HARVARD | BUSINESS | SCHOOL

Missing Trade in Tasks: Employer Outsourcing in the Gig Economy

Christopher T. Stanton
Catherine Thomas

Working Paper 18-096

Missing Trade in Tasks: Employer Outsourcing in the Gig Economy

Christopher T. Stanton
Harvard Business School

Catherine Thomas
The London School of Economics and Political Science

Working Paper 18-096

Missing Trade in Tasks:
Employer Outsourcing in the Gig Economy

Christopher T. Stanton
Harvard, NBER, and CEPR

Catherine Thomas
LSE, CEP, and CEPR

February 2019*

Abstract

Only a small share of employers in high-wage locations procure services from abroad. We document heterogeneous employer adoption of an online labor market that facilitates trade in tasks with global workers. Job vacancies posted by experienced employers who have adopted the market are twice as likely to be filled, and this difference is unrelated to the set of available job applicants or their wages. Instead, hiring demand from experienced employers shifts outward for two reasons that we identify using exogenous variation in workers' wage bids. First, their value for hiring in the market increases—a form of learning-by-doing. Second, experienced employers omit the low-value employers who leave the market. Employers appear to learn their value for online hiring only by trying it out, and new employers' adoption decisions are relatively insensitive to wage rates. Larger firms have lower estimated values for the market. The results suggest that employers' willingness to fragment and outsource production at the task level, rather than the quality of the available online workforce, limits the growth of the gig economy.

*A previous version of this paper was circulated with the title "Experience Markets: An Application to Outsourcing and Hiring." 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 Organizational Economics Working Group, NBER Summer Institute, Queens, Sciences Po, 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, Fabian Lange, 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

Technological change has enlarged the set of feasible relationships between firms and workers, permitting the increase in alternative work arrangements that have come to be known as the "gig" economy. Online labor market platforms are a prominent form of information and communications technology for finding, hiring, and managing task-based contractors (Autor, 2001). Early commentators forecast such platforms would herald the death of distance as a barrier to the global allocation of jobs that do not require co-located production (Cairncross, 1997). However, much like the relatively slow growth of the domestic gig economy (Katz and Krueger, 2019), and despite the technical feasibility of offshoring many jobs, the growth of remote online hiring has fallen far short of early predictions. In the 2012 Survey of U.S. Business Owners, only 1.36% of firms reported offshoring any services or functions abroad.[1] This paper studies demand and supply in one of the largest online labor markets to shed light on why the globalization of labor services has lagged early expectations.

Using data on over 80,000 employers who posted job vacancies between 2008 and 2010, we study employers' decisions to adopt online hiring. Sixty-three percent of new employers never hire a worker after posting at least one job vacancy, 35% hire for one job but do not return to post any future vacancies, and 25% adopt the hiring technology, meaning they hire at least once and return to post additional tasks. The key to our empirical approach is that job-level hiring processes and outcomes differ with employer experience. In particular, employers with prior hiring experience are twice as likely to fill a vacancy compared to employers who have not hired before. Analyzing these differences allows us to distinguish between the possible reasons why only 25% of employers who try it end up adopting the platform. Understanding employers' relatively limited platform adoption is informative about gig economy hiring, and, more broadly, about whether technological improvements in contracting will lead traditional employers to adopt task-level contracting outside the firm (Coase, 1937).

We consider two channels for explaining low adoption that come from employer demand. The first possibility is that employers' value for this platform is lower than anticipated, meaning most employers conclude online hiring does not meet their needs. A second possibility is that the benefits from online hiring accrue only after a period of learning about how to use the platform. In these demand-based explanations, the wage savings that attract employers to online hiring are not sufficient to overcome offsetting factors that make early online hiring costlier than other alternatives.

[1] This is despite the fact that hourly wages differ by up to a factor of 10 for similar workers in high-wage cities in the US and, for example, in India (Payscale.com). The combined annual gross revenue of the two largest online hiring platforms is currently less than $5 billion. While services have long-been viewed as "nontradeables," (Blinder and Krueger, 2013) suggest that around 25% of all services industry jobs can now be performed remotely, enabled for the most part by new technology. Low observed US imports of services trade hence suggests there are "missing" services imports from labor-abundant, low-wage, countries, similar to (Trefler, 1995).

The supply side offers different explanations for low adoption rates, particularly if new employers have difficulty attracting applicants. Prior studies of online labor markets show new workers face market entry barriers due to incomplete information about their quality (Pallais, 2014; Stanton and Thomas, 2016; Agrawal, Lacetera, and Lyons, 2016). We consider whether analogous problems exist for employers, where workers are unwilling to supply labor to new employers. Our analysis reveals whether the platform can increase adoption rates with tweaks to labor supply and by encouraging employer learning-by-doing on the platform, or if demand for gig-level production is low for reasons that are beyond the platform's direct control.

Three institutional features of this market shape our approach to distinguishing between demand- and supply-based explanations for low adoption rates. First, job vacancies are for a single role or position, meaning that the employer hires either one candidate or none (Lazear, Shaw, and Stanton, 2018).[2] Second, job candidates bid an hourly wage when applying for a job, and the bid can be thought of as a take-it-or-leave-it offer. Third, workers can observe whether employers have previously hired, as well as any bilateral feedback scores given to prior workers or left for the employer. Using strategies to exploit exogenous variation in wage bids, we model employer hiring decisions as discrete choices from the set of job applicants and the option to not hire.

To understand demand, we use a simple static discrete choice hiring model in which we allow the parameters to vary with an employer's experience in the market. In our setup, an employer's utility for choosing to hire a worker includes a heterogeneous employer preference for online hiring relative to his outside option. His utility for any particular worker also includes functions of worker characteristics, including wage bids, and job attributes, with parameters that potentially vary based on the employer's prior hiring experience. The demand-driven difference between experienced and inexperienced employers in the propensity to use the platform is revealed in the difference in the utility of online hiring between all new employers and the selected set of employers who become experienced. Variation in wages and candidate characteristics allow us to separate learning-by-doing from the selection of the highest-value employers into becoming experienced.

To account for the potential endogeneity of wage bids, we use Petrin and Train's (2010) control function approach with a set of instrumental variables. The main instrument exploits the fact that all contracts are denominated in US dollars but, because applicants are globally distributed, their offline consumption and employment opportunities are denominated in their local currency. Exogenous exchange rate fluctuations serve as an opportunity cost shock, affecting workers' willingness to supply labor to the platform. A second instrument, based on variation in the expected competitiveness of a job opening, shifts wage bids for all

[2]The technology therefore permits trading in tasks, or fragments of a production process, as discussed in (Grossman and Rossi-Hansberg, 2008).

workers, including those from workers located in countries with dollar-linked currencies. This instrument captures the intensity of applications to jobs in the same job category and time period—information that is available to potential applicants. The combination of these instruments provides plausibly exogenous variation in both the relative bids between job applicants and in the level of bids for all applicants relative to an employer's outside option.

We find the average utility from hiring a worker on the platform differs significantly between inexperienced and experienced employers—a shift in labor demand that explains almost all of the difference in hiring propensity between the two groups. Translating utility differences into a money metric, we show inexperienced employers would require a wage bid reduction of around 30%, which is equivalent to a reduction of about $2.80 per hour, to bring their hiring propensity up to the experienced employer level.

The difference between inexperienced and experienced employers' demand for hiring is due to both selection and learning-by-doing. First, new employers include employers with low valuations. Employers who adopt online hiring are a selected sample of all potential employers—those that have relatively high valuations for finding remote workers. We parameterize the distributions of heterogeneity using a finite mixture model, allowing the type probabilities to differ based on whether the employer is observed to adopt the hiring technology by entering the experienced employer sample. Our estimates of employer type have a higher mean among experienced employers, with substantially less probability weight on the lowest type, consistent with positive selection into adopting the platform.

Because all employers appear in the data by posting jobs, it is their early experiences with the platform that are likely revealing to them whether online hiring is a good fit with their needs, and they select out of the market when it is not. This characteristic is conceptually similar to Nelson's (1970) definition of Experience Goods as goods that consumers can evaluate only through trial. We propose that this platform itself can be viewed analogously as an Experience Market, in which employers learn how the market compares to their outside alternative only after trying it out. We note that online hiring requires an employer to undertake various activities in addition to finding a worker, including fragmenting production to the task level, and coordinating online and offline production. The costs of these requirements are plausibly unknown prior to trial. In our preferred specification, differences in the composition of employers in the market account for around one third of the difference in hiring utility between inexperienced and experienced employers.

The learning-by-doing channel is captured by the extent to which the typical experienced employer's own value for the workers on the platform has shifted out compared to the same employer's value while inexperienced, increasing his propensity to hire. This comparison is a revealed preference approach rather than a direct estimate of productivity improvements with experience as in Levitt, List, and Syverson (2013). The shift in hiring propensity may be driven by employers learning to better manage remote contractors, to tailor their job postings, or to incorporate remote contractors into their existing production process.

Learning-by-doing explains around two-thirds of the change in hiring utility between inexperienced and experienced employers.

The estimates of the hiring model also yield a finding that may be surprising at first glance: Given that experienced employers are more likely to hire for any job, one might expect these employers to have less wage elastic hiring demand. However, we find experienced employers' demand for individual job candidates appears to rotate, making their hiring probability for a given worker more elastic with respect to hourly wage bids. In line with this finding, and interpreting wage bids as equilibrium outcomes in the worker's response to employer experience-specific labor demand, we find that workers include lower markups in their bids to experienced employers and wage bids fall with employer experience. We present summary statistics suggesting that this changing elasticity with experience is due to experienced employers' having accumulated knowledge of a stock of prior applicants that may be recalled on future jobs, increasing the degree of latent competition for a slot.

On the supply side, little evidence suggests that workers avoid applying to inexperienced employers or increase their wage bids to these employers because they anticipate greater costs on the job. The most plausible supply-related costs associated with inexperienced employers relate to their lack of reputation and the potential for receiving poor job-related feedback from them. However, experienced employers who have received good feedback from workers, or who have left good feedback for their prior hires, receive wage bids that are *higher* than the wage bids to other experienced employers, which runs counter to the notion that employer reputation lowers workers' costs. To assess the extent of unobserved changes in workers' expected costs, possibly arising from unobserved changes in how employers describe jobs, we conducted a small field experiment in which we posted identical jobs when inexperienced and experienced employers. Labor supply patterns in these experimental jobs match the observational data.

Armed with our model estimates, we return to ask whether platform policies might encourage additional adoption of this hiring technology for those who try it out. For example, subsidizing hiring may change the extent to which uncertain employers gather information about their own type. We ask whether the platform would find reducing the wage bids that new employers face to be profitable. Our counterfactuals show the platform would have actually earned significantly lower profits had it reduced wage bids to new employers.[3] Although relative subsidies to the inexperienced would increase the share of employers that adopt the platform, because of the relative wage inelasticity of new employers shown in our labor demand estimates, those induced to hire by a subsidy would not go on to hire enough in the future to recover the up-front reduced platform revenue from instituting lower fees.[4] The counterfactuals reveal that the

[3]These conclusions are robust to whether employers are assumed to be forward looking, anticipating changes in future platform fees, or myopic.

[4]In our preferred specification, the probability with which new employers hire has an elasticity with respect to a uniform wage change of -2.93. Failing to account for the greater heterogeneity in employer types among the inexperienced, and selection

challenge of finding the right worker at the right wage cannot explain low adoption rates.

What do our results tell us about the potential size of the market for online labor services? Whereas we find significant learning-by-doing gains from experience in the market, most new potential employers in our sample learn they prefer alternative hiring option to adopting online hiring. Moreover, this preference is relatively insensitive to the wage bids offered.[5] We note that the employers using this platform to contract at arm's length must first fragment production into tasks that an individual worker can complete, and then must coordinate with these workers at arm's length. These costs are employer specific and are likely borne by the employer in local wages. Hence, the subsidies offered by the platform in our counterfactuals do not affect them.

The natural next question is why online hiring works for some employers but not for others. The data include limited information about employers, but we do observe employer size and geography. Projecting our posterior estimate of each employer's heterogeneous value for online hiring on employer size indicators and location fixed effects yields few differences by employer geography. However, our analysis reveals significant heterogeneity in employer value for the market as a function of firm size. Value for online hiring is negatively correlated with the number of internal employees at the firm. Conditional on posting a job, firms with many internal workers are more likely to conclude they prefer to use their alternative labor sources rather than online hiring on this platform.[6] Size may indicate a high cost of fragmenting production to the task level or of integrating and coordinating outsourced work with other parts of a firm's production process, especially if large firms have more complex or bureaucratic processes. Large firms may alternatively be more likely to have substitute internal labor capacity, making the platform relatively less appealing. Although we don't have direct evidence on why large firms are the least likely to adopt this platform after trying it out, the correlation with firm size supports our finding that adoption relates to variation in the value of an employer's outside option, and not to variation in employers' experiences within the platform. We infer that low adoption of globalized online hiring, then, can be attributed to employer-specific costs of fragmenting production and coordinating with individual online contractors at

out, would have led us to overestimate the inexperienced wage elasticity of vacancy fill, at -4.10. We would have then, in turn, overestimated the responsiveness of adoption to new employer subsidies. Our preferred estimate of the vacancy fill rate wage elasticity among experienced employers is -4.03. That is, experienced employers are more elastic than inexperienced employers once the latter group's relative heterogeneity is taken into account.

[5] Furthermore, we expect that our sample encompasses only the higher-tail segment of the platform value distribution because we don't observe employers who don't try out the platform during this time period.

[6] This negative association is in stark contrast to the strong positive association between firm size and propensity to engage in international trade that has been documented in many settings (Bernard, Jensen, Redding, and Schott, 2007; Halpern, Koren, and Szeidl, 2015; Antras, Fort, and Tintelnot, 2017). Recent models of importing typically assume only large firms have the scale to realize sufficient variable cost savings from importing to offset the fixed costs of doing so. In our setting, the single worker-task-level contracting environment doesn't allow any employer scale benefits.

arm's length.

