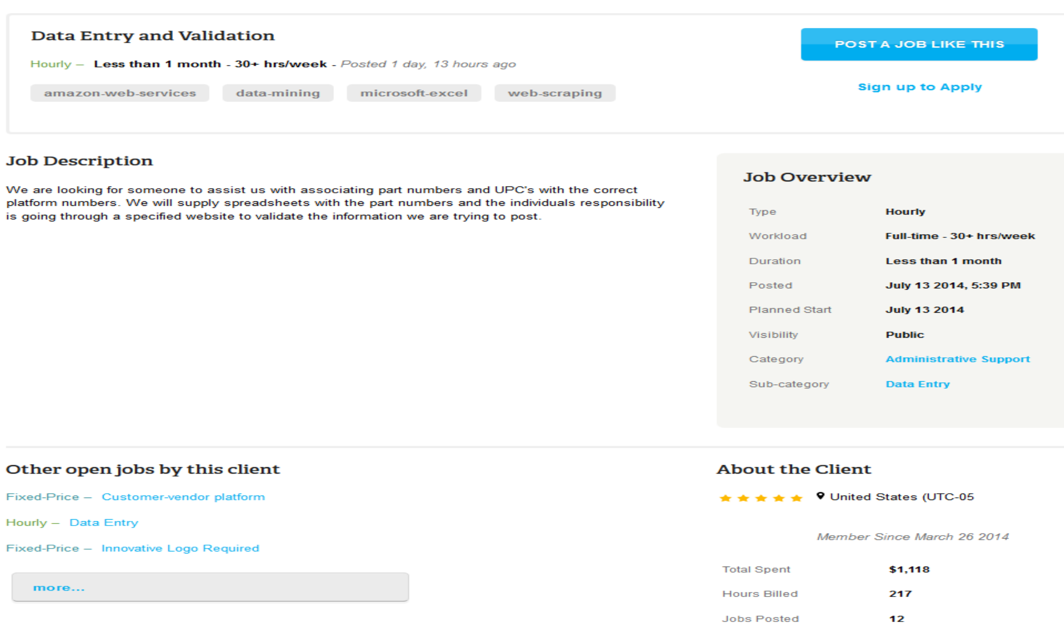

[12] See Horton (2010) for an overview of online labor markets. Other prominent platforms included elance and Guru. eLance merged with oDesk in 2014. The merged company changed its name to Upwork. The data used in this paper pre-date the merger. While several recent papers use data from online labor markets (Agrawal et al., 2013; Ghani et al., 2014; Horton et al., 2017; Lyons, 2017; Pallais, 2014; and Stanton and Thomas, 2016), the majority study information frictions at the worker level. The main exceptions are Horton (2017a) and Horton (2017b).

Interested workers submit applications for the job posting and bid an hourly wage to work on the specific job. Employers also have the option of searching worker profiles directly and inviting applications. Workers' profiles contain information about their skills, education, prior offline work experience, and experience on oDesk (see Figure 2). Workers are located worldwide, and each application clearly displays the applicant's country. The profiles of workers with prior experience on the site show summary feedback scores received from past work on a five-point scale.

Figure 2: A Worker Profile.

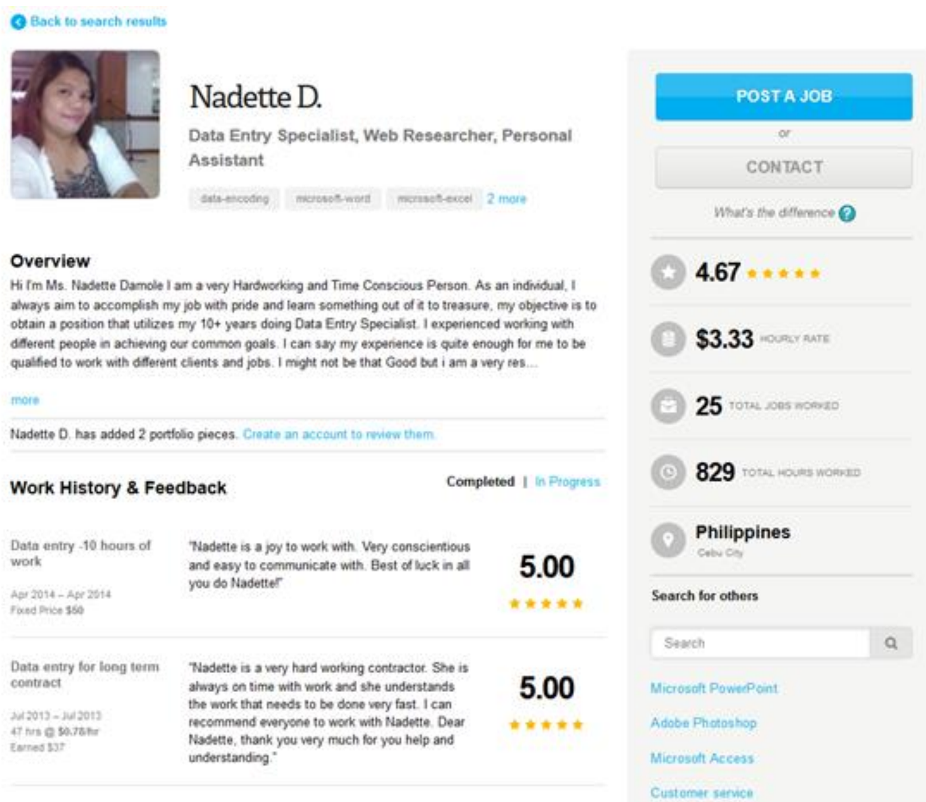

After receiving applications and initiating candidacies, employers can request interviews with any number of applicants for the job. If the applicant agrees, the interview usually takes place via Skype. Whether an interview actually occurs is not recorded in the oDesk database, so the remainder of this paper refers to an accepted interview request as an actual interview.[13] An employer may choose to hire an applicant with or without interviewing her first.

Upon hiring a worker, the employer can monitor work via software provided by oDesk, and oDesk manages all payments for completed work. When a job is complete, the employer is asked for feedback about the worker and vice versa. The employer is also asked whether or not the job was completed successfully. In other contexts, having a market reputation or having received good feedback may be

[13]The data record that an interview takes place whenever an employer sends an interview request to a worker and the worker accepts the request.

negatively associated with the likelihood of a partner absconding without making payment, but oDesk guarantees that workers are paid for the hours billed. Thus, payments and payment risk are unrelated to employer reputation or experience.

The data used in this paper are administrative data from oDesk. In the data, we observe every employer's job-specific search process in each of his successive job postings. For each posting, the data contain information about the entire applicant pool; which candidates, if any, are interviewed; which candidate, if any, is hired; and the feedback and success measures that the employer leaves for the hired worker, and vice versa.

2.2 Summary Statistics by Employer Experience

The data contain job postings and applications between January 2008 and June 2010. They include 82,257 potential employers who posted 322,870 job openings that received more than five million job applications. There are nine job categories. Web Development, the largest technical job category, accounts for 38% of all job postings, followed by Software Development, which accounts for 9%. Administrative Support is the largest non-technical category, with 17% of the postings.[14]

Most employers are located in the United States and include private individuals and those hiring on behalf of firms. Figure 3 presents the distribution of the number of hires per employer throughout the period. 63% of the potential employers posting jobs make no hires at all, while 17% make five or more hires. On each job posting, the number of past hires by the employer and the total hours billed are evident (see Figure 1). Workers can also observe any feedback that employers have received from previous workers.

Table 1 presents summary statistics about job postings, grouped by the number of previous hires made by the relevant employer. Columns 1 to 5 report statistics for all job posts in the sample. Columns 6 to 10 restrict the jobs to what we term the sequential, arms-length sample—a subset of all job postings that we use for the hiring probability estimation in Section 3. Some restrictions are necessary for this estimation to make the hiring processes comparable to employer choice from a set of applicants. Some employers hire multiple workers simultaneously by posting a batch of jobs all at the same time, and, thus, it is not clear how to determine the set of applicants across job posts. Other situations provide difficulties as well, especially if employers bring workers onto the platform or target workers from pre-existing relationships. We restrict the job openings to those that received at least one worker-initiated application, had multiple applicants to consider, and then had a gap of at least one day before or after the same employer posted a different opening.[15]

[14] Other job categories are: Design and Multimedia; Writing and Translation; Sales and Marketing; Business Services; Networking and Information Systems; and Customer Service.

[15] See Appendix 1 and Appendix Table A1 for additional details about sample composition. Of the 119,877 jobs posted by employers without previous hires on the site, 61,160 survive these restrictions. The total number of postings across all employer experience levels falls to 109,764.

Figure 3: Number of Hires per Employer

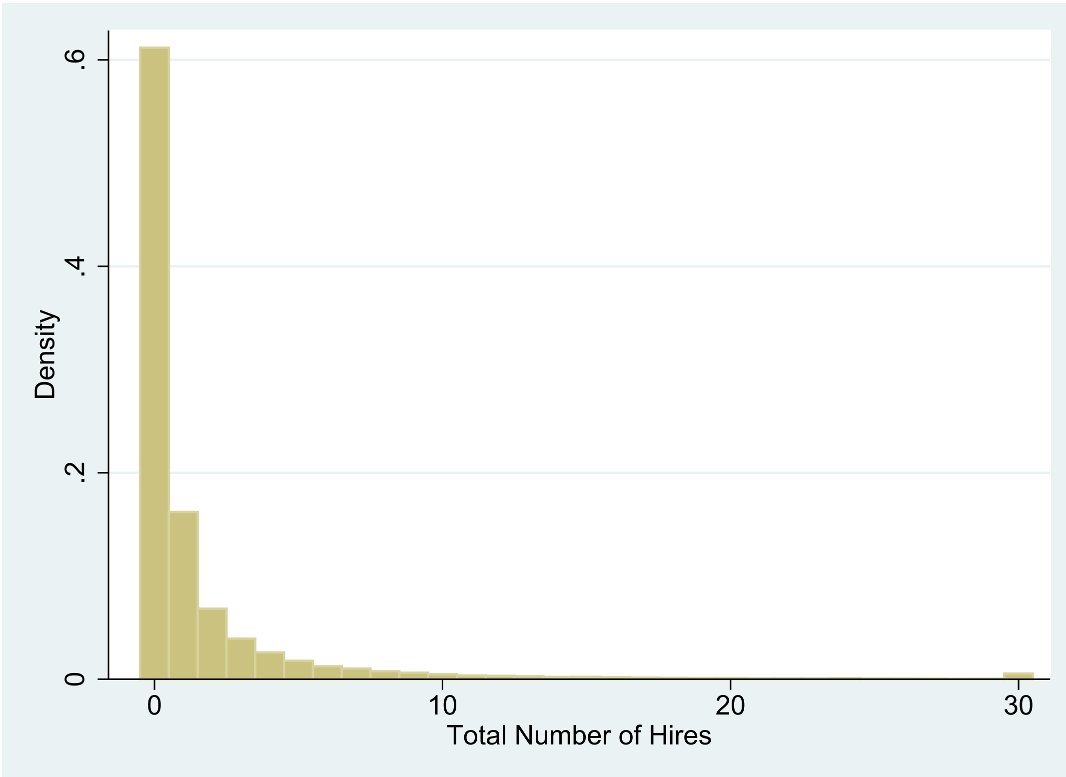

The unit of analysis is an employer. Total hires are censored at 30.

The table provides initial evidence that employer experience in the market is associated with different employer—and job applicant—behavior. The first row of Table 1 summarizes data for jobs posted by inexperienced employers, while subsequent rows summarize data from postings made by experienced employers. The employers in the overall (sequential) sample receive an average of 18.39 (25.47) applications prior to closing the job or making a hire. Only a small fraction of applications are employer-initiated, particularly in the sequential sample, which means that applicants initiate the vast majority. As employers gain experience, the number of applications per job falls in the overall sample, with the largest decline taking place between employers with zero and one prior hire (Column 2), but the number of applications increases with employer experience in the sequential sample.

Only 22% (16%) of openings posted by new employers result in a hire (Columns 5 and 10). Experienced employers are far more likely than inexperienced employers to hire on a given job post. Among employers with at least four prior hires, 57% (28%) hire a worker.

The mean hourly wage bid to employers without prior hiring experience in the overall sample is $10.16. In the sequential jobs sample, the mean bid to inexperienced employers is $10.25. Columns 4 and 9 show that the mean hourly bid declines with employer experience. For those with one prior hire, the average bid falls to $9.77 ($9.85 in the sequential sample). After four or more previous hires, employers receive a mean bid of \$8.71 (\$9.06 in the sequential sample), a decline of 14% (12%).

Table 2 presents evidence that an employer's experience has only limited impact on the resume characteristics of the workers that select into making job applications, in contrast to the observed relationship between employer experience and the wage bids received. It displays data on the average characteristics of applicants for job openings posted by employers with different numbers of past hires. Comparing the means and standard deviations across the first five columns of the table shows the similarity of the applicant pools. The second set of five columns then repeats this exercise with the characteristics of the workers who the employer ultimately hires.

3 Hiring Probabilities by Employer Experience and Implications for Applicants' Bids

Before formally specifying a model of employer demand, we show, in Figure 4, how hiring probabilities at different wage bids differ for new and for experienced employers. Each of the four panels shows a different job category. Overall, both the slopes and the intercepts of the hiring probability functions within a job category differ by employer experience, with experienced employers appearing to have more elastic hiring probabilities over the range of (residualized) bids that are most common in the data.

Other evidence suggests that workers appear differentiated to employers. For example, it is relatively rare for employers to hire the worker who submits the lowest hourly bid.[16] As a result, we consider a differentiated products demand model that captures how employers trade off price and applicant characteristics.

3.1 Employers' Hiring Decisions

To estimate and compare inexperienced and experienced employers' demand, we specify a modified conditional logit hiring probability function with parameters that vary over these two employer segments. For employer i with experience χ, $p_{i\chi j}$ denotes the probability that applicant j will be hired for the job. Our goal is to estimate $p_{i\chi j}$ as a function of the wage bid $w_{i\chi j}$ and worker and job characteristics, X_j. In the model, preferences over these inputs potentially vary with the employer's experience level χ. Consistent with institutional detail about this setting, we assume that the wage bid is a take-it-or-leave-it offer.[17] The

[16]In the Appendix, Figure 8 shows the distribution of the wage bid decile—calculated within each job opening—of the worker who was ultimately hired.

[17]Bargaining between the first offered wage and the starting wage for hired candidates appears to happen infrequently. We investigated this issue explicitly at the beginning of our work with these data, but due to data storage changes over time and the fact that we no longer have access to the company database where the queries were stored, analysis of bargaining using the data we have on hand requires more indirect methods. Appendix 2 provides more detail about this analysis and shows limited explanatory power for bargaining differences to explain the results. Some employers are observed to adjust wages upward, but this tends to happen with a lag after the start of a job.

Figure 4: Residual Hiring Probability as a Function of Residual Bids

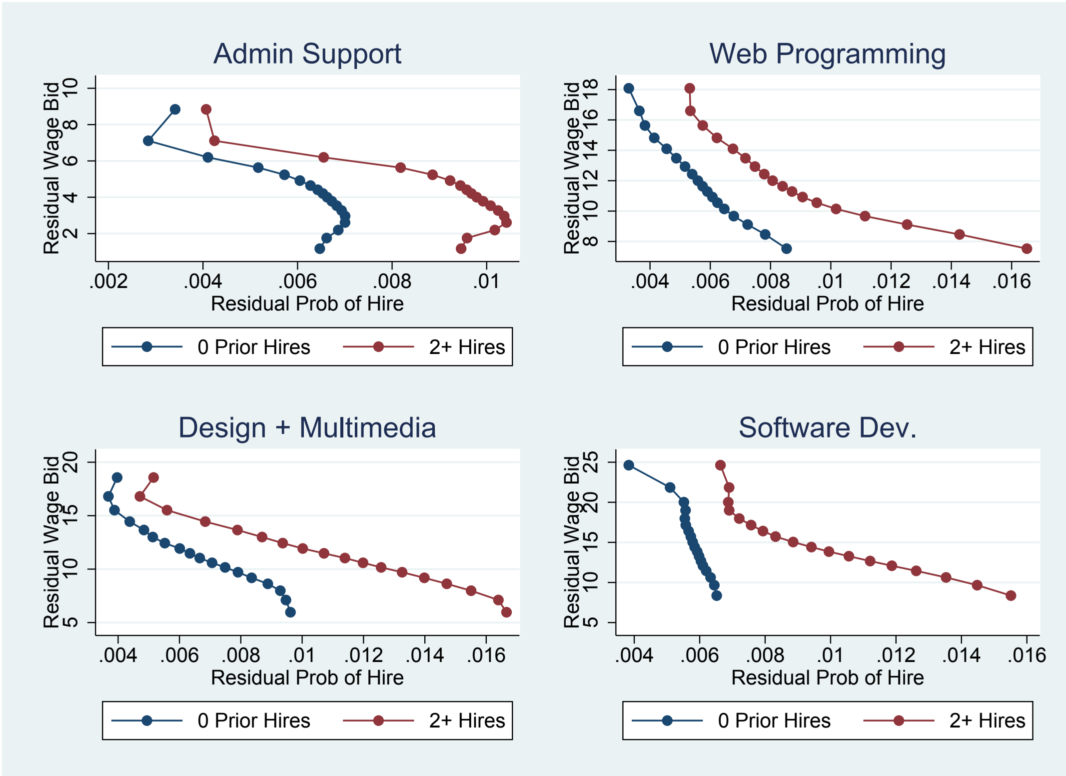

The unit of analysis is a job application. Wage bids and hiring probabilities are residualized within each job category using the worker resume data observed by the employer, a spline for the application number, and a linear time trend. Points are taken from a polynomial smoothing function of the residual hiring indicator on the residual bid.

subscript i is also used here as shorthand for a job opening to underscore that the probability of being hired is related to the set of job applicants for a particular job.

