

NBER WORKING PAPER SERIES

WHO BENEFITS FROM ONLINE GIG ECONOMY PLATFORMS?

Christopher T. Stanton
Catherine Thomas

Working Paper 29477
<http://www.nber.org/papers/w29477>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2021

We thank seminar participants at the AEA Meetings, Boston College, the CEPR Workshop on Incentives, Management and Organisation, Chicago Booth, Columbia, Copenhagen Business School, ESSEC, Harvard, HBS, LMU, LSE, Mannheim, NBER Organizational Economics Working Group, NBER Summer Institute, Queens, Sciences Po, Stanford, Sussex, Utah, and Yale, along with Ajay Agrawal, Ricardo Alonso, Nava Ashraf, Heski Bar-Isaac, Jordi Blanes, Zoe Cullen, Nikhil Datta, David de Meza, Ricard Gil, John Horton, Fabian Lange, Lisa Kahn, Bill Kerr, Bruce Kogut, Ed Lazear, Jin Li, Rocco Macchiavello, Kristof Madarasz, Arnaud Maurel, Ramana Nanda, Amanda Pallais, Luis Rayo, Yona Rubinstein, Scott Schaefer, Kathryn Shaw, Andrei Shleifer, and Nathan Seegert for helpful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Christopher T. Stanton and Catherine Thomas. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Who Benefits from Online Gig Economy Platforms?

Christopher T. Stanton and Catherine Thomas

NBER Working Paper No. 29477

November 2021

JEL No. F66,J23,J8,L24,L51,M5

ABSTRACT

This paper estimates the magnitude and distribution of surplus from the knowledge worker gig economy using data from an online labor market. Labor demand elasticities determine workers' wages, and buyers' past market experience shapes both their job posting frequency and hiring rates. We find that workers on the supply side capture around 40% of the surplus from filled jobs. Under counterfactual policies that resemble traditional employment regulation, buyers post fewer online jobs and fill posted jobs less often, reducing expected surplus for all market participants. We find negligible substitution on the demand side between online and offline jobs by assessing how changes in local offline minimum wages affect online hiring. The results suggest that neither online or offline knowledge workers will benefit from applying traditional employment regulation to the online gig economy.

Christopher T. Stanton
210 Rock Center
Harvard University
Harvard Business School
Boston, MA 02163
and NBER
christopher.t.stanton@gmail.com

Catherine Thomas
Management Department
London School of Economics
Houghton Street
London.
WC2A 2AE
United Kingdom
and CEPR
c.m.thomas@lse.ac.uk

1 Introduction

The growth of ridesharing and delivery platforms over the last decade has motivated regulatory efforts to reclassify independent contractors as employees, with the intent of improving their wellbeing. The most prominent early regulation, AB5 in California (2019), covered independent contractors who have relatively little control over their own terms of work, such as drivers who do not set their own prices. More recent proposals, like the PRO Act that passed the US House of Representatives in early 2021, would expand labor market regulation to the work arrangements of the estimated 1% of the US workforce currently working as independent contractors in the online platform economy (Collins, Garin, Jackson, Koustas, and Payne, 2019). Such regulation would potentially also apply to the remainder of the estimated 160 million service providers on online labor markets worldwide (Kässi and Lehdonvirta, 2018; Kässi, Lehdonvirta, and Stephany, 2021). In the public debate about these proposals, few studies have evaluated the welfare implications of applying the regulatory policies that govern employment in traditional labor markets to the work conducted in the gig economy. Using data from an online platform, we study demand and supply of knowledge work contractors, allowing us to estimate the distribution of surplus between buyers and providers. We then evaluate how surplus would vary under different hypothetical regulatory institutions.

Our approach is based on a model we construct of an online labor market for short-term tasks that features buyers who post jobs (the demand side) and providers who bid hourly wages when applying to posted jobs (the supply side). There are three key features of the model. First, the number of jobs posted by an individual buyer is determined by an arrival process that depends on a buyer’s type and her past history using the market. Second, given a job posting, a buyer determines whom to hire (if anyone) as a function of each applicant’s characteristics, their wage bids, the buyer’s own type, and, potentially, the buyer’s past history of hiring. Third, providers’ hourly wage bids are set strategically based on their individual opportunity cost of work and the buyer’s residual demand elasticity.¹ We estimate the model using data from January 2008 to June 2010 that tracks sequences of activities for individual market participants over time. We account for wage bid endogeneity with two instrumental variables (described below), which allow us to recover credible demand elasticities. From the estimated model, we characterize the magnitude and distribution of static consumer and producer surplus on each job opening and buyers’ present value of using the market based on the arrival rate of jobs.

Our first main finding is that the participants in the online labor market receive significant surplus. Both buyers and providers gain from trade. This result is far from obvious because there are many more providers than buyers active in the market, and most job openings receive numerous applications (Pallais, 2014; Stanton and Thomas, 2016). Despite the relative abundance of providers, their wage bids contain

¹This supply side flexibility is also important because while some platforms curate prices using algorithmic matching, the pricing in this platform is decentralized and equilibrium wage bids reflect changes in providers’ costs and market power.

markups over their outside options that average around 26%. Based on average wages paid, this implies surplus to providers of around \$2 per hour upon landing a job. Hired providers have market power, in part, because of a positioning advantage of being an early applicant to a job, as early applicants are far more likely to be hired.² Patterns in the data from outside the model suggest the estimated provider markups are reasonable.

Based on their estimated demand elasticities, buyers gain an average of \$3 of surplus per hour when they hire for a posting.³ We identify these elasticities by instrumenting for wage bids with two separate instruments. The first instrument exploits the global nature of the platform and uses exchange rate fluctuations that vary relative prices among applicants. This instrument is consistent with recent evidence from [Brinatti, Cavallo, Cravino, and Drenik \(2021\)](#) documenting significant exchange rate passthrough into wages in online labor markets. The second instrument captures variation in prices as a function of market tightness, based on the average number of applicants to similar job postings. These instruments allow us to identify relative demand elasticities across provider applicants and demand elasticities relative to not hiring. We calculate consumer surplus for buyers by integrating demand over changes in wage bids. We can also simulate counterfactuals where only a subset of the applicants would be affected by policy changes, as would be the case if a minimum wage were introduced to the market.