Stepping beyond our specific context, we highlight the general implications of our approach for the study of new markets. If a market is very different from buyers' existing alternatives, it is quite plausible that many users need to engage in trial to learn whether it is valuable to them. Hence, new demand will include potential buyers whose value for the market is revealed to be negative ex-post. Our analysis demonstrates the importance of accounting for how heterogeneity in value for these Experience Markets affects the composition of users. In our application, had we failed to account for changing type heterogeneity as employers gain experience and select out, we would have overestimated the price elasticity on the extensive margin of the market-adoption decision, disguising the real barriers that buyers face when entering a market.

The paper proceeds as follows: Section 2 introduces the data and the empirical context. Section 3 describes the hiring probability estimation strategy. Section 4 presents and interprets the estimation results for inexperienced and experienced employers. Section 5 contains reduced-form results related to wage bids and hiring processes. Section 6 is the counterfactual analysis. Section 7 concludes.

2 The Setting, Data, and Summary Statistics

2.1 oDesk.com: How It Works

oDesk.com is an online platform that allows employers to contract with remote workers who sell labor services.[7] The platform facilitates search and matching, remote task management, and payments. Work includes a range of jobs in which 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% of the transactions in the market span international borders, with the worker typically located in the lower-wage country.

An employer who wants to purchase online labor services creates an account on the platform, with no upfront 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

[7] See Horton (2010) for an overview of how online labor markets work, and Agrawal, Horton, Lacetera, and Lyons (2015) and Horton, Kerr, and Stanton (2017) for stylized facts about patterns of contracting, especially between different countries, in these markets. The market we study has features and characteristics similar to other prominent platforms, which at the time of our sample included eLance and Guru. eLance merged with oDesk in 2014, and the merged company was re-named Upwork.com. The data used in this paper pre-date the merger. Most papers using data from online labor markets focus on workers' careers (Agrawal, Lacetera, and Lyons, 2016; Lyons, 2017; Ghani, Kerr, and Stanton, 2014; Pallais, 2014; Stanton and Thomas, 2016). The main exceptions are Horton (2017a) and Horton (2017b), which examine matching under different platform policies.

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, allowing potential applicants to learn it before submitting their application.

Figure 1: A Job Posting.

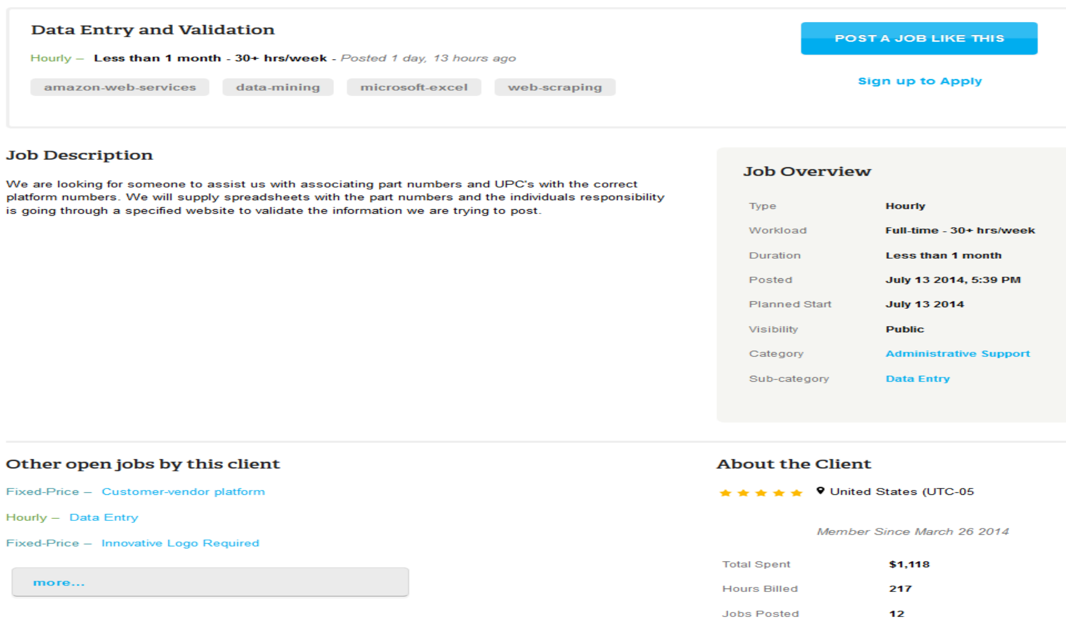

Interested workers submit applications for the job posting and bid an hourly wage to work on the specific job. Employers also have the option to search worker profiles directly and invite applications. Workers' profiles contain information about their skills, education, prior offline work experience, and experience on oDesk (see Figure 2). Workers are located worldwide, and the profiles of workers with prior experience on the site show summary feedback scores received from past work on a 5-point scale.

After receiving applications and inviting workers to apply, employers can request interviews with any number of candidates for the job. If the applicant agrees, the interview usually takes place via Skype.[8] An employer may choose to hire an applicant with or without interviewing her first.

After 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 the job was completed successfully. In other contexts, having a market reputation or having received good feedback mitigates payment risk, but oDesk guarantees workers are paid for the hours billed. Thus, payments and payment risk are unrelated to employer reputation or experience.

[8] The data record interview requests and the applicant's response to the request. 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 interview.

Figure 2: A Worker Profile.

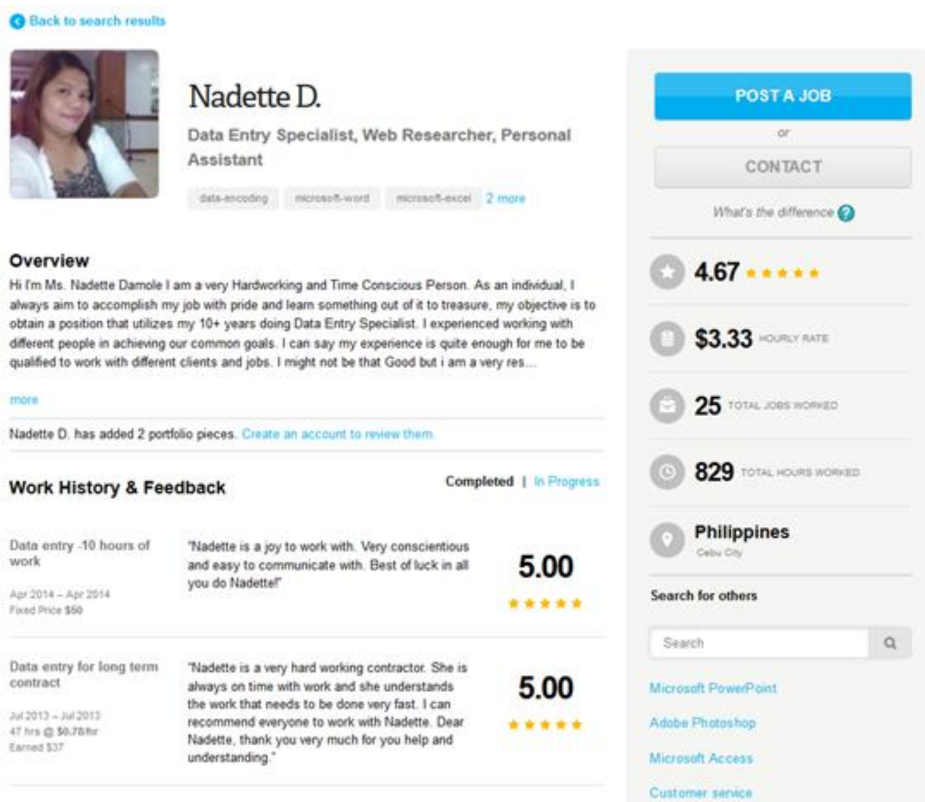

The data used in this paper are administrative data from oDesk. In the data, we observe every employer's job-specific search process in each of his successive job postings. For each posting, the data contain information about the entire applicant pool; which candidates, if any, are interviewed; which 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: 82,257 potential employers posted 322,870 job openings, and more than 5 million job applications were made. There are nine job categories in the data. 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.[9]

Most employers are located in the US and range from private individuals to those hiring within large firms. From the limited information on employers collected by the company, firm size is known for 20,395

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

employers. Of these, 34% are private individuals or sole proprietors, 51% have between two and nine workers, and the remaining 15% have more than 10 in-house employees. Applicants do not see external information about an employer—only the employer's location and past oDesk platform activity are observable when applying for a job, as displayed in Figure 1. On each job posting, workers can observe any feedback that employers have received from previous workers, as well as the number of past hires by the employer and the total hours billed. Figure 3 presents the distribution of the number of hires per employer throughout the period. Sixty-three percent of the potential employers posting jobs hire no applicants at all, whereas 17% make five or more hires.

Figure 3: Number of Hires per Employer

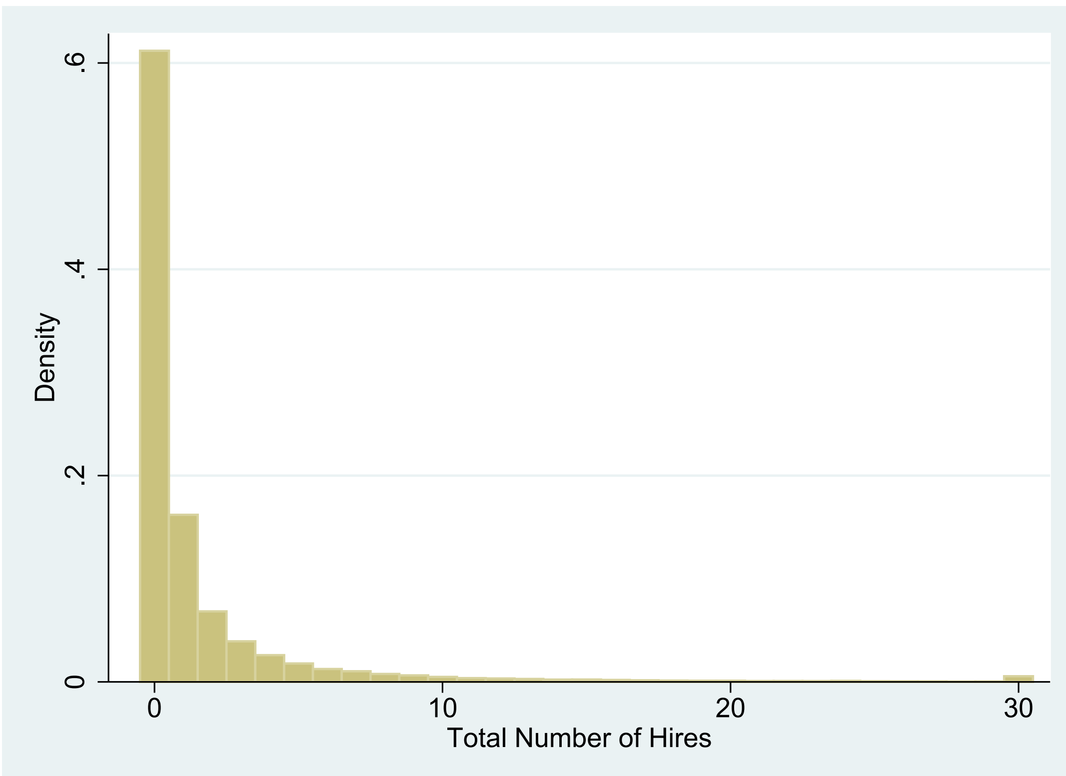

The unit of analysis is an employer. Total hires are censored at 30.

Table 1 summarizes the job openings in the data for two samples. Summary statistics are displayed separately by the number of previous hires made by the employer. Columns 1 to 5 report statistics for all job posts in the sample. Columns 6 to 10 restrict 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. This sample imposes some restrictions that are necessary to ensure features of the hiring processes observed in the data can be mapped into an employer choice problem conditional on a set of applicants. Some employers hire multiple workers simultaneously by posting a batch of jobs all at the same time; therefore, how to determine the set of applicants across job posts is unclear. Other difficulties include employers bringing workers onto the platform or targeting workers whom they have hired in the past. We restrict to job openings with at least one worker-initiated application, which 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.[10]

In both samples in Table 1, hiring behavior differs significantly based on the employer's experience at the time of posting a job. Only 22% (16% among sequential openings) of openings posted by new employers result in a hire (Columns 5 and 10).[11] 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.

Table 1 also provides summary statistics on wage bids and application counts. Columns 4 and 9 indicate wage bids fall with employer experience. We note that wage bids reflect workers' anticipated costs when applying for a job and also capture markups in excess of costs. Differences in bids with employer experience could then be due to cost factors or a consequence of demand changes that alter optimal markups. Table 1 sheds some preliminary light on whether the supply side differs by employer experience by looking at patterns of applications. Raw counts of job candidates decline with employer experience in the overall sample and increase slightly with experience in the sequential sample, suggesting the ease of attracting applicants is unlikely to explain the stark difference in hiring rates. Rates of employer-initiated candidacies are similar across experience levels.

Table 2 shows an employer's experience does not meaningfully alter the characteristics of the candidates who apply. The columns display 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 applicant pools are similar except for the hourly wage bids. The second set of five columns then repeats this exercise with the characteristics of the workers whom employers ultimately hire. Across employer experience levels, we find no major differences in resume characteristics among the workers who are hired. This summary evidence suggests little sorting across applications by employers' experience, casting doubt on the hypothesis that inexperienced employers have difficulty attracting applicants on the extensive margin.

Comparing hired workers with all applicants across columns, Table 2 also shows employers are selective, because the characteristics of hired workers differ from the other applicants. We now turn to estimating how employers evaluate worker characteristics relative to wage bids.

[10] See Appendix 1 and Appendix Table A1 for additional details about sample composition. Of the 119,877 jobs posted by inexperienced employers, 61,160 survive these restrictions, yielding 109,764 total job openings across employer experience levels.

[11] We note some employers post multiple jobs before making any hire, but 63% of employers never hire on any job post.