An employer's objective is to choose one worker out of the set of applicants for an opening, denoted $J_{i\chi}$, with the highest perceived quality per unit of wage.[18] The employer also has the option of hiring no worker, so the best worker's ratio of quality to wage must be greater than the value of the off-platform option, denoted option zero. The payoff to option zero is normalized to 1 in levels (or 0 after taking logs). The employer's objective function takes the form

$$\max_{j \in \{J_{i\chi}, 0\}} \frac{\exp\left(X_j \beta_\chi + \mu_i + \varepsilon_{ij}\right)}{\left(w_{i\chi j}\right)^{\alpha_\chi}}.$$

[18]Workers are assumed to be available when they initiate an application. This assumption is reasonable, requiring only that the probability that a worker will receive two offers over a time interval Δ is small. For example, if, from the worker's perspective, job offers 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 we do not observe declined job offers, this assumption on arrival of offers seems to be reasonable in the data. The observable arrival rate of interview requests fits this assumption. When a worker-day is the unit of analysis, only 3.6% of the worker-days in the sample have more than one interview request arriving; only 0.6% of the worker days have more than two interview requests arriving. A post-candidacy survey also asks employers for reasons that particular workers were not hired and asks workers their reasons for exiting the active candidate set. In some cases, employers or workers explicitly report a realized scheduling conflict. We drop cases of reported scheduling conflicts or when workers refuse invited applications.

To allow for differences in how worker characteristics are evaluated by experience segment, the parameter vector β_χ that specifies the relationship between hiring value and characteristics X_j has an employer experience subscript. This very general structure, allowing parameters to differ by experience, adjusts the scaling of the characteristics relative to the no-hire or outside option as employer experience differs.

The term μ_i is employer-specific heterogeneity and enters the problem to allow for differences across employers that provide variation around the constant term in X_j. That is, μ_i is an employer-specific term that varies the value of hiring on the platform. Because μ_i shifts employer i's value of hiring any applicant, this term does not alter his relative ranking over applicants for a given job, but it does determine the value of hiring an applicant versus not hiring in the market. The term ε_{ij} is assumed to be an idiosyncratic Type-1 extreme value shock for each alternative. Note that flexibility in the parameters by experience segment adjusts scaling relative to the variance of the unobserved shock because the type-1 extreme value shock is iid. We return to this point later, when interpreting how demand differs across segments.

Taking logs of the employer's objective gives a conditional logit function for the hiring probability. Worker j is hired by employer i with experience χ when

$$X_j\beta_\chi + \mu_i + \varepsilon_{ij} - \alpha_\chi \log\left(w_{i\chi j}\right) \geq X_k\beta_\chi + \mu_i + \varepsilon_{ik} - \alpha_\chi \log\left(w_{i\chi k}\right) \tag{1}$$

for all $k \in \{J_i, 0\}$. Conditional on μ_i, the probability that the inequality in (1) holds is given by:

$$\begin{aligned} p_{i\chi j} \quad = \quad & \exp\left(X_j\beta_\chi + \mu_i - \alpha_\chi \log\left(w_{i\chi k}\right)\right) / \\ & \left(1 + \Sigma_j^{J_i} \exp\left(X_j\beta_\chi + \mu_i - \alpha_\chi \log\left(w_{i\chi k}\right)\right)\right). \end{aligned} \tag{2}$$

Allowing uncertainty about an employer's value for the market to affect hiring elasticities is captured through the experience-dependent parameters. The presence of μ_i in the hiring probability function also relaxes a well-known limitation of standard conditional logit models—the independence of irrelevant alternatives or IIA assumption—allowing for different substitution patterns between the no-hire option and the available candidates in the choice set.

Several normalizations allow us to identify the coefficients. First, the variance of ε_{ij} is normalized. Second, if expectations about μ_i were fully observable to the researcher for each employer group, variation in prices, characteristics, and choices would identify the parameters. Because each employer's μ_i is unobserved, only the population distribution of types can be identified. For example, if many employers repeatedly hire when faced with low-quality applicants who submit high bids, while many other employers do not hire when high-quality workers with low bids are available, then the estimated population distribution of types would have a wide dispersion in valuations.

We also must confront selection, in that those employers who return to the market after hiring once are a non-random sample from the population distribution of μ. Because the experienced distribution is likely to truncate the lower tail, assuming symmetry of the type distribution or stability of the distribution over time would not account for employer selection. We build in flexibility to accommodate an arbitrary distribution of types that may change between segments by using a finite mixture model. In our preferred specification, employer types are fixed over time, but we allow the type probabilities to depend on whether the employer is ever observed in the experienced segment. This means that we assume that returning employers are drawn from a distribution with the same latent support as the distribution for the inexperienced employers, but with different probability weights at each support point. Because X_j has a constant term and β_χ is experience-specific, the mean of the distribution is allowed to change as an employer gains experience. Thus, we estimate how the distribution of μ_i varies between the segment of employers that is ever experienced and the segment that is never experienced.

3.2 A Worker's Optimal Bid

Another use of the model is to use the estimated demand elasticities together with data on wage bids to back out the implied markup over workers' costs. A worker j's cost, $c_{i\chi j}$, captures her outside option (opportunity cost of work) and her hassle costs from applying to and/or being hired for the job posted by employer i with experience level χ. When choosing the wage bid, the worker's objective function takes this cost into account along with the hiring probability.[19] It also reflects the ad-valorem platform fees retained by the platform, τ. She chooses the wage bid, $w_{i\chi j}$, that maximizes

$$\underbrace{p_{i\chi j}}_{\Pr(hired)} \times \underbrace{\exp\left(\log w_{i\chi j} - \log(1+\tau)\right)}_{Post-fee\ wage} + (1 - p_{i\chi j}) \times \underbrace{c_{i\chi j}}_{Cost}, \tag{3}$$

where $\log w_{i\chi j}$ is the log of the wage bid inclusive of τ. If she is hired, the worker receives the wage $\frac{w_{i\chi j}}{(1+\tau)} = \exp\left(\log w_{i\chi j} - \log(1+\tau)\right)$ and the employer pays $w_{i\chi j}$. If worker j is not hired, she receives $c_{i\chi j}$, her "net" outside option, which includes the opportunity cost of her alternative use of time, along with the expected direct costs of interviewing for or working on the job.[20]

Each worker's first order condition is given by

$$\frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}} \left(\frac{w_{i\chi j}}{(1+\tau)} - c_{i\chi j} \right) + p_{i\chi j} \frac{w_{i\chi j}}{(1+\tau)} = 0. \tag{4}$$

[19] In our setting, the wage bid is the worker's only strategic variable. We assume that the fit between employer $i\chi$ and worker j is unknown to both workers and employers, and workers cannot choose the productivity signal in their application. This differs from the broader model of quality and price choice in monopolistic competition with experience goods in Riordan (1986).

[20] The objective could alternatively be written:

$$\max_{\log w_{i\chi j}} p_{i\chi j} \times \exp\left(\log w_{i\chi j} - \log(1+\tau) - c_{i\chi jH}\right) + [1 - p_{i\chi j}] \times c_{i\chi jO},$$

where $c_{i\chi jH}$ is a differential hassle cost from on-the-job work associated with being hired for job i, and $c_{i\chi jO}$ is the outside wage for worker j. The first order condition in this case makes clear that only $c_{i\chi j} = c_{i\chi jH} + c_{i\chi jO}$ can be identified.

The system of equations containing the first order condition for each applicant determines equilibrium bids in a Bertrand Nash game in bids. Solving for worker j's optimal wage bid gives

$$w^*_{i\chi j} = c_{i\chi j}\,(1+\tau)\left(1 + p_{i\chi j} / \frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}}\right)^{-1}. \tag{5}$$

This says that the bid is related to three objects: $c_{i\chi j}\,(1+\tau)$, workers' costs and the ad-valorem platform fee; $p_{i\chi j}$, the employer's hiring probability as a function of the bid; and $\frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}}$, the semi-elasticity of the hiring probability with respect to the wage bid. The term $\left(1 + p_{i\chi j} / \frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}}\right)^{-1}$ is the markup over the worker's job-specific costs.

The bid equation can be rearranged to give workers' costs,

$$c_{i\chi j} = \frac{w_{i\chi j}}{(1+\tau)}\left(1 + p_{i\chi j} / \frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}}\right), \tag{6}$$

illustrating how having an estimate of $p_{i\chi j}$ and $\frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}}$, together with the bids and platform fees observed in the data, yields worker-job specific estimates of costs and markups.

3.3 Instruments

It is likely that worker applications contain information other than observed resume characteristics that are relevant to employer demand and also correlated with bids. Correlation of the error and log bids will result in inconsistent estimates of both the price elasticity, α_χ, and the β_χ vectors. Using the workers' first order conditions for bids illustrates potential cost shifters that can be used to account for endogeneity while also defining how demand differences between segments influences workers' bids. We use an instrumental variables strategy based on changes in the dollar-to-local-currency exchange rate for each applicant. Workers are paid in their local currency for offline work but are paid in dollars for their work on oDesk. Frictions limiting exchange rate pass through to local wages mean that offline opportunities are likely to adjust to exchange rates more slowly than online transactions.[21] When the dollar appreciates relative to the local currency, so that one dollar earned on the site provides fewer local currency units, workers' wage bids are predicted to increase.

To see this how this variation affects workers' bids, assume that $c_{i\chi j}$ is denominated in the local currency, while the bids observed by employers are denominated in dollars. Costs in the local currency must be translated into dollars when submitting bids, so the worker's optimal bid becomes $w^*_{i\chi j} = c_{i\chi j}\left(\frac{D}{L}\right)^{\theta}(1+\tau)\left(1 + p_{i\chi j} / \frac{dp_{i\chi j}}{d\log w_{i\chi j}}\right)^{-1}$. The dollar-to-local-currency exchange rate is $\frac{D}{L}$, and the parameter θ captures possible reasons for deviations from complete pass through.[22]

[21]This potential source of variation was revealed in conversations with employers who mentioned the frequency with which exchange rate calculators appear in the screenshots taken by oDesk's monitoring software.

[22]These reasons include: i) some part of a worker's opportunity cost reflects transactions denominated in dollars rather

The worker's optimal log bid can be written as a mapping from local-currency-denominated opportunity costs to dollar-denominated bids, as

$$\log\left(w_{i\chi j}\right) = \theta \log\left(\frac{D}{L}\right) + \log\left(c_{i\chi j}\right) + \log\left(1+\tau\right) - \log\left(1 + p_{i\chi j} / \frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}}\right), \tag{7}$$

which forms the first stage regression for the hiring probability estimation. To control for secular trends and level differences in local exchange rates across countries, each series is detrended, and its time series mean is removed.

Figure 5 illustrates the time-series variation in mean residual log bids and detrended exchange rates that underpins this estimation approach. The top left panel plots the mean residual log bid made by job applicants located in India and the log of the US dollar to Indian rupee exchange rate. The bid and exchange rate time series move together over the time period studied. The other panels in Figure 5 plot the difference between the mean of residualized bids from applicants in India and those from applicants in five other common worker locations (the dots in the figures), along with the difference in the exchange rate between the relevant local currency and the Indian rupee (the crosses in the figures). The figure suggests that the relevant local currencies for the job applicants in the data demonstrate independent dollar exchange rate variation during the time period.[23]

While exchange rate movements are plausibly exogenous to demand on oDesk, there are two additional concerns. First, a subset of job applicants—those based in the US or living in countries with dollar-pegged exchange rates—do not face any cost shocks from exchange rate variation. For these workers, there may be independent variation in relative prices because other applicants face exchange rate shocks, but there is no variation relative to the employer's outside option of not hiring. An additional instrument for worker j's bid that is relevant to all workers, including those with dollar or dollar-pegged local currencies, measures exogenous variation in the intensity of competition for a job posted by employer i with experience χ. It is based on the average number of application arrivals in the first 24 hours for other jobs in the same category in the same week. Job category and week fixed effects are netted out of the application numbers, making this an instrument that varies the extent of competition across job categories, holding fixed the average competition in the category over time and the average competition across all job categories during the week in question.

The second concern about the identification strategy is that there is likely to be sorting on the instruments affecting the composition of workers who apply. For example, an appreciation of a local exchange rate may lead a non-random set of potential applicants to seek work elsewhere. Because many workers

than in the local currency, which may occur if the possibility of receiving an alternative wage comes from searching online; ii) part of a worker's consumption may become cheaper through imports; and iii) the incidence of exchange rate variation is split between workers and employers.

[23] It should be noted that different panels are based on differing numbers of observations; Indians are about 40% of the sample, while Russians and Ukrainians submit under 3% of the total observed bids.

Figure 5: Mean Residual Log Bids and Detrended Exchange Rates.

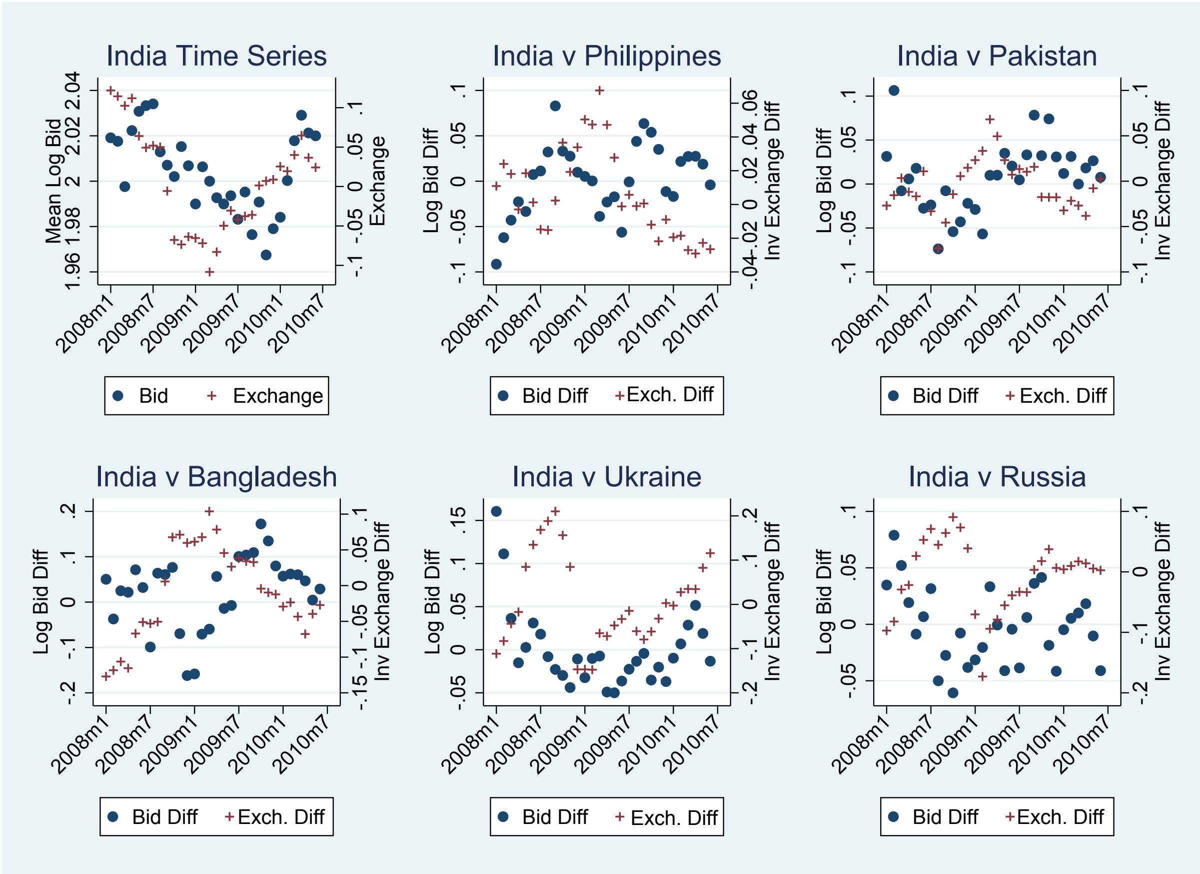

The top left panel plots mean residual log bids against the log of the US Dollar to Indian Rupee exchange rate after removing a time trend. The remaining panels plot log bid differences between India and other countries (left y-axis) and the log of other currency to the Indian exchange rate (right y-axis).

submit only a small number of total applications, it is difficult to assess the extent of selection into application based on unobservable worker characteristics and exchange rate variation. However, assessing the sensitivity of the parameter estimates to the inclusion of different sets of observable worker characteristics offers some insight into whether worker selection into application biases the estimates. For this reason, we discuss a comparison of the estimates with and without worker-level resume data below.