Combining the provider and buyer surplus estimates indicates significant gains from trade on the platform. Over our 30-months of data—from an era when online labor markets were handling much lower volumes of work compared to today—total surplus to providers was about \$6 million and total surplus to buyers was about \$9 million. Averaging over the approximately 36,000 jobs with a hire indicates that each filled post generated total surplus of about \$415 for the participants. These figures represent net welfare gains rather than total revenue or wage bills because they are estimates of buyers’ consumer surplus and providers’ markups over their application costs and their opportunity costs of working. That is, they are estimates of surplus relative to each party’s outside option or reservation utility/wage.⁴

Our second main finding underscores the distinction between static surplus versus lifetime surplus for buyers. We find that lifetime surplus is much higher than initial surplus for buyers who successfully break into the market. As a result, early experiences that deter buyers from successfully hiring serve to reduce the size of the market and limit total surplus.⁵ We arrive at this conclusion by separating buyer lifecycle

²While each posting tends to receive a large number of applications, the probability of being hired falls with applicant order. This reduces the expected surplus from applying for later applicants, such that the estimated expected surplus to the marginal applicant is nearly zero (below \$0.02). Earlier applicants are infra-marginal, with positive expected surplus.

³The expected surplus from a job posting, rather than the surplus conditional on the job being filled, is about \$0.70 per hour before accounting for any sunk costs of posting a job.

⁴The platform collects 10% of all revenues, which, at average wage bids, yields platform revenues of \$1 per hour. Platform marginal costs of servicing transactions are unobserved, and platform surplus is excluded from our surplus calculations.

⁵This is similar to the intuition in [Nosko and Tadelis \(2015\)](#), who show that buyers may draw conclusions about the quality of a platform from their early transactions, deterring or accelerating future use.

dynamics on the platform from unobserved heterogeneity through the use of a finite mixture model that allows an unobserved buyer type to influence both the baseline arrival rate of jobs and the parameters governing the choice to hire. After accounting for this heterogeneity, we allow flexibility in the model to capture changes in the arrival rate of job posting as a function of past hiring and average wage bids on past posted jobs.

Gaining hiring experience in the market is estimated to substantially increase buyer surplus. Greater surplus with experience arises for two reasons: past hiring accelerates the frequency of future job postings and buyers with experience have greater utility conditional on hiring. To illustrate the implications of these dynamics, the ex-ante present value of the market for a new buyer is estimated to be \$412, while an experienced buyer has a present value for the market of \$8,340. Aggregating over all buyers from this time period and assuming a stable job arrival rate into the future, we find that the market generates surplus to buyers totalling \$30 million dollars in present value. Since the time of our data, online labor markets have grown significantly, consistent with buyers, and providers, realizing gains from trade through these markets.

We next turn to several results related to policy experiment counterfactuals that mimic different regulatory environments. In these counterfactuals, we distinguish between the static change in surplus conditional on the set of jobs and applications that we observe in the data and the dynamic change in surplus after allowing the rate of job postings to change. These changes arise through two potential channels: first, fewer buyers gain experience, and, second, because job posting frequency responds to past average wage bids, higher prices deter more intensive market use.⁶ After accounting for dynamics, we consider how the present value of aggregate surplus would change for both buyers and providers under the counterfactuals.⁷

The first policy we evaluate is the introduction of an additional 10% tax paid by buyers when hiring any provider on the platform. This policy is intended to capture how the introduction of payroll taxes, such as the employer portion of Federal Insurance Contributions Act (FICA) contributions for W2 employees, would impact the market. In this scenario, hiring rates on posted jobs fall by around 26%, the static value of buyer surplus falls by around 24%, and the static value of provider surplus falls by 28% because fewer jobs are filled. After accounting for dynamics related to the number of jobs posted, the present value of provider surplus falls by about 50% and buyer surplus falls by 57% for experienced buyers and 69% for inexperienced buyers. The large losses to providers and buyers are mostly due to the fact that the number of jobs posted on the platform falls by 34% in the 30 months of the sample.

⁶Given that buyers are scarce relative to providers, the limiting factor for market size appears to be job openings rather than provider applicants. Our baseline dynamics calculations focus on the buyer side of the market, while sensitivity checks allow provider composition to respond to policy changes.

⁷The structure of the counterfactual analysis does not allow the platform to change any of its own policies. We also do not permit changes to the matching technology that the platform may incorporate under different regulatory environments. Hence, the counterfactual policy analysis should not be used to infer how regulatory changes would impact platform profitability.

The second policy we evaluate is the imposition of a wage floor of \$7.00 per hour. In this scenario, even those providers who would have otherwise bid above \$7.00 increase their optimal bids to reflect reduced low-price competition. Employers respond to the higher wage bids by hiring on posted jobs far less often, while the number of jobs posted falls by 40%. Reductions in the present value of provider and buyer surplus exceed 40% for all parties. The incidence of this surplus reduction is disproportionately borne by the providers who submit low wage bids.

In both counterfactual scenarios, static surplus losses are magnified by reductions in the total number of job postings. When we assume that job posting frequency is unrelated to past wage bids, surplus changes to providers and buyers remain negative, but the magnitude of the dynamic surplus loss shrinks as future job postings are less sensitive to market prices. An additional simulation allows for changes in the composition of applicants, which is most relevant under the hourly wage floor. In this scenario, we replace applicants whose observable characteristics are associated with wage bids below the wage floor with applicants whose characteristics are consistent with wage bids above the wage floor. This change does little to offset the loss in total surplus to providers. Overall, the counterfactuals suggest that attempts to redistribute surplus to the supply side by increasing the costs of hiring would reduce the total surplus generated by the platform and also reduce the surplus earned by each group of market participants.