3 Employer Hiring Decisions

Before specifying a model of hiring, we provide motivating evidence with a figure showing how employers with different experience levels select individual workers at different wage bids. Each of the four panels in Figure 4 contains applicant hiring probabilities for a different job category, plotting the empirical probability of the applicant being chosen as a function of the residual wage bid. In each panel, the experienced employers' hiring probability function is shifted outward, and the slopes also differ by employer experience. Experienced employers tend to be more likely to hire a given worker at any residual wage bid, and they also appear to have more elastic hiring probabilities over the range of residual bids that are most common in the data. This finding suggests both the probability of hire at a residual bid and the sensitivity of the choice probability to the bid vary with employer experience. We now turn to a model that formally captures these relationships in the data.

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

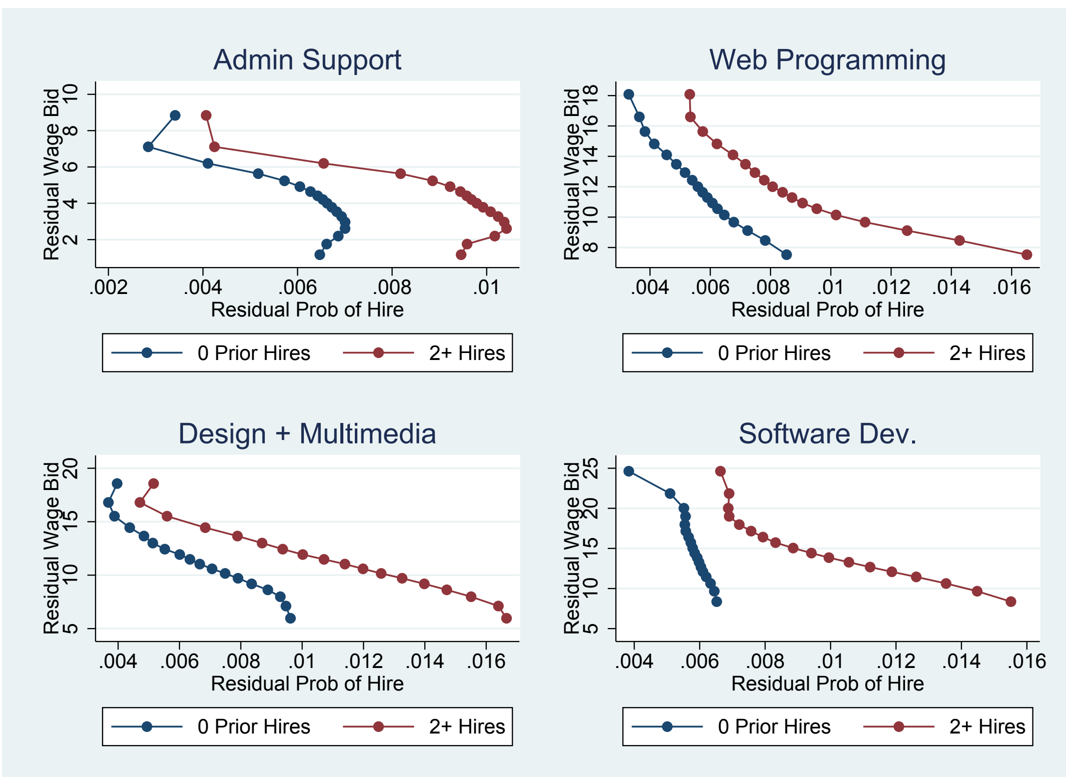

The unit of analysis is a job application. Wage bids and hiring probabilities are residualized within each job category using the worker resume data observed by the employer, a spline for the application number, and a linear time trend. Points are taken from a polynomial smoothing function of the residual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

3.1 The Employer's Problem

To examine shifts in hiring probabilities and different sensitivities to bids by experience, we specify a modified conditional logit hiring probability function with employer-experience-specific parameters. For employer i with experience χ on job opening o, the probability that applicant j is hired for the job is denoted by $p_{i\chi oj}$. Our goal is to estimate $p_{i\chi oj}$ as a function of the worker's wage bid $w_{i\chi oj}$[12] and resume characteristics and job characteristics, X_{oj}.[13] The subscript o for the job opening indicates that the probability that applicant j is hired is related to the set of all applicants for a particular job (Lazear, Shaw, and Stanton, 2018). In addition, we allow $p_{i\chi oj}$ to be a function of employer i's value of hiring any worker on the platform, μ_i, which can be heterogeneous across employers.

An employer's objective is to choose the best applicant out of J_{io}, the set of applicants for an opening, or choose not to hire for this job on the platform.[14] Because the employer is looking for a single worker, he chooses the worker who produces the most or has the highest quality per unit of wage. Quality consists of observable characteristics about the worker in the matrix X_{oj}, an experience-dependent parameter vector β_χ, the employer's heterogeneous value for hiring on the platform, μ_i, and an error term, ε_{oj}. The option of hiring no worker means the best worker's ratio of quality to wage must be greater than the value of the off-platform option, denoted option zero, for the employer to make a hire. The employer's objective function is then

$$\max_{j \in \{J_{io}, 0\}} \left\{ \max_{j \in \{J_{io}\}} \frac{\exp\left(X_{oj}\beta_\chi + \mu_i + \varepsilon_{oj}\right)}{\left(w_{i\chi oj}\right)^{\alpha_\chi}}, \exp\left(\varepsilon_{o0}\right) \right\}.$$

The errors ε_{oj} and ε_{o0} are assumed to be an idiosyncratic type-1 extreme value shock for each alternative worker j and for the outside option, with location and scale normalized. The parameter α_χ scales the importance of wage bids relative to the extreme value errors while partially determining the elasticity of hiring with respect to wages. We note the term μ_i does not appear in the payoff to the outside option, because this term is part of the employer's payoff from hiring any worker on the platform.

[12] Consistent with institutional detail about this setting, we assume the wage bid is a take-it-or-leave-it offer. Bargaining between the first offered wage and the starting wage for hired candidates appears to happen very infrequently. See Appendix 2 for additional detail.

[13] The matrix X_{oj} includes many of the details observable in a worker's profile, including feedback, country, past hours worked, and past compensation history (Barach and Horton, 2017). Job categories and the expected duration of work are among the characteristics included for jobs.

[14] 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 short time interval is small. Although we do not observe declined job offers, only 0.6% of the worker days in the sample have more than two interview requests arriving. A post-candidacy survey also asks employers for reasons 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.

Taking logs of the employer's objective yields an inequality expressing that worker j is hired by employer i with experience χ when

$$X_{oj}\beta_\chi + \mu_i + \varepsilon_{oj} - \alpha_\chi \log\left(w_{i\chi oj}\right) \geq X_{ok}\beta_\chi + \mu_i + \varepsilon_{ik} - \alpha_\chi \log\left(w_{i\chi ok}\right) \tag{1}$$

for all $k \in \{J_{io}, 0\}$. Conditional on μ_i, the probability that inequality (1) holds takes a conditional logit form, with

$$\begin{aligned} p_{i\chi oj} \;=\;& \exp\left(X_{oj}\beta_\chi + \mu_i - \alpha_\chi \log\left(w_{i\chi oj}\right)\right) / \\ & \left(1 + \Sigma_j^{J_{io}} \exp\left(X_{oj}\beta_\chi + \mu_i - \alpha_\chi \log\left(w_{i\chi oj}\right)\right)\right). \end{aligned} \tag{2}$$

The parameter vector β_χ that specifies the relationship between hiring value and job and applicant characteristics X_{oj} has an employer experience subscript, allowing employers with different experience to weigh resume characteristics differently. In addition, α_χ allows the sensitivity of wage bids to vary by experience. These features allow both the baseline hiring probability and the wage bid elasticity to vary with experience. 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 any applicant versus not hiring in the market.[15] The presence of μ_i in the hiring probability function 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.

Identification is straightforward conditional on μ_i. If μ_i were fully observable to the researcher for each employer, variation in wages, applicant characteristics, and choices would identify the parameters up to location and scale using standard arguments. Because each employer's μ_i is unobserved, only the population distribution of types can be identified. For example, if some employers repeatedly hire when faced with low-quality applicants who submit high bids, whereas other employers do not hire when high-quality workers with low bids are available, the estimated population distribution of μ types would have a wide dispersion in valuations.

We require a specification that is sufficiently flexible to capture differences that arise from employer selection versus changes in the underlying parameters that describe how employers assess workers and hiring opportunities. This requirement motivates an estimation approach that measures the extent to which the employers who return to the market after hiring once are a non-random sample from the population distribution of μ. We use a finite mixture model to allow for the flexibility to accommodate an arbitrary distribution of employer types that may change due to selection. Because the distribution among experienced employers is likely to truncate the lower tail of μ, assuming symmetry of the type distribution or stability

[15] Flexibility in the parameters by experience adjusts scaling relative to the variance of the unobserved type-1 extreme value shock.

of the distribution would not account for employer selection. In our preferred specification, employer types are fixed, but we allow the type probabilities to depend on whether the employer is ever observed to post jobs while experienced. Thus, we assume 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.

All employers who gain experience undergo the same shift in preference parameters due to experience, because $\beta\chi$ does not have a type subscript. This allows a simple partition between learning-by-doing and selection.[16] The mean of the distribution of employer valuations can change as employers gain experience, because X_{oj} has a constant term and β_χ is an experience-specific parameter vector. The distribution of μ_i is centered around the constant term in β_χ, and shifts in the constant term with experience reflect all employers' gains from experience. Allowing different type probabilities on μ_i for employers whom we ever observe posting jobs after gaining experience captures employer selection. The parameters are identified using variation in employers' choice sets and exogenous variation in wage bids.

3.2 Instruments and Identification

Worker applications contain unobserved information that may be correlated with wage bids. We address potential correlation of the error and wage bids through an instrumental variables strategy that is based on the assumption that wage bids are determined in part by worker costs, including the opportunity cost of online work. We use changes in the dollar-to-local-currency exchange rate for the applicant's country as an exogenous opportunity cost shifter. Workers are paid in their local currency for offline work, but they are paid in US dollars for their work on oDesk. Frictions limiting exchange rate pass-through to local wages mean offline opportunities are likely to adjust to exchange rates more slowly than online transactions.[17] When the local currency appreciates relative to the dollar, so that one dollar earned on the site provides fewer local currency units, workers' wage bids are predicted to increase.

Figure 5 illustrates the time-series variation in mean residual log wage bids and exchange rates that underpins this instrumental variables approach. We focus on India, which is the largest source country for workers. The top left panel plots the mean residual log wage bid for job applicants located in India and the log of the US dollar to Indian rupee exchange rate. To control for secular trends and level differences in local exchange rates across countries, we detrend each series and remove its time series mean. The mean bid and exchange rate time series move together. The other panels in Figure 5 focus on relative

[16] We also experimented with more complicated models in which the parameters have a more general dependence on employers' type, but this added complexity did little for model fit while adding a large number of additional parameters. These specifications show selection out of market is not related to preferences for worker characteristics or sensitivity to wages.

[17] 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.

bid variation across source countries as a function of local exchange rate differences. These panels 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). The crosses in the remaining panels display differences in the local currency and the Indian rupee exchange rate. These figures indicate that exchange rate variation across applicant countries induces variation across applicants' wage bids.[18]

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

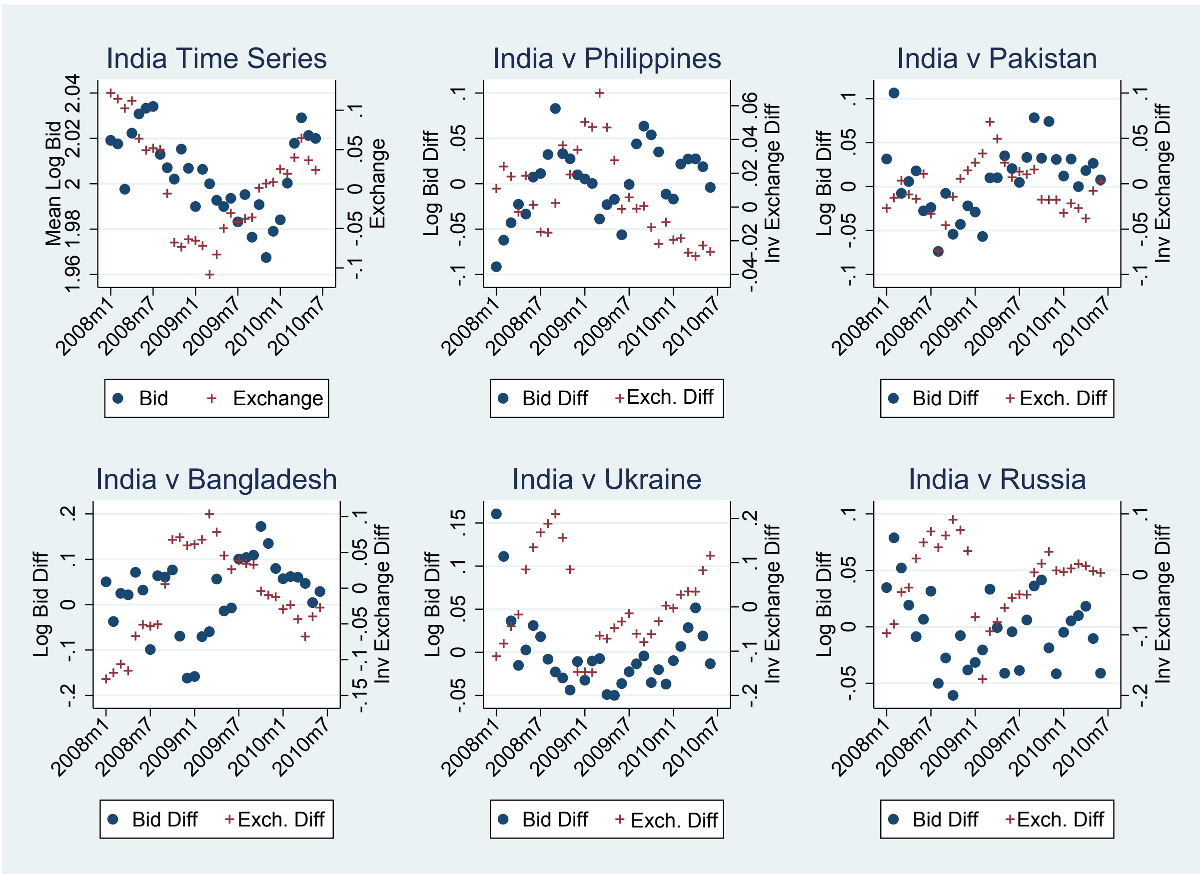

The top left panel plots mean residual log bids against the log of the US Dollar to Indian Rupee exchange rate after removing a time trend. The remaining panels plot log bid differences between India and other countries (left y-axis) and the log of other currency to the Indian exchange rate (right y-axis).