While hiring probabilities and worker costs may be affected by a worker-level unobserved characteristic that affects the employer's perception of quality, the two instruments are plausibly independent of this term other than through the sorting concerns. To make use of the variation in bids induced by the instruments, we use Petrin and Train's (2010) control function approach, putting the two worker-level instruments, Z_j, and worker characteristics, X_j, in a first stage regression of the form

$$\log(w_{ij}) = \gamma_0 + Z_j\gamma_{1\chi} + X_j\gamma_{2\chi} + \nu_{i\chi j}. \tag{8}$$

The coefficients in Equation 8 are estimated separately for the group of new and experienced employers.

The results in Table 3 show that both instruments have a substantive and statistically significant effect on workers' bids. The first column provides estimates for inexperienced employers, including many resume and job controls. Column 2 provides the analogous estimates for experienced employers. The signs of the estimated coefficients are as expected. Bids increase when the local exchange rate increases and decrease with the level of competition on the job. In both cases, the F statistics are extremely large, indicating the strength of the instruments.

Columns 3 and 4 exclude the detailed worker resume data—those columns of X_j in Equation 8 that contain worker characteristics. A comparison of the estimated γ_1 coefficients between those given in the first and second pairs of columns provides some evidence about the extent of sorting on the instrument. Under the null of no sorting, the estimated parameters in Columns 1 and 3 and in Columns 2 and 4 would be statistically indistinguishable. The estimated parameters, while of the same sign in both panels, are larger in absolute magnitude in the columns that exclude the worker characteristics, thus suggesting that there is some sorting into applying in response to exchange rate and/or competition variation.[24] Later, estimation of the hiring probability function with and without worker characteristics will help us to map this sorting into how hiring elasticities change.

3.4 Estimation

This section presents the likelihood over sequences of employer choices across different job posts.[25] The step-by-step approach is as follows. First, the residuals from Equation (8) form control functions for unobserved worker quality, denoted $CF_{i\chi j} = \hat{\nu}_{i\chi j}$. Second, we form choice probabilities conditional on a value of the unobserved term μ_i, taking the form:

$$\begin{aligned} p_{i\chi j} &= \exp\left(X_j\beta_\chi + \mu_i - \alpha_\chi \log\left(w_{ij}\right) + \psi_\chi CF_{i\chi j}\right) / \\ &\quad \left(1 + \Sigma_{k=1}^{J_i} \exp\left(X_k\beta_\chi + \mu_i - \alpha_\chi \log\left(w_{ik}\right) + \psi_\chi CF_{i\chi k}\right)\right). \end{aligned} \tag{9}$$

Third, we assume that μ_i is drawn from a distribution with three distinct types: $\mu_i \in \{\beta_{0\chi}, \beta_{0\chi} + \mu_2, \beta_{0\chi} + \mu_3\}$. For the first type, $\mu_1 = \beta_{0\chi}$ is a constant term that is allowed to vary with employer experience. For the other two types, the deviation from $\beta_{0\chi}$ remains constant with experience. That is, a type 2 employer, who hires on the first job and then posts two additional jobs will have $\beta_{0\chi=I} + \mu_2$ on the first opening and $\beta_{0\chi=E} + \mu_2$ on subsequent openings, where the $\chi = I$ and $\chi = E$ subscripts refer to the parameters for the inexperienced and experienced segment, respectively. Fourth, we then allow

[24] Appendix Table A2 repeats this analysis including worker fixed effects, but the incidental parameters problem means that control functions including the fixed effects are not consistent. The instruments remain strong in these specifications.

[25] In some simple models that omit employer heterogeneity, the likelihood is specified at the job-opening level and μ_i is set to zero.

a flexible pattern of selection into becoming experienced by letting the type probabilities depend on the eventual experience of an employer, using the superscript $S = E$ to denote the ever-experienced group of employers and $S = N$ to denote the never-experienced group. Hence, the type probability vector for an employer who is ever observed in the experienced sample is $\rho^{S=E} = (\rho_1^{S=E}, \rho_2^{S=E}, \rho_3^{S=E})$. On the other hand, the type probability vector for an employer who is never observed in the experienced sample is $\rho^{S=N} = (\rho_1^{S=N}, \rho_2^{S=N}, \rho_3^{S=N})$. The type probabilities are invariant within employer.

We then form the likelihood, which is defined over sequences of employer choices. The probability of a sequence for the employer's choices conditional on μ_i is the product of the choice probabilities for the alternatives selected, $(y = j)$. But, because μ_i is not observed, the marginal likelihood must be used by summing over the likelihoods for different employer types.[26] The marginal likelihood for an employer who ever posts a job in the experienced segment is the sum over the K types weighted by the probability of that type:

$$L = \Sigma_{k=1}^{K} \rho_k^{S=E} \Pi_i \Sigma_j p_{i\chi j} \left(\mu_k\right)^{y=j} .$$

The i subscript on the product term here is a slight abuse of notation; the product of probabilities is taken over all openings posted by the employer. The likelihood contribution for employers who are never observed in the experienced segment is

$$L = \Sigma_{k=1}^{K} \rho_k^{S=N} \Pi_i \Sigma_j p_{i\chi j} \left(\mu_k\right)^{y=j} .$$

3.5 Results

3.5.1 Markups and Wage Bid Sensitivity

Table 4 presents the results for the hiring probability estimation using the sequential openings data sample. There are 61,196 postings by employers with no experience and 48,618 openings by employers with prior experience. The coefficient on the log hourly bid across columns differs by the employer's experience level at the time of posting. The odd-numbered columns show estimates of $\alpha_{\chi=I}$ for employers who have never hired before on the site. The even-numbered columns present deviations from the inexperienced segment (that is, parameter estimates on interactions with experience) for the experienced segment. The important results to note relate to differences in estimated coefficients for the two groups. A comparison across each pair of columns shows that α_χ, the estimated coefficient on the log hourly wage bid, differs substantially with the employer's experience level at the time of posting. It implies significantly larger elasticities for experienced employers across all pairs of columns. Panel B presents some of the implications of these estimates.

Columns 1 and 2 present results without instrumenting for the wage bid. In these models, the likelihood

[26] For further details, see Train (2003).

is at the job-opening level, without taking the product over the sequence of employer job openings. Both the inexperienced and experienced employers in Columns 1 and 2 have inelastic wage bid demand.[27] Columns 3 and 4 include the estimates after including the control function from the first stage estimates. The estimated coefficients and implied wage elasticities become more reasonable: inexperienced employers have an estimated wage elasticity of -5.46, while experienced employers are more wage elastic, with an estimated elasticity of -8.54.

The remaining columns in Table 4 present the full model results. Columns 5 and 6 exclude worker resume characteristics from the first stage (corresponding to Columns 3 and 4 in Table 3), and Columns 7 and 8 include these characteristics (as per Columns 1 and 2 of Table 3). In the specifications in Columns 7 and 8, the mean own-bid elasticity for inexperienced employers is -4.96. The experienced employer segment is more elastic, with an estimated own-bid elasticity that is larger in absolute magnitude by 2.74, at -7.70. These estimated parameters imply differences in markups over workers' opportunity costs. The estimated average markup over cost included in inexperienced employers' bids is 25.2 percent, compared to 14.9 percent for experienced employers. This reflects the different wage elasticities in each segment of employers.

Following the construction of the equilibrium bids in Section 3.2, the observed bids and estimated markups allow for an estimate of worker costs. These estimated costs also differ by employer experience. Referring again to Columns 7 and 8 in Table 4, the mean cost of working for an inexperienced employer is 7.43 USD per hour (before the oDesk fee), compared to 7.35 USD per hour in the experienced sample. This implies that the additional expected hassle cost of applying to an inexperienced employer is limited, at about 0.08 USD per hour. The results suggest that both demand- and supply-side effects play a role in the new employer bid premium, although of differing relative importance. In the preferred estimates in Columns 7 and 8, around 88% of the premium can be attributed to the higher markups set by workers who anticipate that new employers have relatively inelastic hiring probabilities, and only 12% is due to the higher expected costs when applying to new employers.

3.5.2 Employer Type Heterogeneity and Sources of Gains from Experience

The results in Panel A of Table 4 also display the estimated heterogeneity in employer types for never-experienced and ever-experienced employers. Employer types and their corresponding probabilities for the never-experienced group are displayed under columns for inexperienced employers (Columns 5 and 7), while the ever-experienced are in columns for experienced employers (Columns 6 and 8). To help visualize employer heterogeneity, as well as selection out of the market, Figure 6 plots a histogram of employers' types based on whether the employer is in the ever-experienced group or the group that is never observed

[27] The semi-elasticity in the conditional logit model is $(1 - p_{i\chi j})\,\alpha_\chi$. In models with heterogeneity across employers, the mean semi- elasticity is $\Sigma_k \rho_k^S\,(1 - p_{i\chi j}(\mu_k))\,\alpha_\chi$.

in the experienced sample. As is clear, the distribution of employer valuations varies across these two segments, and there is a clear pattern of selection out of the market of type-2 employers, those with low values of μ_i. 34% of experienced employers are type-3, the high μ_i type, compared to only 9% of those that never return as experienced employers.

Figure 6: Employer Types by Ever-Experienced or Never-Experienced Group

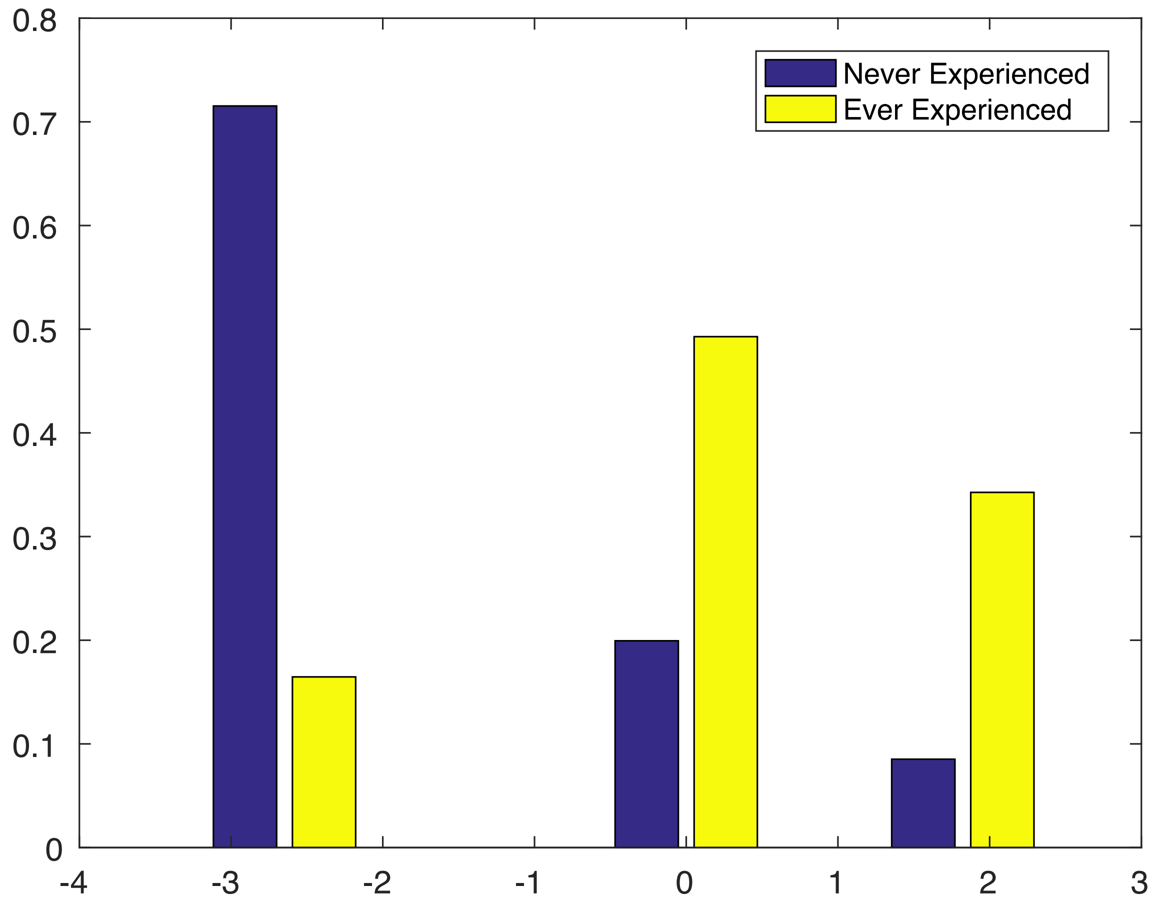

Type probabilities for employers based on whether they are ever observed in the experienced segment. Employers who are inexperienced and eventually transition to the experienced segment are always classified using the "Ever Experienced" type probabilities.

The results in Table 4 also provide estimates of how employer value for the platform changes. We term $X\beta_\chi + E(\mu|\chi)$ the log productive value of hiring and decompose its difference between the inexperienced and experienced into the difference due to changes in X, due to changes in β, and due to changes in μ. Changes in X reflect sorting of workers into jobs or changes in the characteristics of posted jobs. Changes in β suggest changes in employer perceptions of the platform or changing weights on applicant characteristics, and changes in μ suggest employer selection out of the market.

The mean difference in log productive values across segments, while easy to calculate, does not have a natural interpretation on its own. Therefore, we translate the difference into units of log wages. To do this, we define $\tilde{W}$ as the equivalent log wage in the inexperienced segment that would offset the increase in productive value achieved by the experienced segment. This solves the following equation, in expectation:

$$E(X_E\beta_E + \mu_E - \alpha_E\overline{log\,(w_{ijE})}) = E(X_I\beta_I + \mu_I - \alpha_I\tilde{W}). \tag{10}$$

The bottom panel of Table 4 displays the offsetting log-wage bid equivalent reduction to inexperienced employers that would equate hiring probabilities across segments. The results using our preferred specification with heterogeneity across employers suggest that the total change in employer value for the platform with experience is equivalent to reducing bids made to inexperienced employers by about 30 log points, or about \$2.80 per hour.

Using the logic of the Oaxaca-Blinder decomposition (Oaxaca, 1973; Blinder, 1973; Fortin et al., 2011), the difference in log productive values due to differences in X is $\left(\bar{X}_E - \bar{X}_I\right)\beta_I$, which is the difference in characteristics for the experienced and inexperienced segments weighted by the inexperienced parameter estimates. Changes in β provide an estimate of the effect of learning-by-hiring, and are given by $\bar{X}_E\left(\beta_E - \beta_I\right)$. Finally, the difference due to employer composition, changes in μ, calculates the mean of the distribution of μ for the inexperienced and experienced segments (as summarized in Figure 6).[28]

In the models that include worker characteristics, the majority of the difference in employer value for the market across experience levels is due to changes in the estimated β coefficients—that is, the learning-by-hiring effect. Because the worker sorting effect is negative, at around −11% of the log productive value difference, changes in β and in μ explain more than 100% of the difference in the productive value of the market for experienced and inexperienced employers. Overall, 79% of the total increase in log productive value, or a wage shift of about 24%, is due to learning-by-hiring, while 32% of the change in log productive value is due to the employer composition effect, as those employers with low values of μ leave the market. The employer composition change is equivalent to a wage shift of about 9%.

3.5.3 Workers' Sorting and the Instrument

We now assess the sensitivity of the identifying assumptions behind these results. We explore possible worker sorting on the instrument as it relates to *differences* in the estimates between inexperienced and experienced employers. To check whether worker-sorting concerns or omitted variables are driving the results, we compare the estimates in Columns 5 and 6 of Table 5, the estimates without worker resume characteristics as controls, to the estimates given in Table 7 and 8, with these controls. Both specifications suggest a similar shift in employers' log productive value for the platform with experience, but the relative importance of learning-by-hiring versus employer composition is reversed. This is intuitive, as the matrix X includes worker characteristics in Columns 7 and 8, and the variance in outcomes that is correlated with

[28]For experienced employers, this is straightforward and simply uses the type probabilities for the ever-experienced group of employers. For the inexperienced openings, a weighted average of type probabilities is used that corresponds to the fraction of inexperienced employers who eventually become experienced.

these characteristics is attributed to changes in β and μ in Columns 5 and 6.

