Experience Markets:
An Application to Outsourcing and Hiring

Christopher T. Stanton

Harvard, NBER, and CEPR

Catherine Thomas

LSE, Centre for Economic Performance, and CEPR

October 2017*

Abstract

While shopping for a non-standard good or service, a buyer evaluates noisy signals of seller quality. A new buyer, without the benefit of market experience to draw on, puts more weight on each seller's noisy signal of value and therefore perceives alternative sellers as being more differentiated from each other. We call a setting like this an "experience market" and analyze how workers submit wage bids in an online market for global outsourced labor where employers have observably different experience levels. Declining wage bids and changing hiring patterns as employers gain experience suggest that new employers are uncertain about their own value for using this labor market and learn about it when interacting with workers. In some experience markets, platforms may find it profitable to subsidize information acquisition, but, in this market, niche positioning without subsidizing wages for new employers is more profitable because the marginal employer encouraged to hire by the subsidy has a low value for the market. The results offer a possible explanation for limited participation in labor services offshoring relative to predictions based on technical feasibility studies.

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

1 Introduction

An employer learns about the existence of a market for labor services. He considers hiring a worker in this market to complete a specific job, but, having never hired in this market before, he does not know how well the average worker will meet his needs. Recognizing that workers differ in their suitability for his job, he uses information from workers' resumes and interviews to evaluate applicants and also to learn about the overall distribution of workers in the market. A second employer is also seeking a single worker but, unlike the first employer, has hired in the market before. Due to his past experience, the information he receives in applications or interviews does relatively little to change his inference about the pool of workers. Workers in the market observe whether the employer has market experience before submitting job applications. Knowing that the employer will choose, at most, one applicant, each worker makes a strategic wage offer that depends on her forecast of the employer's willingness to pay in the event that she is the most-preferred applicant.

This process characterizes application and hiring decisions in a large online global labor market. Using data from over five million applications for over 300,000 jobs posted by 80,000 employers, we first show that workers' wage bids to employers who have not previously hired in this market are more than seven percent higher than those made to employers who have hired several workers.[1]

Higher wage offers to the inexperienced may have several causes and the first goal of the paper is to disentangle these explanations. We present evidence that the main reason job applicants' make higher wage offers to inexperienced employers is because they expect to be evaluated differently by employers with and without experience, and, in response, mark up wage bids to inexperienced employers. Instrumenting for wage bids in global applications with exchange rate movements and variation in the extent of competition, we identify experience-dependent hiring probabilities as a function of the bids that applicants submit. We show that the relationship between wage bids and employers' hiring probabilities varies with employer experience because of three distinct channels. The overall estimate of the net effect of these changes is that 85% of the observed wage premium to new employers can be attributed to higher markups.

The most important channel causing the hiring probability function for inexperienced and experienced employers to differ lowers equilibrium wage bids to the experienced. Experienced employers perceive less differentiation across applicants as they are better able to separate out the market-level

[1]These bid differences are not bargained down during contract negotiations, so inexperienced employers who hire also pay higher hourly wages.

match quality included in all applicants' signals from individual worker-level match quality. As a result, applicants appear less differentiated to an experienced employer. Therefore, each applicant knows that if she is most preferred by the employer, an experienced employer views the second-most-preferred applicant (or the option of not hiring) as a closer substitute to the most-preferred applicant. We call this the "substitutability" effect. This effect reduces the markup included in the wage bids to the experienced.

The other two channels actually cause bids to increase with experience because experienced employers are more likely to make a hire. For a given employer, having hired in the past may increase the expected value of future hires. We term this second channel a "true experience" effect that serves to increase the probability of hiring because of increasing values for using the market. Third, experienced employers are a selected group, consisting of those employers who have a higher value for the average worker in the market, and who are, hence, more likely to hire any worker, rather than none. We term this third channel an employer "composition" effect. Our estimates show that the true experience effect and the composition effect are substantial, driving the stylized fact that experienced employers are twice as likely to hire. In fact, expressing them in a wage-bid metric suggests that the change with experience is equivalent to cutting wage bids by about 30% to inexperienced employers.[2] While, the substitutability effect appears to dominate for the purposes of applicants' setting wage bids, subsidizing early hiring might be worthwhile if employers do not anticipate the true experience effect.

We highlight that the substitutability effect of sellers appearing more differentiated to new buyers will be present in any market for non-standard goods or services in which (i) buyers are uncertain about their average value for the goods or services offered; and (ii) differentiated sellers send noisy signals of quality that share a common component for the buyer.[3] If, in addition, buyers use market experience to reduce their market-level uncertainty, the extent of perceived differentiation among sellers decreases with buyer experience, and the market appears more competitive. We call a market with these features an Experience Market, in reference to the related concept of an Experience Good (Nelson, 1970), where a buyer's value for a product or service is learned only through consumption.[4] When a buyer's experience is observed, although sellers cannot perfectly

[2]About two-third of the increase comes from the within-employer "true experience" effect, while the remainder is due to employer composition differences.

[3]We have in mind buyers with different idiosyncratic match quality with sellers' goods or services. This allows the analysis to abstract away from sellers' signaling about the quality valued by all buyers through price offers.

[4]Nelson (1970) distinguishes between an experience good and a search good. In the labor market studied here, buyers frequently interview applicants before hiring. We call an interview a partial "experience" with an applicant because it is a costly activity in terms of the employer's time and, unlike the case of a pure search good, interviews

forecast the buyer's perceived match quality for an individual offer, they do know that buyers with less experience are more likely to put additional weight on the noisy signal rather than on their prior about the distribution, thus increasing the optimal bid because of greater perceived differentiation.

We also evaluate other explanations for wage bids that decline with employer experience. Another candidate explanation is the existence of a time-varying factor correlated with experience that affects how employers use the platform. For example, employers may learn how to communicate the requirements of their job by providing different details in a job description. A number of separate analyses, including a field experiment to isolate the effect of observable experience, suggest that concerns about time-varying conditions (other than changes in observable employer experience) are not likely to drive the findings.[5]

Wage bid reductions may also be due to differences in how workers perceive the expected costs of labor supply—perhaps because new employers are more likely to review hired candidates poorly or to create extra hassle for an employee. However, our analysis shows that differences in workers' costs do not explain much of the wage bid premium to new employers.

Additional evidence consistent with the "substitutability" effect—that employers' perceptions change as they learn their own value for using the market—supports the conclusion that this effect explains higher bids to new buyers. First, while bids to experienced employers fall overall, within the set of experienced employers, those who have good feedback from previous hires receive *higher* bids on subsequent jobs than experienced employers who have received poor feedback. This within-experience level result is consistent with the intuition developed above about how the employer composition effect may increase workers' bids when they anticipate that an employer is more likely to hire. In the data, experienced employers with good feedback are 6.7 percentage points more likely to hire than experienced employers who lack good feedback. Good feedback provides a signal that the employer has a high value of the market (due to a combination of the true experience and composition effects) and is less likely to substitute to the outside or no-hire option. This increases the optimal markups charged to these employers.[6] Explanations for the bid premium to inexperienced employers that focus on unobserved worker costs or employer adverse selection would have predicted that highly-rated employers are less costly to work for, driving lower subsequent bids

do not fully reveal the worker's quality.

[5]In the field experiment, wage bids were collected from real jobs posted with identical details but different levels of employer experience. The results show wage bid reductions with experience that are larger than the average reduction estimated from the observational data.

[6]The result is also true for experienced employers with no feedback from workers. In the data, they are 3.7 percentage points more likely to hire than experienced employers with poor feedback (below 4.5 stars), and they also receive higher wage bids than employers with poor feedback.

within experience level.[7]

Further evidence for learning comes from the fact that the average number of interviews conducted on each job, a measure of search effort, declines with employer experience. This pattern may have a number of explanations, such as a declining marginal benefit from learning about the market or the realization that interviews are ineffective. But additional results about reported measures of productivity on the job suggest that interviews do convey information to inexperienced employers. Two measures of productivity that employers report after hiring, public feedback left for the worker and the response to a non-public post-job survey, both vary with the number of interviews conducted on the first job posting. Contrary to most standard intuition about the marginal benefits of search effort, but consistent with a model of sequential sampling with an uncertain distribution of workers (Kohn and Shavell, 1974), the number of interviews is negatively correlated with reported productivity. This is because in the Kohn and Shavell model, the reservation value of hiring declines with experience because additional search conveys less information about the market. Those employers who stop search early do so with a higher reservation threshold, suggesting that they have found a good match early on and are willing to forego learning from subsequent interviews. These results persist into employers' subsequent jobs, suggesting that those who stop early are likely to have a high average value of hiring at the employer-market level.

The second goal of the paper is to examine how our results help us understand market structure in online outsourcing markets and broader offshoring patterns. First, sellers in spot markets, like the one studied here, are likely to fail to internalize the value of hiring experience to the buyer.[8] While there is no reason why far-sighted buyers would not internalize the longer-term benefits of learning through experience, the higher bids to the inexperienced segment may excessively offset experimental hiring from the platform's perspective. This raises the question of whether the platform would find it profitable to subsidize hiring by inexperienced employers. Answering this question sheds light on whether a platform has an incentive to encourage participation by many employers or to target niches of employers with relatively high valuations.

We find that a niche positioning—targeting high value employers—maximizes profits for this platform. In fact, the platform profit-maximizing fees on both experienced and inexperienced employer segments are higher than the 10% ad-valorem rate that was in place at the time of data

[7]This result does not appear to be driven by unobserved employer quality or different sorting of workers to employers with different experience levels, as it remains with employer fixed effects or worker fixed effects.

[8]The uncertainty about the market that drives these results could be expected to encourage new employers to experiment. While our data do not permit an investigation of how different levels of buyer uncertainty change experimental hiring, our data do show that the supply-side response to uncertainty of higher bids to inexperienced buyers reduces the incentive for employers to gather information by hiring workers in the market.

collection.[9] In addition, subsidizing new hiring with a relative fee reduction for the inexperienced would not be profitable. Although some new employers would be induced to gain experience, the reduction in wage offers to new, inframarginal employers (who would have hired without the subsidy) largely outweighs the added long-term volume from the new employers induced to hire.

The estimates show substantial heterogeneity in employer value for the market that appears to be unobserved by employers ex ante. This unobserved heterogeneity can explain both the wage bid premium to inexperienced employers, new employers' relatively high rates exiting without hiring, and a niche market positioning that does not reverse the bid premium to new employers.

What can these findings tell us about recent trends in labor services offshoring more generally? It is clear that relative wage differences across countries offer firms huge cost advantages from the potential wage arbitrage opportunity. In the oDesk data studied here, the mean wage bid for a web programmer in India is $11.89 and in the US the mean bid in the same job category is $20.42. [10] Furthermore, advances in information and communication technology over the last 20 years have allowed for the remote production of a large share of tasks. Blinder and Krueger (2013) estimate the extent to which occupations in the U.S. have the offshorability characteristic.[11] They conclude that a reasonable estimate of the total share of U.S. employment that was offshorable given the technology available at the is around 24%. In information and professional, scientific and professional services, this share increased to around 35%, and to 41% in office and administrative support occupations. Earlier work by Jensen and Kletzer (2010) that focused on services occupations estimated the 93% of computer and mathematical occupations and 64% of office and administrative support could be offshored.

While Firpo et al. (2010) document that some of the increase in U.S. wage inequality during the 1990s and 2000s can be attributed to lower relative wages in offshorable tasks, there has been surprisingly limited study of the associated changes in the number of jobs. Turning to firm-level choices, according to estimates from the US Census Bureau Survey of Business Owners, only 1.36% of firms surveyed in 2012 reported that they outsourced or had transferred a business function outside of the United States. Fort (2017) shows that although 27% of plants fragment production in manufacturing primarily by purchasing manufacturing services from domestic suppliers, only 2%

[9]The platform moved to a different fee structure after the first draft of this paper was written. The new structure has higher average fees than the 10% ad-valorem rate and a second-degree volume discount. The data studied here pre-date this change.

[10]Data available at the website https://www.payscale.com/rccountries.aspx show that Software Developers in India command a median annual salary of $6,000 compared to $69,000 in the U.S.

[11]A job is offshorable if the work can be done at a remote (presumably foreign) location and the work productwhether a good or a servicedelivered to the end user with little or no loss of quality.

of firms report doing so by purchasing manufacturing services offshore.

The low number of firms involved in offshoring, despite its low variable costs and technical feasibility, suggest firms face some barriers to participating in offshore markets. In the oDesk context, we find that new employers are uncertain about the full benefits of using the market before they try it, and need to learn through experimentation.[12] Fort (2017) finds that adopting electronic communication technology at home is more strongly associated with an increase in domestic contracting than in offshore contracting and suggests this is because coordination costs remain relatively high across country borders even with this technology. We view our findings as consistent with this interpretation. Because organizing multinational production requires firms to operate in new labor markets, it presents potential entrants with the challenges of navigating and then operating within an experience market.

The paper proceeds as follows. Section 2 provides intuition for the results and the empirical exercises. In this theoretical motivation, employers gather information and choose workers from a pool of job applicants while being uncertain about the characteristics of the pool.[13] This approach has roots in the early analysis of how employers hire, pioneered by Barron, Bishop, and Dunkelberg (1985). It is also related to the literature on learning in labor markets, with recent contributions by Kahn (2013) and Kahn and Lange (2014). Here, however, rather than having incomplete information about individual workers, an employer has uncertainty about the distribution of workers.

Section 3 introduces the data and the empirical context. This is the first paper, to our knowledge, that examines how an employer's uncertainty affects his hiring problem in an online labor market. Most of the literature using similar data sources, like Pallais (2014) and Stanton and Thomas (2016) analyze the effect of incomplete information about workers.[14]

[12]This is consistent with complementary studies about the value of information in more and less familiar online environments. See Agrawal, Lacetera, and Lyons (2012), Agrawal and Goldfarb (2008), and Agrawal, Catalini, and Goldfarb (2011).