Although exchange rate movements are plausibly exogenous to demand on oDesk, several additional issues are relevant for identification. 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 this instrument. This issue is not a problem for comparisons between workers, because other applicants will have exchange rate variation, shifting relative bids. However, these applicants have no variation relative to the employer's outside option

[18]Each panel is based on differing numbers of observations; Indian workers are about 40% of the sample, whereas Russians and Ukrainians submit under 3% of the total observed bids.

of not hiring. We construct a second instrument for bids that is relevant to all workers, including those with dollar or dollar-pegged local currencies. This instrument captures exogenous variation in the intensity of competition for a vacancy. The ideal variation would alter the extent of perceived competition by other applicants, embodied by situations in which, for example, a subset of potential applicants randomly lost internet access and this supply shock was observable to applicants who remained available. The instrument is based on a similar idea, exploiting the fact that workers can observe counts of other applicants to jobs.

We leverage the extent of competition for jobs in the same category and week as a measure of likely competition for the job in question. We calculate the instrument using the leave-own-opening-out average of the number of job applications in the first 24 hours since posting for the category-by-week cell. This instrument is close to the ideal instrument because potential applicants can observe how many competitors are applying using platform-specific dashboards that summarize openings in a job category. However, the exclusion restriction is vulnerable to demand shocks or time series variation in market conditions. To mitigate these potential violations of the identifying assumptions, we remove job category and time fixed effects. The residuals after removing these fixed effects do not contain time-series variation in the number of jobs or applicants; instead, they capture competition differences across job types in the same week while holding fixed the average competition in the category over time and the average competition across all job categories during the week in question.

An additional concern is that the instruments themselves may influence 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. 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 include a comparison of the first stage and demand estimates with and without worker-level resume data. This approach allows an assessment of whether worker observable characteristics are correlated with the instruments and enables us to comment on the direction of bias in the demand estimates that may result from unobserved correlation.

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_{oj}, and worker characteristics, X_{oj}, in a first stage regression of the form

$$\log(w_{ioj}) = \gamma_0 + Z_{oj}\gamma_{1\chi} + X_{oj}\gamma_{2\chi} + \nu_{i\chi oj}. \tag{3}$$

The coefficients in Equation 3 are estimated separately for the group of new and experienced employers. Later, we derive this first stage as the worker's optimal bid from her supply problem. For additional detail, see Footnote 28.

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_{oj} in Equation 3 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 are larger in absolute magnitude in the columns that exclude the worker characteristics, suggesting some sorting into applying occurs.[19] Later, estimation of the hiring probability function with and without worker characteristics will help us understand how this sorting affects the estimates of demand for each employer segment.

3.3 Estimation

This section presents the likelihood over sequences of employer choices across different job posts.[20] The step-by-step approach is as follows. First, the residuals from Equation (3) form control functions for unobserved worker quality, denoted $CF_{i\chi oj} = \hat{\nu}_{i\chi oj}$. Second, we form choice probabilities conditional on a value of the unobserved term μ_i, taking the form:

$$
\begin{aligned}
p_{i\chi oj} \quad = \quad & \exp\left(X_{oj}\beta_\chi + \mu_i - \alpha_\chi \log\left(w_{ioj}\right) + \psi_\chi CF_{i\chi oj}\right) / \\
& \left(1 + \Sigma_{j=1}^{J_i} \exp\left(X_{oj}\beta_\chi + \mu_i - \alpha_\chi \log\left(w_{ioj}\right) + \psi_\chi CF_{i\chi oj}\right)\right). \qquad (4)
\end{aligned}
$$

Third, we assume 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.[21] 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 $\beta_{0\chi=I} + \mu_2$ prior to hiring 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. The values of μ_2 and μ_3, and the associated type probabilities, tell us the extent of heterogeneity in value for the market among each employer segment.

[19] Because many workers submit only a small number of total applications, including measures of worker heterogeneity in the first stage, or assessing the extent of selection into applications based on exchange rate variation, is hard. Appendix Table A2 contains a sensitivity analysis for the first stage, providing results that from regressions that include job applicant fixed effects. The instruments remain strong in these specifications. However, we note the residuals from the first stage with applicant fixed effects cannot be used in estimation due to the incidental parameters problem that arises because control functions including fixed effects are not the result of a consistent estimator.

[20] In specifications that omit employer heterogeneity, the likelihood is specified at the job-opening level and μ_i is set to zero.

[21] The difference in this term between inexperienced and experienced employers will form our estimate of learning-by-doing, that is, the increase in the average value for the platform that employers gain from having hired on the site before.

Fourth, we allow 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 $\rho^{S=E} = (\rho_1^{S=E}, \rho_2^{S=E}, \rho_3^{S=E})$. 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, and the estimates of the ρ parameters tell us how the distribution of employer types varies with the choice to become experienced.

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.[22] The marginal likelihood for an employer who ever posts a job in the experienced segment is the sum over the three types weighted by the probability of that type:

$$L_i = \Sigma_{k=1}^{3} \rho_k^{S=E} \Pi_o \Sigma_{oj} p_{i\chi oj} \left(\mu_k\right)^{y=j} .$$

The term $\Sigma_{oj} p_{i\chi oj} \left(\mu_k\right)^{y=j}$ is the conditional choice probability for the actual choice made on a particular job opening, summing across the possible chosen alternatives. We then take the product over the sequence of choices across job openings for the employer, as the product is taken over o. The type weights differ depending on whether the employer is observed in the experienced segment, and the likelihood contribution for employers who are never observed as experienced employers is

$$L_i = \Sigma_{k=1}^{3} \rho_k^{S=N} \Pi_o \Sigma_{oj} p_{i\chi oj} \left(\mu_k\right)^{y=j} .$$

4 Results

4.1 Learning-by-Doing and Selection

The estimates indicate significant changes in demand between inexperienced and experienced employers due to both learning-by-doing and selection out of the platform. Learning-by-doing suggests experience is necessary to unlock some benefits of remote hiring, whereas selection out after early experience suggests employers enter with initial uncertainty about their value for using the market. Table 4 presents the results for the hiring probability estimation using the sequential openings sample. Panel A displays the estimated heterogeneity in employer types for never-experienced and ever-experienced employers. Our preferred specification is in Columns 5 and 6, whereas Columns 1 and 2 present a restricted model without employer heterogeneity, and Columns 3 and 4 present the model without worker resume characteristics. Employer types and their corresponding probabilities for the group of employers who don't return as

[22] For further details, see Train (2009).

experienced employers are displayed under columns for inexperienced employers (Columns 3 and 5) and the ever-experienced are in the columns for experienced employers (Columns 4 and 6).

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 in the experienced sample, and helps to visualize employer heterogeneity and selection into posting further jobs. This figure uses the parameter estimates in Panel A of Table 4 as well as the estimated share of employers of each type in each employer group, also given in Panel A. The μ distribution of employer valuations varies by experience, and a clear pattern of selection exists, in which type-2 employers, those with low values of μ_i, are much less likely to be observed in the ever-experienced group. Thirty-four percent 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

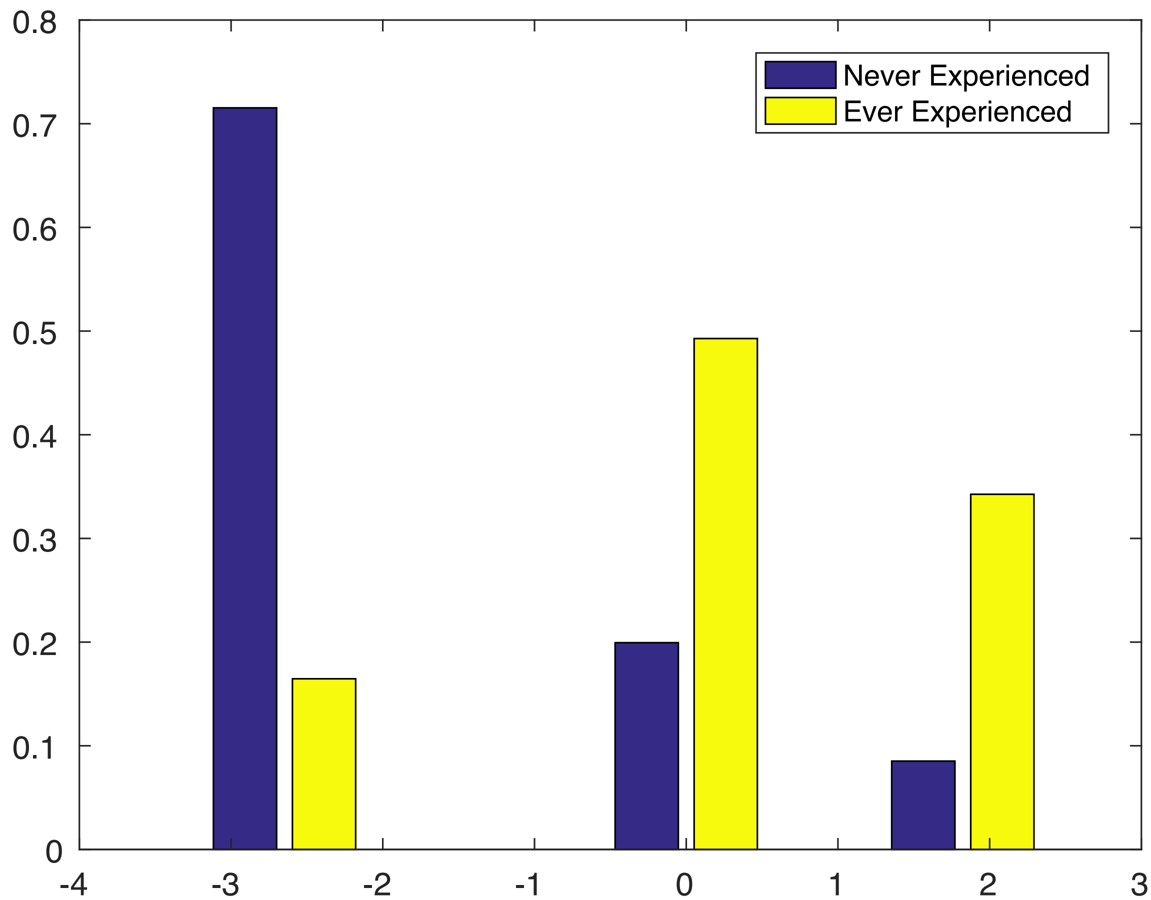

Type probabilities for employers based on whether they are ever observed in the experienced segment. Employers who are inexperienced and eventually transition to the experienced segment are always classified using the "Ever Experienced" type probabilities.

Quantifying the extent of both learning-by-doing and selection, Panel B of Table 4 provides estimates of how employer valuations change with experience through a decomposition. We term $E(X_{oj}|\chi)\beta_\chi + E(\mu|\chi)$

the log productive value of hiring an applicant with characteristics X_oj for employers with experience χ and decompose its difference between experienced and inexperienced employers. Differences in the log productive value of hiring arise due to changes in X, due to changes in β, and due to changes in μ between job postings by the inexperienced and by the experienced. Changes in X reflect sorting of workers into jobs or changes in the characteristics of job postings. Changes in β suggest changes in employer perceptions of the platform or changing weights on applicant characteristics, and changes in μ indicate employer selection out of the market based on their value of hiring online.

Using the logic of decomposition estimators (Oaxaca, 1973; Blinder, 1973; Fortin, Lemieux, and Firpo, 2011), let I and E subscripts denote the X matrix and parameter vector β for inexperienced employers and experienced employers, respectively. The difference in log productive value 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-doing, 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).[23]

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

$$E(X_E\beta_E + \mu_E - \alpha_E\overline{log\left(w_{ijE}\right)}) = E(X_I\beta_I + \mu_I - \alpha_I\tilde{W}). \tag{5}$$

This measure captures the log-wage bid reduction to inexperienced employers that would equate the probabilities of hiring a given applicant for experienced and inexperienced employers. It shows substantial bid-equivalent increases in the value of the market for the experienced. The specification that omits employer-type heterogeneity (Columns 1 and 2) yields an 11-log-point bid-equivalent increase in employer value for the platform from experience. When allowing for heterogeneity across employers, we estimate that inexperienced employers would have to receive bids that were around 30 log points lower to equate their value for the platform with that of employers with experience. This amount is an average reduction of about \$2.80 per hour, a substantial percentage of average bid levels.

In the models that include worker characteristics, the majority of the change by experience level is due

[23]For experienced employers, this calculation is straightforward and simply uses the type probabilities for the ever-experienced group of employers. The openings posted by inexperienced employers include those posted by employers who never become experienced and those that do. So, for the inexperienced openings, both distributions in Figure 6 are used, along with a weighted average of type probabilities that corresponds to the fraction of inexperienced employers who eventually become experienced.

to changes in the estimated β coefficients—that is, the learning-by-doing 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. Seventy-nine percent of the total increase in log productive value, or a wage shift of about 24%, is due to learning-by-doing. Thirty-two percent of the change in log productive value is due to the employer composition effect of employers with low values of μ selecting out of the market. The employer composition change is equivalent to a wage shift of about 9%.

Do employer characteristics, unobserved by applicants but visible in the administrative data, explain the estimated heterogeneity in market value? To answer this question, we first compute the posterior type of an employer, denoted $\hat{\mu}_i$, using Bayes' rule.[24] We then regress $\hat{\mu}_i$ on observable characteristics of the employer. Some observable employer characteristics, such as geographic location, are populated in the database for most employers, but other characteristics are contained only in sparse employer surveys. We, hence, explored several different models: (1) We regressed $\hat{\mu}_i$ on state fixed effects for US-based employers, asking whether geographic variation contributes to value for the platform. (2) We then reduced the sample by adding survey information about the headcount of regular workers at the employer's firm. State fixed effects lead to a large number of parameters, so we discipline the analysis of across-state heterogeneity using a LASSO approach with a cross-validated penalty term. Most state fixed effects do not survive the LASSO; when they do, small cell sizes are responsible. Firm-size dummies remain.