More relevant to our understanding of employer demand is the impact of the omission of workers' resume characteristics for the elasticity estimates produced. Omitting these characteristics results in a smaller estimated markup difference between inexperienced and experienced employers. Including worker characteristics in the model, the estimated markup difference across segments is 10.3% (25.2% - 14.9%). Without worker characteristics, the estimated markup difference across segments is 4.3%. This comparison suggests that including worker characteristics that are unobserved to us would lead to estimated markup differences that are even larger than those in Columns 7 and 8. Unfortunately, we cannot estimate a model that includes worker fixed effects in the first stage due to the incidental parameters problem.[29] If sorting into applying on worker fixed effects goes in the same direction as the sorting based on observable worker characteristics, for which we can control, our estimate of the difference in markups to inexperienced and experienced employers may be downward biased.

4 Intuition About Expected Willingness To Pay in Experience Markets

This section provides intuition for how employer experience affects willingness to pay for individual workers. Market experience reduces employer uncertainty, which, in turn, affects employer willingness to pay for each worker relative to other applicants. We use the context of employer hiring due to our empirical setting, but the logic applies to situations in which buyers navigate unfamiliar markets for differentiated products or services.

4.1 The Setup

Employers post jobs and evaluate applicants over multiple periods. Employers are heterogeneous in their average fit with using the services provided by the applicants in this market. To reflect this heterogeneity, we use a notion similar to that used in the formal estimation of demand set out in Section 3.1.[30] We specify that the value of the average worker to employer i is an employer-specific parameter μ_i. Workers are also differentiated, and an employer's goal is to choose the worker who presents the best match out of the set of applicants.

To make this concrete, assume that employer i observes a (non-strategic) quality signal from each of J applicants, denoted q_{ij}. These signals are

$$q_{ij} = \mu_i + \eta_{ij} + \xi_{ij}.$$

[29]Many of the workers observed in the sample make only a single bid. Small numbers of bids mean that sampling error in the estimates of the worker effects make them inconsistent. Including these inconsistently estimated effects in a non-linear transformation would bias the estimates of other parameters.

[30]The exposition here uses a more restrictive setup for heterogeneity across employers than that which is estimated earlier; the setting is purely to provide intuition.

Each quality signal is centered on μ_i, the value of the average worker to the employer. The signal also includes a match productivity component with the employer, denoted η_{ij}, for the employer-worker pair. Noise in the signal, ξ_{ij}, means that the employer cannot perfectly distinguish between these components. The random variables η_{ij} and ξ_{ij} are both assumed to be normally distributed with mean zero and variances σ_η^2 and σ_ξ^2. All market participants know the parameters.

A new employer knows that the distribution of μ_i is normal, with mean and variance μ_{i0} and σ_μ^2, but he does not know his own μ_i draw from this distribution.[31] We explore how the employer evaluates his own willingness to pay for each of the job applicants at varying levels of uncertainty about μ_i.[32]

4.2 Evaluating Applicants and Willingness to Pay

When evaluating his willingness to pay for each applicant, employer i attempts to distinguish the productive component of each signal, $(\mu_i + \eta_{ij})$, from the noise term, ξ_{ij}. An employer with the prior $\mu_i \sim N\left(\mu_{i0}, \sigma_{\mu 0}^2\right)$ estimates that the productive component in signal q_{ij} is:

$$\hat{q}_{ij} = \frac{\left(\sigma_\eta^2 + \sigma_{\mu 0}^2\right) q_{ij} + \mu_{i0}\sigma_\xi^2}{\sigma_\eta^2 + \sigma_{\mu 0}^2 + \sigma_\xi^2}. \tag{11}$$

The gap between the estimated productive component in the highest and second highest signals determines the employer's willingness to pay for the most-preferred candidate. Note that the prior beliefs about μ_i do not affect the ordering of applicants, so the identity of the most preferred applicant among the set of workers is independent of μ_{i0} and σ_μ^2. There are, hence, three alternative cases: 1) the employer's assessment of the two highest ranked applicants exceeds his outside option of hiring no worker; 2) only his most preferred applicant is assessed as exceeding his outside option; 3) the outside option is more valuable than his assessment of any applicant.

Comparative statics related to equation (11) for different priors about the distribution of μ_i generate two propositions:

Proposition 1 *Consider two employers facing the same set of applicants with the same prior belief about the mean value of match quality, μ_{i0}. The employer with the greater $\sigma_{\mu 0}^2$ has a weakly higher willingness to pay for the most preferred applicant whenever $\mu_{i0} \leq 0$, that is, whenever the employer believes that the average applicant is less preferred than his alternative source of labor. When $\mu_{i0} > 0$, the employer with the greater $\sigma_{\mu 0}^2$ will more frequently have a higher willingness to pay for the most preferred applicant.*

Proof. Take the two applicants who have sent the highest signals, labeled q_{i1} and q_{i2}, with $q_{i1} > q_{i2}$,

[31] This normal-normal model provides a convenient way to parameterize uncertainty and to describe how an employer weights the signal and prior to evaluate the expected quality of each applicant (Farber and Gibbons, 1996; Altonji and Pierret, 2001; Kahn and Lange, 2014).

[32] We also abstract away from explicit consideration of different numbers of applicants, focusing instead on the surplus available for a fixed number of applicants. In the previous implementation, we account for the number of applicants.

and call them applicant 1 and 2, respectively. In case (1), both applicants exceed the no-hire option, with $\hat{q_{i1}} > \hat{q_{i2}} > 0$. The employer's estimate of the relative productivity difference between applicants 1 and 2 is:

$$\hat{q}_{i1} - \hat{q}_{i2} = \frac{\left(\sigma_\eta^2 + \sigma_{\mu 0}^2\right)(q_1 - q_2)}{\sigma_\eta^2 + \sigma_{\mu 0}^2 + \sigma_\xi^2}. \tag{12}$$

This difference is increasing with $\sigma_{\mu 0}^2$. In other words, as employers become more certain about their own average value for workers in the market, μ_i, the two highest-ranked applicants appear to be closer substitutes. To see this, differentiate equation (12) with respect to σ_μ^2, which, because $q_1 > q_2$, gives

$$\frac{\partial(\hat{q}_{i1} - \hat{q}_{i2})}{\partial \sigma_{\mu 0}^2} = \frac{(q_1 - q_2)\,\sigma_\xi^2}{(\sigma_\eta^2 + \sigma_{\mu 0}^2 + \sigma_\xi^2)^2} > 0. \tag{13}$$

In case (2), the employer's estimate of his willingness to pay for his most preferred applicant is the only estimate that exceeds his value for the outside option, so $\hat{q_1} > 0 > \hat{q_2}$. In this case, the outside option (the alternative source of labor, normalized to zero) pins down willingness to pay and Equation (11) provides the expected value relative to the outside option. The change with respect to uncertainty is

$$\frac{\partial \hat{q}_{i1}}{\partial \sigma_{\mu 0}^2} = \frac{(q_{i1} - \mu_{i0})\sigma_\xi^2}{(\sigma_\eta^2 + \sigma_{\mu 0}^2 + \sigma_\xi^2)^2}.$$

When $\mu_{i0} \leq 0$, $\frac{\partial \hat{q}_{i1}}{\partial \sigma_{\mu 0}^2} > 0$ because $q_{i1} > 0 > \mu_{i0}$.

When $\mu_{i0} > 0$ the sign depends on $(q_{i1} - \mu_{i0})$. In this sub-case, it is possible that willingness to pay is increasing or decreasing with $\sigma_{\mu 0}^2$. All that can be said is that willingness to pay is more frequently increasing in $\sigma_{\mu 0}^2$. To see this, note that for employers with rational expectations, such that $E(\mu_i) = \mu_{i0}$, then $E(q_{ij}) = \mu_{i0}$. Because $\mu_{i0} > 0$, case (1), where $\hat{q_{i1}} > \hat{q_{i2}} > 0$, occurs more frequently than case (2) where $\hat{q_{i1}} > 0 > \hat{q_{i2}}$. Because willingness to pay is unambiguously increasing in case (1), in the parameter range where $\mu_{i0} > 0$, employers are more likely to be in case (1) where willingness to pay for the most preferred applicant is unambiguously increasing with uncertainty. Case (3) is irrelevant, as no applicant is hired. ∎

Proposition 2 *For two employers facing the same set of applicants with the same level of uncertainty $\sigma_{\mu 0}^2$, the employer with higher μ_{i0} will have a weakly higher willingness to pay for the most preferred applicant.*

Proof. Suppose that the employer has two applicants whose expected values exceed his outside option. Because μ_{i0} affects his estimated value for both applicants similarly, its level cancels out of the expression for his relative willingness to pay for the most preferred applicant. In this case (case 1), the level of μ_{i0} is irrelevant. In case (2), however, with the next best option being the outside option, differentiating equation (11) gives $\frac{\partial \hat{q}_{ij}}{\partial \mu_{i0}} > 0$, so that the willingness to pay for the most preferred applicant is increasing in μ_{i0}. Again, case (3) is irrelevant, as no applicant is hired. ∎

4.3 Experience, Uncertainty, and Signals of Willingness to Pay

Propositions 1 and 2 relate an employer's beliefs about his own value for the average worker in the market to his willingness to pay for his most preferred applicant. We propose that an employer has greater certainty about his own μ_i when he has gained market experience—that is, when he has observed many worker quality signals.

First, then, experience in the market reduces $\sigma^2_{\mu 0}$. Each interaction with a worker provides a quality signal that allows an employer to update his beliefs, increasing the precision of his estimate of μ_{i0}. Under Proposition 1, inexperienced (uncertain) employers have larger $\sigma^2_{\mu 0}$ estimates, so that, more frequently, they have a higher relative willingness to pay for their most preferred job applicant.[33]

When taking Proposition 2 to the data, we note that, among experienced employers, there is variation in the feedback score that they have given to, and received from, previously hired workers. Employers who have good feedback are revealed to have had positive past outcomes in the market, and good employer feedback is positively associated with hiring one of the applicants for an opening. This feedback variation possibly captures differences in the level of μ_{i0} among experienced employers, along with reputation, hassle to work for, etc. Under Proposition 2, experienced employers with good feedback scores (higher values of μ_{i0}) have a higher relative willingness to pay for their most preferred job applicant than do experienced employers with low feedback.[34]

The fact that employers with low values of μ_{i0} are likely to select out of the market creates some tension between the results in Propositions 1 and 2. In this case, experienced employers will have a higher average level of μ_{i0}. Proposition 2 suggests that the higher average fit with the market in the experienced segment that comes from selection of employers will make it harder to detect support for Proposition 1.

We look for evidence consistent with Proposition 1 using variation in workers' bids to employers with different experience levels. We look for evidence consistent with Proposition 2 using variation in feedback within the set of experienced employers. The following sections present reduced form analysis consistent with the predictions of both propositions.

[33]This result depends on the distribution of the top two signals being equal across employers of different experience levels. Sorting into job applications and competitive entry has the potential to alter this distribution. As shown in the Appendix, in the oDesk market, an increase in the number of applications does not appear to offset this difference in expected willingness to pay.

[34]If higher levels of feedback were associated with lower expected costs for workers—for example, due to a lower risk of receiving bad feedback on the job or of higher costs of doing business—we would expect a negative correlation between experienced employer feedback and worker bids. This is in direct contrast to the prediction of Proposition 2, which relates to a demand-side effect.

5 Reduced form empirical analysis

5.1 Applicants' Wage Bids Decline with Employer Experience

Table 5 provides regressions analyzing the log of applicants' hourly wage bids. The bids studied are the offers submitted to employers. As detailed in Appendix 2, there is very little change from an applicant's initial bid to the wage paid upon hiring. The coefficients capture the percentage difference between the hourly wage bids received by inexperienced employers and the bids received by employers after hiring one, two, three, four, or five-plus workers.

The first column presents differences in wage bids received by employers of different experience levels after including fixed effects for job categories and calendar time. Employers receive bids that are approximately 3% lower after making one hire. Bids continue to fall after subsequent hires, and employers who have made five or more previous hires receive bids that are 7.6% lower than those received by employers who have not previously hired. Column 2 adds controls for detailed worker resume characteristics, controls for who initiated the application, and a third-order polynomial in the number of characters included in the job description. These controls are intended to allow for worker sorting into making an application and for the possibility that the employer is learning how to write an effective job posting. The observed decline in average wage bids received by employers with five or more prior hires is 5.9%.

Subsequent columns of Table 5 provide results within employer and within worker. This addresses concerns about sorting and selection into application. Columns 3 and 4 add employer fixed effects. A given employer receives bids that are 4.7% lower when he has five or more prior hires than when he has no experience. Adding, again, the controls for worker and job characteristics that were included in Column 2 reduces the effect to 3.7%, indicating some sorting on observable characteristics and employer experience rather than sorting simply on the identity of the employer. Still, a large bid reduction with employer experience remains. Columns 5 and 6 remove employer fixed effects and add worker fixed effects. For a given worker, bids to employers who have previously made five or more hires are, on average, 3.9% lower than bids to employers with no observable experience.

5.1.1 Different Expected Job Outcomes and Different Employer Hiring Practices with Experience

The data allow us to rule out several plausible reasons for the observed within-worker and within-employer decline in wage bids with employer experience. We first consider the role played by actual or anticipated long-term employer-employee relationships. It is possible that workers view a job application to an experienced employer as having better or worse long-term prospects than an application to an inexperienced employer. For this effect to explain higher bids to inexperienced employers, workers must anticipate a greater chance of a long term relationship with an experienced employer and be willing to take a short term wage reduction to secure this relationship. The distribution of contract duration in the data allows us to evaluate whether these worker expectations would be reasonable. Table 6 shows that experienced employers are not more likely than the inexperienced to start a long-term relationship. This table displays regressions in which the unit of analysis is the first time an employer hires a worker, while the outcomes are calculated over the entire duration of the future employer-worker match, including any subsequent contracts. The first two columns regress the log number of hours worked over the first and all future contracts on indicators for employer experience the first time an employer-worker match occurs. Odd-numbered columns are simple OLS regressions, while even-numbered columns include fixed effects for the employers' stated duration of a job. In each of the first two columns, the coefficients on employer experience tend to decline, meaning that inexperienced employers' new relationships tend to last longer than those started by experienced employers.[35]

Contracting patterns are, therefore, inconsistent with workers expecting a tradeoff in which more hours make up for lower hourly wages. Columns 3 and 4 evaluate whether applicants prefer working for experienced employers because they are more likely to provide raises after the contract begins. There is no evidence that this is the case. Finally, a worker might expect that the feedback she receives would differ by employer experience. Columns 5 and 6 of Table 6 reveal a positive correlation between such feedback and employer experience. However, the magnitude of these differences is small. The feedback score given by employers with up to four previous hires is, on average, less than one-half of one percent higher than the score given by employers with no previous hires.[36] The increase in the feedback score is 1.5% for employers with five or more previous hires compared to those with none. Unlike wages, patterns of feedback differences do not display a monotonic relationship with employer experience, suggesting that this factor is unlikely to explain the monotone decline in wage bids with employer experience shown in Table 5. We later show that increases in the feedback score that employers leave for past hires actually raise rather than lower wage bids.

We next explore whether employer experience is correlated with some aspect of the job posting or hiring process that affects workers' bid decisions. Appendix 2 shows that the bid premium to inexperienced employers does not arise because of differences in the number of job applicants.[37] Additionally, the bid premium does not appear to be driven by time-varying unobservables relating to how employers post jobs. A field experiment designed to isolate the effect of observable employer experience from the effects of environmental changes—detailed in Appendix 2 and with results in Appendix Table A4—shows large

[35]Inexperienced employers state in job posts that their jobs are likely to be of shorter duration, which is not surprising if small trial jobs are used to resolve uncertainty. Column 2 conditions on the expected duration of the job as stated in the job post. Within expected-duration categories, the results are stronger.

[36]For employers with one prior hire, this calculation is (0.0243/4.303) and (0.0312/4.285); similar calculations can be done for other experience levels.

[37]The premium remains when controlling for the number of job applicants in the first twenty four hours after job posting, as shown in Appendix Table A3.

wage-bid declines to experienced employers.