[13]Much of the literature on experience goods has learning about the characteristics of an individual product or service in the monopoly case (Shapiro, 1983), but some versions of the problem with monopolistic competition allow prices to signal underlying product quality that is commonly valued (Riordan, 1986). Here, because the match with an employer is unknown to the seller ex-ante, bids respond to the expected surplus available. This has a similar intuition to Benabou and Gertner (1993), who consider learning from prices with the goal of resolving aggregate uncertainty, but different predictions are derived here because the uncertainty is idiosyncratic. Bergemann and Valimaki (2006) analyze the problem of a monopoly provider of an experience good when the firm knows the fraction of buyers who have experience. Empirical applications leveraging consumer learning about product or service characteristics include Erdem and Keane (1996) and Israel (2005). Empirical work that leverages learning about an agent's own match with a sector or firm often extends from the Miller (1984) model. A recent application is Arcidiacono, Aucejo, Maurel, and Ransom (2016) who estimate students' learning about job sector and educational fit.

[14]See Horton (2010) for an overview of how online labor markets work and Horton, Kerr, and Stanton (2017) for a review of literature on these markets. Most of the literature is in line with the analysis of incomplete information in other markets (for example, Cabral and Hortacsu (2010)). Other papers that evaluate employer hiring, but consider different empirical variation or different mechanisms, are Ghani, Kerr, and Stanton (2014) who document ethnic

Section 4 presents the problem of estimating hiring probabilities. We derive a conditional logit model of hiring applicants. Using the Petrin and Train (2010) control function approach, exogenous variation in exchange rates and competition, both of which are passed through to bids, identifies the elasticity of hiring with respect to wage bids. The model also permits estimation of the degree of latent employer heterogeneity for using the market.

Section 5 conducts robustness and supporting analysis, building in the sequential nature of how employers gather information. Section 6 considers counterfactuals that vary platform fees, taking as input the parameter estimates and hiring probabilities from Section 4. Section 7 concludes.

2 Intuition About Offers in Experience Markets

Before turning to the empirical application, we provide some intuition about the results described in the introduction. We use the context of an employer hiring, but the logic applies to many situations in which buyers are in unfamiliar markets. A job opening arises because an employer has a task that requires a single worker. Employers are heterogeneous in that some types of employers are better suited than others to using the market overall.[15] To reflect this, the value of the average worker to employer i is modeled as an employer-specific parameter, μ_i, from a known distribution.

After posting a job, applications arrive. In each application, the employer observes a signal of worker quality and a wage offer.[16] Before considering how workers choose what wage to bid, we focus on conditions that lead to variation in the distribution of employer willingness to pay to hire a given worker. Optimal wage offers will take the expected employer willingness to pay into account. Because the employer is looking for just a single worker, his willingness to pay for his most preferred worker upon hiring is pinned down by the perceived difference in his expected payoff between the best applicant and the next best alternative. His next best alternative may be hiring another applicant or not hiring anyone in the market.

matching between Indians and Horton (2017a) and Horton (2017b). Horton (2017a) analyzes a field experiment that provided algorithmic recommendations to employers, documenting an increase in hiring. Horton (2017b) studies how employers react to the imposition of a minimum wage in the market. Using this source of variation, and consistent with the results here, he finds employers trade off wage and productivity attributes, substituting to more productive workers after the imposition of a minimum wage. Both the algorithmic recommendations and the minimum wage imposition took place after our sample ends.

[15]This may reflect differences in the ability to communicate with workers abroad, to screen candidates in a new setting, or to find applicants with skills that match the specific needs of the firm.

[16]For now, we abstract away from explicit consideration of different numbers of applicants, instead focusing on the surplus available for a fixed number of applicants. In the empirical implementation, we will also account for the number of applicants.

To make this concrete, assume that employer i observes signals from each of J applicants, denoted q_{ij}. These signals are

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

Each of the J quality signals is assumed to be centered on μ_i, the value of the average worker to the employer. The signal also includes a component that reflects match productivity with the employer, denoted η_{ij}, for the employer-worker pair. When μ_i is unknown, noise in the signal, ξ_{ij}, means that the employer cannot perfectly distinguish between the market-level component, the individual worker-employer-level component, and overall productivity. When μ_i is known, noise means that the employer cannot distinguish between η_{ij}, the worker-employer-specific component, and ξ_{ij}, the noise term. In both cases, the employer's task is to filter the productive components of the signal, $(\mu_i + \eta_{ij})$, from the noise, ξ_{ij}.

A convenient way to parameterize employer uncertainty about his own value of μ_i and to describe how he may filter the productive component, $(\eta_{ij} + \mu_i)$, of each signal q_{ij} is a normal-normal model (Farber and Gibbons, 1996; Altonji and Pierret,2001; Kahn and Lange, 2014). An uncertain employer believes that the distribution of μ_i is normal, with prior mean and variance μ_{i0} and $\sigma^2_{\mu 0}$.[17] The random variable η_{ij} is also assumed to be normally distributed with mean 0 and variance σ^2_η. On receiving a set of signals , an employer uses his prior about μ_i, and the signals, to evaluate each applicant's employer-specific match quality.

Suppose that employer A is unaware of his μ_i. On receiving a set of signals q_{ij}, employer A uses his prior about μ_i, and the signals, to evaluate each applicant's employer-specific match quality. Given his imprecise prior, employer A's best estimate of the expected productivity of the applicant is:

$$\hat{q}_{ij} = \frac{\left(\sigma^2_\eta + \sigma^2_{\mu 0}\right) q_{ij} + \mu_{i0}\sigma^2_\xi}{\sigma^2_\eta + \sigma^2_{\mu 0} + \sigma^2_\xi}. \tag{1}$$

This expression shows that the weight placed on the signal q_{ij} in the estimate of worker quality is increasing with the prior variance $\sigma^2_{\mu 0}$.[18] Since the application-specific noise term ξ_{ij} is a component of the signal, expression (1), therefore, also implies that the weight placed on application-specific noise is increasing in $\sigma^2_{\mu 0}$.

Now consider a second employer, labeled B, who knows with certainty that $\mu_i = \mu_B$. Employer

[17] These may be the priors after seeing all applicants; what is important is the relative variance of the prior across different experience levels.

[18] This is because $\frac{\partial\left(\sigma^2_{\mu 0}+\sigma^2_\eta\right)/\left(\sigma^2_\eta+\sigma^2_\xi+\sigma^2_{\mu 0}\right)}{\partial \sigma^2_{\mu 0}} = \frac{\sigma^2_\xi}{\left(\sigma^2_\eta+\sigma^2_\xi+\sigma^2_{\mu 0}\right)^2} > 0.$

B's expectation of quality given the same signal q_{ij} is

$$\hat{q}_{ij} = \frac{(q_{ij} - \mu_B)\,\sigma_\eta^2}{\sigma_\eta^2 + \sigma_\xi^2} + \mu_B.$$

Assume that employers A and B are the same in every way other than the precision of their beliefs about μ_i, and assume that they receive the same set of signals. We are interested in each employer's willingness to pay for the candidate j, who they judge to be of the highest expected quality among all applicants. Their willingness to pay for the highest-ranked applicant depends on how much they prefer this applicant to the second-highest-ranked applicant, or to the outside option of no hire if this exceeds the expected payoff of the second-highest-ranked applicant.

Because they see the same data, the employers rank applicants in the same order. Take the two best applicants, $j = 1$ and $j = 2$, who have signals q_1 and q_2, with $q_1 > q_2$. Employer A's assessment of the relative quality of applicant 1 is:

$$\hat{q}_{A1} - \hat{q}_{A2} = \frac{\left(\sigma_\eta^2 + \sigma_{\mu 0}^2\right)(q_1 - q_2)}{\sigma_\eta^2 + \sigma_{\mu 0}^2 + \sigma_\xi^2}. \tag{2}$$

Employer B's assessment of the relative quality of worker 1, however, is:

$$\hat{q}_{B1} - \hat{q}_{B2} = \frac{\sigma_\eta^2\,(q_1 - q_2)}{\sigma_\eta^2 + \sigma_\xi^2}. \tag{3}$$

The employers' assessments of the relative quality of applicants 1 and 2, $\hat{q}_{A1} - \hat{q}_{A2} > \hat{q}_{B1} - \hat{q}_{B2}$, are increasing in $\sigma_{\mu 0}^2$. That is, comparing equations (2) and (3) shows that any two applicants appear more differentiated when the employer has uncertainty about his own μ_i. Because the data generating process for the signals is invariant to employer experience, the distribution of signals and the probability of being the top ranked applicant are invariant to $\sigma_{\mu 0}^2$. The resolution of uncertainty affects only the difference in an employer's relative preference for applicant 1.

What about the case when the next best alternative to the preferred applicant is the outside option of not hiring? To analyze how the outside option affects willingness to pay for the most preferred applicant, we normalize the payoff from the no-hire option to zero and assume that $\mu_A = \mu_B$, meaning that employer A's signals are centered on μ_B,. We can compare how much more employer A is willing to pay for applicant 1 versus not hiring by examining the difference in $\hat{q}_{A1}$

and $\hat{q}_{B1}$:

$$\begin{aligned}
&\frac{(\mu_B + \eta_{ij} + \xi_{ij})\left(\sigma_\eta^2 + \sigma_{\mu 0}^2\right) + \mu_0 \sigma_\xi^2}{\sigma_\eta^2 + \sigma_{\mu 0}^2 + \sigma_\xi^2} - \left(\frac{(\eta_{ij} + \xi_{ij})\,\sigma_\eta^2}{\sigma_\eta^2 + \sigma_\xi^2} + \mu_B\right) \\
= &\left(\frac{\sigma_\eta^2 + \sigma_{\mu 0}^2}{\sigma_\xi^2 + \sigma_\eta^2 + \sigma_{\mu 0}^2} - \frac{\sigma_\eta^2}{\sigma_\xi^2 + \sigma_\eta^2}\right)(\xi_{ij} + \eta_{ij}) + \frac{\left(\sigma_\eta^2 + \sigma_{\mu 0}^2\right)\mu_B + \mu_0 \sigma_\xi^2}{\sigma_\xi^2 + \sigma_\eta^2 + \sigma_{\mu 0}^2} - \mu_B. \qquad (4)
\end{aligned}$$

The overall sign of this expression is ambiguous, but comparative statics with respect to μ_B are informative. [19] The comparative static piece of equation (4), when comparing two identical employers but with different experience levels, is based on whether the realized μ_B is greater or less than the population mean, which forms the prior for employer A. That is, the latter (deterministic) term, $\frac{\left(\sigma_\eta^2+\sigma_{\mu 0}^2\right)\mu_B+\mu_0\sigma_\xi^2}{\sigma_\xi^2+\sigma_\eta^2+\sigma_{\mu 0}^2} - \mu_B \geq 0$ if $\mu_0 \geq \mu_B$. The expected surplus for hiring a worker compared to the outside option is increasing with μ_B and is greater for employers with positive ex-post realizations relative to the population mean value.

This analysis focuses on how surplus versus the outside option may change as employers' valuations stay fixed but uncertainty is reduced. The sign is ambiguous and depends on whether the employers' valuation is above or below the population mean valuation. Later, we will also build in the possibility of a "true experience" effect, as discussed in the introduction. With a true experience effect, an employer's μ_i will vary as he gains experience through learning by doing or due to other factors that change the attractiveness of the market to him. Any actual shift in μ_i must be compared against the mean μ_i for inexperienced employers. For experienced employers, a positive shift gives increased surplus relative to the outside option.

When employer experience reduces uncertainty, we also need to allow for the possibility that market experience will change the composition of employers in the market. It is likely that employers who update sufficiently negatively on their own draw of μ_i after early interactions with sellers will opt to leave the market altogether. The employers who return to post subsequent jobs are those more likely to have higher values of μ_i. When bidding to work for experienced employers, applicants face a distribution of market valuations that is likely to be missing density in the left tail of the range of the original μ_i distribution.

[19] With a large number of applicants, the expression can be signed in expectation in some cases. The sign of the first term depends on the sum of random variables $(\xi_{ij} + \eta_{ij})$, since $\left(\frac{\sigma_\eta^2+\sigma_{\mu 0}^2}{\sigma_\xi^2+\sigma_\eta^2+\sigma_{\mu 0}^2} - \frac{\sigma_\eta^2}{\sigma_\xi^2+\sigma_\eta^2}\right)$ is always positive. The term $(\xi_{ij} + \eta_{ij})$ has mean zero in expectation and ambiguous sign. However, when applicant 1 is the most preferred applicant (the only relevant case for the decision of interest here of whether to hire or not), the expected sign of $(\xi_{ij} + \eta_{ij})$ is positive. With normal random variables and two applicants, the maximum order statistic is positive with probability 0.25. With six applicants, the maximum order statistic is positive with probability greater than 0.98.

In Section 4, we will examine in more detail the relationship between the optimal price and the seller's expectation of μ_i conditional on both buyer composition and any true experience effect. However, to outline the intuition here: a higher expected μ_i does not affect the relative willingness to pay for the most preferred worker and the next best applicant, but it does increase an employer's relative preference to hire his most preferred worker versus not hiring. The composition effect and true experience effect *may*, therefore, both lead to higher bids to experienced employers in equilibrium, as every worker knows that there is some chance of her being the most preferred worker competing with the outside option of not hiring.

This relatively stark model, which allows employers to evaluate applicants differently after learning about the market through experience has several implications for the relationship between employer experience and wage bids:

1. If wage bids rise with employer experience, then the composition of employers remaining in the market and the true experience effect for those employers together outweigh the substitutability effect as employers learn about their market value.

2. If wage bids fall with employer experience, then the substitutability effect of decreased perceived differentiation among applicants and the outside option outweighs the sum of the composition and true experience effects.

3. Further, among experienced employers in the market, if those that have received good feedback on prior hires receive higher bids, then their feedback has revealed to applicants that these employers have a high value for the market and are less likely to select the outside option.

We note that allowing employers to gather information from interviews as well as from hiring leads to a richer set of predictions related to employer experience, some of which are examined in later sections of the paper. The next section discusses the setting and data used to evaluate these implications.

3 The Setting: oDesk.com

3.1 How it Works

oDesk.com is an online platform that allows employers to contract with remote workers who are selling online labor services.[20] The platform facilitates search and matching, remote task and project management, and payments. Work includes a range of jobs in which output can be delivered electronically, with the most frequently observed job categories being Web Development and Administrative Support. Jobs tend to be short-term spot transactions, and the majority of postings require less than three months of work. Around 85 percent of the transactions in the market span international borders and, therefore, constitute international labor services trade.[21]

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

Interested workers submit applications for the job posting and bid an hourly wage to work on the specific job. Employers also have the option of searching worker profiles directly and inviting applications from individual workers. Workers' profiles, visible to potential employers, contain information about their skills, education, prior offline work experience, and experience on oDesk (see Figure 2). The country where they are located is also displayed prominently on the profile. The profiles of workers that have prior experience on the site also show summary feedback scores based on a five-point scale.