The analysis described so far has only considered the loss of surplus from online opportunities, but it is possible that implementing these policies would lead to substitution to offline alternatives, increasing offline economic activity.⁸ At one extreme, if online contract work and traditional work were perfect substitutes, increasing the price of online tasks would result in employment flows to local contract workers or W2 employees. At the other extreme, online labor markets may simply expand the set of tasks done by others so that, in the absence of these work arrangements, a buyer would undertake the task him or herself or would not do it at all.

Although our data do not reveal buyers' alternative sources of labor—that is, their options outside the platform—it is possible to infer whether platform use responds to regulations that impact offline labor markets. If buyers considered online and offline sources of labor to be substitutes, we would expect changes in the price of local offline labor, in the form of exogenous changes to offline minimum wages, to affect online demand from locations with increases in the local minimum wage. To address this question, we study how online transactions respond in US states that did and did not increase the minimum wage in July of 2009, which is about half way through the time period we study. The July 2009 date corresponds to the effective date of a US Federal minimum wage increase of 10.6% (to \$7.25 per hour), which caused

⁸The literature on alternative work arrangements tends to assume implicitly that traditional and alternative arrangements are considered to be substitutes by workers. For example, recent papers estimate workers' preferences for flexibility versus more rigid traditional jobs (Mas and Pallais, 2017; Datta, Giupponi, and Machin, 2018), comparing the costs and benefits of each arrangement. There is little evidence on how buyers substitute between contract and traditional employment (or online and offline labor) and how the costs and benefits of each compare.

some states to raise their minimum wage at the same time as the Federal increase (even if the Federal minimum did not bind) or to raise the effective minimum wage because the Federal minimum was binding.

We estimate difference-in-differences regressions of platform use for states that raised the minimum wage compared to those that did not. We find no differential platform use when outside labor costs increase. The number of jobs posted and the hiring rate rate for states with and without minimum wage increases are no different for affected and unaffected states around the event date. There are also no differences in the number of non-technical job postings or hiring rates, for which offline wages are more likely affected by local minimum wages, or in specifications with buyer fixed effects.

The evidence shown here suggests the substitution elasticity between online and offline tasks is relatively limited. The fact that online job postings fall in response to online hiring costs, hence, implies that the “lost” postings would not be completed, rather than shifted to offline labor markets. The flexibility the platform offers buyers appears to fit their heterogeneous and idiosyncratic needs, and the majority of missing jobs in the counterfactual appear to be for one-off gigs or tasks that likely would not be staffed offline. A large fraction of the job postings and accompanying surplus that are lost online in the counterfactuals would therefore be net losses.

Our work is related to a growing literature about the foundations of the gig economy and the new work arrangements that it allows. Despite the ongoing debate about how to measure the number of workers in the gig economy (Katz and Krueger, 2019; Collins, Garin, Jackson, Koustas, and Payne, 2020), relatively little is known about the surplus generated from these work arrangements or how surplus is distributed among market participants. The closest work focuses on ridesharing, where recent papers quantify demand and supply (Hall, Horton, and Knoepfle, 2021), analyze changes in surplus from surge pricing (Castillo, 2020), and assess offline benefits (Gorback, 2020) and the value of flexibility to drivers (Chen, Rossi, Chevalier, and Oehlsen, 2019). More general labor economics research underscores how different workers have heterogeneous demands for different working arrangements, including flexibility, which can influence job choice (Mas and Pallais (2017)).

In the public discourse about the degree of surplus generated for the supply side by the online platform economy, a core feature of many curated markets is that the platform itself can greatly influence who receives value from transactions and who instead may be nearly indifferent between participating or not.⁹ Some common themes that emerge from the literature on online labor markets include that there are many more suppliers than buyers, that contracting flows are predictable, and that information frictions can be significant in these markets (Pallais, 2014; Agrawal, Horton, Lacetera, and Lyons, 2015; Horton, Kerr, and Stanton, 2017; Kässi and Lehdonvirta, 2018). Our paper is among the first to examine how these features interact to explain how surplus is distributed among market participants when pricing is decentralized

⁹See Jin, Kominers, and Shroff (2021).

rather than set by a platform.

Although the alternative work settings where contractors set their own prices are currently outside the remit of recent regulation, proposed extensions would apply to this part of the gig economy. Our work is thus related to studies that examine the welfare consequence of labor market regulation by considering a range of employer response margins (Clemens, Kahn, and Meer (2018)). One such paper is Horton (2017b), who studies hiring choices in an online labor market under a \$3 wage floor. His experiment was carefully designed to measure quantity responses to higher wages conditional on job posting and on applicant composition, which he was able to do because the treatment was a surprise to buyers after a job posting went live. In contrast, our study combines data on buyers’ job posting decisions over time as well as their hiring decisions in order to explore how buyers’ past experience affects their future actions.

Considering the estimated surplus losses shown in the regulatory counterfactuals and together with the limited substitutability to offline workers, our estimates suggest that policies designed to help workers will instead raise the cost of online hiring and will likely fail to meaningfully redistribute economic surplus to either online providers or to offline workers. This conclusion reflects the facts that providers are able to extract surplus when hired, and that the task-based, or gig, nature of labor demand implies that increases in hiring costs shrink future market size significantly.

The rest of the paper proceeds as follows. Section 2 describes the data and presents descriptive statistics that guide our modeling choices. Section 3 presents the model of hiring probability, job-posting frequency, the provider’s bidding problem, model identification, estimation, and surplus calculations. Section 4 presents the estimation results and the calculations of buyer and provider surplus. Section 5 is the counterfactual analysis. Section 6 examines online-offline wage elasticities. Section 7 concludes.