5.2 Experienced Employers with Good Feedback Receive Higher Wage Bids

In this subsection, bids are shown to increase with the interaction of employer experience and feedback left for the employer. This is consistent with the experience market comparative statics. These results are also important for understanding whether the risk of receiving poor feedback motivates worker bid differences.[38]

In contrast to predictions about feedback risk, applicants actually make higher bids to experienced employers with good feedback than to experienced employers with poor feedback. Experienced employers who have not received feedback also receive higher bids than employers who have poor past feedback. In line with Section 4, we interpret these results as consistent with poor employer feedback signaling a reduced likelihood to hire that lowers the optimal bid to these employers. Consistent with this interpretation, relative to an experienced employer who has poor feedback, the likelihood of hiring an applicant increases by 6.7 percentage points (about 13%) for experienced employers with good feedback and 3.7 percentage points (roughly 8%) for experienced employers with no feedback.[39]

Table 7, Panel A, in which the log bid is the dependent variable, presents two sets of results on how bids vary with employer experience and employer feedback. First, even controlling for the feedback received from past workers, wage bids decline robustly with employer experience. Only employers with hiring experience can have a feedback score, but many experienced employers have no feedback from past hires. Indicators for the employer having no observable feedback and for the employer having observable feedback of 4.5 or higher are interacted with an indicator for having observable prior hiring experience. Thus, the baseline group in the regression is experienced employers who have feedback scores lower than 4.5. The interaction terms capture deviations from the baseline for experienced employers with good feedback and for those with no feedback. The baseline point estimates in Columns 2 through 6, which present the effect of experience for employers who have bad visible feedback, show that these employers receive bids that are significantly lower than those made to inexperienced employers. The positive coefficients on good feedback and no observable feedback are inconsistent with an adverse selection channel that would predict lower bids for employers with good feedback.

We interpret these results as showing that workers tailor their wage bids to the information revealed

[38]For example, in an experiment on the provision of worker-level information, Pallais (2014) shows that the revelation of public feedback about workers is beneficial to their later careers on this platform. Stanton and Thomas (2016) show that the effect of worker feedback is concentrated among workers for whom employers have the least information. Similar issues may be at play among employers—that is, feedback about how an employer treated previous hires may allow future applicants to tailor their bids based on variation in the expected cost of working for that employer.

[39]These results are precisely estimated and are significant at the 1% level after clustering standard errors by employer. The dependent variable in the regression is an indicator for hiring any applicant. The independent variables are indicators for employer experience, the employers' feedback, job category, expected duration, and time, along with a cubic polynomial for the characters in the opening description.

about employers through their feedback scores because these measures are informative of employer type. In particular, they are likely to be positively correlated with the employer's value for the market. In this case, the findings are consistent with workers playing Bertrand Nash strategies when bidding in the presence of an outside option of not hiring. In a differentiated workers' bidding game, equilibrium bids are increasing in applicants' assessment that the employer is likely to hire any worker rather than not hire.

Table 7, Panel B assesses whether feedback risk may drive these results by allowing heterogeneity in bids based on the feedback an employer left for prior workers. Although it took some effort for applicants to observe this at different times over the platform's evolution, applicants could navigate to view the history of feedback that employers left for workers on prior jobs. Employers who left low feedback (possibly because of miscalibration with the distribution or because of a bad initial match) did not receive higher wages. In fact, employers who left good feedback were the ones who received higher wage offers later, inconsistent with poor feedback risk driving the wage-bid results.

5.3 Employers' Number of Interviews Falls with Experience

If employers are using interviews to both look for the best applicant for the job and to learn about their own value for the market, then the marginal benefit of an interview is higher when employers know relatively little about their value for the market. Therefore, it is optimal for inexperienced employers to interview more applicants. Table 8 presents the results of a regression of (1+ the log of) the number of interviews conducted for each job by the number of previous hires made. The first column does not include employer fixed effects, and 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. This table presents the overall sample, and the results are similar for the sequential sample. Even with different levels of controls, in all specifications with employer fixed effects, the number of interviews decreases, and at a decreasing rate, on successive jobs. The predicted number of interviews falls by 67% after five prior hires, and a large share of that decline happens between the first and second jobs.

5.4 Employer Outcomes and Search Effort are Consistent with the Resolution of Uncertainty

If an employer's search process can be modeled as an optimal stopping problem in which the employer hires the first interviewed applicant whose expected value exceeds a threshold, the relevant stopping threshold will be higher when the employer is also using interviews to learn about the market. This is because the more information that an interview conveys about the market, the greater the benefit of a marginal interview. An implication of using interviews to learn is that, as an employer interviews more candidates and gains a more precise assessment of the market, the threshold stopping value for hiring falls because the marginal learning value declines (Kohn and Shavell, 1974). A further implication is that a new employer who hires after conducting a small number of interviews must have found an applicant with a very high expected value early on in his search process. An employer who interviews many applicants before hiring will likely end up hiring an applicant with a lower expected value to him. These findings are borne out in Appendix Table A6.

5.5 Omitted Employer Characteristics

The results so far are consistent with the hypothesis that there is an employer-level market value term that employers learn about during early experiences. This section assesses whether some other employer-level omitted variable determines these outcomes. Because we observe the order of interview requests in the data, it is possible to examine whether inexperienced employers' early actions on a job posting are correlated with later choices and outcomes. If this were the case, some of the results could be forecast by employer characteristics or actions that determine their search processes rather than by employers' experience and learning in the market. If ex-ante differences were driving, say, the number of interviews conducted or whether the employer eventually hires for the opening, one might expect these differences to be reflected in the characteristics of the applicants whom the employer interviews.

Building on the analysis in Table 9, which controls for early actions in looking at job outcomes, Figure 7, Panel A, examines the distribution of the hourly wage bids of the worker selected for the first interview, based on the employer's eventual action. The split in Panel A is based on whether the employer does more or fewer than five total interviews. 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 by employers who go on to interview few or many applicants.

Panel B of Figure 7 repeats this exercise by whether the employer hires for the first job and again shows that early interviewing choices are unrelated to whether or not an employer makes a hire. These comparisons cast doubt on the hypothesis that inexperienced employers' eventual differences in interviewing or hiring behavior are related to unobservable differences in information, preference differences at the time of the initial job posting, or preference differences at the time of selecting the first interview candidate.

Even more important, the overlap in wage bids on the initial interview request suggests that workers are not able to segment inexperienced employers based on the eventual number of interviews that they conduct or by the probability that they will hire. These results are consistent with the hypothesis that actions are shaped by what is learned rather than by another employer type variable that the employer knows before entering the market.

Figure 7: Residual Hourly Bids by First Applicant Selected for Interview

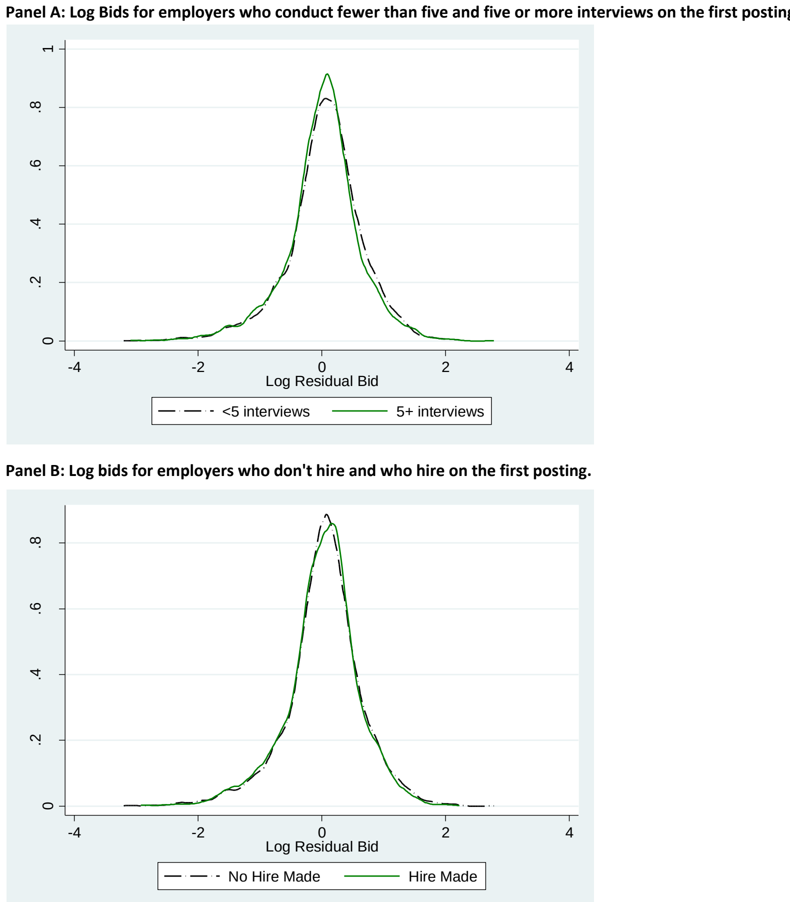

6 Counterfactual Analysis

In experience markets, buyers learn about their own value for the market from early interactions with sellers. Whether they continue in the market depends on what they infer from the signals they receive. For some buyers, these signals will lead to false negatives, suggesting a role for the platform to encourage further experimentation. Focusing on the employers who have not previously hired, Table 1 tells us that 78% of this group post a job opening and receive applications, but leave the market without hiring. That is, a large share of employers receive signals about their own market value from applications that lead them to negatively update their prior about how valuable they will find the market.

We investigate whether, from the platform's perspective, false negatives for employers result in too much exit. To do so, we assess whether a different fee structure would alter the mix of employers hiring in the market and, hence, platform profits. Since its founding through the end of the sample period, oDesk's fee was constant, at 10% of wages.[40] In the absence of variation in fees altering bids for all workers, we use

[40]After the sample period ended, the platform raised baseline fees and implemented quantity discounts, but we do not have

a model-based assessment to analyze how changes in platform fees would affect transitions of inexperienced employers to the experienced segment.

Appendix 3 provides intuition about the exercise by analyzing the problem using specific (fixed) fees. The important intuition from the fixed-fee case is simple and translates to thinking about ad-valorem fees (but with substantial additional algebra): a fee has revenue implications but also governs selection into becoming experienced. Optimal fees on inexperienced employers balance revenue today with the reduced probability of hiring and, hence, providing future expected revenue after transitioning into the experienced segment.

To simulate how employer hiring evolves under different fees and how the associated market revenues and platform profits change, we need to make an additional assumption about how many jobs an experienced employer goes on to post after becoming informed about his own value for the market through early experience.[41] Assuming a large number of subsequent job postings presents an extreme case in which it is very valuable for the platform to induce an employer to stay in the market. We choose this extreme case for illustrative purposes and find that offering lower fees to the inexperienced employer becomes optimal for the platform only when the employer posts eight or more later jobs. If the employer were to remain in the market for fewer than eight more jobs, the optimal fee would not be lower for inexperienced employers. Still, even if we assume that an experienced employer posts eight successive jobs, we find that the optimal fee on inexperienced employers is higher than the existing fee of 10%, and the optimal fee on experienced employers is even higher.

Table 9 presents platform profits and transactions volumes at different fee levels for the case in which an experienced employer's long-term value is high because he posts eight jobs once he is relatively informed about his own value for the market. The simulations make use of the hiring probability parameters estimated in Section 3. The estimated distribution of the employer-specific heterogeneity characterizes selection into continuation in the market given the estimated parameters. We simulate employer hiring using a grid of different fees—the fees to the inexperienced vary across rows, and the fees to the experienced vary across columns. Appendix 3 details the steps used in this simulation.

Panel A of Table 9 analyzes the change in platform profits relative to the current fee structure, while Panel B provides percentage changes in the number of employers transitioning to the experienced segment relative to the baseline 10% fee. The main result is that the optimal fee for both segments is higher than the 10% fee charged by oDesk at the time of the data. Again, because employers in the data typically

data from this period.

[41] 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. In the simulated calculations, the estimated parameters governing hiring and transition rates for each type are used. Different types may be affected by changes in fees, and transition rates are calibrated from the empirical rate of transitioning to the experienced segment for each type.

post fewer than eight jobs, the assumption that each posts eight future jobs leads to an overestimate of profits that may be had from inducing more employers to gain experience. Nonetheless, platform profits are maximized when a smaller share of employers than under the current fee structure are induced to return. With eight subsequent hires assumed for employers who elect to remain in the market, platform profits are maximized with a 15% fee on the inexperienced segment and a 20% fee on the experienced segment of employers. The profits at this fee structure are estimated to increase by about 21%, and the number of returning employers is estimated to fall by 22%. That is, platform size would be reduced, but per-transaction profitability would increase enough to more than offset the reduction in transactions volume.[42] With fewer than eight future jobs, optimal fees are approximately equal across experience levels, at 20%.

7 Conclusion

We characterize online labor platforms as experience markets in which differentiated sellers vary in their fit with an individual buyer's needs. Buyers who are new to the market are uncertain about their own value for what sellers offer. The experience market concept is not limited to digital platforms, and many new markets are likely to share its defining features. The characterization emphasizes changes in how buyers assess products over time, with implications for demand differences between inexperienced and experienced buyers. In this online labor market, the buyers are employers and the sellers are job applicants offering to work.

We estimate demand for job applicants flexibly, allowing it to vary based on whether the employer has hiring experience in the market. Instrumenting for wages with exchange rate fluctuations and competition among workers for jobs, we estimate significant changes in demand between employer segments. Employers appear to use interactions with job applicants to learn how the workers in the market fit with their own needs, while new employers' demand for labor reflects uncertainty about their own valuations. This demand difference also explains why new employers receive higher wage bids: it is an equilibrium response to their relative inelasticity of demand. There is limited evidence that workers submit higher wage bids to these employers because they are less desirable employers who would impose higher costs on job applicants.

In this market, the majority of workers looking for jobs and the majority of employers looking to hire never participate in a market transaction, but this paper shows that the reason for the low participation rate is very different for employers and for workers. Inexperienced workers are penalized because their quality is hidden from potential employers, and the literature has established that hiring rates for inexperienced

[42] One limitation of this analysis is that it does not account for other platforms' competitive response or entry of competing marketplaces. These considerations may reduce optimal fees. In addition, the analysis abstracts away from tailored offers that deviate from this fee structure. Here, wage offers to employers are assumed to be based on only observed historical use of the platform, but additional segmentation or non-linear schemes may make it possible for the platform to price discriminate on other observable employer attributes.

employees are inefficiently low (Pallais, 2014; Stanton and Thomas, 2016).[43] In contrast, we show that information about new employers that is hidden prior to experience encourages their market participation to learn about their type. Low hiring rates by new employers reflect the decisions that employers make after some information has been revealed.

The estimated differences in demand by employer experience segment have implications for understanding platform policies to attract new employers. In a counterfactual exercise, using the estimated hiring demand functions and hiring rates among employers with different experience levels, we ask whether the platform should use different fees to entice new employers to hire. Market revenues and, hence, platform profits would be maximized at higher fees than those that are currently used. This finding requires an understanding of how the elasticity of demand varies as employers gain experience, and market counterfactuals that do not allow for demand changes with experience would conclude that early hiring decisions were more responsive to wage subsidies. Our analysis shows that early subsidies would increase new employer entry but would fail to recover their cost.

This platform provides employers with the opportunity to engage in offshoring labor services, with the potential to realize large hourly wage differences across country borders. Other gains from collaboration might also be substantial (see Agrawal and Goldfarb (2008)). Our analysis shows that the majority of potential employers discover that the market is far less valuable to them than these wage differences would suggest.[44] Heterogeneity in the ability of potential employers to offshore appears to suggest that some firms face broader organizational costs when fragmenting production across country borders.

[43] The focus on workers is mirrored in results from other contexts, like crowdfunding (Agrawal, Catalini, and Goldfarb, 2011).

[44] These findings also suggest that technical offshorability studies at the job or task level may overstate the feasibility of offshoring by not accounting for the associated employer-level costs. Jensen and Kletzer (2010) and Blinder and Krueger (2013) suggest that many jobs are able to be offshored, but the pace of actual offshoring does not match these projections. In the 2012 Survey of Business Owners (SBO) conducted by the US Census Bureau, only 1.36% of firms responded that they offshore services or functions abroad.

Appendix

Appendix 1: Data Details and Cleaning

Data cleaning details

Appendix Table 1 gives details about the resume data used in the full sample and in the sample used for hiring probability estimation. The following restrictions are used to clean job openings for the purposes of estimating hiring probabilities. First, to be able to characterize the set of applicants for individual jobs, the sample is restricted to openings that have at least one day of elapsed time between the current job posting and the next job posting and at least one day of elapsed time between the previous posting and the current posting. This allows for at least a single cycle of applicants from different time zones to apply to the different jobs, eliminating batched hiring for which available applicants may blend across jobs. The sample for estimating hiring probabilities also drops jobs for which the employer hires a worker from a previous engagement. Many jobs also appear to originate from bringing an offline relationship onto the platform. Filtering these jobs requires that at least one application be worker-initiated, and the total number of candidates must be greater than five. The restriction to five total applicants eliminates most obvious cases in which an employer posts a job as publicly visible but with the intention of hiring a pre-selected candidate or set of candidates. Any job from an employer who sends over 100 interview requests on the first job or who sends over 60 interview requests on a subsequent posting is omitted. These postings are likely to be fake jobs posted by spammers. Finally, any job that is later declared to have been posted by mistake is dropped. The following restrictions are used to clean applications: first, applications from invited workers who later report they are unavailable are dropped; second, applications are dropped if the employer reports obvious spam.