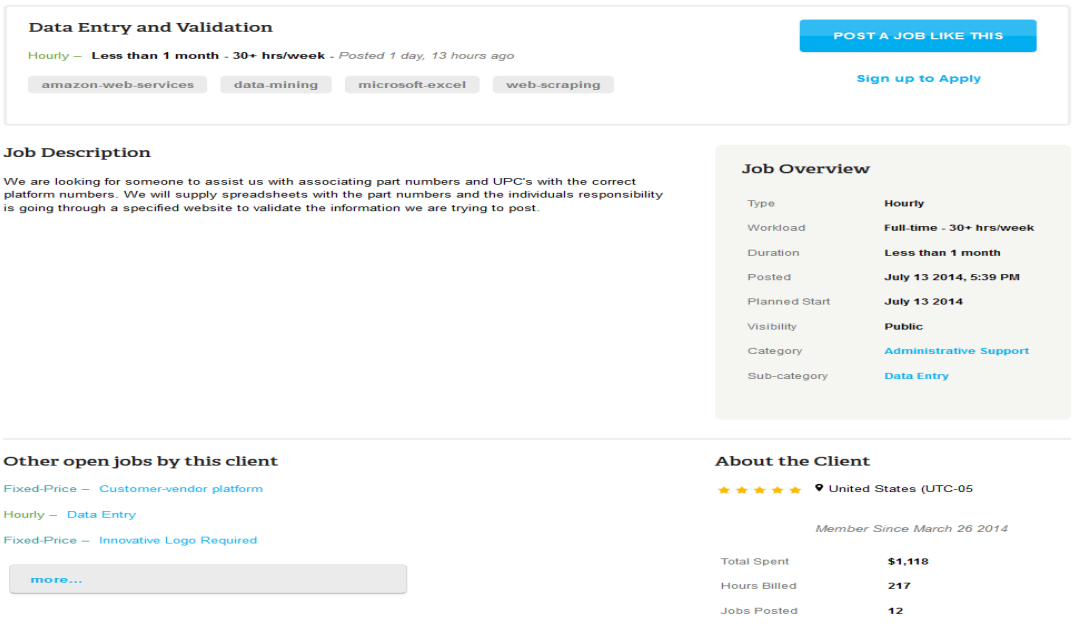
After receiving applications and initiating candidacies, employers can request interviews with any number of applicants for the job. If applicants agree, the interview usually takes place via Skype. Whether an interview actually occurs is not recorded in the oDesk database, so the remainder of the paper will refer to an interview request as an actual interview.[22] An employer may choose to hire an

[20] Other prominent platforms included elance and Guru, the first of which merged with oDesk in 2014. The merged company has recently changed its name to Upwork. The data used in this paper predate the merger.

[21] While several recent papers use data from online labor markets (Agrawal et al., 2013a and 2013b; Ghani, Kerr, and Stanton, 2014; Horton, Kerr, and Stanton, 2017; Lyons, 2017; Pallais, 2014; and Stanton and Thomas, 2016), many focus on information frictions at the worker level. The main exceptions are Horton (2017a) and Horton (2017b).

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

Figure 1: A Job Posting.



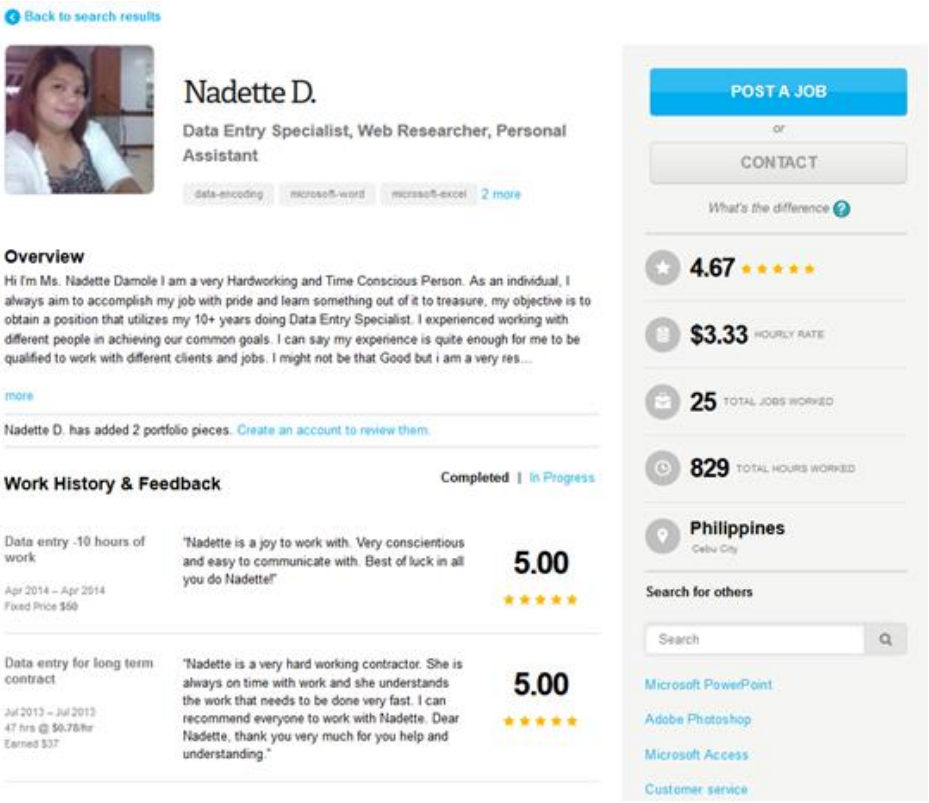
applicant with or without interviewing her first. Upon hiring a worker, the employer can monitor the work as it is performed via software provided by oDesk, and then oDesk manages all payments for completed work. When a job is complete, the employer is asked for feedback about the worker and vice versa. The employer is also asked whether or not the job was completed successfully.

The data used in this paper are administrative data from the company, which means that every employer's job-specific search process is observed on each of his successive job postings. For each posting, the data contain information about the entire applicant pool; which candidates, if any, are interviewed; and which candidate, if any, is hired. This paper does not use information about on-the-job monitoring, but does exploit information on feedback after job completion and indicators of job success.

3.2 Summary Statistics by Employer Experience

The sample consists of job postings and applications between January 2008 and June 2010. It contains 82,257 unique potential and actual employers who posted 322,870 jobs over the time period studied and received more than 5 million job applications. There are nine job categories in the data. 38% of all job postings are in Web Development, the largest technical job category. The next largest technical job category, Software Development, contains 9% of all the job postings.

Figure 2: A Worker Profile.



Administrative Support is the largest non-technical category, with 17% of the postings.[23]

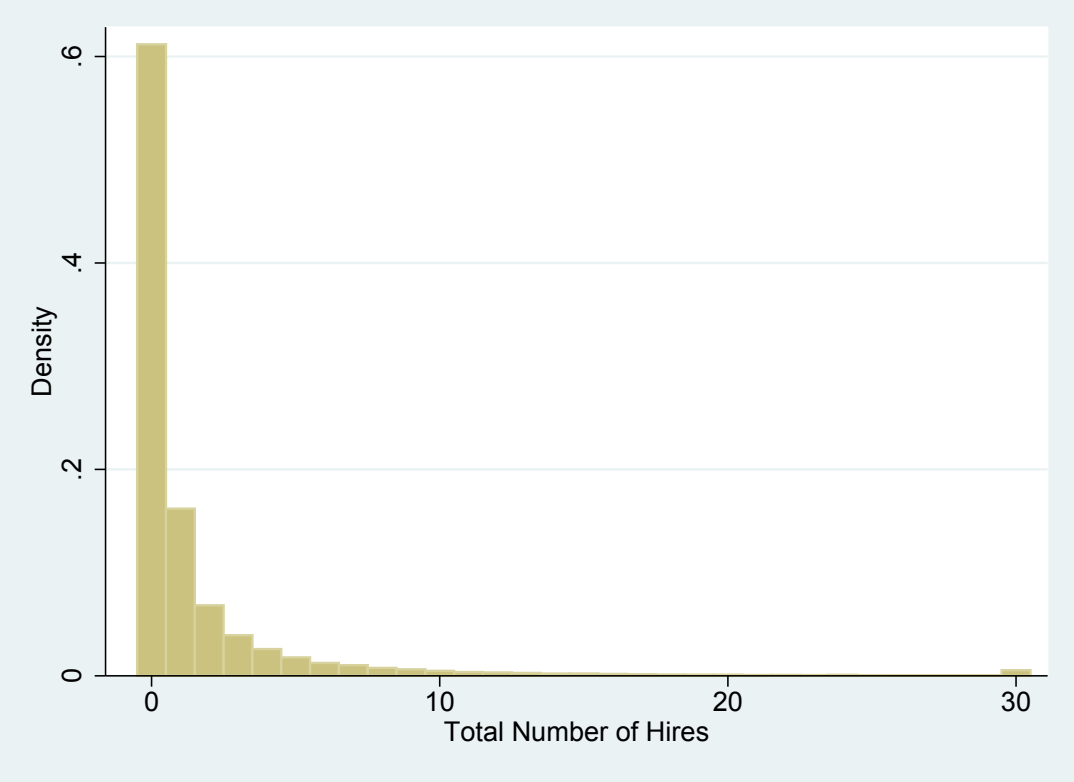
Most employers are located in the United States and include private individuals as well as individuals hiring on behalf of established firms. Figure 3 presents the distribution of the number of hires per employer throughout the period. 63% of employers make no hire at all, and 17% of employers make five or more hires. On each employer's job posting page, the number of past hires and the total hours are prominently displayed (see Figure 1). Workers can also observe any feedback that employers have received from previously employed workers.

For the hiring probability estimation that follows in Section 4, we use a subset of all job postings. Some employers hire multiple workers simultaneously by posting a batch of jobs at the same time, thus making it impossible to observe the choice set for each opening. When estimating hiring probabilities, we restrict the job openings in the sample to those that received at least one worker-initiated application, had multiple applicants to consider, and then had a gap of at least one day before or after the same employer posted a different opening.[24] Of the 119,877 jobs posted by em-

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

[24]We also apply filters that remove job openings that appear to violate "arms-length" hiring. These violations occur when employers hire workers from a previous contract, appear to be bringing workers onto the platform from

Figure 3: Number of Hires per Employer



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

ployers without previous hires on the site, 61,160 survived these restrictions; and the total number of postings across all employer experience levels fell from from 322,870 to 109,764. Throughout the paper, this subset of jobs is referred to as the non-overlapping or sequential sample. The following analysis presents summary statistics for both the overall and sequential samples.

Tables 1A and 1B present summary statistics about job postings, grouped by the number of previous hires made by the relevant employer. Table 1A includes all job posts in the sample, while Table 1B includes only the sequential, arms-length job sample. These tables provide initial evidence that employer—and worker—behavior differs according to the number of previous hires that the employer has made on the site.

The first row of Table 1A summarizes new employers' behavior on the platform and how workers interact with these employers. The mean hourly wage bid to employers without prior hiring experience in the overall sample is 10.16 USD. In the sequential jobs sample, summarized in Table 1B, the mean bid is 10.25 USD. Inexperienced employers in the overall (sequential) sample receive an average of 18.39 (25.47) applications prior to closing the job or making a hire. Only a small

offline, or have posted a publicly visible job but intend to hire a known worker.

fraction of applications are employer-initiated, particularly in the sequential sample, which means that the vast majority are initiated by the applicant. 34% (31%) of applications are from workers who have a current feedback score above 4.5, termed good feedback, at the time of application, and 44% (47%) of applicants have no recorded feedback.[25] New employers conduct an average of 2.96 (3.64) interviews in the full (sequential) sample.

Column 8 shows that only 22% (16%) of openings posted by inexperienced employers result in a hire. Columns 9 through 11 present data on jobs for which a hire was made. The mean hourly wage paid (Column 9) is slightly lower than the mean bid received; the hired worker is more likely to have received good past feedback and is less likely to have missing feedback than the typical applicant (Columns 10 and 11).

The subsequent rows in Tables 1A and 1B summarize data on applicants and the hiring processes of employers with at least one prior hire. As employers gain experience, the number of applications per job falls in the overall sample, with the largest decline taking place between having hired zero and one prior worker (Column 2), but the number of applications increases with employer experience in the sequential sample.

Column 4 presents the first fact that motivates this paper: the mean hourly bid declines with employer experience. For those with one prior hire, the average bids fall to 9.77 USD (9.85 USD in the sequential sample). After four or more previous hires, employers receive a mean bid of 8.71 USD (9.06 USD in the sequential sample), a decline of 14% (12%). Among employers with prior hires, similar shares of applications come from workers who have good feedback (Column 5) or no observable feedback (Column 6), regardless of the number of past hires. Experienced employers typically conduct fewer interviews than the inexperienced (Column 7) in the overall sample, and employers with at least four prior hires conduct an average of 2.06 interviews. In the sequential sample, the reverse is true; the number of interviews increases with experience.[26]

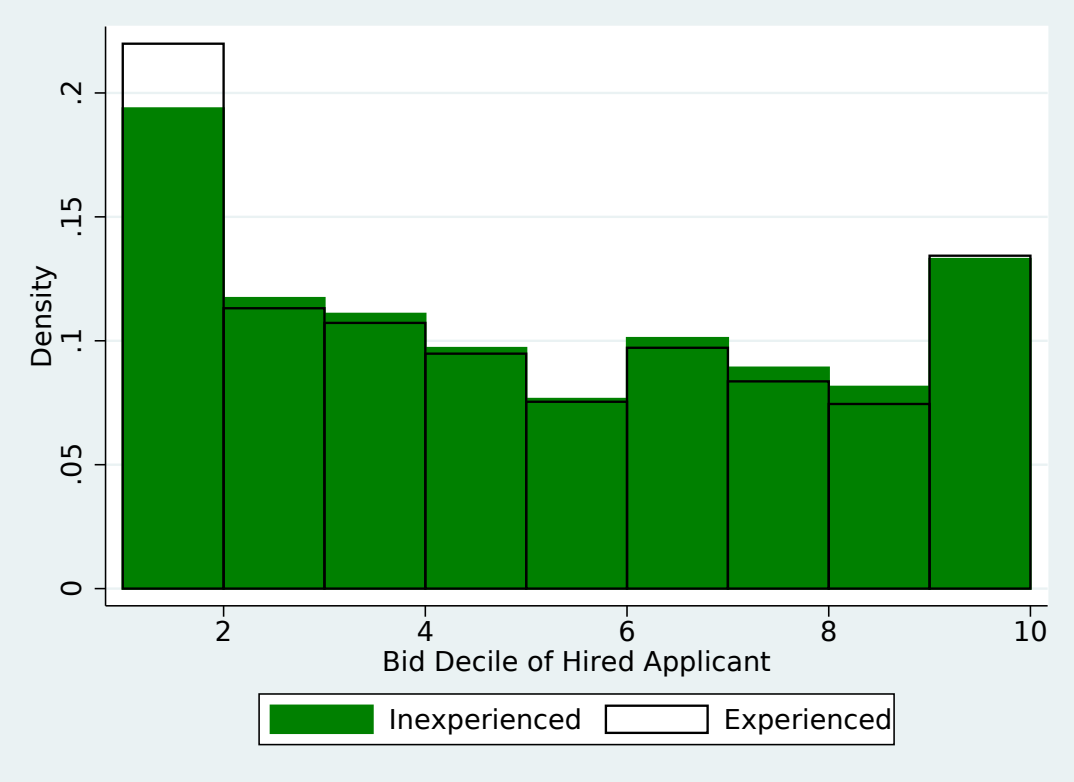
Column 8 presents the second motivating fact for this paper: experienced employers are far more likely than inexperienced employers to make a hire. Among employers with at least four prior hires, 57% (28%) hire a worker. The mean wage paid to the hired worker also declines with employer experience (Column 9); on average, the wage paid by employers with four or more prior hires is 10% (8%) lower than the wage paid by the inexperienced.