Table 5 presents the results of regressions that capture the spirit of the LASSO exploration, but here the approach allows for the computation of standard errors. Columns 1 and 2 report a linear probability model where the dependent variable is that the employer has adopted the market, as indicated by posting jobs after gaining experience. The coefficients capture how the probability of market adoption changes with employer size. In Column 1, employers with 10 or more internal workers are 7 percentage points (25 percent) less likely to adopt compared to sole proprietors and individuals (the excluded category). It is not the existence of internal workers that drives the negative result, as employers with between 2 and 9 workers have similar estimates to individuals. The estimates in Column 2 change little with geography fixed effects. This negative pattern between market and adoption and firm size appears to be driven by a lower value for the market, as captured from our estimate of $\hat{\mu}_i$. Columns 3 and 4 change the dependent variable to $\hat{\mu}_i$ and show that it is the baseline value for the market (estimated from exogenous wage variation) that is lower for larger employers. Relative to employers with no internal employees, the average $\hat{\mu}$ is over 28 log points lower for employers with 10 or more internal workers.

The negative relationship between an employer's value for the market and firm size may be due to several underlying reasons. Large employers may have difficulty integrating online hiring into their existing

[24] This calculation is $\hat{\mu}_i = \Sigma_{k=1}^{3} \mu_k \frac{\rho_k \Pi_{io} \Sigma_j p_{i\chi oj}(\mu_k)^{y=j}}{\Sigma_{k=1}^{K} \rho_k \Pi_{io} \Sigma_j p_{i\chi oj}(\mu_k)^{y=j}}$.

processes. There is substantial literature showing there is a firm size wage premium (Oi and Idson, 1999), which may put pressure on large firms to save wages. Our results suggest that jobs at large firms are less likely to be outsourced online even in the face of mounting domestic outsourcing (Goldschmidt and Schmieder, 2017; Bloom, Guvenen, Smith, Song, and von Wachter, 2018). This suggests that the overhead costs of specifying and hiring for individual tasks is less likely to be worthwhile when tasks are frequent and somewhat standardized. Alternatively, large employers may have easily accessible sources of internal labor resources that are unavailable to smaller employers or single individuals.

4.2 Wage Elasticities and Interpreting Workers' Bids

This section presents a framework for interpreting how hiring elasticities influence applicant bidding behavior. Elasticities are reported in Table 4, Panel C. Inexperienced employers are less elastic than experienced employers, indicating a greater sensitivity of the hiring probability with respect to wage bids.

Our starting point is to model worker bidding using tools for analyzing competition between heterogeneous producers that have some local market power. This assumption is reasonable because the data suggest workers appear differentiated to employers. For example, hiring the worker who submits the lowest hourly bid is relatively rare for employers.[25] As a result, we view workers as akin to "differentiated products," and the demand estimates allow us to characterize how employers trade off wage bids and other characteristics. Implicit is that workers with different attributes compete to be hired through their hourly wage bid choices. Several resulting implications result for the supply side of this labor market. Recall that the summary statistics provided in Tables 1 and 2 show that wage bids decline with employer experience. We examine the extent to which lower bids are due to differing hiring elasticities.

4.2.1 The worker's problem

We interpret observed wage bids as optimal functions of workers' costs and the demand elasticity for each worker in a Betrand-Nash oligopoly. We then use the estimated elasticities together with data on wage bids to back out the implied markups included in the bids over workers' costs. Worker j's cost, $c_{i\chi j}$, captures her outside option and her expected costs of 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.[26] Her wage bid also takes into account the ad-valorem fees

[25] In the Appendix, Figure 9 shows the distribution of the wage bid decile, calculated within each job opening, of the worker who was ultimately hired.

[26] In our setting, the wage bid is the worker's only strategic variable. We assume the fit between the employer/job and the worker is unknown at the time of application. This partially differs from the broader model of quality and price choice in monopolistic competition with experience goods in Riordan (1986).

retained by the platform, τ. The worker chooses the wage bid, $w_{i\chi oj}$, that maximizes

$$\underbrace{p_{i\chi oj}}_{\Pr(hired)} \times \underbrace{\exp\left(\log w_{i\chi oj} - \log(1+\tau)\right)}_{Post-fee\ wage} + (1 - p_{i\chi oj}) \times \underbrace{c_{i\chi j}}_{Cost}, \tag{6}$$

where $\log w_{i\chi oj}$ is the log of the wage bid inclusive of τ. If she is hired, the worker receives the wage $\frac{w_{i\chi oj}}{(1+\tau)} = \exp\left(\log w_{i\chi oj} - \log(1+\tau)\right)$ and the employer pays $w_{i\chi oj}$. 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 or working on the job.[27] The worker's first order condition is given by

$$\frac{\partial p_{i\chi oj}}{\partial \log w_{i\chi oj}} \left(\frac{w_{i\chi j}}{(1+\tau)} - c_{i\chi j} \right) + p_{i\chi oj} \frac{w_{i\chi oj}}{(1+\tau)} = 0. \tag{7}$$

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

$$w^*_{i\chi oj} = c_{i\chi j}\,(1+\tau) \left(1 + p_{i\chi oj} / \frac{\partial p_{i\chi oj}}{\partial \log w_{i\chi oj}} \right)^{-1}. \tag{8}$$

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

$$c_{i\chi j} = \frac{w_{i\chi oj}}{(1+\tau)} \left(1 + p_{i\chi oj} / \frac{\partial p_{i\chi oj}}{\partial \log w_{i\chi oj}} \right), \tag{9}$$

illustrating how having an estimate of $p_{i\chi oj}$ and $\frac{\partial p_{i\chi oj}}{\partial \log w_{i\chi oj}}$, together with the bids and platform fees observed in the data, yields worker-job specific estimates of costs and markups.[28]

---

[27] The objective could alternatively be written: $\max_{\log w_{i\chi oj}} p_{i\chi oj} \times \exp\left(\log w_{i\chi oj} - \log(1+\tau) - c_{i\chi jH}\right) + [1 - p_{i\chi oj}] \times c_{i\chi jO}$, where $c_{i\chi jH}$ is a 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.

[28] This optimal bid equation offers more insight as to how the exchange rate instrument described in Section 3 affects workers' bids. Assume that $c_{i\chi j}$ is denominated in the local currency, whereas 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 $w^*_{i\chi oj} = c_{i\chi j} \left(\frac{D}{L}\right)^{\theta} (1+\tau) \left(1 + p_{i\chi oj} / \frac{dp_{i\chi oj}}{d \log w_{i\chi oj}} \right)^{-1}$. The dollar-to-local-currency exchange rate is $\frac{D}{L}$, and the parameter θ captures possible reasons for deviations from complete pass through. These reasons include include the following: (i) Some part of a worker's opportunity cost reflects transactions denominated in dollars rather 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. The worker's optimal log bid can be written as a mapping from local-currency-denominated opportunity costs to dollar-denominated bids, as

$$\log\left(w_{i\chi oj}\right) = \theta \log\left(\frac{D}{L}\right) + \log\left(c_{i\chi j}\right) + \log\left(1+\tau\right) - \log\left(1 + p_{i\chi j} / \frac{\partial p_{i\chi oj}}{\partial \log w_{i\chi oj}} \right), \tag{10}$$

which forms the first stage regression for the hiring probability estimation.

4.2.2 Elasticity estimates

We find the wage bid decline with employer experience can primarily be attributed to differences in demand—that is, differences in the optimal bid and markup arising from elasticities. The coefficient on the log hourly bid, α_χ, differs by employer experience, as shown in Panel C of Table 4. 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 additive deviations for experienced employers relative to the baseline estimate for inexperienced employers. That is, the experienced employer columns contain parameter estimates on interactions with experience. A comparison across each pair of columns shows α_χ is significantly larger in magnitude (more negative) for experienced employers, and these employers have significantly larger hiring wage elasticities.[29]

Columns 1 and 2 include the estimates after including the control function from the first stage estimates but without allowing for employer heterogeneity in type. We comment first on estimates of elasticities with respect to individual worker bids. Inexperienced employers have an estimated wage elasticity of -5.46, whereas experienced employers are more wage elastic, with an estimated elasticity of -8.54. Columns 3 and 4 exclude worker resume characteristics from the first stage (corresponding to Columns 3 and 4 of Table 3), and Columns 5 and 6 include these characteristics (as per Columns 1 and 2 of Table 3). In the specifications in Columns 5 and 6, 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 elasticities imply differences in markups over workers' opportunity costs. The estimated average markup over cost included in inexperienced employers' bids is 25.2%, compared to 14.9% for experienced employers. Referring again to Columns 5 and 6 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. Thus, the estimates show only a small additional expected cost when applying to an inexperienced employer. The results suggest 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 5 and 6, around 88% of the observed wage-bid premium to inexperienced employers can be attributed to the higher markups set by workers who anticipate that new employers have relatively inelastic hiring probabilities. Only 12% is due to the higher expected costs when applying to new employers.

If costs of applying to an inexperienced employer are similar but markups are higher, why doesn't entry of additional applicants drive down markups? We contend application order affects whether employers consider an applicant, with late arrivals much less likely to be hired. Included in the X_{oj} matrix is a

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

flexible spline for the applicant's position in the order of candidates who arrive, with knots that correspond to the pagination order that would require an employer to click on a new page. The elasticity estimates control for arrival position of an applicant, which is known to an applicant when making a bid. The parameter estimates show that later applicants, after the first 30, are much less likely to be hired. This finding is consistent with the relationship in Figure 7, which shows employers do not tend to consider late-arriving applicants, as measured by interviews or hires. Employers' probability of choosing later applicants falls faster than the prospective benefits of competition from more applicants, allowing positive markups.

Figure 7: Probability of Being Interviewed or Hired, by Applicant Order

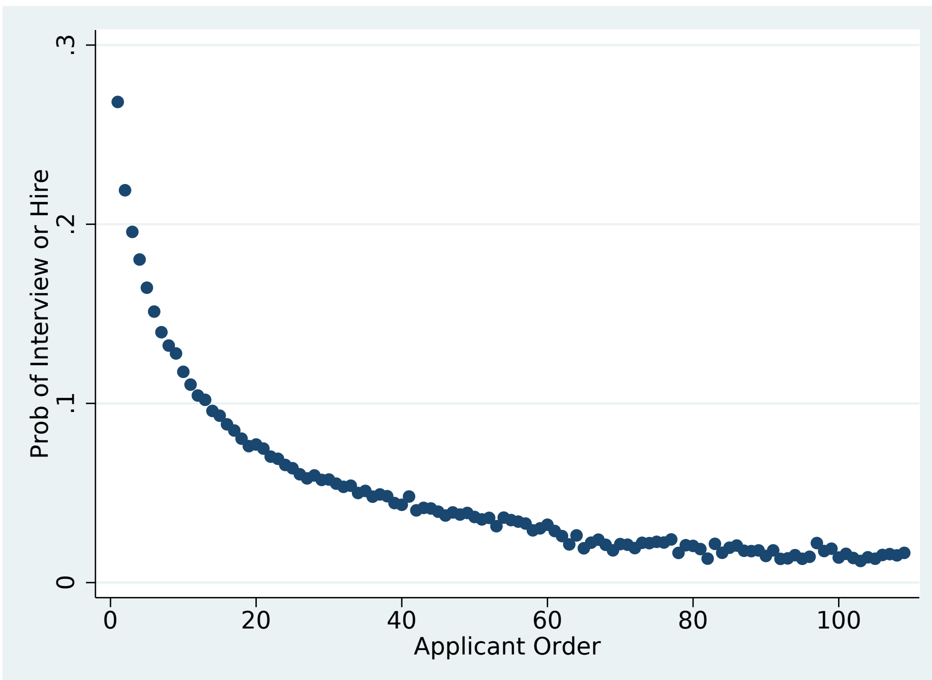

This figure plots the probability an employer either interviews or hires a worker as a function of the order in which she applies. Earlier applicants are much more likely to land interviews or jobs.

The final row of Table 4 presents estimates of the elasticity of vacancy fill with respect to a change in wage bids for all job applicants. These estimates are calculations of the elasticity that the employer hires any worker for this job posting with respect to a percentage increase in wage bids from all workers. Comparing Columns 1 and 5 demonstrates the role of employer heterogeneity in estimating the elasticity of the hiring decision for inexperienced employers. Controlling for the increased variance of employer type among the inexperienced reveals new employers' vacancy fill rate is relatively inelastic, at 2.93. Importantly, comparing the odd- and even-numbered columns in the table, experienced employers are found to be more elastic than the inexperienced on the extensive margin of hiring any worker across all specifications. The difference in vacancy-fill wage elasticity is greatest in Columns 5 and 6, which is the specification that controls for both varying employer heterogeneity with experience and the full set of worker resume characteristics.

This relative inelasticity of new employer hiring will become relevant later when discussing counterfactuals related to changes in platform fees.

4.2.3 Search effort and wage elasticity differences

Why are experienced employers' hiring probabilities relatively wage elastic? We consider explanations based on experienced employers' ability to recall prior job applicants. We note that experienced employers interview fewer applicants before hiring, reducing their search effort. We then show experienced employers have a stock of prior applicants to whom they can turn rather than evaluating new applicants. This stock provides latent competition for a job and is a potential reason for the decline in observed search effort with experience. A related possible explanation behind the changing elasticity of hiring is that declining interviews suggest employers have a different perception of how to evaluate workers because of differences in the precision of signals received from job applicants.[30] Some models, as discussed in (Englmaier, Schmöller, and Stowasser, 2017), focus on inattention or limited information processing. In our context, the weights on resume characteristics are allowed to change, capturing some information processing differences.

Consistent with an Experience Market, if employers are using interviews to learn about their own value for the market as a whole, the marginal benefit of an interview is higher when employers know relatively little about their value for the market. It is therefore optimal for inexperienced employers to interview more applicants. Appendix Table A3 presents the results of a regression of (1+ the log of) the number of interviews conducted for each job on the number of previous hires made, along with various controls. In all specifications with employer fixed effects, the number of interviews decreases, and at a decreasing rate, as the cumulative number of prior hires increases. 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.