Bids of hired workers relative to other applicants

The solid bars of the histogram in Figure 8 present the bid decile among all applicants of the worker who was hired when applying to jobs posted by employers with no prior hiring experience. Around 13% of all employers hire a worker whose bid is in the top decile of the distribution of bids for the job. Less than 20% of inexperienced employers who hire choose a worker whose bid is in the lowest decile of the job-specific bid distribution. Experienced employers, shown in the histogram with the outlined bars, are somewhat more likely to hire workers in the lowest wage bid decile.

Figure 8: Bid Decile of Hired Worker

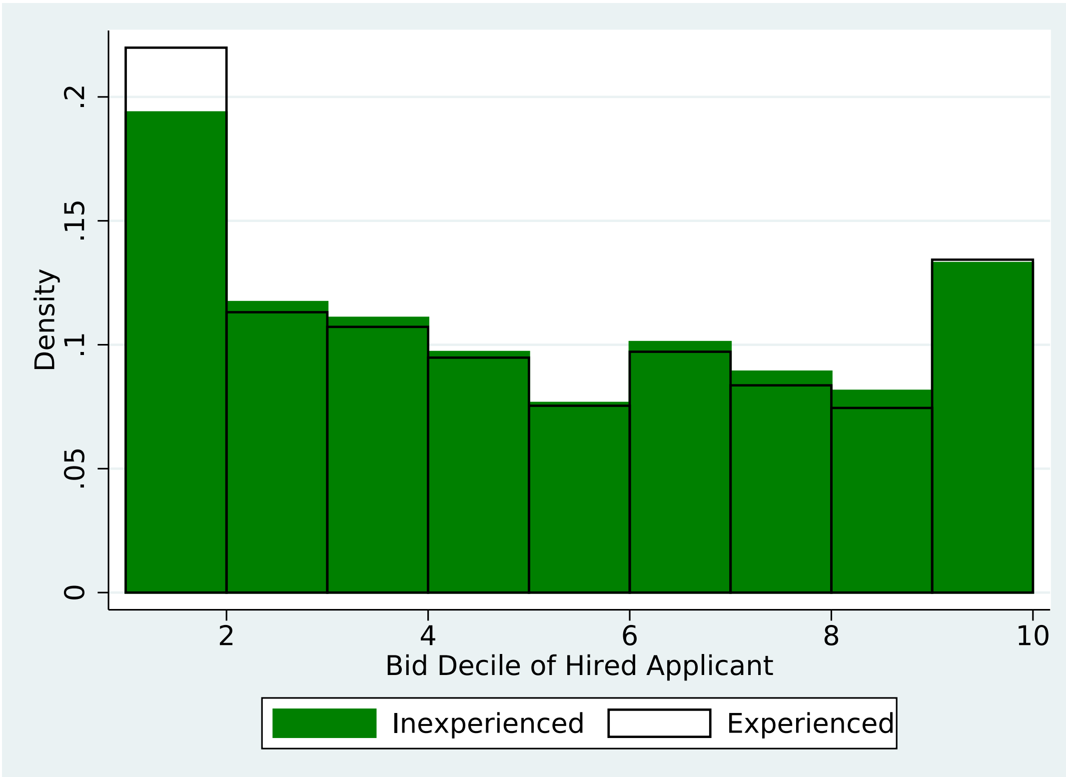

The figure shows the bid decile of the worker who is hired from the set of applicants. For each job opening, we find the decile of wage bid for each applicant. We then take the decile of the applicant who was hired and plot the histogram of wage bid deciles for applicants selected by inexperienced employers (solid bars) and by experienced employers (outlined bars).

Appendix 2: Alternative Explanations for Declining Wage Bids

Different Application Rates

It is possible that the extent of competitionfor a job posting changes with employer experience, and workers might submit lower wage bids when they anticipate a more competitive market. For variation in anticipated market competitiveness to explain the bid premium to inexperienced employers, workers must anticipate that the job postings by experienced employers are more competitive. Table 1 (Columns 2 and 7) showed that inexperienced employers in the sequential sample receive a smaller number of applicants in total, suggesting that, on average, competition might indeed be greater for employers' later jobs.

To examine this possibility, Table A3 repeats the analysis from Table 5, but the estimations include the log arrival rate of applicants within the first 24 hours of posting the job as an additional control. Note that the regressions already include a spline in the application number, and bidders can observe the number of prior applicants when making their bid. This additional regressor removes the effect of expected future competition on bids. The faster the rate, the lower are all bids received by the employer. However, including this control does not change the main finding from Table 5 that experienced employers receive significantly lower bids.

Results from Experimental Job Postings

To guard against time-varying unobservables at the job opening level that might change workers' bids, we ran a small field experiment to isolate the effect of employer experience alone.[45] We posted identical jobs from the accounts of employers with different levels of experience. Employer 1 had no experience, while employer 2 had prior hiring experience and a good feedback score. Each employer posted a short, identically worded job description in the "Data Entry" job sub-category. The task description read "I need you to take data from a website and put it into excel." No additional detail was provided.

Two dependent variables are of interest in regressions using the experimental data to estimate the causal effect of experience on bids. The first is the actual log bid submitted to the job. The second is the difference in the log bid and the log hourly rate posted in the worker's profile. This latter measure helps to pick up unobserved heterogeneity about workers who may sort to jobs. These measures are regressed on an indicator that the job was posted by the experienced employer. Some specifications also control for the number of hours the applicant has previously worked on the platform or application order fixed effects.

Table A4 contains the results. In each specification, there is a significant, negative point estimate on the experienced employer indicator variable. The effect sizes are larger in magnitude than those estimated in Table 5. In addition, pairwise comparisons of Columns 1-3 and Columns 4-6 indicate very similar point estimates when the dependent variable is the log of the actual bid versus the difference in bid from the profile rate. This exercises isolates the effect of observable experience in driving lower bids.

Wage Bargaining

Our wage bid data contain the final offer to an employer. When an applicant is not hired, the final and first offers coincide: when she is, the first and final offers may differ. It is possible that applicants' first bids vary with employer experience because they expect employers to have become more or less skilled negotiators.

Early on in our previous project (Stanton and Thomas, 2016), we investigated the extent to which offers changed between first and final offers; unfortunately, those queries were not pulled down from the company's servers. Those early queries found, however, very limited bargaining from initial offer to final wage for either employer segment. To construct an analysis of the extent to which bargaining may affect our results, we use the insight that rejected applications will have the same initial and final wage offer. Then, for each worker who is hired, we take the last rejected wage offer in the same job category to employers of

[45] Unfortunately, we are unable to conduct this experiment on a larger scale because of the Upwork terms of use relating to creating employer accounts.

the same experience level. We then compare the final wage offer on hires to the last wage offer on rejected applications.

Using these measures to assess the extent of bargaining is imperfect, as they likely overstate the extent to which bargaining occurs. This is because i) wages tend to decline over time if an applicant hasn't landed a job and continues to apply; and ii) the parties may set up side payments off the platform to avoid the platform fee, especially if they have prior experience working together. With declining wages over time, accepted wages on hires that are below the last rejected wage bid will inflate the extent to which employers bargain over wages. However, under the assumption that this measurement problem does not differ by employer experience, regressions of the difference in the final log wage when hired and the last log wage bid on a rejected application are suggestive of whether expected differences in bargaining may change the interpretation of the results. The second problem, of payments off-platform, are likely to be small due to difficulties in transferring funds across banking systems. However, to the extent that they do exist, we expect they will show up for experienced employers who have reputations.

Table A5 contains the results. In OLS regressions, there may be some small reductions in wages due to the appearance of bargaining for employers with five or more hires. The point estimate is a reduction of about 2%. These differences do not, however, reflect the immediate reduction in bids with experience seen in other tables. Bargaining differences are not significant with the addition of employer fixed effects, meaning this channel is unlikely to drive the pattern of results documented elsewhere.

Employer Outcomes, Search Effort, and the Resolution of Uncertainty

We consider how employer search effort is linked to outcomes in a model in which the employer is searching for a good match while trying to learn something about the distribution of applicants. Under the standard model in which the employer is not using interviews to learn about the distribution of worker value, the threshold value to hire a worker remains constant with the number of interviews. In this case, the expected stopping value of the hired applicant is independent of the number of interviews conducted before hiring. Under the hypothesized learning process, according to Kohn and Shavell (1974), an inexperienced employer who uses interviews to learn about the distribution of matches will have worse expected outcomes as the number of interviews increases. An employer who hires after a small number of interviews is more likely to hire a worker of higher value to him, since this worker exceeds a higher reservation value.

Table A6 shows that inexperienced 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. For example, Columns 4 and 5 show that inexperienced employers who hire after one interview are 6% more likely than inexperienced employers who conduct 11 or more interviews to report success and give good feedback after controlling for the hourly wage paid to the hired worker. These results are consistent with the hypothesis that employers who hired after one interview, thus forgoing the learning value of additional interviews, must have been lucky in finding an interviewee with a high expected value (and a high actual value) to them. Columns 6 through 9 then consider whether these results transfer over to outcomes on the second job. These columns consider whether employers who stop interviewing earlier on the first job have better future outcomes later on. These regressions are, thus, second job outcomes as a function of interviews on the first job post. The results are consistent with the model, suggesting that those who stop interviewing more quickly find better average matches across job posts.

Appendix 3: Platform Profits

Fixed Fee Differences by Employer Experience

The platform's objective is to maximize total profits, which is equivalent to maximizing the total value of transactions in the market, and it can do this by setting different fees for inexperienced and experienced employers. To denote specific fees, we call the fee to inexperienced employers t_I and the fee to experienced employers t_E. Let H_I be an indicator for an employer hiring while inexperienced and H_E be an indicator for hiring while experienced. Wages for the inexperienced and experienced segment are w_I and w_E, respectively.[46] The platform's problem is

$$\max_{t_I, t_E} \Pr\left(H_I|w_I\right) \times \left[t_I + t_E \times \Pr\left(H_E|H_I\left(w_I\right), w_E\right)\right],$$

where $\Pr\left(H_I|w_I\right)$ is the probability that an inexperienced employer will hire given wages w_I, and $\Pr\left(H_E|H_I\left(w_I\right), w_E\right)$ is the probability an experienced employer will hire as a function of wages w_E conditional on the first hire, $H_I\left(w_I\right)$. Notice that the platform does not set wages, only fees, but wages that employers face will vary with platform fees because they are passed through.

Adding uncertainty and selection makes the fee-setting problem more interesting. When employers are uncertain about platform valuation and some uncertainty is resolved through hiring, experienced employer hiring probabilties, $\Pr\left(H_E|H_I\left(w_I\right), w_E\right)$, may depend on the evolution of employers' beliefs about the platform as a result of hiring. That $\Pr\left(H_E|H_I\left(w_I\right), w_E\right)$ specifically conditions on $H_I\left(w_I\right)$ and the wage paid 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.

Thus far, this formulation says nothing about how beliefs evolve with employer experience. This leaves the learning process free, allowing for models with myopic or anticipated learning.

Using H_I as shorthand for $\Pr\left(H_I|w_I\right)$ and H_E as shorthand for $\Pr\left(H_E|H_I\left(w_I\right), w_E\right)$, the first order

[46] In this setup, we assume that employers have the opportunity to hire in the experienced segment only after they have hired while inexperienced.

conditions for the optimal fee levels 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:

$$t_E^* = -\frac{H_E}{\frac{\partial H_E}{\partial w_E}\frac{\partial w_E}{\partial t_E}}. \tag{14}$$

The fee for the inexperienced is:

$$t_I^* = -\frac{H_I}{\frac{\partial H_I}{dw_I}\frac{\partial w_I}{\partial t_I}} - t_E^* H_E - t_E^* H_I \frac{\partial H_E}{\partial H_I}. \tag{15}$$

The first term in t_I^* is the standard static markup for the segment of inexperienced employers. However, this markup is reduced by the latter two terms. The second term includes the future value of fees for those hiring in the experienced segment, adjusting t_I downward to account for the spillover to future demand. The final term, which accounts for composition effects, is of particular interest. The expression $\frac{\partial H_E}{\partial H_I}$ incorporates how the marginal employer induced to hire by the fee set for the inexperienced will change the likelihood of future hiring.

Algorithm for Simulating Platform Fees

The following steps are used in the simulation. First, inexperienced employers are assigned draws from the types according to the population fraction of types in the inexperienced segment.[47] These types are assigned independently from the applicant set. Then, for each ad-valorem fee pair, (τ_I, τ_E), simulated profits are constructed according to the following procedure: 1) Log wage bids to inexperienced employers are calculated, where pass-through of the fee is computed according to the worker's first order condition for setting bids. 2) Inexperienced employers receive a random uniform draw and choose whether or not to hire based on the computed choice probability. 3) For those inexperienced employers who hire, we iterate the following steps until convergence: a) a candidate set of employers posts additional jobs; b) given the openings posted, elasticities are calculated, and log wage bids including fees and workers' markups are determined; c) given wage bids, the expected surplus from posting additional jobs is computed; d) employers' choices to transition are updated. Employers rationally choose to post additional jobs if the expected surplus from an opening, accounting for wage bids and fees, exceeds the value of exiting the

[47] These types are a weighted-average of the "ever experienced" and "never experienced" segments.

market. This is modeled as a probabilistic function of the expected surplus and is calibrated from the transition probability between the inexperienced and experienced sample;[48] e) if the set of employers is stable, the loop is terminated and, if not, we return to step (a). The loop in 3) involves re-calculating markups and log wage bids conditional on the set of employers that transition into posting successive jobs. 4) Profits are then based on hiring probabilities from the model and the fee-rate associated with the chosen bid.

[48] The parameters are calibrated to minimize the distance between the predicted transition probability to the experienced segment given hiring and the actual transition rate. The transition rate is allowed to depend on a constant and the expected surplus in the experienced segment. The expected surplus is a function of the employer's type and the expected value of hiring a worker, computed from the well-known surplus formula for the conditional logit.

References

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

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

[3] Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons. 2016. ”Does standardized information in online markets disproportionately benefit job applicants from less developed countries?” textitJournal of international Economics, 103, 1-12.

[4] Agrawal, Ajay, John Horton, Nicola Lacetera, and Elizabeth Lyons. 2013. ”Digitization and the Contract Labor Market: A Research Agenda.” in “The Economic Analysis of the Digital Economy’ edited by Avi Goldfarb, Shane Greenstein and Catherine Tucker.

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

[6] Arcidiacono, Peter, Esteban Aucejo, Arnaud Maurel, and Tyler Ransom. 2016. ”College Attrition and the dynamics of information revelation.” Working Paper

[7] Barron, John M, John Bishop, and William C. Dunkelberg. 1985. ”Employer Search: The Interviewing and Hiring of New Employees.” *The Review of Economics and Statistics*, 67(1), 43-52.

[8] Benabou, Roland and Robert Gertner. 1993. ”Search with Learning from Prices: Does Increased Inflationary Uncertainty Lead to Higher Markups.” *Review of Economic Studies*, 60(1), 69-94.

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

[10] Blinder, Alan. 1973. Wage Discrimination: Reduced Form and Structural Estimates. *The Journal of Human Resources*, 8(4), 436-455.

[11] Blinder, Alan S., and Alan B. Krueger. 2013. ”Alternative measures of offshorability: a survey approach.” Journal of Labor Economics 31.

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

[13] Cairncross, Frances. 2001. ”The death of distance: 2.0: How the communications revolution will change our lives”. Texere, 2001.

[14] Clemens, Michael A. 2011. ”Economics and emigration: Trillion-dollar bills on the sidewalk?.” *Journal of Economic Perspectives,* 25(3), 83-106.

[15] Cremer, Jacques. 1984. ”On the economics of repeat buying”. *The RAND Journal of Economics*, 396-403.

[16] 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.

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

[18] Fort, Theresa C. 2017. Technology and production fragmentation: Domestic versus foreign sourcing. The *Review of Economic Studies*, 84(2), 650-687.

[19] Fortin, Nicole, Thomas Lemieux, and Sergio Firpo. 2011. Decomposition Methods in Economics. *Handbook of Labor Economics*, 4(1), 1-102.

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

[21] Horton, John. 2010. ”Online Labor Markets.” *Internet and Network Economics*: 515-522.

[22] Horton, John. 2017a. ”The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment.” *Journal of Labor Economics,* 35(2).

[23] Horton, John. 2017b. ”Price Floors and Employer Preferences: Evidence from a Minimum Wage Experiment.” Working Paper.

[24] Horton, John, and Joseph Golden. 2015. Reputation inflation: Evidence from an online labor market. Working Paper, NYU.

[25] Horton, John, William R. Kerr, and Christopher Stanton. 2017. ”Digital Labor Markets and Global Talent Flows.” NBER Working Paper No. 23398.