[25]The feedback distribution is skewed, and the cutoff of 4.5 is between the mean (4.40) and the median (4.68) of feedback scores observed in the data among active workers with past feedback. See Horton and Golden (2015) for an analysis of how reputation scores have evolved over time in this market.

[26]Later, we show that this is due to employer selection into making repeat postings. Within-employer, interviews fall in the overall and sequential samples.

The data also provide some evidence that workers appear differentiated to employers. It is relatively rare for employers to hire the worker who submits the lowest hourly bid. Figure (4) shows the distribution of the wage bid decile, calculated within each job opening, of the worker who was ultimately hired. The solid bars of the histogram present these findings for employers with no prior hiring experience. Around 13% of employers hire a worker whose bid is in the top decile of the distribution of all bids received 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 4: 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, the decile of wage bid for each applicant is computed. 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).

3.3 Experienced Employers Receive Lower Wage Bids

Table 2 presents more detail about the decline in hourly bids with employer experience. These bids are the final offers submitted to employers and there is very little change from an applicant's initial

bid to the wage actually paid at the start of a job if she is hired.[27] The dependent variable in these regressions is the log of an applicant's bid, meaning that the coefficients capture the percentage difference between the hourly wage bids received by inexperienced employers and the bids received by employers after hiring one, two, three, four, or five-plus workers. The first column presents differences in wage bids received by employers of different experience levels after including fixed effects for job categories and calendar time. Employers receive bids that are approximately 3% lower after making one hire. Bids continue to fall after subsequent hires, and employers who have made five or more previous hires receive bids that are 7.6% lower than those received by employers who have not previously hired.

Column 2 adds controls for detailed worker resume characteristics, controls for who initiated the application, and a third-order polynomial in the number of characters included in the job description. These controls are intended to allow for worker sorting and for the possibility that the employer is learning how to write an effective job posting. The observed decline in average wage bids received by employers with five or more prior hires is 5.9%.

Subsequent columns of Table 2 add further controls to account for other plausible explanations for the decline in wage bids with employer experience. Columns 3 and 4 add employer fixed effects. A given employer receives bids that are 4.7% lower when he has five or more prior hires than when he has no experience. Adding again the controls for worker and job characteristics that were included in Column 2 reduces the effect to 3.7%, indicating sorting on observable characteristics and employer experience rather than sorting simply on the identity of the employer. 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.

Appendix 2 discusses evidence on some possible alternative explanations for these findings. First, the bid premium to inexperienced employers does not arise because of a positive correlation between employer experience and the number of job applicants. The bid premium remains when controlling for the number of job applicants in the first 24 hours after job posting, as shown in Table A2. Second, the bid premium does not appear to be driven by time-varying unobservables about

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

how employers post jobs. A field experiment designed to isolate the effect of observable employer experience from environmental changes—detailed in Appendix 2 and with results in Appendix Table 3—shows large wage-bid declines to experienced employers.

3.4 Among Experienced Employers, Those with Good Feedback Receive Higher Wage Bids

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

Somewhat surprisingly, applicants actually make higher bids to experienced employers with good feedback than to experienced employers with poor feedback. Experienced employers who have not received feedback also receive higher bids than employers who have poor past feedback. In line with Section 2, we interpret these results as consistent with poor employer feedback signaling a reduced likelihood to hire that lowers the optimal bid to make to these employers. Relative to an experienced employer who has received poor feedback, the likelihood of hiring an applicant on a subsequent job increases by 6.7 percentage points (about 13%) for experienced employers with good feedback and 3.7 percentage points (roughly 8%) for experienced employers with no feedback.[28]

Table 3, in which the log bid is the dependent variable, presents two sets of results about employer experience and employer feedback. First, even controlling for the feedback received from past workers, wage bids decline robustly with employer experience. Only employers with hiring experience can have a feedback score, but many experienced employers have no feedback from past hires. Indicators for the employer having no observable feedback and for the employer having observable feedback of 4.5 or higher are interacted with an indicator for having observable prior hiring experience. Thus, the baseline group in the regression is experienced employers who have feedback scores lower than 4.5. The interaction terms capture deviations from the baseline for experienced employers with good feedback and for those with no feedback. The baseline point

[28]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.

estimates in Columns 2 through 6, which present the effect of experience for employers who have bad visible feedback, show that these employers receive bids that are significantly lower than those made to inexperienced employers. The positive coefficients on good feedback and no observable feedback are inconsistent with an adverse selection channel that would predict lower bids for employers with good feedback.[29]

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

4 Estimating New and Experienced Employers' Hiring Probabilities

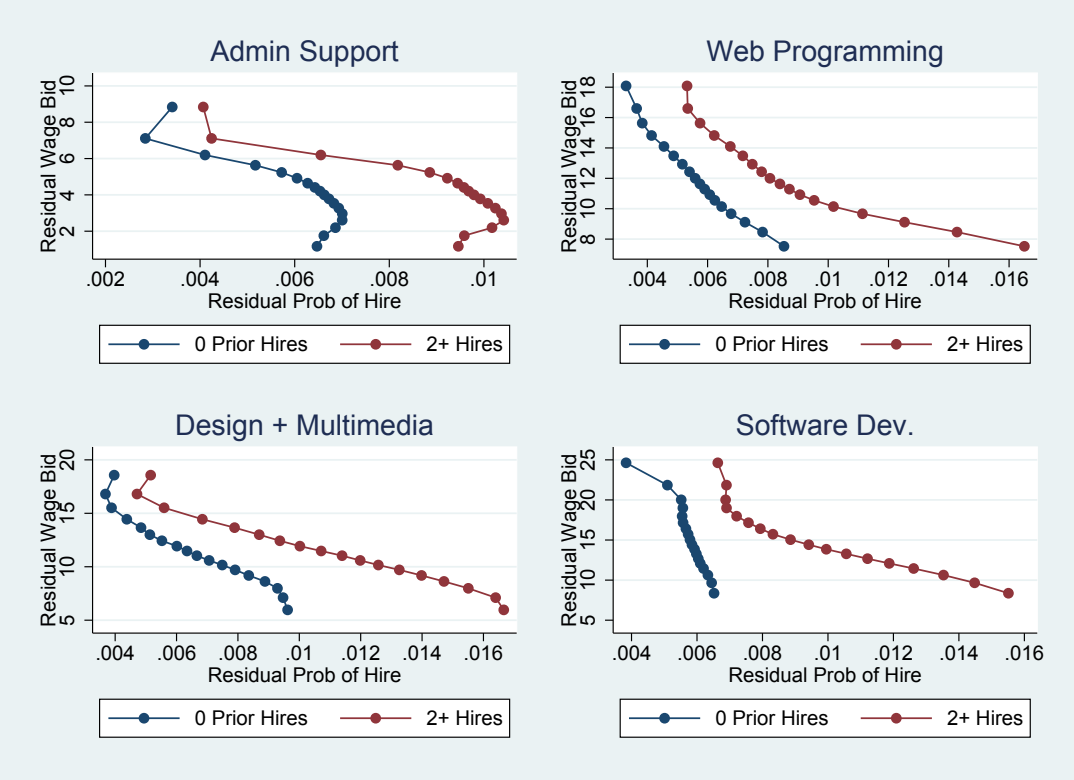
Figure 5 plots hiring probabilities at different wage bids for new and for experienced employers. Each of the four panels shows a different job category. Overall, both the slopes and the intercepts of the hiring probability functions differ by employer experience within a job category, with experienced employers appearing to have more elastic hiring probabilities over the range of (residualized) bids that are most common in the data.

To evaluate these differences systematically, we specify a hiring probability function with parameters that can vary with whether the job was posted by a new or by an experienced employer. For a job posting by employer i with experience χ, $p_{i\chi j}$ denotes the probability that applicant j is hired for that job. The function $p_{i\chi j}$ depends on observable characteristics of the applicant, her wage bid, the characteristics of the job, and the composition of the applicant pool. The subscript χ allows the function to differ by employer experience segment.

Using data on wage bids and the estimates of employer hiring probabilities allows us to decompose the bids made to experienced and inexperienced employers into workers' costs and their markup over marginal costs for each segment.

[29] We also repeated this exercise with the feedback that employers have left for previously-hired workers (although this information is harder for future applicants to observe). The results have similar signs, and the feedback received and the feedback left by employers are positively correlated.

Figure 5: 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.

4.1 A Worker's Optimal Bid

A worker j's cost, $c_{i\chi j}$, captures her outside option (opportunity cost of work) or hassle costs from applying to and/or being hired for the job posted by employer i with experience level χ. When choosing the wage bid, her objective function takes this cost into account along with the hiring probability. She chooses the wage bid, $w_{i\chi j}$, that maximizes

$$\underbrace{p_{i\chi j}}_{\Pr(hired)} \times \underbrace{\exp\left(\log w_{i\chi j} - \log(1+\tau)\right)}_{Post-fee\ wage} + (1 - p_{i\chi j}) \times \underbrace{c_{i\chi j}}_{Opportunity\ cost} \tag{5}$$

where $\log w_{i\chi j}$ is the log of the wage bid inclusive of ad-valorem platform fees, τ. Accounting for oDesk's fee, if she is hired, the worker receives the wage $\frac{w_{i\chi j}}{(1+\tau)} = \exp\left(\log w_{i\chi j} - \log(1+\tau)\right)$ and the employer pays $w_{i\chi j}$. If worker j is not hired, she receives $c_{i\chi j}$, her "net" outside option, which includes the opportunity cost of her alternative use of time, along with the expected direct costs of

interviewing or working on the job.[30] When choosing her wage bid, the worker trades off a higher wage and a lower probability of being hired.

Each worker's first order condition is given by

$$\frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}} \left(\frac{w_{i\chi j}}{(1+\tau)} - c_{i\chi j} \right) + p_{i\chi j} \frac{w_{i\chi j}}{(1+\tau)} = 0. \tag{6}$$

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

$$w^*_{i\chi j} = c_{i\chi j} (1+\tau) \left(1 + p_{i\chi j} / \frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}} \right)^{-1}. \tag{7}$$

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

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

$$c_{i\chi j} = \frac{w_{i\chi j}}{(1+\tau)} \left(1 + p_{i\chi j} / \frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}} \right), \tag{8}$$

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

4.2 Employer Hiring Probabilities

The set of applicants to an opening is denoted $J_{i\chi}$, which is taken as given, and workers are assumed to be available when they initiate an application.[31] We also assume that the employer

[30] The objective could alternatively be written:

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

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

[31] This assumption is reasonable, requiring only that the probability that a worker receives two offers over a time interval Δ is small. For example, if, from the worker's perspective, job offers follow a Poisson process, then the probability of receiving a job offer in the interval $(t, t+\Delta)$ is $\lambda\Delta + o(\Delta)$, and the probability of two offer arrivals in $(t, t+\Delta)$ is $o(\Delta)$.

Although we do not observe declined job offers, this assumption on arrival of offers seems to be reasonable in the data. The observable arrival rate of interview requests fits this assumption. When a worker-day is the unit of analysis, only 3.6% of the worker-days in the sample have more than one interview request arriving; only 0.6% of the

observes signals of applicant quality q_{ij}. We adjust the structure of q_{ij} from that used in Section 2 by including a χ subscript to account for worker characteristics and changing perceptions of the utility of using the platform, introducing an experience subscript, χ. We assume that $q_{ij\chi}$ is centered on an experience-specific mean, with an employer-specific offset, μ_i, which varies the value of hiring on the platform.[32] Because μ_i shifts employer i's value of hiring any applicant, this term does not alter his relative ranking over applicants for a given job, but it does determine the value of hiring an applicant compared to not hiring in the market.

An employer's objective is to choose the worker with the highest perceived quality per unit of wage, subject to this being greater than the value of the off-platform option (denoted option zero, with payoff normalized to 1 in levels and 0 in logs). His objective function takes the form

$$\max_{j \in \{J_{i\chi}, 0\}} \frac{q_{ij\chi} \exp(\mu_i) \exp(\varepsilon_{ij})}{(w_{i\chi j})^{\alpha_\chi}}.$$

If we were to use the setup in Section 2 with normally distributed stochastic signals, this objective could be viewed as the foundation of a multinomial probit model. However, with many alternatives to consider per job opening, the numerical integration required in the multinomial probit model is not feasible for estimation. Instead, the term ε_{ij} is assumed to be an idiosyncratic Type-1 extreme value shock for each alternative. Because the type-1 extreme value shock is iid, we approximate the structure motivating the exercise by allowing all model parameters to depend on the employer experience level, χ.[33]

This setup for the employer's objective, after taking logs, gives a conditional logit function for the hiring probability. The probability that worker j is hired by employer i with experience χ is

$$\log(q_{ij\chi}) + \mu_i + \varepsilon_{ij} - \alpha_\chi \log(w_{i\chi j}) \geq \log(q_{ik\chi}) + \mu_i + \varepsilon_{ik} - \alpha_\chi \log(w_{i\chi k}) \tag{9}$$

for all $k \in \{J_i, 0\}$. Conditional on μ_i and $q_{ij\chi}$, the employer's probability of hiring applicant j is

worker days have more than two interview requests arriving. A post-candidacy survey also asks employers for reasons why particular workers were not hired and workers for 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.