2 The Setting, Data, and Descriptive Statistics

2.1 Empirical Setting

The labor market studied in this paper is an online platform that allows buyers to contract with remote providers who sell labor services. The platform facilitates search and matching, remote task management, and payments. Job types vary, but a key feature is that they include tasks for which the output can be delivered electronically. The most frequently observed task categories during the time period studied are Web Development and Administrative Support. Job postings tend to be short-term spot transactions, and the typical job requires 17 hours of work per week, with 58% of postings lasting less than three months. Filled jobs average around 75 total hours worked. 89% of the transactions in the market span international borders, with the provider typically located in the lower-wage country.¹⁰

¹⁰See Horton (2010) for an overview of how online labor markets work, and Agrawal et al. (2015) and Horton, Kerr, and Stanton (2017) for stylized facts about patterns of contracting, especially between different countries, in these markets. This

To purchase online labor services, a buyer first needs to create an account on the platform. She must then post a job, which requires selecting the task’s work category and its expected duration, giving the job a title, and describing the work to be done and required skills. Once the posting is live, potential applicants learn about it by searching on the site or through email notifications about new jobs. Like the example in Figure 1, the postings contain information about the buyer and the job that providers can observe before applying.

Interested providers submit applications to an opening, bidding an hourly wage to work on the specific task. Buyers also have the option of searching provider profiles directly and inviting applications. Provider profiles contain information about their skills, education, prior offline work experience, and experience on the platform. Providers are located worldwide, and the profiles of those with prior experience on the platform include summary feedback scores received from past work.

After receiving applications and inviting providers to apply, buyers can request interviews with any number of candidates for the job. During the time period studied, if the applicant agreed to an interview, it would usually take place via an off-platform messaging or conferencing system.¹¹ A buyer may choose to hire a provider with or without interviewing her first.

After hiring a provider for the posting, the buyer can monitor work via software provided by the platform, and the platform also manages all payments for completed work. When a task is complete, the buyer is asked for feedback about the provider and about whether the task was completed successfully. The provider is also asked for feedback about the buyer. The platform guarantees providers are paid for the hours billed and so payments and payment risk are unrelated to buyer reputation or experience.

The data used in this paper are administrative data obtained from the platform. Crucially, for incorporating buyer dynamics in the model, every buyer’s posting-specific hiring process is observed in each of her earlier successive postings. For each vacancy, the data contain information about the entire applicant pool; which candidate, if any, is hired; and the feedback and success measures that the buyer leaves for the hired provider, and vice versa.

2.2 Descriptive Statistics

The data cover the 30 months from January 2008 to June 2010. During this time, 67,292 potential buyers posted 169,578 jobs, and received more than 4.4 million applications. There are nine task categories in the data. Web Development, the largest category, makes up 40% of all job postings; Software Development, the

market has similar features and characteristics to other prominent platforms. Most papers about online labor markets focus on workers’ careers (Agrawal, Lacetera, and Lyons, 2016; Lyons, 2017; Ghani, Kerr, and Stanton, 2014; Pallais, 2014; Stanton and Thomas, 2016). The main exceptions are Horton (2017a) and Horton (2017b), which examine matching under different platform policies. See Stanton and Thomas (2020) for a discussion on the evolution of these markets over time.

¹¹The data record interview requests and the applicant’s response to the request. Whether an interview actually occurs is not recorded in the data.

next largest technical category, makes up 9%. Administrative Support is the largest non-technical category, with 15% of total postings.¹² The sample includes all job postings for hourly-billed work after filtering for spam job postings, which are defined as those where the buyer sent over 60 interview requests. Any job that was later declared to have been posted by mistake was also excluded. We also drop applications that buyers themselves flag as obvious spam.

Table 1 displays descriptive statistics about buyers. The first three rows show that most buyers in the data are from the US, while 75% are located in G10 countries with high levels of income per capita. Buyers range from private individuals and those working at small firms, to the occasional buyer who hires on behalf of a large enterprise. Providers observe the country where the buyer is located and her past platform activity when applying, as displayed in Figure 1.¹³ Only 6% of buyers are located in India or the Philippines, which are two of the largest source countries for providers. The majority of buyers post their first job in a technical job category.

The average buyer posts 2.52 jobs during the 30 months of data, but the standard deviation is large. That is, there is substantial heterogeneity in buyer engagement with the platform. Later, our model will attempt to distinguish between whether differences in engagement arise from innate type heterogeneity or from different experiences while hiring. This distinction is relevant for understanding how transaction volumes would change under different counterfactuals.

Table 2 presents descriptive statistics at the job posting level. Column 1 includes all postings, while subsequent columns condition on buyers' prior hiring experience. The first row displays the number of months between job postings and is right-censored at the end of the data period. Buyers post an opening every 1.15 months, on average. Average lags between job postings are shorter for experienced buyers, since those with no prior hires post every 1.53 months and those with three or more prior hires post jobs every 0.67 months on average.

Postings receive an average of 26 applications, and this ranges from an average of 24 applications for inexperienced buyers to 30 applicants for those with three or more hires. Applicant characteristics other than wage bids vary little with buyer experience. Across all jobs, only 11% of applicants are based in the US. Around 34% have good feedback scores from past work. The mean hourly wage bid is \$10.88, but bids are higher for those without experience (averaging \$11.28 per hour) and lower for highly experienced buyers (averaging \$10.31 per hour). This suggests that similar providers submit lower hourly wage bids to the jobs posted by experienced buyers. Later we will explore whether this is because providers face lower costs when working for these buyers or because their wage bids include lower markups over their costs of

¹²Other task categories are, in order of frequency: Sales and Marketing, Writing and Translation, Design and Multimedia, Networking and Information Systems, Business Services, and Customer Service.

¹³Specifically, they observe any feedback that the buyer has received from previous hires, as well as the number of past hires she has made and the total hours billed.

providing services.

In order to provide some context for counterfactual imposing a \$7.00 per hour price floor, we note that around one third of applicants to the typical posting submit an hourly wage bid of less than \$7.00 per hour. This statistic is relevant because in July 2009, around the middle of the time period studied, the Federal minimum wage increased from \$6.55 to \$7.25 per hour. While this minimum wage did not apply to online work, and providers were free to submit any non-zero hourly wage bid throughout our sample, in the counterfactual analysis, we simulate the impact on equilibrium wage bids of a minimum wage bid threshold of \$7.00 to investigate the impact of such a wage floor on platform surplus.