[26] Israel, Mark. 2005. ”Services as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance.” *The American Economic Review*, 95(5), 1444-1463.

[27] Jensen, J. Bradford and Lori G. Kletzer. 2010. ”Measuring Tradeable Services and the Task Content of Offshorable Services Jobs.” in *Labor in the New Economy* eds Katharine Abraham, James Spletzer, and Michael Harper.

[28] Kahn, Lisa. 2013. "Asymmetric Information between Employers." *American Economic Journal: Applied Economics*. 5(4): 165-205.

[29] Kahn, Lisa and Fabian Lange. 2014. "Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures." *Review of Economic Studies* , 81.4: 1575-1613.

[30] Katz, Lawrence, and Alan Krueger. 2016. The rise and nature of alternative work arrangements in the United States, 1995-2015 (No. w22667). *National Bureau of Economic Research.*

[31] Kohn, Meir G. and Steven Shavell. 1974. "The Theory of Search." *Journal of Economic Theory*, 9(2): 93-123.

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

[33] Lyons, Elizabeth. 2017. "Team Production in International Labor Markets: Experimental Evidence from the Field." *American Economic Journal: Applied Economics*, forthcoming July, 2017.

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

[35] Moreno, Antonio, and Christian Terwiesch. 2014. "Doing business with strangers: Reputation in online service marketplaces." *Information Systems Research* 25(4), 865-886.

[36] Nelson, Phillip. 1970. Information and consumer behavior. *Journal of Political Economy*, 78(2), 311-329.

[37] Oaxaca, Ronald. 1973. Male-Female Wage Differentials in Urban Labor Markets. *International Economic Review*, 14(3), 693-709.

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

[39] Petrin, Amil and Kenneth Train. 2010. "A Control Function Approach to Endogeneity in Consumer Choice Models." *Journal of Marketing Research,* 47(1): 3-13.

[40] Riordan, Michael H. 1986. "Monopolistic Competition with Experience Goods." *Quarterly Journal of Economics,* 101(2): 265-280.

[41] Stanton, Christopher and Catherine Thomas. 2016. "Landing the First Job: The Value of Intermediaries in Online Hiring." *The Review of Economic Studies*, 83(2), 810-854.

[42] Shapiro, Carl. 1983. "Optimal Pricing of Experience Goods." *The Bell Journal of Economics*, 14(2), 497-507.

[43] Tadelis, Steven. 2002. "The Market for Reputation as an Incentive Mechanism." *Journal of Political Economy*, 110(4), 854-882.

[44] Train, Kenneth. 2003. "Discrete Choice Methods with Simulation." *Cambridge University Press*.

Table 1: Job Openings and Hiring Probabilities by Employer Hiring Experience

	All Job Openings					Sequential, Arms-Length Openings				
Employers' Previous Hires	Number of Job Openings	Number of Candidates	Share of Candidates Initiated by Employer	Mean Wage Bid	Probability a Hire is Made	Number of Job Openings	Number of Candidates	Share of Candidates Initiated by Employer	Mean Wage Bid	Probability a Hire is Made
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
0	119877	18.39	7.8%	10.16	22%	61160	25.47	2.1%	10.25	16%
		(25.15)		(7.20)			(27.25)		(7.00)	
1	32526	14.67	7.9%	9.77	49%	10173	27.73	2.0%	9.85	31%
		(24.68)		(7.03)			(29.65)		(6.87)	
2	22269	14.11	8.1%	9.33	52%	6220	29.40	1.8%	9.35	30%
		(27.83)		(6.97)			(36.97)		(6.82)	
3	16820	13.91	7.7%	9.33	54%	4525	28.28	1.8%	9.60	31%
		(26.68)		(7.03)			(31.31)		(6.89)	
4+	131378	13.74	7.4%	8.71	57%	27686	31.15	1.7%	9.06	28%
		(30.37)		(7.02)			(39.82)		(6.91)	

Notes: Sample period is from January 2008 to June 2010. For details on sample composition, see Appendix 1. Numbers in parentheses are standard deviations.

Table 2: Worker Characteristics in the Applicant Pool and For Those Hired in the Sequential, Arms-Length Sample

	Job Applicants					Workers hired				
Employers' Previous Hires	0	1	2	3	4+	0	1	2	3	4+
Mean Hourly Wage Bid	10.25	9.85	9.35	9.60	9.05	10.67	10.03	9.95	9.98	9.84
	(7.004)	(6.868)	(6.817)	(6.887)	(6.914)	(7.376)	(6.973)	(6.863)	(7.027)	(7.199)
Good Worker Feedback	0.314	0.338	0.335	0.338	0.333	0.476	0.476	0.440	0.472	0.451
	(0.464)	(0.473)	(0.472)	(0.473)	(0.471)	(0.499)	(0.5)	(0.497)	(0.499)	(0.498)
No Feedback for Worker	0.473	0.427	0.437	0.428	0.443	0.304	0.327	0.337	0.301	0.343
	(0.499)	(0.495)	(0.496)	(0.495)	(0.497)	(0.46)	(0.469)	(0.473)	(0.459)	(0.475)
Number of Prior Jobs	5.01	5.83	5.74	5.81	5.67	9.68	9.51	9.25	9.76	8.99
	(9.982)	(11.619)	(11.745)	(11.905)	(11.675)	(14.83)	(15.018)	(16.396)	(14.89)	(13.976)
BA or Higher Degree	0.348	0.350	0.353	0.351	0.354	0.371	0.362	0.366	0.374	0.359
	(0.476)	(0.477)	(0.478)	(0.477)	(0.478)	(0.483)	(0.481)	(0.482)	(0.484)	(0.48)
Agency Affiliation	0.354	0.350	0.331	0.340	0.315	0.313	0.307	0.304	0.294	0.285
	(0.478)	(0.477)	(0.471)	(0.474)	(0.465)	(0.464)	(0.461)	(0.46)	(0.456)	(0.452)
Number of Observations	1,558,429	282,229	182,975	128,030	862,898	10,009	3,139	1,889	1,400	7,733

Notes: Means of Application Characteristics for all job applicants and workers who are hired. Standard deviations in parentheses.

Table 3: First Stage Regression of Log Hourly Bids on Exchange Rate and Arrivals Instruments

Sample	Inexperienced Employers	Experienced Employers	Inexperienced Employers, Excluding Resume Characteristics	Experienced Employers, Excluding Resume Characteristics
	(1)	(2)	(3)	(4)
Log Dollar to Local Currency Exchange Rate, de-trended	0.0899***	0.0995***	0.0940***	0.100***
	(0.00676)	(0.00832)	(0.00735)	(0.00913)
Residual Log Applicants per Job Opening	-0.0681***	-0.0695***	-0.0805***	-0.0847***
	(0.00357)	(0.00383)	(0.00381)	(0.00409)
Number of Observations	1,558,429	1,456,132	1,558,429	1,456,132
R-Squared	0.612	0.644	0.545	0.580
F Statistic on Excluded Instruments	166.9	131.5	182.4	151.7

Notes: First stage regression coefficients with robust standard errors in parentheses. The inexperienced sample is employers on their first job post. The experienced sample is employers who have hired a previous worker. The first instrument is the log of the average monthly dollar to local currency exchange rate after removing a currency-specific linear trend. The second instrument uses the average number of applications arriving per job opening in the first 24 hours for other jobs in that week and job category cell. After taking logs, the instrument is what remains after removing week and job category fixed effects. Indicators that each instrument is missing or invariant within country are also included. All models contain a calendar time trend, a separate trend for technical categories, job category fixed effects, a spline with four knots for applicant order (knots correspond to pagination after sorting by arrival time), an indicator that the application was employer-initiated, and eight country-group fixed effects. The last country group includes many countries with small application shares. Models in Columns 1 and 2 also include the following applicant characteristics: a dummy for good reported English skills; a dummy for a BA degree or higher; a dummy for having no prior work experience, a dummy for agency affiliation and its interaction with having no prior work experience; the number of prior jobs; and the log of the wage on the last hourly job. See Appendix Table 1 for details and summary statistics on the resume data.

Table 4: Demand Model Estimates, Elasticities, Costs, and Markups

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Employer Experience	Inexperienced	Experienced	Inexperienced	Experienced	Inexperienced	Experienced	Inexperienced	Experienced
Resume Characteristics	Yes		Yes		No		Yes	
Control Function for Bids	No		Yes		Yes		Yes	
Multiple Types	No		No		Yes		Yes	
Panel A: Parameter Estimates from Demand Models								
Note: Experienced Employer Columns Contain Additive Interaction Terms Relative to Inexperienced Employer Baseline								
Log Hourly Bid	-0.46***	-0.11***	-5.50***	-3.13***	-5.08***	-0.87***	-5.10***	-2.78***
	(0.02)	(0.03)	(0.42)	(0.72)	(0.15)	(0.06)	(0.26)	(0.28)
Type-2 Intercept					-2.78***		-2.65***	
					(0.81)		(0.27)	
Type-3 Intercept					1.85***		1.82***	
					(0.38)		(0.06)	
Fraction Type 1					0.19	0.49	0.20	0.49
Fraction Type 2					0.72	0.16	0.72	0.16
Fraction Type 3					0.09	0.35	0.09	0.34
Panel B: Valuations, Elasticities, and Costs								
Mean Own-Price Elasticity	-0.46	-0.56	-5.46	-8.54	-5.04	-5.89	-4.96	-7.70
Mean % Markup, (Pre oDesk-Fee)	NA	NA	22.4%	13.3%	24.7%	20.5%	25.2%	14.9%
Mean Implied Cost (USD, Pre oDesk-Fee)	NA	NA	$7.61	$7.46	$7.46	$7.01	$7.43	$7.35
Mean Wage Bid (USD, Pre oDesk-Fee)	$9.31	$8.45	$9.31	$8.45	$9.31	$8.45	$9.31	$8.45
Percentage of Bid Difference Due to Markups	NA	NA	79%		35%		88%	
Log-bid Equivalent Change in Productive Value:	NA	NA	0.11		0.30		0.30	
Fraction of Change Due to Worker Characteristics (X)			-14%		-27%		-11%	
Fraction Due to Coefficients (β)			-114%		45%		79%	
Fraction Due to Heterogeneity (μ)					82%		32%	

Notes: Standard errors computed using the sandwich form are in parentheses below estimated coefficients. There are 109,814 job openings in the sample used for estimation, with 61,196 postings by inexperienced employers and 48,168 postings by experienced employers. The type-probabilities in odd-numbered columns are for employers in the "Never experienced" group, and those in even-numbered columns are for employers in the "Ever experienced" group. The likelihood in columns 1-4 is over job openings, while the likelihood in columns 5-8 is over sequences of job openings within employer. See the text for details about the estimation procedure. In Panel B, when the model has employer heterogeneity, own price elasticities are type-weighted averages of the individual elasticities, and mean markups are computed from the type-weighted average. Columns 1 and 2 do not have markup estimates because these are undefined on the inelastic segment of the demand curve. The log productive value decomposition is described in the text.

Table 5: Log Wage Bids Decline with Observable Employer Experience

	OLS	OLS	Employer Fixed Effects	Employer Fixed Effects	Worker Fixed Effects	Worker Fixed Effects
	(1)	(2)	(3)	(4)	(5)	(6)
On posts after making 1 hire	-0.0281***	-0.0197***	-0.0177***	-0.0130***	-0.0101***	-0.0113***
	(0.00342)	(0.00293)	(0.00385)	(0.00324)	(0.000639)	(0.000627)
2 hires	-0.0407***	-0.0292***	-0.0283***	-0.0211***	-0.0170***	-0.0178***
	(0.00477)	(0.00416)	(0.00510)	(0.00436)	(0.000787)	(0.000773)
3 hires	-0.0381***	-0.0266***	-0.0272***	-0.0200***	-0.0176***	-0.0184***
	(0.00478)	(0.00407)	(0.00524)	(0.00441)	(0.000893)	(0.000882)
4 hires	-0.0608***	-0.0447***	-0.0437***	-0.0323***	-0.0208***	-0.0215***
	(0.00577)	(0.00500)	(0.00617)	(0.00523)	(0.000995)	(0.000977)
5+ hires	-0.0756***	-0.0591***	-0.0468***	-0.0372***	-0.0386***	-0.0391***
	(0.00569)	(0.00518)	(0.00541)	(0.00458)	(0.000633)	(0.000619)
Detailed Worker and Job Controls	No	Yes	No	Yes	No	Yes
Observations	5,040,655	5,040,628	5,040,655	5,040,628	5,040,655	5,040,628
R-Squared	0.466	0.562	0.540	0.613	0.841	0.845

Notes: The dependent variable is the log of the hourly wage bid. Robust standard errors are clustered by employer. All specifications contain a spline for the applicant's arrival order, detailed job category fixed effects, calendar time fixed effects at the monthly level, and expected duration of the job by required hours-per-week fixed effects. Specifications with detailed worker and job controls also include the following about the worker: a third-order polynomial of the worker's feedback score; a dummy for good reported English skills in the worker's resume; a dummy for a BA degree or higher; a dummy for having no prior work experience; a dummy for agency affiliation and for its interaction with having no prior work experience; the number of prior jobs; the log of the wage received on the last hourly job; and an indicator that no last wage is displayed when the worker is experienced. The detailed job controls in these specifications include: a third-order polynomial in the number of characters in the job description; and an indicator that the employer initiated contact with the worker.

Table 6: Outcomes over Entire Relationship (Including Future Contracts) by Employer Experience on First Hire

	Log Total Hours Over all Employer-Worker Contracts		Indicator that Worker Ever Receives Higher Compensation on Any Contract		Maximum Feedback Score over the Relationship if Feedback is Given	
	(1)	(2)	(3)	(4)	(5)	(6)
On relationships where the employer has made 1 hire	-0.0674***	-0.0982***	-0.00105	-0.000239	0.0243*	0.0312**
	(0.0175)	(0.0169)	(0.00200)	(0.00199)	(0.0131)	(0.0131)
2 hires	-0.117***	-0.170***	-0.00286	-0.00135	-0.000618	0.0126
	(0.0200)	(0.0194)	(0.00224)	(0.00224)	(0.0152)	(0.0152)
3 hires	-0.110***	-0.168***	-0.00214	-0.000374	0.0157	0.0304*
	(0.0222)	(0.0215)	(0.00247)	(0.00247)	(0.0166)	(0.0166)
4 hires	-0.116***	-0.176***	-0.00519**	-0.00323	0.0399**	0.0562***
	(0.0246)	(0.0240)	(0.00263)	(0.00263)	(0.0180)	(0.0179)
5+ hires	-0.0992***	-0.205***	-0.00523***	-0.00156	0.120***	0.152***
	(0.0128)	(0.0125)	(0.00141)	(0.00142)	(0.00943)	(0.00948)
Constant	3.033***	3.094***	0.0425***	0.0404***	4.303***	4.285***
	(0.00993)	(0.00982)	(0.00116)	(0.00115)	(0.00781)	(0.00785)
Expected Job Duration Category Fixed Effects	No	Yes	No	Yes	No	Yes
Observations	118,772	118,772	127,120	127,120	87,179	87,179
R-squared	0.001	0.057	0.000	0.003	0.003	0.011

Notes: Robust standard errors in parentheses. The unit of analysis is a worker-employer pair at the time of first hire. The dependent variable in each column includes all future contracts for the pair. Observation counts differ between columns 1-2 and 3-4 due to some contracts lacking recorded hours. Counts differ in latter columns because ongoing jobs and jobs where no feedback was given are excluded from the feedback calculation.

Table 7: Log Wage Bids and Observable Employer Feedback

	OLS	OLS	Employer Fixed Effects	Employer Fixed Effects	Worker Fixed Effects	Worker Fixed Effects
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Feedback Left For The Employer						
On posts after making 1+ hires	-0.0563***	-0.0868***	-0.0549***	-0.0432***	-0.0523***	-0.0521***
	(0.00342)	(0.0112)	(0.00737)	(0.00645)	(0.00713)	(0.00701)
1+ hires and no observable employer feedback		0.0549***	0.0267***	0.0215***	0.0343***	0.0334***
		(0.0110)	(0.00726)	(0.00640)	(0.00705)	(0.00692)
1+ hires and good observable employer feedback		0.0501***	0.0289***	0.0249***	0.0287***	0.0275***
		(0.0111)	(0.00774)	(0.00695)	(0.00711)	(0.00698)
Detailed Worker and Job Controls	No	Yes	No	Yes	No	Yes
Observations	5,019,235	5,019,208	5,019,235	5,019,208	5,019,235	5,019,208
R-Squared	0.491	0.591	0.564	0.640	0.857	0.861
Panel B: Feedback The Employer Left for Workers						
On posts after making 1+ hires	-0.0563***	-0.0817***	-0.0805***	-0.0636***	-0.0453***	-0.0447***
	(0.00342)	(0.00671)	(0.00519)	(0.00436)	(0.00402)	(0.00397)
1+ hires and no observable feedback left		0.0448***	0.0595***	0.0481***	0.0221***	0.0206***
		(0.00588)	(0.00446)	(0.00377)	(0.00352)	(0.00347)
1+ hires and good observable feedback left		0.0242***	0.0245***	0.0212***	0.0106**	0.0103**
		(0.00765)	(0.00574)	(0.00487)	(0.00473)	(0.00467)
Detailed Worker and Job Controls	No	Yes	No	Yes	No	Yes
Observations	5,019,235	5,019,235	5,019,235	5,019,208	5,019,235	5,019,208
R-Squared	0.491	0.591	0.564	0.640	0.857	0.861

Notes: The dependent variable is the log of the hourly wage bid. Robust standard errors are clustered by employer. All specifications mirror those in Table 2 and contain the same controls as detailed in the notes to Table 2. Observation counts differ from Table 2 because, for some observations, the timing of the initial feedback cannot be classified as occurring before or after later job postings.