[32] None of the other the stochastic assumptions about the structure of quality signals made in Section 2 for illustrative purposes will be imposed in estimation.

[33] Allowing the parameters to vary flexibly with experience approximates the changing variance with experience that would be captured explicitly in the multinomial probit model. This is because the scale of the error is normalized, so a change in a parameter with experience reflects a changing scale relative to the variance normalization.

given by:

$$\begin{aligned} p_{ij\chi} &= \exp\left(\log\left(q_{ij\chi}\right) + \mu_i - \alpha_\chi \log\left(w_{i\chi j}\right)\right) / \\ &\quad \left(1 + \Sigma_j^{J_i} \exp\left(\log\left(q_{ij\chi}\right) + \mu_i - \alpha_\chi \log\left(w_{i\chi j}\right)\right)\right). \end{aligned} \tag{10}$$

4.3 Estimation

Our objective is to estimate the hiring probability function for two segments of employers—inexperienced and experienced—with flexibility in the parameter estimates for observed worker and job characteristics. It is also necessary to account for worker heterogeneity. We allow worker quality $q_{ij\chi}$ to be a function of the worker-level data that the employer observes in the application, so that

$$\log\left(q_{ij\chi}\right) = X_j \beta_\chi.$$

The matrix X_j contains a rich set of resume, worker, and job characteristics, including a constant term. The subscript χ on β_χ allows the relationship between values of X_j and perceived worker quality to vary with employer experience. This flexibility also applies to the estimate of employers' sensitivity to wage bids, α_χ.

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

Several normalizations allow us to identify the coefficients. First, the variance of the error term ε_{ij} is normalized so that the estimated coefficients reflect a scaling by this error variance. Second, if expectations about μ_i (and $q_{ij\chi}$) were fully observable to the researcher for each employer group, variation in prices, characteristics, and choices would identify the parameters. Because we don't observe levels of μ_i, only the population distribution of types can be identified. For example, if many employers persistently hire when faced with low-quality workers submitting high bids while many other employers do not hire when high-quality workers with low bids are available, then the estimated population distribution of types would have wide dispersion in valuations.

In the exposition in Section 2, the distribution of μ_i was assumed to be normal, and other assumptions were imposed on the nature of applicant signals, including the noise term. In estimation,

we confront the selection issue in that those employers who return to the market after hiring once are a non-random sample from the population distribution of μ_i. To account for this, we assume that returning employers are drawn from a different distribution, and we estimate how the distribution of μ_i varies between the segment that is ever experienced and the segment that is never experienced. Because the experienced distribution may truncate the lower tail, assuming symmetry for each distribution may not accurately characterize selection. Instead, we use a finite mixture model in which employer types are allowed to depend on whether the employer is ever observed in the experienced segment.

The remaining difficulty is handling any aspect of worker-level quality that the researcher does not observe but that is relevant in the employer's choice problem. Suppose that log quality in the data is $\log(q_{ij\chi}) = X_j\beta_\chi$, but the employer observes $\log(q_{ij\chi}) = X_j\beta_\chi + \upsilon_j$. If υ_j is correlated with log bids, this will result in inconsistent estimates of both the price elasticity, α_χ and the β_χ vectors. The following subsection addresses this endogeneity concern and discusses the instrumentation strategy.

4.4 Instruments

To account for the endogeneity of wage bids arising from worker characteristics that the researcher does not observe, we use an instrumental variables strategy based on changes in the dollar-to-local-currency exchange rate for each applicant. Workers are paid in their local currency for offline work but are paid in dollars for their work on oDesk. Frictions limiting exchange rate pass through to local wages mean that offline opportunities are likely to adjust to exchange rates more slowly than online transactions.[34] When the dollar appreciates relative to the local currency, so that one dollar earned on the site provides fewer local currency units, workers' wage bids are predicted to increase.

To see this how this variation affects workers' bids, assume that $c_{i\chi j}$ is denominated in the local currency, while the bids observed by employers are denominated in dollars. Costs in the local currency must be translated into dollars when submitting bids, so the worker's optimal bid becomes $w^*_{i\chi j} = c_{i\chi j}\left(\frac{D}{L}\right)^\theta (1+\tau)\left(1 + p_{i\chi j}/\frac{dp_{i\chi j}}{d\log w_{i\chi j}}\right)^{-1}$. The dollar-to-local-currency exchange rate is $\frac{D}{L}$, and the parameter θ captures possible reasons for deviations from complete pass through.[35]

The worker's optimal log bid can be written as a mapping from local-currency-denominated

[34] 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.

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

opportunity costs to dollar-denominated bids as

$$\log(w_{i\chi j}) = \theta \log\left(\frac{D}{L}\right) + \log(c_{i\chi j}) + \log(1+\tau) - \log\left(1 + p_{i\chi j}/\frac{\partial p_{i\chi j}}{\partial \log w_{i\chi j}}\right), \tag{11}$$

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

Figure 6 illustrates the time-series variation in mean residual log bids and detrended exchange rates that underpins this estimation approach. The top left panel plots the mean residual log bid made by job applicants located in India and the log of the US dollar to Indian rupee exchange rate. The bid and exchange rate time series move together over the time period studied. The other panels in Figure 6 plot the difference between the mean of residualized bids from applicants in India and those from applicants in five other common worker locations (the dots in the figures), along with the difference in the exchange rate between the relevant local currency and the Indian rupee (the crosses in the figures). The figure suggests that the relevant local currencies for the job applicants in the data demonstrate independent dollar exchange rate variation during the time period.[36]

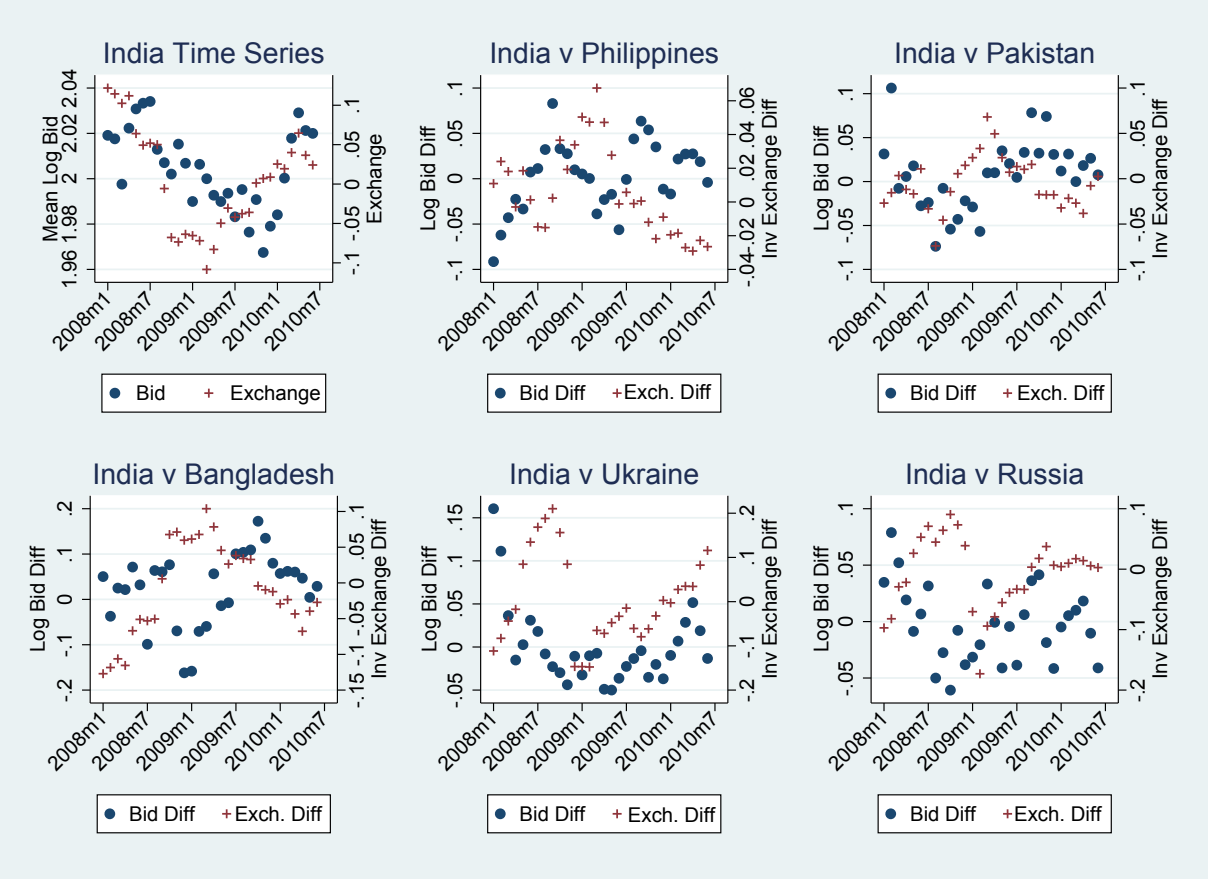
While exchange rate movements are plausibly exogenous to demand on oDesk, there are two additional concerns. First, a subset of job applicants—those based in the US—do not face any dollar cost shocks resulting from exchange rate variation. This does not pose a problem for comparing relative bid differences across workers as long as some are subject to exchange rate variability. However, for American applicants or those from countries with currencies pegged to the dollar, there is no wage bid variation relative to the employer's outside option of not hiring.

An additional instrument for worker j's bid that is relevant to all workers, including those with dollar or dollar-pegged local currencies, captures exogenous variation in the intensity of competition for employer i with experience χ. It is based on the average number of application arrivals in the first 24 hours after posting to other jobs in the same category in the same week. Job category and week fixed effects are netted out of the application numbers, making this an instrument that varies the extent of competition across job categories holding fixed the average competition in the category over time and the average competition across all job categories during the week in question.

The other concern about the identification strategy is that there is likely to be sorting on the instruments that influences which workers apply to employer i with experience χ. For example, an

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

Figure 6: 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).

appreciation of a local exchange rate may lead a non-random set of potential applicants to seek work elsewhere or to use the platform. Because many workers submit only a small number of total applications, it is difficult to assess the extent of selection into application based on unobservable worker characteristics and exchange rate variation. However, assessing the sensitivity of the parameter estimates to the inclusion of different sets of observable worker characteristics offers some insight about whether worker selection into application biases the estimates. Below, we discuss a comparison of the estimates with and without worker-level resume data for this reason.

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

the form

$$\log(w_{ij}) = \gamma_0 + Z_j\gamma_{1\chi} + X_j\gamma_{2\chi} + \nu_{ij\chi}. \tag{12}$$

The coefficients in Equation 12 are estimated separately for the group of new and experienced employers.

The results in Table 4 show that both instruments have a substantive and statistically significant effect on workers' bids. The first column provides estimates for inexperienced employers, including many resume and job controls. Column 2 provides the analogous estimates for experienced employers. The signs of the estimated coefficients are as expected. Bids increase when the local exchange rate increases and decrease with the level of competition on the job. In both cases, the F statistics are extremely large, indicating the strength of the instruments.

Columns 3 and 4 exclude the detailed worker resume data—those columns of X_j in Equation (12) that contain worker characteristics. A comparison of the estimated γ_1 coefficients between those given in the first two columns provides some evidence about the extent of sorting on the instrument. Under the null of no sorting, the estimated parameters in Columns 1 and 3 and in Columns 2 and 4 would be statistically indistinguishable. The estimated parameters, while of the same sign in both panels, are larger in absolute magnitude in the columns that exclude the worker characteristics suggesting that there is some sorting into applying in response to exchange rate and/or competition variation.[37] Later, estimation of the hiring probability function with and without worker characteristics will help us to map this sorting into sensitivity over hiring elasticity.

4.5 Estimating Choice Parameters

This section presents the estimation approach for the full model with employer heterogeneity that is unobserved by both inexperienced employers and the econometrician. This approach requires forming a likelihood over sequences of employer choices. We will present some estimates from simpler models estimated without heterogeneity; for these models, only the likelihood at the job-opening level is relevant. The step-by-step approach is as follows.

First, the residuals from Equation (12) form control functions for unobserved worker quality, denoted $CF_{ij} = \hat{\nu}_{ij}$. Second, we form choice probabilities conditional on a value of the unobserved

[37] Appendix Table 5 repeats this analysis including worker fixed effects, but the incidental parameters problem means that control functions including the fixed effects are not consistent. The instruments remain strong in these specifications.

term μ_i, taking the form[38]

$$\begin{aligned} p_{i\chi j} &= \exp\left(X_j\beta_\chi + \mu_i - \alpha_\chi \log\left(w_{ij}\right) + \psi_\chi CF_{ij}\right) / \\ &\quad \left(1 + \Sigma_j^{J_i} \exp\left(X_j\beta_\chi + \mu_i - \alpha_\chi \log\left(w_{ij}\right) + \psi_\chi CF_{ij}\right)\right). \end{aligned} \tag{13}$$

Third, we assume that μ_i is drawn from a distribution with three distinct types: $\mu_i \in \{\beta_{0\chi}, \beta_{0\chi} + \mu_2, \beta_{0\chi} + \mu_3\}$. For the first type, $\mu_1 = \beta_{0\chi}$ is a constant term that may vary with employer experience. For the other two types, the deviation from $\beta_{0\chi}$ remains constant with experience. That is, a type 2 employer, who hires on the first job and then posts two additional jobs will have $\beta_{0\chi=I} + \mu_2$ on the first opening and $\beta_{0\chi=E} + \mu_2$ on subsequent openings. Fourth, we then allow an extremely flexible pattern of selection by letting the type probabilities depend on the eventual experience of an employer, using superscript $S = E$ to denote the ever-experienced group of employers and $S = N$ to denote the never-experienced group. Hence, for an employer who is ever observed in the experienced sample, the type probability vector for that employer is $\rho^{S=N} = (\rho_1^{S=E}, \rho_2^{S=E}, \rho_3^{S=E})$. On the other hand, for an employer is never observed in the experienced sample, the type probability vector for that employer is $\rho^{S=N} = (\rho_1^{S=N}, \rho_2^{S=N}, \rho_3^{S=N})$.[39]

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

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

The i subscript on the product term is a slight abuse of notation; the product of probabilities is taken over all openings posted by the employer. The likelihood contribution for employers who are never observed in the experienced segment is $L = \Sigma_{k=1}^K \rho_k^{S=N} \Pi_i \Sigma_j p_{ij\chi}\left(\mu_k\right)^{y=j}$.