The final row of Table 2 shows that an average of 21% of posted openings are filled on the platform, and job fill rates for buyers with past hiring experience. The next section will consider buyers’ hiring decisions more formally. However, a preliminary look at the data shows that hiring the provider who submits the lowest hourly bid is relatively rare. Figure A.1 in the Appendix shows the distribution of the wage bid decile within each posting of the provider who was ultimately hired. Only around 20% of hired providers come from the lowest bid decile for the opening. As a result, it will be important to account for characteristic differences across providers and other sources of perceived differentiation.

The descriptive statistics illustrate that a relatively small number of buyers generate a disproportionately large share of total platform volume, both by posting jobs more often and being more likely to hire on those postings. In seeking to quantify the magnitude of the platform’s surplus and its distribution, it is therefore important to understand how buyers make posting and hiring decisions at each experience level, as well as the dynamics relating to their gaining experience. The following section models these processes.

3 A Model of the Demand and Supply of Online Labor Services

3.1 Demand: Hiring Decisions for Posted Jobs

On a job opening, denoted o , a buyer with heterogeneous type k and prior experience with hiring $\chi \in \{0, 1\}$ evaluates applicants based on their characteristics relative to their wage bids. The buyer’s indirect utility from hiring provider j is:

$$u_{k\chi oj} = \frac{\exp(X_{o\chi j}\beta_{k\chi} + \mu_{k\chi} + \varepsilon_{oj})}{(w_{\chi oj})^{\alpha_\chi}}, \quad (1)$$

where $X_{o\chi j}$ contains observable provider and job posting characteristics.¹⁴ $\beta_{k\chi}$ is a vector of buyer type- and experience-specific preferences for provider characteristics. The term ε_{oj} is an idiosyncratic utility shock for provider j that is assumed to be drawn from a location and scale normalized type-1 extreme value distribution. The other parameter in the numerator of the buyer’s indirect utility, $\mu_{k\chi}$, captures the

¹⁴The vector X_{oj} includes many of the details observable in a provider’s profile that are shown to be relevant to hiring decisions in Barach and Horton (2017), including feedback score, country, past hours worked, and past compensation history. Job categories and expected duration are among the characteristics included for postings.

value that a buyer of type k with experience χ gains from hiring any provider on the platform relative to their normalized outside option. The denominator contains provider j 's job-specific wage bid, $w_{\chi oj}$, which is made after observing the job posting, buyer experience, and possibly signals of buyer type. The term α_χ measures buyers' experience-specific wage disutility.¹⁵ The possibility that providers tailor their wage bids to particular openings or types of buyers necessitates an instrumental variables strategy, which we describe in Section 3.4.

The buyer's decision problem is to choose the best applicant out of the set of applicants for opening o , denoted J_o , or to choose not to hire on the platform.¹⁶ If the buyer chooses not to hire, she receives indirect utility $\exp(\varepsilon_{o0})$, where ε_{o0} is another idiosyncratic type-1 extreme value shock. The payoff from the outside option excludes the term $\mu_{k\chi}$.

Because the buyer is looking for a single provider, if she chooses to hire, she selects the provider who creates the highest utility per unit of wage. The option of not hiring means the best provider's ratio of quality to wage must be greater than the value of the off-platform option, denoted option zero, for the buyer to make a hire. The buyer's objective function is then:

$$\max_{j \in \{J_o, 0\}} \left\{ \max_{j \in \{J_o\}} \frac{\exp(X_{o\chi j} \beta_{k\chi} + \mu_{k\chi} + \varepsilon_{oj})}{(w_{\chi oj})^{\alpha_\chi}}, \exp(\varepsilon_{o0}) \right\}.$$

Taking logs of the buyer's objective function yields an inequality stating that provider j is hired by buyer of type k with experience χ when

$$X_{o\chi j} \beta_{k\chi} + \mu_{k\chi} + \varepsilon_{oj} - \alpha_\chi \log(w_{\chi oj}) \geq X_{o\chi l} \beta_{k\chi} + \mu_{k\chi} + \varepsilon_{ol} - \alpha_\chi \log(w_{\chi ol}) \quad (2)$$

for all $l \in \{J_o, 0\}$. Note that the right-hand-side in inequality (2) contains only the term ε_{o0} when comparing applicant j to the outside option.

Conditional on $\mu_{k\chi}$ and the error structure, the probability that provider j is the preferred choice, $p_{k\chi oj}$, takes a conditional logit form,

$$p_{k\chi oj} = \frac{\exp(X_{o\chi j} \beta_{k\chi} + \mu_{k\chi} - \alpha_\chi \log(w_{\chi oj}))}{\left(1 + \sum_j^{J_o} \exp(X_{o\chi j} \beta_{k\chi} + \mu_{k\chi} - \alpha_\chi \log(w_{\chi oj}))\right)}. \quad (3)$$

Because $\mu_{k\chi}$ shifts the buyer's value of hiring any provider on the platform, this term does not alter how

¹⁵ Allowing the parameters to have a more general dependence on buyers' type, including allowing wage disutility α_χ to vary with k , does little for model fit.

¹⁶ Providers are assumed to be available when they initiate an application. This assumption requires only that the probability that a provider will receive two offers over a short time interval is small. Although declined job offers cannot be distinguished in the data, only 0.6% of the provider days in the sample have more than two interview requests. A post-candidacy survey also asks buyers for reasons particular providers were not hired and asks providers their reasons for exiting the active candidate set. In the few cases where buyers or providers explicitly report a realized scheduling conflict or where providers refuse invited applications, the application is excluded from the sample.

the buyer ranks applicants for a given job posting except relative to the outside option. It does, however, relax the independence of irrelevant alternatives, or IIA, assumption in standard conditional logit models, allowing for different substitution patterns between the no-hire option and the available candidates in the choice set.

Identification is intuitive conditional on buyer type k and experience. If a buyer’s type were fully observable, variation in wages, applicant characteristics, and hiring decisions would identify the parameters, up to location and scale. However, because each buyer’s type is unobserved, only the population distribution of types can be identified overall. This population distribution is identified up to the random utility assumptions by buyers who take different actions in the face of similar choice sets. For example, if some buyers repeatedly hire when faced with low-quality applicants who submit high bids, whereas other buyers do not hire when high-quality providers with low bids are available, the estimated population distribution of μ types would have a wide dispersion in valuations.