Table 8: 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	-0.0515***	-0.292***	-0.289***	-0.288***	-0.286***
	(0.00576)	(0.00773)	(0.00772)	(0.00774)	(0.00773)
Two previous hires	-0.0918***	-0.362***	-0.357***	-0.357***	-0.352***
	(0.00657)	(0.00928)	(0.00929)	(0.00928)	(0.00929)
Three previous hires	-0.114***	-0.406***	-0.401***	-0.400***	-0.396***
	(0.00726)	(0.0106)	(0.0106)	(0.0106)	(0.0106)
Four previous hires	-0.135***	-0.442***	-0.436***	-0.437***	-0.431***
	(0.00786)	(0.0114)	(0.0114)	(0.0115)	(0.0115)
Five or more previous hires	-0.177***	-0.520***	-0.513***	-0.513***	-0.506***
	(0.00704)	(0.0123)	(0.0123)	(0.0123)	(0.0123)
Mean Log Bid				0.0805***	0.0807***
				(0.00474)	(0.00409)
Constant	0.893***	1.121***	1.110***	0.910***	0.898***
	(0.121)	(0.172)	(0.177)	(0.168)	(0.173)
Includes job duration fixed effects and third order polynomial of job description length	No	No	Yes	No	Yes
Observations	322,333	322,333	322,332	322,333	322,332
R-Squared	0.021	0.414	0.416	0.416	0.418

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

Table 9: Platform Profits and Repeat Job Posting for Different Fee Schedules Assuming a High Future Value of Experience

Panel A: Percent Change in Profits Relative to 10% Uniform Fee

Inexperienced \ Experienced Fee	5%	10%	15%	20%	25%	30%
5%	-37%	-13%	1%	4%	1%	-4%
10%	-24%	0%	12%	15%	12%	8%
15%	-14%	8%	18%	21%	18%	14%
20%	-11%	8%	18%	20%	18%	13%
25%	-9%	8%	17%	19%	18%	14%
30%	-11%	5%	13%	14%	14%	10%

Panel B: Percent Change in Number of Employers Becoming Experienced

Inexperienced \ Experienced Fee	5%	10%	15%	20%	25%	30%
5%	21%	12%	4%	-3%	-11%	-18%
10%	7%	0%	-7%	-14%	-21%	-27%
15%	-3%	-9%	-16%	-22%	-28%	-34%
20%	-15%	-21%	-27%	-33%	-38%	-43%
25%	-25%	-30%	-35%	-40%	-44%	-49%
30%	-36%	-40%	-44%	-49%	-52%	-56%

Notes: Simulations use the parameters from the last two columns of Table 6. Inexperienced employers are first assigned types from the weighted distribution of types observed overall in the inexperienced sample. The weights for these types match the population fraction of employers who begin as inexperienced and transition to the experience sample and the population fraction who never transition to the experienced sample. Simulated profits are then constructed according to the following procedure for each pair of fees: 1) Log wage bids to inexperienced employers are calculated, where pass-through of the fee is computed according to the worker's first order condition for setting bids. 2) Inexperienced employers hire or not based on the choice probabilities calculated from the demand model under the bids that reflect the inexperienced fee. 3) Iterating until convergence, log wage bids including fees and the worker's markup are calculated for experienced employers who post jobs. The set of experienced employers posting subsequent jobs is calculated using the types who hire, the expected surplus given wage bids, and a probability model for transitioning to post experienced jobs as a function of the expected surplus given bids that is calibrated off of the data with fixed 10% fees. Markups and log wage bids are recomputed given the set of experienced employers. After step 3) has converged, profits are calculated based on hiring probabilities from the model and the fee-rate associated with the chosen bids. Employers who become experienced are assumed to have a present value of eight jobs while in the experienced sample. This number doesn't match the data, but is the smallest number of future jobs for which a fee on the inexperienced lower than 20% is profitable for the platform.

Table A1: Details about Resume Data

Variable	Full Sample Mean	Std. Dev.	Min	Max	Sequential, Arms-Length Sample Mean	Std. Dev.	Min	Max
Log of Hourly Rate on Last Job (Zero for Missing)	1.11	1.19	-4.61	7.01	1.12	1.20	-4.61	6.62
No Last Rate Displayed	0.42	0.49	0	1	0.43	0.49	0	1
Self-reported Good English Skills	0.90	0.30	0	1	0.90	0.30	0	1
BA or Higher Degree	0.35	0.48	0	1	0.35	0.48	0	1
Number of Prior Jobs	5.44	10.99	0	266	5.35	10.85	0	266
Indicator for No Prior Jobs	0.39	0.49	0	1	0.39	0.49	0	1
Feedback Score (Including Zeros)	2.44	2.27	0	5	2.39	2.26	0	5
Feedback Score Squared	11.53	10.84	0	25	10.81	10.76	0	25
Feedback Score Cubed	53.43	52.80	0	125	49.84	52.06	0	125
Prior Experience and Zero Feedback	0.06	0.24	0	1	0.06	0.24	0	1
Agency Affiliate	0.32	0.47	0	1	0.34	0.47	0	1
Agency Affiliate x No Prior Jobs	0.08	0.27	0	1	0.09	0.28	0	1

Notes: This table provides summary measures for the detailed resume data used in estimation. The full sample has 5,040,791 observations, while the sequential, arms-length sample has 3,014,561 observations. In the hiring probability estimation, fixed effects for country groups, job category, a spline for applicant order, an indicator for an employer-initiated application, and a time trend are also included. The log of the hourly rate on the last job is set to zero for applicants who have never been hired for hourly jobs.

Table A2: First Stage Regression of Log Hourly Bids on Exchange Rate and Arrivals Instruments Including Worker Fixed Effects

Sample	Inexperienced Employers	Experienced Employers	Inexperienced Employers, Excluding Resume Characteristics	Experienced Employers, Excluding Resume Characteristics
	(1)	(2)	(3)	(4)
Log Dollar to Local Currency Exchange Rate, de-trended	0.119***	0.135***	0.131***	0.175***
	(0.00577)	(0.00744)	(0.00588)	(0.00766)
Residual Log Applicants per Job Opening	-0.0294***	-0.0205***	-0.0325***	-0.0205***
	(0.00271)	(0.00304)	(0.00275)	(0.00309)
Number of Observations	1,558,429	1,456,132	1,558,429	1,456,132
R-Squared	0.868	0.868	0.864	0.864
F Statistic on Excluded Instruments	153.1	103.1	178.8	154

Notes: This table replicates the first stage regression table but includes applicant fixed effects. Columns 1 and 2, therefore, differ from 3 and 4 only by including time-varying applicant characteristics. See notes for Table 3.

Table A3: Log Wage Bids controlling for the arrival rate of job applicants

	OLS	OLS	Employer Fixed Effects	Employer Fixed Effects	Worker Fixed Effects	Worker Fixed Effects
	(1)	(2)	(3)	(4)	(5)	(6)
On posts after making 1 hire	-0.0247***	-0.0161***	-0.0180***	-0.0135***	-0.00961***	-0.0106***
	(0.00339)	(0.00290)	(0.00382)	(0.00320)	(0.000638)	(0.000626)
2 hires	-0.0371***	-0.0252***	-0.0287***	-0.0217***	-0.0166***	-0.0170***
	(0.00482)	(0.00423)	(0.00508)	(0.00435)	(0.000786)	(0.000772)
3 hires	-0.0345***	-0.0228***	-0.0279***	-0.0212***	-0.0171***	-0.0176***
	(0.00474)	(0.00404)	(0.00519)	(0.00435)	(0.000891)	(0.000881)
4 hires	-0.0563***	-0.0399***	-0.0438***	-0.0328***	-0.0202***	-0.0206***
	(0.00572)	(0.00494)	(0.00611)	(0.00516)	(0.000993)	(0.000976)
5+ hires	-0.0718***	-0.0551***	-0.0473***	-0.0384***	-0.0380***	-0.0383***
	(0.00567)	(0.00516)	(0.00536)	(0.00452)	(0.000630)	(0.000617)
Log Applicant Arrivals in First 24 Hours	-0.0428***	-0.0483***	-0.0306***	-0.0363***	-0.00854***	-0.0123***
	(0.00219)	(0.00178)	(0.00201)	(0.00160)	(0.000480)	(0.000444)
Detailed Worker and Job Controls	No	Yes	No	Yes	No	Yes
Observations	5,040,077	5,040,050	5,040,077	5,040,050	5,040,077	5,040,050
R-Squared	0.467	0.563	0.540	0.613	0.841	0.845

Notes: The dependent variable is the log of the hourly wage bid. The sample is limited to worker-initiated applications on sequential job openings. Robust standard errors are clustered by employer. All specifications contain a spline for the applicant's arrival order, detailed job category fixed effects, monthly time fixed effects, and expected duration by hours-per-week fixed effects. Specifications with detailed worker and job controls also include the following: a third-order polynomial in the number of characters in the job descriptio;, a dummy for good reported English skills; a dummy for a BA degree or higher, a dummy for having no prior work experience; a dummy for agency affiliation and its interaction with having no prior work experience; the number of prior jobs; the log of the wage on the last hourly job; and an indicator that no last wage is displayed when the worker is experienced.

Table A4: Field Experimental Evidence on Log Bids by Employer Experience

	Log Bid			Difference in Log Bid and Log Profile Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
Experienced Employer	-0.222***	-0.205***	-0.305**	-0.213***	-0.208**	-0.246**
	(0.0623)	(0.0677)	(0.127)	(0.0785)	(0.0828)	(0.109)
Constant	1.406***	1.396***	1.425***	-0.0689	-0.0671	-0.0154
	(0.0606)	(0.0659)	(0.0685)	(0.0545)	(0.0546)	(0.0607)
Hours Control Included		Y	Y		Y	Y
Applicant Order Fixed Effects			Y			Y
Observations	128	97	97	97	97	97
R-squared	0.138	0.111	0.622	0.069	0.070	0.826

Notes: Robust standard errors in parentheses below. Experimental jobs were posted in March of 2015. Experienced employer refers to a job posted by an employer with past hiring experience in data entry and with good employer feedback. Otherwise, the job was posted by an inexperienced employer. The number of hours worked previously by each candidate, termed "hours control," was hand-collected by an RA. This was not available for applicants who did not have publicly visible resumes, reducing the sample size.

Table A5: Analysis of Bargaining Proxy by Employer Experience

	OLS	OLS	Employer Fixed Effects	Employer Fixed Effects	Worker Fixed Effects	Worker Fixed Effects
	(1)	(2)	(3)	(4)	(5)	(6)
On posts after making 1 hire	0.00114	0.000378	0.000232	-0.00105	0.00525	0.00438
	(0.00357)	(0.00357)	(0.00598)	(0.00598)	(0.00470)	(0.00471)
2 hires	0.00319	0.00201	0.00878	0.00720	0.00472	0.00402
	(0.00400)	(0.00400)	(0.00666)	(0.00666)	(0.00528)	(0.00527)
3 hires	0.00216	0.000731	0.000334	-0.00163	0.00148	0.000442
	(0.00445)	(0.00445)	(0.00733)	(0.00734)	(0.00599)	(0.00601)
4 hires	-0.00642	-0.00764	-0.00490	-0.00654	-0.00219	-0.00342
	(0.00487)	(0.00485)	(0.00792)	(0.00791)	(0.00652)	(0.00650)
5+ hires	-0.0220***	-0.0221***	-0.00804	-0.0104	-0.0149***	-0.0159***
	(0.00337)	(0.00334)	(0.00687)	(0.00689)	(0.00377)	(0.00378)
Detailed Worker and Job Controls	No	Yes	No	Yes	No	Yes
Observations	61,958	61,958	61,958	61,958	61,958	61,958
R-Squared	0.018	0.024	0.402	0.405	0.353	0.357

Notes: Dependent variable is the log of the final hourly wage bid when the applicant is hired, less the last log hourly wage bid on a rejected application in the same job category and employer experience segment. Robust standard errors are clustered by employer. All specifications contain a spline for the applicant's arrival order, detailed job category fixed effects, calendar time fixed effects at the monthly level, and expected duration of the job by required hours-per-week fixed effects. Specifications with detailed worker and job controls also include the following about the worker: a third-order polynomial of the worker's feedback score; a dummy for good reported English skills in the worker's resume; a dummy for a BA degree or higher; a dummy for having no prior work experience; a dummy for agency affiliation and its interaction with having no prior work experience; the number of prior jobs; the log of the wage received on the last hourly job; and an indicator that no last wage is displayed when the worker is experienced. The detailed job controls in these specifications include: a third-order polynomial in the number of characters in the job description and a dummy that the employer initiated contact with the worker.

Table A6: Productivity and Search Effort on the First and Second Jobs, for Employers who Interview on First Job

	First Job Outcomes					Second Job Outcomes, for Employers who Hire on Second Job			
	Hires a Worker	Hires and Reports Success	Hires with Good Feedback	Hires and Reports Success, Wage Control	Hires with Good Feedback, Wage Control	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)	(7)	(8)	(9)
Panel A: Specifications Without Buyer Group Fixed Effects									
2-5 Interviews	-0.021***	-0.014***	-0.012***	-0.029***	-0.048***	-0.033***	-0.050***	-0.032***	-0.050***
	(0.005)	(0.004)	(0.004)	(0.010)	(0.011)	(0.012)	(0.015)	(0.012)	(0.015)
6-10 Interviews	-0.013*	-0.016***	-0.009*	-0.054***	-0.051***	-0.053***	-0.068***	-0.049***	-0.063***
	(0.007)	(0.005)	(0.005)	(0.014)	(0.016)	(0.016)	(0.018)	(0.016)	(0.018)
11+ Interviews	-0.079***	-0.054***	-0.044***	-0.064***	-0.057**	-0.115***	-0.072***	-0.107***	-0.062***
	(0.007)	(0.005)	(0.005)	(0.020)	(0.024)	(0.019)	(0.019)	(0.019)	(0.020)
Includes Buyer Group Fixed Effects	No	No	No	No	No	No	No	No	No
Mean of DV	0.238	0.137	0.119	0.677	0.674	0.639	0.659	0.639	0.659
Observations	53,687	51,254	49,612	10,362	8,724	8,766	7,269	8,766	7,269
R-Squared	0.048	0.056	0.043	0.090	0.047	0.073	0.038	0.079	0.047
Panel B: Including Buyer Group Fixed Effects									
2-5 Interviews	-0.029***	-0.020***	-0.016***	-0.038***	-0.045***	-0.047***	-0.043**	-0.046***	-0.044**
	(0.005)	(0.004)	(0.004)	(0.014)	(0.017)	(0.017)	(0.020)	(0.017)	(0.020)
6-10 Interviews	-0.021***	-0.024***	-0.015***	-0.069***	-0.049**	-0.059***	-0.051**	-0.056***	-0.047*
	(0.007)	(0.006)	(0.005)	(0.019)	(0.023)	(0.021)	(0.024)	(0.021)	(0.024)
11+ Interviews	-0.088***	-0.061***	-0.050***	-0.068***	-0.046	-0.120***	-0.059**	-0.115***	-0.050*
	(0.008)	(0.006)	(0.005)	(0.026)	(0.032)	(0.026)	(0.027)	(0.026)	(0.027)
Includes Buyer Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of DV	0.238	0.137	0.118	0.677	0.674	0.639	0.659	0.639	0.659
Observations	53,641	51,221	49,587	10,348	8,718	8,762	7,266	8,762	7,266
R-Squared	0.131	0.137	0.128	0.298	0.291	0.304	0.318	0.308	0.324

Notes: Column headings display the dependent variable. Robust standard errors clustered by employer. The interview counts parameter estimates are interviews on the employer's first job opening and are carried over to examine outcomes on the second job opening (Columns 6 - 9). Panel B includes fixed effects for groups of employers who interview workers with similar characteristics on the first job. All specifications contain fixed effects for expected duration of job, time, and detailed job category.