[38] In specifications without employer heterogeneity, μ_i is set to zero.

[39] In the presentation of the results, employer type 1 has a value of μ_i, while type 2 employers have the lowest type and type 3 the highest. Hence, μ_2 can be thought of as negative and μ_3 as positive.

4.6 Results

4.6.1 Markups and Wage Bid Sensitivity

Table 5 presents the results for the hiring probability estimation on the sequential openings sample. There are 61,196 postings by employers with no experience and 48,618 openings by employers with prior experience. The coefficient on the log hourly bid across columns differs by the employer's experience level at the time of posting. The odd-numbered columns show estimates of $\alpha_{\chi=I}$ for employers who have never hired before on the site. The even-numbered columns present deviations from the inexperienced segment (that is, parameter estimates on interactions with experience) for the experienced segment. The important results to note relate to differences in estimated coefficients for the two groups. A comparison across each pair of columns shows that α_χ, the estimated price coefficient, differs substantially and significantly by the employer's experience level at the time of posting. Panel B presents some of the implications of these estimates.

Columns 1 and 2 present results without instrumenting for the wage bid. In these models, the likelihood is at the job-opening level, without taking the product over the sequence of employer job openings. Both the inexperienced and experienced employers in Columns 1 and 2 have inelastic wage bid demand.[40]

Columns 3 and 4 include the estimates after including the control function from the first stage estimates. The estimated coefficients and implied wage elasticities become more reasonable: inexperienced employers have an estimated wage elasticity of -5.46, while experienced employers are more wage elastic, with an estimated elasticity of -8.54.

The remaining columns of Table 5 present the full model results. Columns 5 and 6 exclude worker resume characteristics from the first stage (corresponding to Columns 3 and 4 in Table 4), and Columns 7 and 8 include these characteristics (as per Columns 1 and 2 of Table 4). Columns 6 and 8, which present additive interaction terms relevant for the experience employers relative to the inexperienced), show in each case that the experienced employers are more wage elastic than the inexperienced. In the specifications in Columns 7 and 8, the mean own-bid elasticity for inexperienced employers is -4.996. The experienced employer segment is more elastic, with an estimated own-bid elasticity that is larger in absolute magnitude by 2.776, at -7.772.

These estimated parameters imply differences in the markups. Despite inexperienced employers' lower probability of hiring, the estimated average markup over cost included in inexperienced

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

employers' bids is 25.2 percent, compared to 14.9 percent for experienced employers. This reflects the different wage elasticities in each segment of employers.

Following the construction of the equilibrium bids in Section 4.1, the observed bids and estimated markups allow for an estimate of worker costs. These estimated costs also differ by employer experience. Referring again to Columns 7 and 8 in Table 5, the mean cost of working for an inexperienced employer is 7.43 USD per hour (before the oDesk fee), compared to 7.35 USD per hour in the experienced sample. This implies that the additional expected hassle cost of applying to an inexperienced employer is limited, at about 0.08 USD per hour. The results suggest that both demand and supply side effects play a role in the new employer bid premium, although of differing relative importance. In the preferred estimates in Columns 7 and 8, around 88% of the premium can be attributed to the higher markups set by workers who anticipate that new employers have relatively inelastic hiring probabilities, meaning that only 12% is due to the higher expected costs of applying to new employers.

4.6.2 Workers' Sorting and the Instrument

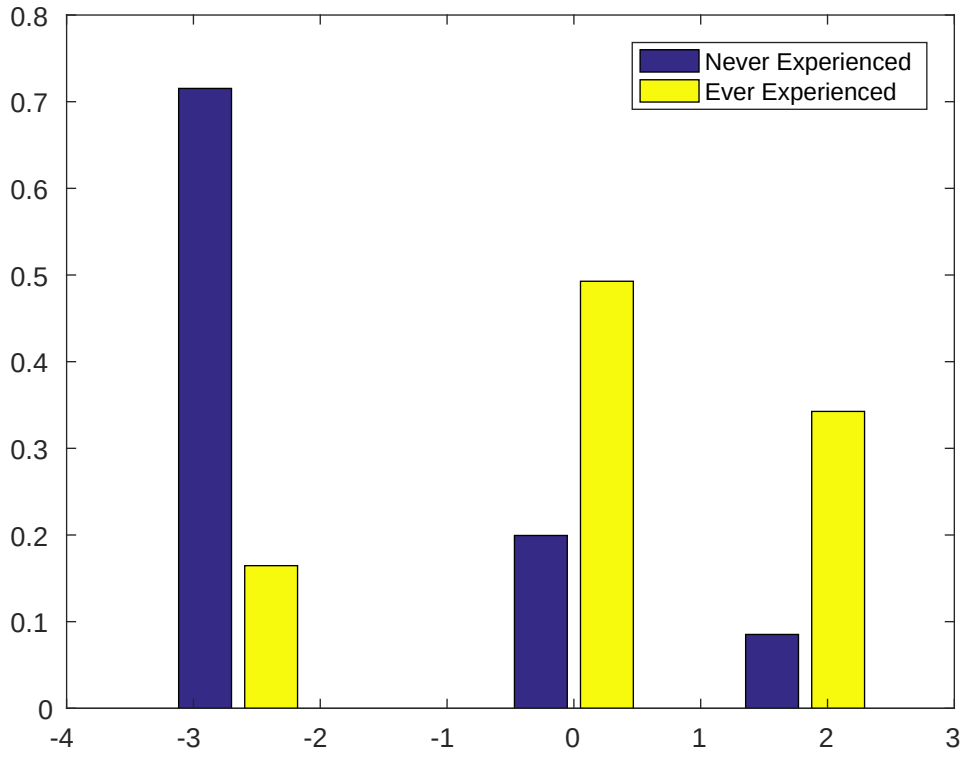
We return to explore the possible issue of worker sorting on the instrument as it relates to *differences* in the estimates between inexperienced and experienced employers. To check whether worker-sorting concerns or omitted variables are driving the results, we note that omitting workers' resume characteristics, as shown in Columns 5 and 6, results in a smaller difference in estimated markups between inexperienced and experienced employers than is estimated in Columns 7 and 8, which include worker characteristics. Including worker characteristics in the model, the estimated markup is 10.3%. Without worker characteristics, the estimated markup difference is 4.3%. This comparison is useful in estimating any possible bias from worker selection. We cannot estimate a model that includes worker fixed effects in the first stage due to the incidental parameters problem.[41] If sorting into applying on worker fixed effects goes in the same direction as the sorting based on observable worker characteristics, for which we can control, our estimate of the difference in markups to inexperienced and experienced employers may be downward biased.

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

4.6.3 Employer Type Heterogeneity and Sources of Gains from Experience

The results in Table 5 display heterogeneity in employer types for never-experienced and ever-experienced employers. Types and their corresponding probabilities for the never-experienced group are displayed under columns for inexperienced employers, while the ever-experienced are in columns for experienced employers. To help visualize employer heterogeneity, as well as selection out of the market, Figure 7 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. As is clear, the distribution of employer valuations varies across these two segments, and there is a clear pattern of selection out of the market of type-2 employers, those with low values of μ_i. 34% of experienced employers are type-3, the high μ_i type, compared to only 9% of those that remain inexperienced.

Figure 7: Employer Types by Ever-Experienced or Never-Experienced



In addition to showing the "substitutability" effect of lower markups to the experienced, the results in Table 5 also permit estimates of the "composition" and "true experience" effects on wage bids, while controlling for worker sorting on worker-level observables. We term $X\beta_\chi + E(\mu|\chi)$ the log

productive value of hiring and decompose it for experienced and inexperienced employers. The mean difference in log productive values, while easy to calculate across experienced and inexperienced buyers, does not have a natural interpretation on its own. We therefore translate the difference into units of log wages. To do so, we define $\tilde{W}$ as the equivalent log wage in the inexperienced segment that would offset the increase in productive value achieved by the experienced segment:

$$E(X_E\beta_E + \mu_E - \alpha_E\overline{log\left(w_{ij\chi}\right)}) = E(X_I\beta_I + \mu_I - \alpha_I\tilde{W}). \tag{14}$$

The bottom panel of Table 5 displays the log-wage bid equivalent change to inexperienced employers that would offset the increase in productive value with experience. The results using our preferred specification with heterogeneity across employers suggests that the total change with experience is equivalent to reducing bids made to inexperienced employers by about 30 log points, or about \$3 per hour.

We then decompose this increase in employers' mean log productive value for the platform into the true experience, composition and worker-sorting effects. Using the logic of the Oaxaca-Blinder decomposition (Oaxaca, 1973; Blinder, 1973; Fortin et al., 2011), the difference in log productive values due to differences in the characteristics of applicants and jobs (worker sorting) 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. The difference in log productive values due to different parameters, or the true experience effect, is $\bar{X}_E\left(\beta_E - \beta_I\right)$. Finally, the difference due to employer composition calculates the mean of the distribution of μ for the inexperienced and experienced segments (as summarized in Figure 7).[42]

In the models that include worker characteristics, the majority of the difference across experience levels is due to the change in coefficients, that is, the true experience effect. Because the worker sorting effect is negative, at around -11% of the log productive value difference, the true experience and composition effects explain more than 100% of the difference in the productive value of the market for experienced and inexperienced employers. Overall, 79% of the total increase in log productive value is due to a true experience effect. 32% is due to the composition effect as those employers with low values of μ leave the market.

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

5 Employer Learning about Value for the Market

This section presents several analyses that support the hypothesis that new employers undertake a learning process during early market experience that resembles the model outlined in Section 2.

5.1 Employers' Number of Interviews Falls with Experience

Table 1 showed that the number of interviews falls, on average, with successive job postings. If employers are using interviews to both look for the best applicant for the job and to learn about their own value for the market, then the marginal benefit of an interview is higher when employers know relatively little about their value for the market. Therefore, it is optimal for inexperienced employers to interview more applicants. Table 6 presents the results of a regression of (1+ the log of) the number of interviews conducted on each job by the number of previous hires made. The first column does not include employer fixed effects, and subsequent columns include combinations of employer fixed effects, controls for qualitative opening features and fixed effects for expected job duration, and controls for the mean log bid on the opening. This table presents the overall sample, and the results are similar for the sequential sample. Even with different levels of controls, in all specifications with employer fixed effects, the number of interviews decreases, and at a decreasing rate, on successive jobs. The predicted number of interviews falls by 67% after five prior hires, and a large share of that decline happens between the first and second jobs.

5.2 Employer Outcomes and Search Effort

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

Under the standard model in which the employer is not using interviews to learn about the distribution of worker value, the threshold value remains constant with the number of interviews. In this case, the expected stopping value of the hired applicant is independent of the number of interviews conducted before hiring. Under this alternative, an employer who had to interview more applicants before hiring was simply unlucky in his interviewee choices, compared to an employer who found a worker whose expected value exceeded the threshold level early on in the interview process.

Under the hypothesized learning process, then, an inexperienced employer who hires after a small number of interviews is more likely to hire a worker of higher value to him, since this worker exceeds a higher reservation value. Table 7, Panel A, shows that inexperienced employers who interview fewer candidates are more likely to hire, more likely to report having had a successful hire, and more likely to give good feedback to the employed worker. Columns 4 and 5 show that inexperienced employers who hire after one interview are 6% more likely to report success and give good feedback after controlling for the hourly wage paid to the hired worker than inexperienced employers who conduct 11 or more interviews. These results are consistent with the hypothesis that employers who hired after one interview, thus forgoing the learning value of additional interviews, must have been lucky in finding an interviewee with a high expected value (and a high actual value) to them.

Columns 6 to 9 of Table 7, Panel A, examine a similar prediction for employers who go on to post a second job. Using the variation in interviews on the first job, these results show that employers who conduct more interviews on the first job are also less likely to report success or give workers good feedback on the second job. The interpretation of these results is that, because new employers cannot distinguish market-level from worker-level signals, those who stop interviewing and hire early do so because of either a superior match with a worker or a favorable match with the market. However, those who conduct a larger number of interviews on the first job are less likely to have discovered a favorable match with the market. The results show that a larger number of interviews on the first job is indicative of a lower average assessment of productivity on later jobs. While one might think that early search effort would allow employers to better screen applicants, measures of search effort do not predict better matches. Instead, a large number of early interviews appears to indicate an employer's lack of overall fit with the market.

Table 7, Panel B repeats the analysis in Panel A but includes an additional control for unobserved employer characteristics. This control is a group fixed effect for all employers who take

the same action on the first interview.[43] The negative relationship between the number of initial interviews and job outcomes that was seen in Panel A is largely robust to the inclusion of this control. Overall, the evidence supports the model prediction that inexperienced employers who conduct fewer interviews have better first- and second-job outcomes.

5.3 Omitted Employer Characteristics

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

Building on the analysis in Table 7 that controls for early actions in looking at job outcomes, Figure 8, Panel A, examines the distribution of the hourly wage bids of the worker selected for the first interview based on the employer's eventual action. The split in Panel A is based on whether the employer does more or fewer than five total interviews. The figure plots the residual of a regression of the log of the hourly wage on job category and year-month fixed effects. A comparison of the two distributions shows very little difference in the choice of first interviewee by employers who go on to interview few or many applicants.