From the population distribution, we calculate individual buyers’ posterior type probabilities via Bayes’ rule. Individual posterior types put more weight on the lowest type if a buyer does not hire with attractive choice sets, while a buyer who repeatedly hires despite having poor alternatives would have a posterior reflecting a high type.

3.2 Demand: Extensive Margin Job Postings

So far we have focused only on the static hiring choice problem, but the full model also accounts for the extensive margin decision to post a job as a function of buyer type and past transaction history. This is because understanding how changes in market institutions alter the total surplus available requires a specification that is versatile enough to capture differences in demand arising from (i) changes in the underlying parameters describing how buyers assess providers at different experience levels, and (ii) buyer selection in or out of the market at different experience levels. This requirement motivates an estimation approach allowing the buyers who return to the market at different levels of experience to be a non-random sample from the population distribution of types. This flexibility will help to match patterns from Table 2, where the distribution of buyer job posting frequency is related to buyer platform experience. Table 2 also suggests that there is a thick tail of postings coming from a small number of buyers.

We accomplish these goals by specifying a job posting arrival rate as a stochastic process that is allowed to depend on a buyer’s history of market use and type k . Postings are modeled as an exponential arrival process, $\lambda_{k\chi}$, that depends on buyer type, past hiring experience, and the past wage bids the buyer has observed for a particular type of job. Allowing the job posting process to vary with buyer type k captures non-random selection into using the market or into posting jobs frequently.

To consider how changes in hiring costs affect market use, the job posting process is allowed to depend

on the wage bids encountered on past job posts. This dependency could arise if the bids received in the past shaped expectations about future bids and buyers' decisions to post future jobs depended, in turn, on expected bids. The dependency of future job posting rates on past wage bids is modeled in reduced form, using variation in market-level average log wage bids at the time of past job posts. This is to guard against idiosyncratic or buyer-specific unobservables from biasing this estimated relationship. Specifically, we define $\widetilde{\log(w_{oj})} = \frac{1}{O-1} \sum_1^{O-1} (\overline{\log(w_{o,t,Category})} - t_{Category} - 1_{Category})$ as the average log wage residual for each buyer on all prior openings 1 to $O - 1$ after de-trending monthly log wage bid averages for each job category. Buyers who post jobs in time periods where wages are relatively high (compared to trend) will have higher individual values of $\widetilde{\log(w_{oj})}$.

The job posting arrival process $\lambda_{k\chi}$ is specified as follows:

$$\lambda_{k\chi} = \delta_{1k} + \delta_2 1(\chi > 0) + \delta_3 \widetilde{\log(w_{oj})}. \quad (4)$$

The vector δ_{1k} contains buyer type-specific constants. The other two terms reflect buyer hiring experience and the previous wage bids received. The parameters δ_2 and δ_3 are common across all buyer types and shift the mean hazard for experienced buyers. Buyers posting their first job have not received any prior wage bids, have no experience, and their waiting time is undefined. The arrival rate for new buyers is therefore not modeled.

3.3 Supply: Provider Wage Bids

Provider bidding is modeled using tools for analyzing heterogeneous products competition with local market power. Wage bids are characterized as optimal functions of providers' costs and the application-specific demand elasticity in a Bertrand-Nash oligopoly. Provider j 's cost, $c_{\chi oj}$, captures her outside option and her expected costs of applying to and/or being hired for the job opening o posted by a buyer with experience level χ . Thus, $c_{\chi oj}$ captures both the opportunity cost of work through the provider's alternative use of time and the direct cost of applying to a job. When choosing the wage bid, the provider's objective function takes this cost into account along with the hiring probability.¹⁷ Her wage bid also takes into account the ad-valorem fees retained by the platform, τ .

The provider chooses the wage bid, $w_{\chi oj}$, that maximizes

$$\underbrace{p_{\chi oj}}_{\text{Pr(hired)}} \times \underbrace{\exp(\log w_{\chi oj} - \log(1 + \tau))}_{\text{Post-fee wage}} + (1 - p_{\chi oj}) \times \underbrace{c_{\chi oj}}_{\text{Cost}}, \quad (5)$$

¹⁷In our setting, the wage bid is the provider's only strategic variable. The fit between the buyer or job and the worker is assumed to be unknown at the time of application. This differs from the broader model of quality and price choice in monopolistic competition with experience goods in Riordan (1986).

where $\log w_{\chi oj}$ is the log of the wage bid that would be paid by the buyer, meaning it is inclusive of τ . If she is hired, provider j receives the wage $\frac{w_{\chi oj}}{(1+\tau)} = \exp(\log w_{\chi oj} - \log(1 + \tau))$ and the buyer pays $w_{\chi oj}$. If provider j is not hired, she receives $c_{\chi j}$, her “net” outside option.¹⁸ The provider’s first order condition is given by

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

Taking the first order condition in equation (6) for each applicant to a job opening and forming a system of equations characterizes Bertrand-Nash equilibrium bids. For an individual provider, note that the elasticity in equation (6) is a function of other applicants’ bids. Solving for provider j ’s optimal wage bid yields

$$w_{\chi oj}^* = c_{\chi oj} (1 + \tau) \left(1 + p_{\chi oj} / \frac{\partial p_{\chi oj}}{\partial \log w_{\chi oj}} \right)^{-1}. \quad (7)$$

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

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

which shows how having an estimate of $p_{\chi oj}$ and $\frac{\partial p_{\chi oj}}{\partial \log w_{\chi oj}}$, together with the observed bids and platform fees, yields applicant-job-specific estimates of costs and markups. Section 3.6 details how these objects are used to calculate provider surplus.