At first glance, these results might suggest that experienced employers are sampling or considering fewer applicants, which would arguably make them less elastic when considering a worker. However, a different interpretation is that employers simultaneously interview fewer applicants and become more elastic with experience because they have built up a stock of prior applicants who might be latent competitors for the job. An experienced employer is likely to know the productive value of previously hired workers, while also having a relatively precise signal about each previously interviewed worker. If new applicants to later jobs understand they are competing not only with all other new applicants, but also with the workers that are already known to the employer, they will act as if they are part of a more competitive applicant pool. Our

[30] However, attributing differences in the elasticity to employer learning would require a relatively complicated setup compared to the seminal papers on learning and information acquisition in labor economics. Standard learning models feature normal signals and normal noise (Lange, 2007; Kahn, 2013; Kahn and Lange, 2014). If employers are learning about the market as a whole, any noise in the difference between workers (which is relevant for pinning down workers' market power) can be filtered by the employer. We thank Fabian Lange for pointing this out.

sampling approach shuts down repeat hiring, but the sequential sample does include hires of applicants who were interviewed for a previous job opening.

We look for direct evidence that employers retain access to the workers that they interacted with on prior job postings. The data show that in the 38,342 cases in which a previously interviewed applicant is an applicant for a subsequent job posting, 49% of these applications were initiated by the employer. This finding contrasts with the 8% of applications that were employer-initiated among applicants that were not previously interviewed. Experienced employers thus appear to "know where to look" and do return to access their stock of previously interviewed candidates. Of the 19,670 experienced employer hires, 1,408 of the hires are workers who were interviewed for a previous job posting. Access to this latent supply of workers serves as a further benefit to being experienced in the market—subsequent jobs are perceived by new applicants as being more competitive.

4.3 Worker Sorting and the Instrument

We now assess the sensitivity of the identifying assumptions behind the demand results. We explore possible worker sorting on the instrument as it relates to differences in the estimates for inexperienced and experienced employers. To check whether worker sorting concerns or omitted variables are driving the results, we compare the estimates in Columns 3 and 4 of Table 4, the estimates without worker resume characteristics, with the estimates given in Columns 5 and 6 that include these characteristics. 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 finding is intuitive because the matrix X_{oj} includes worker characteristics in Columns 5 and 6; with the additional characteristics, changes in β_χ have more weight, reducing the weight placed on changes in the distribution of μ_i between employer groups.

Also relevant to our understanding of employer demand is the impact of the omission of workers' resume characteristics for the elasticity estimates produced. Including worker characteristics in the model, the estimated markups for the inexperienced and experienced are 25.2% and 14.9%, respectively. Omitting these characteristics results in a smaller estimated markup difference between inexperienced and experienced employers. This comparison suggests that had we been able to include worker characteristics that are unobserved to us, the estimated markup differences would have been even larger than those in Columns 5 and 6. Unfortunately, we cannot estimate a model that includes worker fixed effects in the first stage, due to the incidental parameters problem.[31] If sorting into applying based on worker fixed effects goes in the same direction as the sorting based on observable worker characteristics, for which we can control,

[31] As discussed previously, many of the workers observed in the sample make only a single bid. Small numbers of bids mean 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.

our estimate of the difference in markups to inexperienced and experienced employers may be downward biased.

5 Analysis of Supply-Based Explanations Using Variation in Workers' Bids

The analysis so far attributes almost all the changes in outcomes between inexperienced and experienced employers to demand-side explanations. Our data permit further investigation of alternative supply side explanations, using, primarily, information about wage bids. We assume worker-level wage bids reflect workers' costs and benefits from applying for a job, and our approach asks whether workers are willing to submit different bids because employer experience itself is correlated with attributes that change workers' expected costs. These attributes might include relationship duration, the propensity of employers to leave good feedback, or the risk due to lack of employer reputation. Several other potential explanations are unobservable in the data. For example, how an employer writes a job description may change with experience in a way that affects workers' anticipated costs of applying. To address the role of unobservables, we validate our estimates with evidence from a small-scale field experiment in which we varied observable employer experience in job postings.

5.1 Wage Bids Decline with Employer Experience

Table 6 motivates the supply-side analysis by presenting regressions of the log of applicants' hourly wage bids as a function of employer experience. Consistent with Tables 1 and 2, experienced employers receive lower bids. 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 includes 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. The observed decline in average wage bids received by employers with five or more prior hires is 5.9%. The monotone decline in bids in both specifications is consistent with more experienced employers having sampled a larger stock of prior workers who may be latent competitors for the job.

The subsequent columns of Table 6 show these results remain within employer and within worker. Exercises with employer fixed effects further indicate whether sorting or unobserved heterogeneity (to the econometrician, but not to workers) drives workers' bidding decisions. In Columns 3 and 4 with employer fixed effects, bids are 4.7% lower when an employer has five or more prior hires than when he has no experience. This decline suggests within-employer change, rather than permanent differences in what workers observe about job postings, drives the declining bids. The bid reduction remains when adding worker and job characteristics. 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. We now evaluate and rule out several plausible supply-side reasons for the observed within-worker and within-employer decline in wage bids with employer experience.

5.2 Job Characteristics, Employer Attributes, and Variation in Wage Bids

We first consider the possibility that workers view a job application to an experienced employer as having better long-term prospects than an application to an inexperienced employer. For this effect to explain higher bids to inexperienced employers, workers must be willing to take a short-term wage reduction to secure this relationship. Table 7, which presents data on the empirical distribution of contract duration, shows that experienced employers are not more likely 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, whereas 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, and 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 jobs started with inexperienced employers do not result in shorter relationships than those started with more experienced employers.[32] We also find no evidence for ex-post differences in bargaining power, performance-based pay, or renegotiation after a job starts. Columns 3 and 4 evaluate whether applicants prefer working for experienced employers because they are more likely to provide raises after the contract begins, and the findings show no differences by employer experience.

Alternatively, an applicant might expect the feedback she receives to differ by employer experience. Receiving good feedback is helpful to workers' careers, and could therefore motivate worker bid differences.[33]

[32]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.

[33]For example, in an experiment on the provision of worker-level information, Pallais (2014) shows the revelation of public feedback about workers is beneficial to their later careers on this platform. Stanton and Thomas (2016) show the effect of worker feedback is concentrated among workers for whom employers have the least information. Similar issues may be at play

Columns 5 and 6 of Table 7 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 1% higher than the score given by employers with no previous hires.[34] 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 6.

Table 8, Panel A, 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 would have required some effort from applicants at different times over the platform's evolution, they could have navigated to the history of feedback that employers left for workers on prior jobs. We show that employers who left low feedback (possibly because of miscalibration with the distribution or because of a bad initial match) did not receive higher subsequent wage bids. In fact, employers who left good feedback for workers received higher wage offers later, inconsistent with poor feedback risk driving the wage-bid results.

Third, we show that bids actually *increase* with the interaction of employer experience and the feedback an employers has received from past workers. After observing good bilateral feedback about an employer, future 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. Table 8, Panel B, 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, 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 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 a risk-premium channel that would predict lower bids for employers with good feedback.

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.

[34]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.

These results offer further evidence that workers' bids are sensitive to the likely elasticity of hiring probability, because experienced employers with good feedback are more likely to hire, and are thus less wage elastic than other experienced employers. Poor employer feedback appears to signal a reduced likelihood of hiring, lowering the optimal bid to these employers. 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 by 3.7 percentage points (roughly 8%) for experienced employers with no feedback.[35] Workers appear to tailor their wage bids to the information revealed about employers through their feedback scores, because these measures are informative about employer future hiring decisions rather than about the costs of working for these employers.

5.3 Unobservable Employer Attributes, Wage Bids, and Interview Decisions

We next assess whether employer experience is correlated with some aspect of the job posting or hiring process that affects workers' application decisions. Appendix 2 shows the bid premium to inexperienced employers does not arise due to differences in the number of *observed* job applicants.[36] However, the number of latent competitors through the stock of prior experience may cause applicants to anticipate additional competition controlling for the extent of organic applications to the jobs.

Bid differences with experience also do 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 unobservable changes shows that wage bids do decline with employer experience in a controlled setting with identical job descriptions. Appendix 2 provides more details about the experiment, and Appendix Table A5 presents the results.

We interpret the evolution of hiring and selection as arising from information that arrives after employers enter the platform rather than something that is known to employers ex ante. Supporting this interpretation, Figure 8 asks whether the identity of the first candidate interviewed differs based on whether the inexperienced employer subsequently conducts many or few interviews or goes onto hire. If the candidates who are evaluated first look similar but subsequent choices differ across employers, it suggests intermediate information is allowing employers to learn about their value for the market. The figure plots the distribution of the hourly wage bids of the worker selected for the first interview, based on an employer's eventual action. The split in Panel A is based on whether the employer does more or fewer than five total interviews

[35]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.

[36]The premium remains when controlling for the number of job applicants in the first 24 hours after job posting, as shown in Appendix Table A4.

on the first job. Little difference exists in wage bids of the first interviewee for employers who go on to interview few or many applicants. Panel B of Figure 8 plots residual bids as a function of whether the employer hires for the first job, again showing the wage bid for early interviewees is unrelated to whether an employer makes a hire. These results are consistent with the hypothesis that employer actions are shaped by what they learn during the hiring process rather than by unobserved information that the employer knows before entering the market.

Figure 8: Residual Hourly Bids by First Applicant Selected for Interview

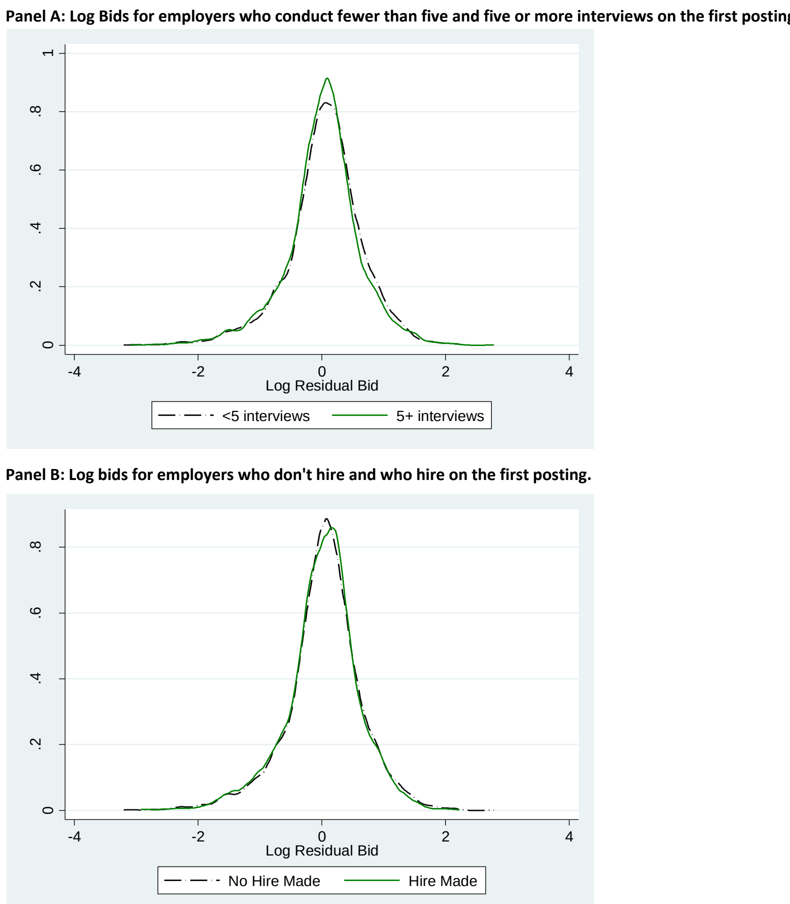

The figure plots the density of residual hourly wage bids of the worker selected for the first interview, based on an employer's eventual action. Plotted wage bids are net of job category and calendar time fixed effects. The split in Panel A is based on whether the employer does more or fewer than five total interviews on the first job. The split in Panel B is based on whether the employer hires.

6 Counterfactual Analysis

In Experience Markets, where buyers learn about their own value for the market through interactions with sellers, continued market participation depends on what buyers infer from these early interactions. Focusing on the employers who have not previously hired, Table 1 shows 78% of the jobs posted by inexperienced employers go unfilled. Thirty-five percent of employers hire once, and only 25% of all employers who enter the market hire and then post jobs as experienced employers. These statistics suggest a large share of employers receive signals that lead them to negatively update their beliefs about their own value for using the market. At the same time, inexperienced employers receive higher wage bids than their experienced counterparts, and these higher bids may deter the potential realization of future value that accrues with experience. In this section, we explore how the platform could encourage initial hiring by altering the platform fee schedule to depend on the employers' experience with hiring. We assess how platform profits change for a variety of different fees to inexperienced and experienced employers.

From its founding through the end of the sample period, the oDesk fee was constant, at 10% of wages.[37] We implement counterfactual analysis of different fees using a model-based assessment of how changes in these fees would affect the number of employers who become experienced and their propensity to hire initially and thereafter. We take the estimates of demand by employer experience as a sufficient statistic for how employers hire, and then we adjust the parameters based on various different assumptions about what employers anticipate in the future.

An important aspect of this exercise is based on the premise that information is revealed with experience. Although our estimates do not model information acquisition, we take the type heterogeneity among inexperienced employers as an approximation of the distribution of signals that arrive to these employers. We model how inexperienced hiring and selection into becoming experienced changes with different pairs of fees charged to the inexperienced and experienced. To do so, we calculate a baseline transition rate conditional on hiring into becoming an experienced employer using the posterior distribution of types in the data. We then adjust these transition probabilities based on how surplus from using the market changes as a function of (i) the fees imposed and (ii) workers' endogenous markups in response to these fees and the changing composition of employers.