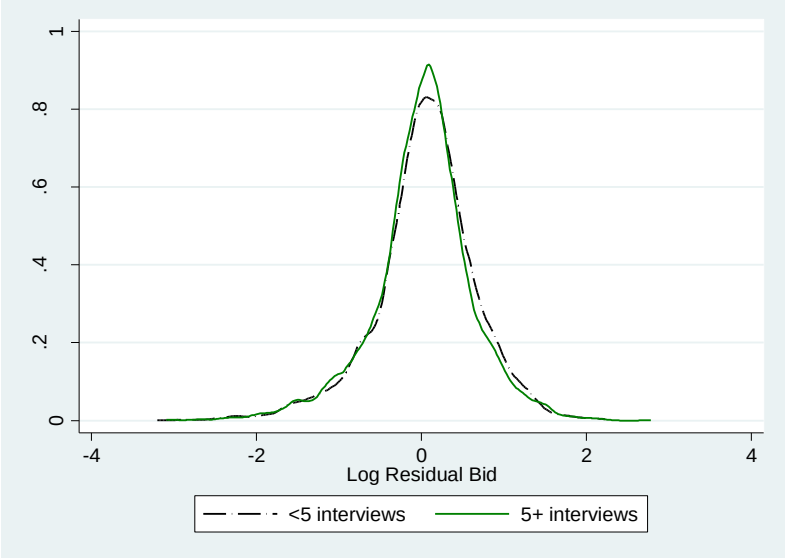
Panel B of Figure 8 repeats this exercise by whether the employer hires for the first job and again shows that early interviewing choices are unrelated to whether or not an employer makes a hire. These comparisons cast doubt on the hypothesis that inexperienced employers' eventual differences in interviewing or hiring behavior are related to unobservable differences in information, preference

[43] These actions are: they post in the same detailed job category; they either all make one first interview request or all make simultaneous first interview requests to similar applicants; their first interview is segmented based on whether it is an applicant- or employer-initiated candidacy; and the residual bid made by the first interviewee is within the same 5% bucket of the residual bid distribution.

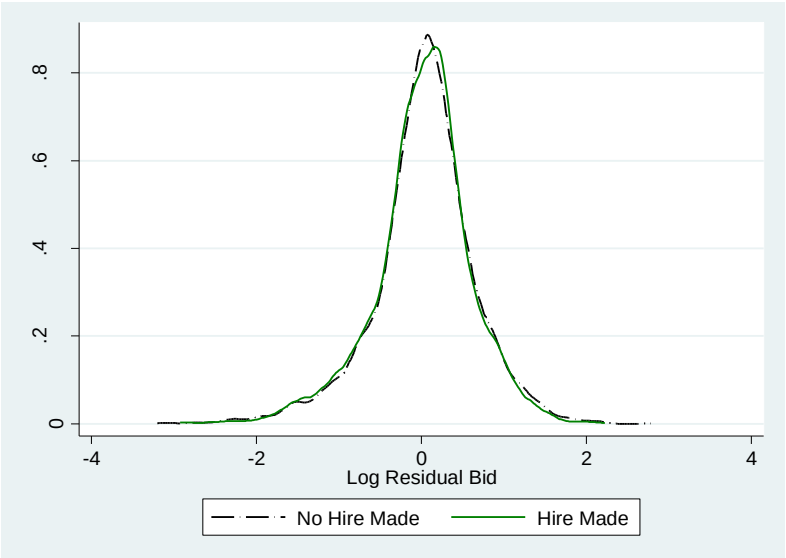
differences at the time of the initial job posting, or preference differences at the time of selecting the first interview candidate. Even more important, the overlap in wage bid on the initial interview request sent by employers suggests that workers are not able to segment inexperienced employers based on the eventual number of interviews that they conduct or by the probability that they will hire. These results are consistent with the hypothesis that actions are shaped by what is learned rather than by another employer type variable that is known by the employer before entering the market.

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

Panel A: Log Bids for employers who conduct fewer than five and five or more interviews on the first posting



Panel B: Log bids for employers who don't hire and who hire on the first posting.



6 Counterfactual Analysis and Discussion of Outsourcing

In experience markets, buyers learn about their own value for the market from early interactions with sellers. Whether they continue in the market depends on what they learn. In the oDesk setting, data on exit rates at various stages of employer experience suggest that employers learn from viewing applications, conducting interviews, and hiring. Focusing on the employers who have not previously hired, Table 1 tells us that 78% of this group post a job opening and receive applications, but leave the market without hiring. That is, a large share of employers receive signals about their own market value from applications, which leads them to negatively update their prior about how valuable they will find the market. This group would need very low wage bids to be persuaded to hire one employee or post further jobs.

To investigate whether a different fee structure would alter the mix of employers in the market and, hence, platform profits, we turn to the issue of platform fees. Since its founding up through the end of the period, the oDesk fee was constant, at 10% of wages.[44] We provide intuition for the exercise in Appendix 3 by analyzing the problem using specific (fixed) fees that do not depend on the actual wage. The important intuition from the fixed-fee case is simple: a fixed fee has revenue implications but also governs selection into becoming experienced based on an employer's expected value for using the market. We next perform simulations of platform profitability with different ad-valorem fees accounting for employer selection into returning to the market having gained market experience.

6.1 Market Size and Platform Profits under Different Ad-Valorem Fees

We simulate how employer hiring evolves under different fees and then compute the associated market revenues and platform profits. To do this, we need to make an additional assumption about how many jobs an experienced employer goes on to post after becoming informed about his own value for the market through early experience.[45] Assuming a large number of subsequent job postings presents an extreme case in which it is very valuable for the platform to induce an employer

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

[45] If employers anticipate that they will learn about their individual valuations for the platform, initial hiring may reflect employers' recognition of an option to no longer use the platform if the initial experience is unsuccessful. On the other hand, myopic employers who have low expectations of platform valuation may require inducement if the option value of gaining information is not recognized. In the simulated calculations, the estimated parameters governing hiring and transition rates for each type are used. Different types may be affected by changes in fees, and transition rates are calibrated from the empirical rate of transitioning to the experienced segment for each type.

to stay in the market. We choose this extreme case for illustrative purposes and find that offering lower fees to the inexperienced becomes optimal for the platform only when an employer posts eight or more later jobs. If the employer were to remain in the market for fewer than eight more jobs, the optimal fee would not be lower for inexperienced employers. Still, even if we assume that an experienced employer posts eight successive jobs, we find that the optimal fee on inexperienced employers is higher than the existing fee of 10%, and the optimal fee on experienced employers is even higher.

Table 8 presents platform profits and transactions volumes at different fee levels for the case in which an experienced employer's long-term value is high because he posts eight jobs once he is perfectly informed about his own value for the market. The simulations make use of the hiring probability parameters estimated in Section 4, and the estimated distribution of the employer-specific heterogeneity characterizes selection into continuation in the market given the estimated parameters. We simulate employer hiring using a grid of different fees—the fees to the inexperienced vary across rows, and the fees to the experienced vary across columns. Appendix 3 details the steps used in this simulation.

Panel A of Table 8 analyzes the change in platform profits relative to the current fee structure, while Panel B provides percentage changes in the number of employers transitioning to the experienced segment, relative to the baseline 10% fee. The main result is that the optimal fee for both segments is higher than the 10% fee charged on oDesk at the time of the data. Again, because employers in the data typically post fewer than eight jobs, the assumption that each posts eight future jobs leads to an overestimate of profits that may be had from inducing more employers to gain experience. Nonetheless, platform profits are maximized when a lower share of employers than that under the current fee structure are induced to return. Platform profits are maximized with a 15% fee on the inexperienced segment and a 20% fee on the experienced segment of employers. The profits at this fee structure are estimated to increase by about 21%, and the number of returning employers is estimated to fall by 22%. That is, platform size would be reduced, but per-transaction profitability would increase enough to more than offset the reduction in transactions volume.[46] With fewer than eight future jobs, optimal fees are approximately equal across experience levels, at 20%.

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

6.2 Linking the Results to Outsourcing

This analysis reveals that the unobserved heterogeneity in employer value for the market is large in this setting. Even among those employers who are willing to explore its value by creating a profile and posting a job, for many, the early information they receive convinces them that the value of the market to them is lower than their outside option, and remains so even with low initial wages. These findings for the platform–that employers appear to learn and that it is more profitable for the platform to target the high-value niche rather than the mass-market with lower fees–may help to explain the low rate of offshoring observed in administrative data. Employers seeking to offshore services likely confront the same problems as employers looking to outsource online: they often must evaluate candidate suppliers from an unknown pool, with unfamiliar credentials, and then learn about how to organize production in a new way. Under this interpretation, the relative cost of learning and instituting new processes, compared to the available wage savings, appears high. In the 2012 Survey of Business Owners (SBO) conducted by the US Census Bureau, only 1.36% of firms responded that they offshore services or functions abroad. While that rate was up from 1.14% in the 2007 SBO, it does not come close to matching the extent of offshoring predicted by technical estimates of feasibility. [47]

7 Conclusion

This paper provides evidence that new employers enter the largest global online labor market with incomplete information about the value of the market to them. They use early market experiences to learn how workers in the market fit with their own needs. Their initial uncertainty affects hiring probabilities and, because job applicants observe employer inexperience, can explain the observed supply-side response of higher wage bids to these employers.

We argue that this empirical setting can be characterized as an experience market. Nelson (1970) described an experience good as one in which a consumer finds it less costly to learn about product quality through consumption experience than through search. He showed that missing information about product quality can have profound implications for market structure when consumers are informed about only a few brands. His paper was one of the first studies to show how consumer behavior can affect market power.

An experience market is related to an experience good, but has two specific distinguishing

[47] These figures resemble the findings in Fort (2017) who reports low levels of foreign contracted manufacturing services at the plant level using the 2007 Census of Manufactures in the U.S.

features: first, a buyer's willingness to pay for all goods contains an individual-specific, market-level component that the buyer may learn about through consumption. Second, new buyers are unable to fully separate the consumption value due to overall market value from that due to variety value.

The implication of these two features is that when sellers are ex ante differentiated in some dimension of quality, each seller knows that inexperienced buyers faced with a variety of products will attribute a larger part of the differences in quality signals to differences in product quality. The buyer will, therefore, view the products as less close substitutes, and, depending on his value of not using the market, may have a higher willingness to pay for his most preferred variety. Market inexperience thereby conveys local market power to the most preferred product in an experience market even when there are many competing varieties. In equilibrium, suppliers who are uninformed about their own appeal relative to other suppliers, but are aware of a buyers inexperience, will set a higher price.

Once a buyer has amassed experience that allows him to increase the precision of his belief about his own market value, the same set of noisy signals of product variety permit a more precise estimate of differences in variety-specific quality. The most preferred applicant will appear less differentiated, and equilibrium bids may fall for all suppliers.

We are able to characterize the empirical setting studied in this paper as an experience market because it has a third feature: suppliers observe buyer/employer experience and are able to price discriminate by employer group by submitting wage bids that reflect a buyers willingness to pay to hire the most preferred job applicant. The data suggest that applicants understand that new employers are less able to accurately evaluate their job fit. Employers appear to have less elastic hiring probabilities when hiring for the first time. Viewing the phenomenon of services outsourcing as an experience market—where there is an initially unobserved component of market-level value—provides an explanation for another striking feature of market outcomes for inexperienced employers: the majority set up profiles, post job openings, and view applications, but leave the market without hiring. Experienced employers returning to the market are over twice as likely to hire for a job. This pattern is consistent with employers updating their priors about their own market value from early experiences posting jobs and viewing applications, before even hiring.

In a counterfactual exercise, using the estimated hiring probability functions and hiring rates among employers with different experience levels, we ask whether the platform should use different fees to entice new employers to hire. Market revenues and, hence, platform profits are maximized

at higher fees than are currently observed. Moreover, inducing the platform to offer lower wages to the inexperienced, with the intention of subsidizing learning, requires an assumption that all employers who hire once find it worthwhile to hire many more times than currently observed in the data. This finding suggests that lower fees to induce early hiring would not be profitable for the platform.

We suggest that many empirical settings, particularly digital platforms or offshoring markets, have the defining features of an experience market. These insights may shed light on the relatively slow growth of offshoring. In this particular labor services offshoring setting, the heterogeneity in employer value for the market is surprising, given the potential wage arbitrage from contracting with workers in developing countries. However, employer heterogeneity that is unknown to individual employers ex ante appears to explain the lack of hiring on the extensive margin. The majority of potential employers discover that the market is far less valuable to them than they had initially expected when posting their first job. That is, for many potential offshore outsourcers, dipping a toe in the water serves to convince them to stay firmly grounded onshore.

Appendix

Appendix 1: Data Details and Cleaning

Appendix Table 1 gives details about the resume data used in the full sample and 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 to individual jobs, the sample is restricted to openings that have at least one day of elapsed time between the next job posting and the current job posting and at least one day of elapsed time between the last posting and the current posting. This allows for at least a single cycle of applicants from different time zones to apply to the different jobs, eliminating batched hiring for which available applicants may blend across jobs. The sample for estimating hiring probabilities also drops jobs on which the employer hires a worker from a previous engagement. Many jobs also appear to originate from bringing an offline relationship onto the platform. Filtering these jobs involves requiring that at least one application be worker-initiated, while 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 to hire a pre-selected candidate or set of candidates. Any job from an employer who sends over 100 interview requests on the first job or who sends over 60 interview requests on a subsequent posting is omitted. These postings are likely to be fake jobs posted by spammers. Finally, any job that is later declared to have been posted by mistake is dropped.

The following restrictions are used to clean applications: first, applications from workers invited who later report they are unavailable are dropped; applications are also dropped if the employer reports obvious spam.

Appendix 2: Alternative Explanations for Declining Wage Bids

Different Application Rates

It is possible that the extent of competition on a job posting changes with employer experience, and workers might submit lower wage bids when they anticipate a more competitive market. For variation in anticipated market competitiveness to explain the bid premium to inexperienced employers, workers must anticipate that the job postings by experienced employers are more competitive. Table 1B showed that inexperienced employers in the sequential sample received 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 A2 repeats the analysis from Table 2 but includes the log arrival rate of applicants within the first 24 hours of posting the job as an additional control. Note that the regressions already include a spline in the application number and bidders could observe the number of prior applicants when making their bid. This additional regressor removes the effect of expected future competition on bids. The faster the rate, the lower are all bids received by the employer. However, including this control does not change the main finding from Table 2 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.[48] We posted identical jobs from the accounts of employers with different levels of experience. Employer 1 had no experience while employer 2 had prior experience hiring past workers 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 A3 contains the results. In each specification, there is a significant, negative point estimate on the experienced employer point estimate. The effects sizes are larger in magnitude than those estimated in Table 2. In addition, pairwise comparisons of Columns 1-3 and Columns 4-6 indicate very similar point estimates when the dependent variable is the log of the actual bid versus the difference in bid from the profile rate. This exercises isolates the effect of observable experience in driving lower bids.