3.4 Instruments and Identification of Hiring Probabilities

Job applications potentially contain unobserved characteristics that may be correlated with providers’ wage bids. We account for potential correlation of the error and wage bids with an instrumental variables strategy. The first instrument is based on the notion that wage bids are determined in part by provider costs, including the opportunity cost of online work, as shown in equation (7). We use changes in the dollar-to-local-currency exchange rate for a provider’s country as an exogenous opportunity cost shifter. Providers are paid in their local currency for offline work, but they are paid in US dollars for their work on the platform. Hence, frictions limiting exchange rate pass-through to local wages mean offline opportunities are likely to adjust to exchange rates more slowly than online transactions.¹⁹ When the local currency

¹⁸The objective could alternatively be written: $\max_{\log w_{\chi oj}} p_{\chi oj} \times \exp(\log w_{\chi oj} - \log(1 + \tau) - c_{\chi oHj}) + [1 - p_{\chi oj}] \times c_{\chi Oj}$, where $c_{\chi oHj}$ is a cost from on-the-job work associated with being hired for job o , and $c_{\chi Oj}$ is the outside wage for provider j . The first order condition in this case makes clear that only $c_{\chi oj} = c_{\chi oHj} + c_{\chi Oj}$ can be identified.

¹⁹This potential source of variation was revealed in conversations with buyers who mentioned the frequency with which exchange rate calculators appear in the screenshots taken by the platform’s monitoring software.

appreciates relative to the dollar, so that one dollar earned on the site provides fewer local currency units, applicants' wage bids are predicted to increase.

Figure 2 focuses on India, the largest non-US provider source country, to illustrate the time-series variation in mean residual log wage bids and exchange rates that underpins this instrumental variables approach. The blue squares represent the median residual log wage bid, net of job category fixed effects, for applicants located in India. The red circles display the normalized log US dollar to Indian rupee exchange rate.²⁰ The mean bid and exchange rate time series move together.

Although exchange rate movements are plausibly exogenous to demand on the platform, applicants based in the US or living in countries with dollar-pegged exchange rates do not face any cost shocks from this instrument. This does not pose a problem for comparisons between applicants, because others will have exchange rate variation, shifting relative bids, but it leaves these applicants' wage bids with no variation relative to the buyer's outside option of not hiring.

We therefore use a second instrument that is relevant to all providers, including those with dollar or dollar-pegged local currencies. This instrument captures variation in the intensity of competition that applicants face for an individual job posting, which affects the optimal mark up included in the bid, as shown in equation (7). The ideal variation would alter the extent of perceived competition from other providers, embodied by situations in which, for example, a subset of potential applicants randomly lost internet access and this supply shock was observable to providers who remained available. The instrument we construct is based on a similar idea because it exploits the fact that providers can observe the count of other applications to each job posting at each point in time, which, for an individual bid decision, allows providers to evaluate the extent of competition across a portfolio of job openings and to tailor their bid accordingly. On part of the exclusion restriction for this instrument entails an assumption that buyers do not base their job posting decisions on the thickness of the market or on how many providers are competing. Because buyers and providers view the site through different entry points and have different user experiences, such that buyers tend to see provider profiles rather than other job posts or statistics about application rates to competing jobs, this part of the exclusion restriction seems reasonable.

We calculate the instrument using the log of the leave-own-opening-out average of the number of applications in the first 24 hours since posting for all other jobs posted in the same job category that week. While the instrument captures aggregate competition, the exclusion restriction is vulnerable to demand shocks that affect market conditions or the intertemporal mix of providers who supply labor. To mitigate this potential violation of the identifying assumptions, we net out job category and time fixed effects from the instrument so that these sources of variation are not used in estimating the first stage relationship. The residuals after removing these fixed effects capture competition differences across job

²⁰To control for secular trends and level differences in local exchange rates across countries, the exchange-rate series is normalized and de-trended.

categories in the same week while holding fixed the average competition in the category over time and the average competition across all job categories during the week in question. Later, for both instruments, we test how the first stage relationship and the model parameters change when we control for the number of applications an individual provider submits in a given month. We find that this control, which proxies for exclusion restriction violations on the supply side, does little to alter the estimates.²¹

To make use of the variation in bids induced by the instruments, we utilize [Petrin and Train’s \(2010\)](#) control function approach, including the two instruments, Z_{oj} , and provider characteristics, X_{oj} , in a first stage regression of the form

$$\log(w_{\chi oj}) = \gamma_0 + Z_{oj}\gamma_1 + Z_{oj} \times 1(\chi > 0)\gamma_2 + X_{oj}\beta_1 + X_{oj} \times 1(\chi > 0)\beta_2 + \nu_{oj}. \quad (9)$$

The notation in equation (9) makes clear that separate interactions are estimated for jobs posted by buyers who have no hiring experience and by those who have made prior hires ($\chi > 0$).²²

Table 3 shows that both instruments have a substantive and statistically significant effect on applicants’ bids, with the effects in the expected direction. Log bids increase when the local currency appreciates (i.e. when the log dollar-to-local-currency exchange rate increases) and log bids decrease when the log number of applicants to similar jobs increases. The first two rows in each column present the coefficient estimates for all buyers. Columns 1 and 2 contain the instruments alone, without experienced buyer interactions. The second two rows in columns 3 and 4 present interaction terms between the instruments and an indicator that the buyer is experienced (has hired at least once). These interaction terms are small and insignificant. Two F-statistics based on different levels of clustering are reported to assess the strength of the instruments. The first captures variation at the job opening level, which is the relevant source of variation compared to the buyer’s off-platform option. The second clustering level is at the provider level, which would account for within-provider correlation in when applications are submitted and wages are bid. In each case, the

²¹One potential exclusion restriction violation is that the instruments change the mix of providers available. In related analysis, [Horton \(2021\)](#) finds that the collapse of the Russian ruble (outside of our sample period) drove large changes in the quantity of applications. In a larger panel from more countries, however, [Brinatti et al. \(2021\)](#) find larger rates of pass through of exchange rates to wages.