To provide intuition about the tradeoff faced by the platform, Appendix 3 presents a simple version of the platform's problem for the case of specific (fixed) fees. The insight 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 for the platform but also governs selection into becoming experienced. Optimal fees for inexperienced employers balance revenue today with the reduced probability of hiring and, hence, future

[37] After the sample period ended, the platform raised baseline fees and implemented quantity discounts, but we do not have data from this period.

expected revenue after transitioning into the experienced segment.

This description characterizes the platform's tradeoff when employers have no foresight about the distribution of future wage bids. We also consider a different scenario where we assume that employers perfectly anticipate future equilibrium wage bids when they become experienced. When employers anticipate future wage bids, our solution requires iterating until convergence over: (i) hiring decisions among inexperienced employers that depend on the future discounted value of using the platform as a function of changes in experienced employer wage bids, (ii) inexperienced employers' decisions to post jobs after initial hiring, adjusting for any change in expected future surplus due to the fee change on experienced employers, and (iii) the wage-bid markups that result from the elasticity of demand and fees among experienced employers who post additional jobs. Additional computational details are provided in Appendix 3.

After implementing these steps, we also must determine a long-run value of experienced employers to the platform.[38] This value is akin to a discount rate for the platform in order to compute the present value of future hires. We choose a present value of four future jobs per experienced employer, which is a larger number than the average number of future jobs for those employers who gain experience; choosing this larger number, however, biases us towards concluding the platform should induce new employers to gain experience.

Despite an optimistic view of the long-run value of an experienced employer, we find the optimal fee on new employers is higher than the existing fee of 10%, whereas the optimal fee on experienced employers depends on the assumption made about how inexperienced employers anticipate future wages. Table 9 presents the results under each of two assumptions. Panels A and B present percentage changes in profits relative to the current fee structure for each case, whereas Panel C provides percentage changes in the number of employers transitioning to the experienced segment when employers do not anticipate future wages.

The main result is that the profit-maximizing fee for inexperienced employers is always at least as high as the profit-maximizing fee for experienced employers. Profit-maximizing fees for both experience levels are weakly higher than the 10% fees in place at the time of the sample. When employers anticipate future bids, however, the platform shades down fees on experienced employers relative to the case in Panel B where future wage bids do not have an effect on the hiring decision. Panel C shows that under the profit-maximizing fee for myopic employers, a dramatic decline occurs in the number of experienced employers on the platform, with the share of new employers who become experienced falling by about 38%. When inexperienced employers anticipate future wage bid changes (unreported), the share of employers who transition to become experienced falls by over 30%. Platform profits are maximized when a lower share of

[38]Finding this value requires potential extrapolation because future hiring behavior for many employers is only observed over a potentially short horizon.

employers than that under the current fee structure is induced to gain experience. Platform size would be reduced, but per-transaction profitability would more than offset the reduction in transaction volumes.[39]

7 Conclusion

Much of the recent interest in alternative employment arrangements has focused on the domestic labor market (Katz and Krueger, 2019), but the underlying technology behind gig economy platforms also enables remote contracting with distant workers. The potential wage savings from outsourcing service tasks domestically or abroad are enormous (Clemens, 2011), but relatively little is known about employers' adoption of new technologies to access this hiring channel. This paper sheds light on barriers to employers' use of online outsourcing for service tasks by characterizing how labor demand and workers' supply change as a function of an employer's experience with online hiring. To do so, we use data from a large online labor market and estimate employers' propensity to hire individual workers as a function of worker characteristics, job characteristics, and wages. We allow the model to vary with an employer's experience with hiring, while incorporating unobserved employer heterogeneity and non-random selection into gaining experience. We account for the endogeneity of wages using an instrumental variables strategy and Petrin and Train's (2010) control function estimator.

The majority of new employers who try out online hiring through this platform never end up using it to fill a vacancy. Employers who have gained experience, however, are more than twice likely to fill a vacancy than inexperienced employers. There are two main factors that contribute to this difference, and both relate to employer demand. First, hiring experience shifts out employers' value for remote workers, reflecting gains from learning-by-doing. Although we do not observe the source of learning-by-doing, we document that experienced employers expend less search effort during the online hiring process. They are likely to have gained familiarity with tailoring their communications and coordination of work with remote employees.

Second, the rising vacancy fill rate with employer experience is also due to exit from the market of employers with low valuations for online work, despite the anticipated benefits of learning-by-doing. Given that these low value employers created accounts and posted jobs online, we infer that employers need to try out the platform to determine their value for online hiring. The presence of buyers who don't know their own value for the market relative to their outside option is likely to be a characteristic of all unfamiliar

[39] 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 profit-maximizing 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 allow the platform to price discriminate on other observable employer attributes.

markets for procuring goods or services, similar to the concept of experience goods in Nelson (1970). An important implication is that accurate measurement of new buyers' price elasticity requires researchers to account for the larger heterogeneity in market value among these buyers compared to the selected group of buyers who have purchased in the market before. The extent of heterogeneous employer value for online hiring is substantial. The average change in employers' value for online hiring due to the exit of low value types is equivalent to about a 9% reduction in workers' wages.

The employer characteristic that correlates most strongly with our estimates for the value of online hiring is firm size. Large employers have relatively low average valuations for using the market. We note that these valuations are all relative to the individual employer's alternative hiring options. One possible explanation for the negative association between platform value and employer size is that the ease with which remote hiring can be coordinated with other production activities decreases with the size of the firm. This inference is consistent with (Fort, 2017), which shows that technology enables more fragmentation and outsourcing in industries whose production specifications are easy to codify. Alternative explanations also relate to the employer's production process, as a large firm may find it more cumbersome to manage the internal bureaucracy of contracting at arm's length.

On the applicant side of the market, workers' supply to new employers does not prevent employers from adopting the market. In contrast to the importance of positive feedback for workers in improving their future outcomes on the platform (Pallais, 2014; Stanton and Thomas, 2016), good feedback for experienced employers actually results in worker's marking up their wage bids. Workers appear to account for the positive correlation between feedback and the probability of filling a vacancy when bidding for jobs.

The ability to hire globally has not led to the forecast death of distance in the distribution of work, despite the large potential wage savings from online hiring. Our results suggest the challenge of outsourcing at the task-level is what creates most of the distance cost, rather than the geographic distance between employers and workers. We use the estimates to consider how the size and profitability of the platform would change after encouraging new employer entry through subsidies. Subsidizing new employers would not be profitable for the platform. Instead, pursuing a niche strategy, targeting the small share of employers who are able to coordinate outsourced task-level online work with their other production activities is most profitable for platform operators.

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 restriction allows for a full application cycle of workers from different time zones and eliminates batched hiring for which applicants may blend across jobs. The sample for estimating hiring probabilities also drops jobs for which the employer hires a previously hired worker, because whether these jobs are new or the continuation of a pre-existing contract under different terms is unclear.

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. We omit any job from an employer who sends over 100 interview requests for the first job or who sends over 60 interview requests for a subsequent posting. These postings are likely to be fake jobs posted by spammers. Finally, we drop any job that is later declared to have been posted by mistake. The following restrictions are used to clean applications: First, we drop applications from invited workers who later report they are unavailable; second, we drop applications if the employer reports obvious spam.

Appendix 2: Alternative Explanations for Declining Wage Bids

Bids of hired workers relative to other applicants

The solid bars of the histogram in Figure 9 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 9: Bid Decile of Hired Worker

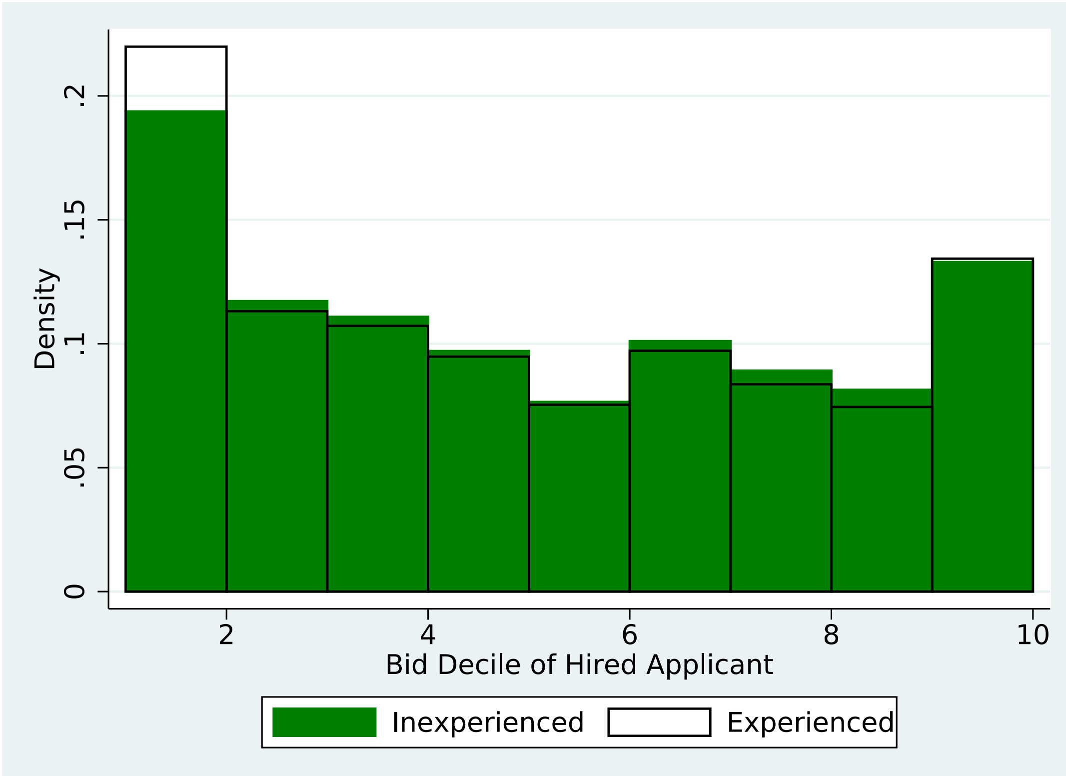

The figure shows the bid decile of the worker who is hired from the set of applicants. For each job opening, we find the decile of wage bid for each applicant. We then take the decile of the applicant who was hired and plot the histogram of wage bid deciles for applicants selected by inexperienced employers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

Different Application Rates

The extent of competition for a job posting might change 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) shows 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 A4 repeats the analysis from Table 6, but the estimations include the log arrival rate of applicants within the first 24 hours of posting the job as an additional control. Note 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 all bids received by the employer. However, including this control does not change the main finding from Table 6 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.[40] We posted identical jobs from the accounts of employers with different levels of experience. Employer 1 had no experience, whereas 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 A5 contains the results. In each specification, we see a significant, negative point estimate on the experienced employer indicator variable. The effect sizes are larger in magnitude than those estimated in Table 6. 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 as opposed to the log of the difference in the actual bid from the rate posted in the applicant's profile. 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. Applicants' first bids might 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. Rejected wage offers for employers with similar experience are highly serially correlated. To construct

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

an initial offer for workers who are hired, we take the last rejected wage offer to employers of the same experience level in the same job category. We then compare the observed wage on hires made to the last wage offer received on rejected applications.

Using these measures to assess the extent of bargaining is imperfect, because they likely overstate the extent to which bargaining occurs, 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, is likely to be small due to difficulties in transferring funds across banking systems. However, to the extent that they do exist, we expect them to show up for experienced employers who have reputations.

Table A6 contains regression results analyzing bargaining. In OLS regressions, some small reductions may exist in final wages 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.

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 so 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 let H_E be an indicator for hiring while experienced. Wages for the inexperienced and experienced segment are w_I and w_E, respectively.[41] 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,

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

$H_I(w_I)$. Notice 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, $\Pr(H_E|H_I(w_I), w_E)$, may depend on the evolution of employers' beliefs about the platform as a result of hiring. That $\Pr(H_E|H_I(w_I), w_E)$ specifically conditions on $H_I(w_I)$ and the wage paid captures the possibility that the identity of the marginal inexperienced employer may affect experienced hiring. 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 $\Pr(H_I|w_I)$ and H_E as shorthand for $\Pr(H_E|H_I(w_I), w_E)$, the first order 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:

$$t_E^* = -\frac{H_E}{\frac{\partial H_E}{\partial w_E}\frac{\partial w_E}{\partial t_E}}. \tag{11}$$

The fee for the inexperienced is:

$$t_I^* = -\frac{H_I}{\frac{\partial H_I}{dw_I}\frac{\partial w_I}{\partial t_I}} - t_E^* H_E - t_E^* H_I \frac{\partial H_E}{\partial H_I}. \tag{12}$$

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.[42] These types are assigned independently from the applicant set. Then, for each ad-valorem fee pair, (τ_I, τ_E), simulated

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

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. (1a) If the inexperienced employer is forward looking and anticipates future wage bids, the inside utility of any hire made while inexperienced changes (after looping over bids to the experienced employers until convergence in step (3) by $Pr(EverExperienced|\mu_i) \times \alpha_E \overline{\Delta log w_E}$, where $\alpha_E \overline{\Delta log w_E}$ is the coefficient on future wages and $Pr(EverExperienced|\mu_i)$ is the probability that this change will become relevant for an employer of type μ_i becoming experienced. (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 market. This surplus is modeled as a probabilistic function of the expected surplus and is calibrated from the transition probability between the inexperienced and experienced sample;[43] (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) If inexperienced employers are forward looking, we return to step 1a and re-compute hiring decisions adjusting $\overline{\Delta log w_E}$ by future wage bids received by experienced employers. Otherwise, end. (5) Profits are then based on hiring probabilities from the model and the fee-rate associated with the chosen bid.

[43]The parameters are calibrated to minimize the distance between the predicted transition probability and 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

Agrawal, Ajay, John Horton, Nicola Lacetera, and Elizabeth Lyons. 2015. "Digitization and the Contract Labor Market: A Research Agenda." In *Economic Analysis of the Digital Economy*. University of Chicago Press, 219–250.

Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons. 2016. "Does Standardized Information in Online Markets Disproportionately Benefit Job Applicants from Less Developed Countries?" *Journal of International Economics* 103:1–12.

Antras, Pol, Teresa C Fort, and Felix Tintelnot. 2017. "The Margins of Global Sourcing: Theory and Evidence from US Firms." *American Economic Review* 107 (9):2514–64.