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

Wage Bargaining

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

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

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

Appendix Table 4 contains the results. In OLS regressions, there may be some small reductions in wages due to the appearance of bargaining for employers with five or more hires. The point estimate is a reduction of about 2%. These differences do not, however, reflect the immediate reduction in bids with experience seen in other tables. Bargaining differences are not significant with the addition of employer fixed effects, meaning this channel is unlikely to drive the pattern of results documented elsewhere.

Appendix 3: Platform Profits

Fixed Fee Differences by Employer Experience

The platform's objective is to maximize total profits, which is equivalent to maximizing the total value of transactions in the market, and it can do this by setting different fees for inexperienced and experienced employers. To denote specific fees, we call the fee on inexperienced employers t_I and the fee on experienced employers t_E. Let H_I be an indicator for an employer hiring while inexperienced and H_E be an indicator for hiring while experienced. Wages for the inexperienced and experienced segment are w_I and w_E.[49] 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 hires given wages w_I and $\Pr\left(H_E | H_I\left(w_I\right), w_E\right)$ is the probability an experienced employer hires as a function of wages w_E conditional on the first hire, $H_I\left(w_I\right)$. Notice that the platform does not set wages, only fees, but wages that employers face will vary with platform fees because they are passed through.

Adding uncertainty and selection makes the fee-setting problem more interesting. When employers are uncertain about platform valuation and some uncertainty is resolved through hiring, experienced employer hiring probabilties, $\Pr\left(H_E | H_I\left(w_I\right), w_E\right)$, may depend on the evolution of employers' beliefs about the platform as a result of hiring. That $\Pr\left(H_E | H_I\left(w_I\right), w_E\right)$ specifically conditions on $H_I\left(w_I\right)$ and the wage paid captures the possibility that experienced hiring may be affected by the identity of the marginal inexperienced employer. Variation in wages, induced by different platform fees, induces variation in the identity of the marginal employer.

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

Using H_I as shorthand for $\Pr\left(H_I | w_I\right)$ and H_E as shorthand for $\Pr\left(H_E | H_I\left(w_I\right), w_E\right)$, the first order conditions for the optimal fee levels are:

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

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

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:

$$t_E^* = -\frac{H_E}{\frac{\partial H_E}{\partial w_E}\frac{\partial w_E}{\partial t_E}}. \tag{15}$$

The fee for the inexperienced is:

$$t_I^* = -\frac{H_I}{\frac{\partial H_I}{dw_I}\frac{\partial w_I}{\partial t_I}} - t_E^* H_E - t_E^* H_I \frac{\partial H_E}{\partial H_I}. \tag{16}$$

The first term in t_I^* is the standard static markup for the segment of inexperienced employers. But, 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 is of particular interest, which accounts for composition effects. 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.[50] These types are assigned independently from the applicant set. Then, for each ad-valorem fee pair, (τ_I, τ_E), simulated profits are constructed according to the following procedure: 1) Log wage bids to inexperienced employers are calculated, where pass-through of the fee is computed according to the worker's first order condition for setting bids. 2) Inexperienced employers receive a random uniform draw and choose to hire or not 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 is modeled as a probabilistic function of the expected surplus and is calibrated from the transition probability between the inexperienced

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

and experienced sample;[51] e) if the set of employers is stable, the loop is terminated and, if not, we return to step (a). The loop in 3) involves re-calculating markups and log wage bids conditional on the set of employers that transition into posting successive jobs. 4) Profits are then based on hiring probabilities from the model and the fee-rate associated with the chosen bid.

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

References

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

[16] Firpo, S., Fortin, N. M., and Lemieux, T. 2010. "Occupational tasks and changes in the wage structure." Working paper.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Table 1A: Summary Statistics

									Postings where a Hire is made		
Previous Hires	Number of Job Openings	Mean Wage Bid	Number of Applicants	Share of Employer-Initiated Candidates	Share with Good Worker Feedback	Share with Missing Worker FB	Number of Interviews	Probability a Hire is Made	Mean Wage Bid of Hired Worker	Share with Good Worker Feedback	Share with Missing Worker FB
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
0	119877	10.16	18.39	7.8%	34%	44%	2.96	22%	9.55	38%	39%
		(7.20)	(25.15)				(5.39)		(6.90)		
1	32526	9.77	14.67	7.9%	37%	40%	2.43	49%	9.16	40%	36%
		(7.03)	(24.68)				(4.36)		(6.80)		
2	22269	9.33	14.11	8.1%	37%	40%	2.28	52%	9.04	40%	36%
		(6.97)	(27.83)				(4.29)		(6.75)		
3	16820	9.33	13.91	7.7%	37%	40%	2.15	54%	8.89	40%	36%
		(7.03)	(26.68)				(3.93)		(6.75)		
4+	131378	8.71	13.74	7.4%	36%	42%	2.06	57%	8.56	40%	37%
		(7.02)	(30.37)				(4.30)		(6.82)		

Table 1B: Summary Statistics for Sequential, Arms-Length Openings

									Postings where a Hire is made		
Previous Hires	Number of Job Openings	Mean Wage Bid	Number of Applicants	Share of Employer-Initiated Candidates	Share with Good Worker Feedback	Share with Missing Worker FB	Number of Interviews	Probability a Hire is Made	Mean Wage Bid of Hired Worker	Share with Good Worker Feedback	Share with Missing Worker FB
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
0	61160	10.25	25.47	2.1%	31%	47%	3.64	16%	9.51	35%	43%
		(7.00)	(27.25)				(5.84)		(6.70)		
1	10173	9.85	27.73	2.0%	34%	43%	4.04	31%	9.17	36%	39%
		(6.87)	(29.65)				(5.65)		(6.57)		
2	6220	9.35	29.40	1.8%	33%	44%	3.95	30%	8.90	36%	40%
		(6.82)	(36.97)				(5.74)		(6.51)		
3	4525	9.60	28.28	1.8%	34%	43%	3.88	31%	8.88	36%	39%
		(6.89)	(31.31)				(5.26)		(6.47)		
4+	27686	9.06	31.15	1.7%	33%	44%	4.06	28%	8.72	36%	40%
		(6.91)	(39.82)				(6.00)		(6.59)		

Notes: Sample period is from January 2008 to June 2010. Wage bids are winsorized at the 99th percentile. For details on sample composition, see Appendix 1. Numbers in parentheses are standard deviations.

Table 2: 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 or higher degree, 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 3: 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)
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 2 and contain the same controls as detailed in the notes to Table 2. Observation counts differ from Table 2 because, for some observations, the timing of the initial feedback cannot be classified as occurring before or after later job postings.

Table 4: First Stage Regression of Log Hourly Bids on Exchange Rate and Arrivals Instruments

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

Notes: First stage regression coefficients with robust standard errors in parentheses. The inexperienced sample is employers on their first job post. The experienced sample is employers who have hired a previous worker. The first instrument is the log of the average monthly dollar to local currency exchange rate after removing a currency-specific linear trend. The second instrument uses the average number of applications arriving per job opening in the first 24 hours for other jobs in that week and job category cell. After taking logs, the instrument is what remains after removing week and job category fixed effects. Indicators that each instrument is missing or invariant within country are also included. All models contain a calendar time trend, a separate trend for technical categories, job category fixed effects, a spline with 4 knots for applicant order (knots correspond to pagination after sorting by arrival time), an indicator that the application was employer initiated, and 8 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 or higher degree, 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 5: Demand Model Estimates, Elasticities, Costs, and Markups

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

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

Table 6: Log interviews per job opening fall with hiring experience

DV: Log Number of Interviews +1	OLS	Employer Effects	Employer Effects	Employer Effects	Employer Effects
	(1)	(2)	(3)	(4)	(5)
One previous hire	-0.0515***	-0.292***	-0.289***	-0.288***	-0.286***
	(0.00576)	(0.00773)	(0.00772)	(0.00774)	(0.00773)
Two previous hires	-0.0918***	-0.362***	-0.357***	-0.357***	-0.352***
	(0.00657)	(0.00928)	(0.00929)	(0.00928)	(0.00929)
Three previous hires	-0.114***	-0.406***	-0.401***	-0.400***	-0.396***
	(0.00726)	(0.0106)	(0.0106)	(0.0106)	(0.0106)
Four previous hires	-0.135***	-0.442***	-0.436***	-0.437***	-0.431***
	(0.00786)	(0.0114)	(0.0114)	(0.0115)	(0.0115)
Five or more previous hires	-0.177***	-0.520***	-0.513***	-0.513***	-0.506***
	(0.00704)	(0.0123)	(0.0123)	(0.0123)	(0.0123)
Mean Log Bid				0.0805***	0.0807***
				(0.00474)	(0.00409)
Constant	0.893***	1.121***	1.110***	0.910***	0.898***
	(0.121)	(0.172)	(0.177)	(0.168)	(0.173)
Includes job duration fixed effects and third order polynomial of job description length	No	No	Yes	No	Yes
Observations	322,333	322,333	322,332	322,333	322,332
R-Squared	0.021	0.414	0.416	0.416	0.418

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

Table 7: Productivity and Search Effort on the First and Second Jobs, for Employers who Interview on First Job

	First Job Outcomes					Second Job Outcomes, for Employers who Hire on Second Job			
	Hires a Worker	Hires and Reports Success	Hires with Good Feedback	Hires and Reports Success, Wage Control	Hires with Good Feedback, Wage Control	Reports Success, Sequential Openings	Gives Good Feedback, Sequential Openings	Reports Success, Sequential Openings and Wage Control	Gives Good Feedback, Sequential Openings and Wage Control
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Specifications Without Buyer Group Fixed Effects									
2-5 Interviews	-0.021***	-0.014***	-0.012***	-0.029***	-0.048***	-0.033***	-0.050***	-0.032***	-0.050***
	(0.005)	(0.004)	(0.004)	(0.010)	(0.011)	(0.012)	(0.015)	(0.012)	(0.015)
6-10 Interviews	-0.013*	-0.016***	-0.009*	-0.054***	-0.051***	-0.053***	-0.068***	-0.049***	-0.063***
	(0.007)	(0.005)	(0.005)	(0.014)	(0.016)	(0.016)	(0.018)	(0.016)	(0.018)
11+ Interviews	-0.079***	-0.054***	-0.044***	-0.064***	-0.057**	-0.115***	-0.072***	-0.107***	-0.062***
	(0.007)	(0.005)	(0.005)	(0.020)	(0.024)	(0.019)	(0.019)	(0.019)	(0.020)
Includes Buyer Group Fixed Effects	No	No	No	No	No	No	No	No	No
Mean of DV	0.238	0.137	0.119	0.677	0.674	0.639	0.659	0.639	0.659
Observations	53,687	51,254	49,612	10,362	8,724	8,766	7,269	8,766	7,269
R-Squared	0.048	0.056	0.043	0.090	0.047	0.073	0.038	0.079	0.047
Panel B: Including Buyer Group Fixed Effects									
2-5 Interviews	-0.029***	-0.020***	-0.016***	-0.038***	-0.045***	-0.047***	-0.043**	-0.046***	-0.044**
	(0.005)	(0.004)	(0.004)	(0.014)	(0.017)	(0.017)	(0.020)	(0.017)	(0.020)
6-10 Interviews	-0.021***	-0.024***	-0.015***	-0.069***	-0.049**	-0.059***	-0.051**	-0.056***	-0.047*
	(0.007)	(0.006)	(0.005)	(0.019)	(0.023)	(0.021)	(0.024)	(0.021)	(0.024)
11+ Interviews	-0.088***	-0.061***	-0.050***	-0.068***	-0.046	-0.120***	-0.059**	-0.115***	-0.050*
	(0.008)	(0.006)	(0.005)	(0.026)	(0.032)	(0.026)	(0.027)	(0.026)	(0.027)
Includes Buyer Group Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean of DV	0.238	0.137	0.118	0.677	0.674	0.639	0.659	0.639	0.659
Observations	53,641	51,221	49,587	10,348	8,718	8,762	7,266	8,762	7,266
R-Squared	0.131	0.137	0.128	0.298	0.291	0.304	0.318	0.308	0.324

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

Table 8: Platform Profits and Repeat Job Posting for Different Fee Schedules Assuming a High Future Value of Experience

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

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

Panel B: Percent Change in Number of Employers Becoming Experienced

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

Notes: Simulations use the parameters from the last 2 columns of Table 5. Inexperienced employers are first assigned types from the weighted distribution of types observed overall in the inexperienced sample. The weights for these types match the population fraction of employers who begin as inexperienced and transition to the experience sample and the population fraction who never transition to the experienced sample. Simulated profits are then constructed according to the following procedure for each pair of fees: 1) Log wage bids to inexperienced employers are calculated, where pass-through of the fee is computed according to the worker's first order condition for setting bids. 2) Inexperienced employers hire or not based on the choice probabilities calculated from the demand model under the bids that reflect the inexperienced fee. 3) Iterating until convergence, log wage bids including fees and the worker's markup are calculated for experienced employers who post jobs. The set of experienced employers posting subsequent jobs is calculated using the types who hire, the expected surplus given wage bids, and a probability model for transitioning to post experienced jobs as a function of the expected surplus given bids that is calibrated off of the data with fixed 10% fees. Markups and log wage bids are recomputed given the set of experienced employers. After step 3) has converged, profits are calculated based on hiring probabilities from the model and the fee-rate associated with the chosen bids. Employers who become experienced are assumed to have a present value of 8 jobs while in the experienced sample. This number doesn't match the data, but is the smallest number of future jobs for which a fee on the inexperienced lower than 20% is profitable for the platform.

Table A1: Details about Resume Data

	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 while the sequential, arms-length sample has 3,014,561 observations. In the hiring probability estimation, fixed effects for country groups, job category, a spline for applicant order, an indicator for an employer-initiated application, and a time trend are also included. The log of the hourly rate on the last job is set to zero for applicants who have never been hired for hourly jobs.

Appendix Table 2: 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 or higher degree, 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.

Appendix Table 3: Field Experimental Evidence on Log Bids by Employer Experience

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

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

Table A4: Analysis of Bargaining Proxy by Employer Experience

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

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

Table A5: 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 in including time-varying applicant characteristics. See notes for Table 4.