Autor, David. 2001. "Wiring the Labor Market." *Journal of Economic Perspectives* 15 (1):25–40.

Barach, Moshe and John J Horton. 2017. "How Do Employers Use Compensation History?: Evidence from a Field Experiment." Tech. rep.

Bernard, Andrew B, J Bradford Jensen, Stephen J Redding, and Peter K Schott. 2007. "Firms in International Trade." *Journal of Economic perspectives* 21 (3):105–130.

Blinder, Alan S. 1973. "Wage Discrimination: Reduced Form and Structural Estimates." *Journal of Human Resources* :436–455.

Blinder, Alan S and Alan B Krueger. 2013. "Alternative Measures of Fffshorability: A Survey Approach." *Journal of Labor Economics* 31 (S1):S97–S128.

Bloom, Nicholas, Fatih Guvenen, Benjamin S Smith, Jae Song, and Till von Wachter. 2018. "The Disappearing Large-Firm Wage Premium." In *AEA Papers and Proceedings*, vol. 108. 317–22.

Cairncross, Frances. 1997. *The Death of Distance: How the Communications Revolution Will Change Our Lives.* {Harvard Business School Press}.

Clemens, Michael A. 2011. "Economics and Emigration: Trillion-dollar Bills on the Sidewalk?" *Journal of Economic perspectives* 25 (3):83–106.

Coase, Ronald H. 1937. "The Nature of the Firm." *Economica* 4 (16):386–405.

Englmaier, Florian, Arno Schmöller, and Till Stowasser. 2017. "Price Discontinuities in an Online Market for Used Cars." *Management Science* 64 (6):2754–2766.

Fort, Teresa C. 2017. "Technology and production fragmentation: Domestic versus foreign sourcing." *The Review of Economic Studies* 84 (2):650–687.

Fortin, Nicole, Thomas Lemieux, and Sergio Firpo. 2011. "Decomposition Methods in Economics." In *Handbook of Labor Economics*, vol. 4. Elsevier, 1–102.

Ghani, Ejaz, William R Kerr, and Christopher Stanton. 2014. "Diasporas and Outsourcing: Evidence from oDesk and India." *Management Science* 60 (7):1677–1697.

Goldschmidt, Deborah and Johannes F Schmieder. 2017. "The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure." *The Quarterly Journal of Economics* 132 (3):1165–1217.

Grossman, Gene M and Esteban Rossi-Hansberg. 2008. "Trading Tasks: A Simple Theory of Offshoring." *American Economic Review* 98 (5):1978–97.

Halpern, László, Miklós Koren, and Adam Szeidl. 2015. "Imported Inputs and Productivity." *American Economic Review* 105 (12):3660–3703.

Horton, John, William R Kerr, and Christopher Stanton. 2017. "Digital labor markets and global talent flows." Tech. rep., National Bureau of Economic Research.

Horton, John J. 2010. "Online Labor Markets." In *International Workshop on Internet and Network Economics*. Springer, 515–522.

———. 2017a. "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment." *Journal of Labor Economics* 35 (2):345–385.

———. 2017b. "Price Floors and Employer Preferences: Evidence from a Minimum Wage Experiment." Tech. rep.

Kahn, Lisa B. 2013. "Asymmetric Information Between Employers." *American Economic Journal: Applied Economics* 5 (4):165–205.

Kahn, Lisa B and Fabian Lange. 2014. "Employer Learning, Productivity, and the Earnings Distribution: Evidence from Performance Measures." *The Review of Economic Studies* 81 (4):1575–1613.

Katz, Lawrence F and Alan B Krueger. 2019. "Understanding Trends in Alternative Work Arrangements in the United States." Tech. rep., National Bureau of Economic Research.

Lange, Fabian. 2007. "The Speed of Employer Learning." *Journal of Labor Economics* 25 (1):1–35.

Lazear, Edward P, Kathryn L Shaw, and Christopher T Stanton. 2018. "Who Gets Hired? The Importance of Competition Among Applicants." *Journal of Labor Economics* 36 (S1):S133–S181.

Levitt, Steven D, John A List, and Chad Syverson. 2013. "Toward an Understanding of Learning by Doing: Evidence from an Automobile Assembly Plant." *Journal of Political Economy* 121 (4):643–681.

Lyons, Elizabeth. 2017. "Team Production in International Labor Markets: Experimental Evidence from the Field." *American Economic Journal: Applied Economics* 9 (3):70–104.

Nelson, Phillip. 1970. "Information and Consumer Behavior." *Journal of Political Economy* 78 (2):311–329.

Oaxaca, Ronald. 1973. "Male-female Wage Differentials in Urban Labor Markets." *International Economic Review* :693–709.

Oi, Walter Y and Todd L Idson. 1999. "Firm Size and Wages." *Handbook of labor economics* 3:2165–2214.

Pallais, Amanda. 2014. "Inefficient Hiring in Entry-Level Labor Markets." *American Economic Review* 104 (11):3565–99.

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

Riordan, Michael H. 1986. "Monopolistic Competition with Experience Goods." *The Quarterly Journal of Economics* 101 (2):265–279.

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

Train, Kenneth E. 2009. *Discrete Choice Methods with Simulation.* Cambridge University Press.

Trefler, Daniel. 1995. "The Case of the Missing Trade and Other Mysteries." *The American Economic Review* :1029–1046.

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: Characteristics of Applicants and Hired Workers in the Sequential, Arms-Length Sample

	Job Applicants					Hired Workers				
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	5.007	5.833	5.736	5.812	5.674	9.682	9.509	9.248	9.758	8.985
	(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 Affiliate	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 applicant resume characteristics for all job applicants and workers who are hired. The good worker feedback dummy is an indicator that the applicant's feedback score is greater than 4.5 out of 5 stars and is set to missing for workers who lack feedback. Agency affiliate means the applicant is a member of a third-party outsourcing agency that operates on the platform, and the agency's collective reputation score is observable to employers. 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 Local Currency to Dollar 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 who have never hired previously. The experienced sample is employers who have hired. The first instrument is the log of the average monthly local currency to US dollar 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. After taking logs, the second instrument is purged of variation due to week and job category fixed effects. Regressions include indicators if an instrument is missing or, for the first instrument, is invariant within country. 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

Note: Experienced Employer Parameter Estimate Columns Contain Additive Interaction Terms Relative to Inexperienced Employer Baseline

	(1)	(2)	(3)	(4)	(5)	(6)
Employer Experience	Inexperienced	Experienced	Inexperienced	Experienced	Inexperienced	Experienced
Resume Characteristics	Yes		No		Yes	
Control Function for Bids	Yes		Yes		Yes	
Multiple Employer Types	No		Yes		Yes	
Panel A: Parameter Estimates for Employer Types in Demand Models						
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						
Log-bid Equivalent Change in Productive Value:	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 Employer Heterogeneity (*μ*)			82%		32%	
Panel C: Parameter Estimates, Elasticities, and Costs						
Log Hourly Bid	-5.50***	-3.13***	-5.08***	-0.87***	-5.10***	-2.78***
	(0.42)	(0.72)	(0.15)	(0.06)	(0.26)	(0.28)
Mean Own-Price Elasticity	-5.46	-8.54	-5.04	-5.89	-4.96	-7.70
Mean % Markup, (Pre oDesk-Fee)	22.4%	13.3%	24.7%	20.5%	25.2%	14.9%
Mean Implied Cost (USD, Pre oDesk-Fee)	$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
Percentage of Bid Difference Due to Markups	79%		35%		88%	
Elasticity of Vacancy Fill With Respect to All Bids	-4.1	-5.08	-2.9	-3.28	-2.93	-4.03

Notes: Robust 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 1-2 is over job openings, whereas the likelihood in Columns 3-6 is over sequences of job openings within employer. See the text for details about the estimation procedure. In Panel B, own price elasticities are type-weighted averages of the individual elasticities when the model has employer heterogeneity. The log productive value decomposition is described in the text. The elasticity of vacancy fill with respect to all bids is the elasticity of the employer making any hire with respect to a percentage increase in the bids of all applicants.

Table 5: Market Adoption and Estimated Valuations are Declining with Firm Size

Dependent Variable:	Adoption		Estimated Employer Value for Market	
	(1)	(2)	(3)	(4)
Firm Size:				
2-9 Internal Workers	0.0283**	0.0105	-0.0137	-0.0691***
	(0.0109)	(0.0081)	(0.032)	(0.0184)
10+ Internal Workers	-0.0736***	-0.0676***	-0.300***	-0.279***
	(0.0118)	(0.0106)	(0.0354)	(0.0301)
Constant	0.286***	0.294***	-1.104***	-1.079***
	(0.0166)	(0.00497)	(0.0515)	(0.0123)
Geography Fixed Effects		X		X
Observations	20,395	20,395	20,395	20,395
R-squared	0.006	0.043	0.007	0.054

Notes: The unit of analysis is an employer. The sample includes surveyed employers who report information about themselves or their business. The excluded category in all regressions is sole proprietors or individuals. Robust standard errors are clustered by employer geography. The dependent variable in Columns 1 and 2 is an indicator that the employer posts a job after gaining hiring experience--our definition of market adoption. The dependent variable in Columns 3 and 4 is the posterior estimate of the employer's individual value for the market, μ_i, which can be interpreted on a scale relative to log wage bids.

Table 6: 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 7: 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 8: 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 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
Panel B: 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

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 6 and contain the same controls as detailed in the notes to Table 6. Observation counts differ from Table 5 because, for some observations, the timing of the initial feedback cannot be classified as occurring before or after later job postings.

Table 9: Platform Profits and Repeat Job Posting for Different Fee Schedules

Panel A: Percent Change in Profits Relative to 10% Uniform Fee When Employers Anticipate Future Fee Changes

Inexperienced \ Experienced Fee	5%	10%	15%	20%	25%	30%
5%	-35.2%	-19.1%	-14.1%	-17.8%	-23.8%	-33.3%
10%	-13.7%	0.0%	2.0%	-1.8%	-9.8%	-20.3%
15%	1.3%	11.4%	13.0%	8.0%	-0.8%	-10.7%
20%	9.9%	17.7%	17.5%	11.3%	3.8%	-5.9%
25%	13.4%	20.4%	19.4%	12.6%	4.6%	-5.4%
30%	13.1%	17.7%	15.7%	10.8%	2.6%	-7.1%

Panel B: Percent Change in Profits Relative to 10% Uniform Fee When Employers Do Not Anticipate Future Fees

Inexperienced \ Experienced Fee	5%	10%	15%	20%	25%	30%
5%	-38.5%	-19.6%	-8.0%	-6.1%	-5.9%	-10.0%
10%	-18.3%	0.0%	9.6%	13.8%	12.2%	7.9%
15%	-4.2%	12.9%	22.1%	25.2%	23.9%	20.9%
20%	3.5%	18.1%	26.4%	29.6%	28.6%	25.3%
25%	7.6%	21.0%	27.8%	30.9%	30.5%	27.2%
30%	6.9%	18.6%	25.0%	27.8%	27.5%	25.0%

Panel C: Percent Change in Number of Employers Becoming Experienced When Employers Do Not Anticipate Future Fees

Inexperienced \ Experienced Fee	5%	10%	15%	20%	25%	30%
5%	19.3%	11.9%	5.2%	-1.1%	-6.6%	-12.6%
10%	5.6%	0.0%	-6.1%	-11.7%	-17.0%	-22.2%
15%	-4.8%	-10.0%	-15.3%	-19.7%	-24.7%	-29.5%
20%	-16.8%	-21.2%	-25.8%	-30.0%	-34.2%	-38.9%
25%	-27.0%	-30.7%	-34.9%	-38.7%	-42.2%	-46.0%
30%	-37.3%	-40.6%	-44.1%	-46.9%	-49.9%	-53.2%

Notes: Simulations use the parameters from the last two columns of Table 4 and consider a different grid of fees that vary by employer experience. Panel A displays percent changes in total platform profits when inexperienced employers anticipate future wage bids as a function of experienced fees and the resulting markups induced by selection into job posting in the experienced segment. Panel B displays percent changes in total platform profits when inexperienced employers do not anticipate future fee and markup changes even though they are realized in equilibrium and affect the contribution to platform profits by experienced employers. Panel C displays the percent change in the fraction of employers who go onto the experienced sample as a function of different fees. Selection into or out of experience happens because (i) inexperienced employers are more or less likely to hire in the first place as a function of inexperienced fees and (ii) are more or less likely to realize future surplus after hiring as a function of experienced fees. See Appendix 4 for details behind these calculations and how transitions between inexperienced and experienced employers are calibrated.

Table A1: Details about Resume Data

	Full Sample				Sequential, Arms-Length Sample			
Variable	Mean	Std. Dev.	Min	Max	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, whereas 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 interviews per job opening fall with hiring experience

DV: Log Number of Interviews +1	OLS	Employer Effects	Employer Effects
	(1)	(2)	(3)
One previous hire	-0.0515***	-0.292***	-0.289***
	(0.00576)	(0.00773)	(0.00772)
Two previous hires	-0.0918***	-0.362***	-0.357***
	(0.00657)	(0.00928)	(0.00929)
Three previous hires	-0.114***	-0.406***	-0.401***
	(0.00726)	(0.0106)	(0.0106)
Four previous hires	-0.135***	-0.442***	-0.436***
	(0.00786)	(0.0114)	(0.0114)
Five or more previous hires	-0.177***	-0.520***	-0.513***
	(0.00704)	(0.0123)	(0.0123)
Constant	0.893***	1.121***	1.110***
	(0.121)	(0.172)	(0.177)
Includes job duration fixed effects and third order polynomial of job description length	No	No	Yes
Observations	322,333	322,333	322,332
R-Squared	0.021	0.414	0.416

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 A4: 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 description, 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 A5: 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 information was not available for applicants who did not have publicly visible resumes, reducing the sample size.

Table A6: Analysis of Bargaining 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 a proxy for the extent of bargaining, calculated as the log of the final hourly wage bid when the applicant is hired less the last log hourly wage bid on a rejected application for that candidate in the same job category to an employer with the same experience. 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.