²² The first-stage regression in equation (9) is a transformation of the optimal bid equation, equation (7), derived by mapping local-currency-denominated opportunity costs to dollar-denominated bids. Assume that $c_{\chi oj}$ is denominated in the local currency, whereas the bids that buyers observe are denominated in dollars. Costs in the local currency must be translated into dollars when submitting bids, so the provider’s optimal bid becomes $w_{\chi oj}^* = c_{\chi oj} \left(\frac{D}{L}\right)^\theta (1 + \tau) \left(1 + p_{\chi oj} / \frac{\partial p_{\chi oj}}{\partial \log w_{\chi oj}}\right)^{-1}$, where $\left(\frac{D}{L}\right)$ is the dollar-to-local-currency exchange rate and the parameter θ captures possible reasons for deviations from complete pass through. These reasons include the following: (i) some part of a provider’s opportunity cost reflects transactions denominated in dollars rather than in the local currency, which may occur if the possibility of receiving an alternative wage comes from searching online; (ii) part of a provider’s consumption may become cheaper through imports; and (iii) the incidence of exchange rate variation is split between providers and buyers. Taking logs of this optimal bid yields the first-stage regression:

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

instruments are strongly jointly significant.

As mentioned previously, an additional concern is that the instruments may influence the composition of workers who apply. For example, an appreciation of a local exchange rate may lead a selected group of potential applicants to seek work elsewhere. The second and fourth columns of Table 3 show the first stage results when including the number of applications made by the provider in the relevant month as a control variable. If a non-random subset of applicants increase or decrease their applications in response to exchange rate and competition variation, including this control will absorb some of the variation in wage bids that arises because of the instruments. The parameter estimate on the competition instrument is nearly identical with and without this control, while the estimate on the exchange rate instrument declines roughly ten percent, from 0.087 to 0.077. Overall, the similarity of estimates suggests that extensive-margin provider sorting is unlikely to affect inference about wage bids, and hence markups and surplus. We later demonstrate that including this control in the choice model does little to alter our conclusions about surplus.

3.5 Estimation

We estimate the model by maximizing the likelihood of the observed sequences of buyer job posting and hiring decisions. The step-by-step estimation approach is as follows: First, we calculate the residuals from equation (9) to form control functions that account for unobserved provider quality, denoted $CF_{\chi oj} = \hat{v}_{\chi oj}$. Second, we calculate choice probabilities conditional on a value of the unobserved type k , as in equation (3), while including $CF_{\chi oj}$, which take the form:

$$p_{k\chi oj} = \frac{\exp(X_{oj}\beta_{k\chi} + \mu_{k\chi} - \alpha_{\chi} \log(w_{\chi oj}) + \psi_{\chi} CF_{\chi oj})}{\left(1 + \sum_j^{J_o} \exp(X_{oj}\beta_{k\chi} + \mu_{k\chi} - \alpha_{\chi} \log(w_{\chi oj}) + \psi_{\chi} CF_{\chi oj})\right)}. \quad (10)$$

Third, we assume there are $K = 3$ three distinct types, and for each type, the parameters can vary with experience. These parameters are estimated along with the overall share of buyers of each type k , ρ_k .

We then form the likelihood conditional on buyer type k , which is defined over sequences of buyer choices across different job postings, as:

$$L_k = \Pi_o(\sum_{oj}(p_{k\chi oj})^{y_o=j})(\lambda_{k\chi} \exp(-\lambda_{k\chi} t))^{o \in (2,O)} (\exp(-\lambda_{k\chi} T))^{o=O}. \quad (11)$$

The term $\sum_{oj}(p_{k\chi oj})^{y_o=j}$ in equation (11) is the conditional choice probability for the provider selected on posting o after accounting for all of the possible alternatives j in the summation operator. $\lambda_{k\chi} \exp(-\lambda_{k\chi} t)^{o \in (2,O)}$ is the density of a random variable with exponential distribution and parameter $\lambda_{k\chi}$ that captures the probability of waiting t months to post job o from the time when the previous job, $o - 1$, was posted. The first posting for every buyer, opening 1, does not have a lag between openings, while the last posting must

account for the fact that the data are right censored by the end of the time period studied. The final term, $\exp(-\lambda_{k\chi}T)^{o=O}$, accounts for right censoring of the last posting in June 2010, which is the last month in the data. Each of these components for the arrival process is computed at the job opening level, and the product over openings o —that is, over the sequence of choice probabilities and densities of posting times for the buyer—is then calculated. Finally, because type k is not observed, the marginal likelihood is the type-weighted sum over the likelihoods for different buyer types:²³

$$L = \sum_{k=1}^3 \rho_k L_k \quad (12)$$

3.6 Buyer and Provider Surplus

Estimation yields parameters of the buyer’s indirect utility function given in equation (3) and the job posting arrival process given in equation (4), as well as the share of buyer types, ρ_k , which can be used with the arrival process to recover type shares by buyer experience. The elasticity estimates enable inference of provider-job-specific costs, based on equation (8). With these objects, the total gains from trade and the division of surplus between buyers and providers can be calculated.

3.6.1 Buyer Surplus

Buyer surplus on a given job opening is given by integrating over the product of quantities (hiring probabilities) and prices (wages paid) starting from the base wage observed, w_0 . A buyer of type k with experience χ has per-opening surplus for each hour of work given by:

$$Surplus_{k\chi} = \int_{w_0}^{\infty} \sum_{oj} (p_{oj}(w)|k) \times (w - w_0) dw. \quad (13)$$

In this expression, the probability of hiring each applicant, given applicant characteristics, at small wage bid intervals dw , is computed conditional on the buyer being type k . This expression makes clear that if $p_{oj}(w)$ changes little with wages (i.e. demand is very inelastic), buyers enjoy substantial surplus because they would be willing to pay amounts in excess of w_0 . If instead demand across all providers (because of the summation) is relatively elastic, buyer surplus on an opening is lower. Estimates of type-specific opening surplus are weighted by the posterior distribution of buyer types, recovered by Bayes’ rule, to give total surplus of $\sum \hat{\rho}_k Surplus_{k\chi}$.²⁴

To aggregate expected surplus across multiple job openings for an individual buyer requires treating inexperienced and experienced buyers slightly differently. For those that are experienced, the present

²³For further details, see Train (2009).

²⁴We compute the posterior types for each buyer as $\hat{\rho}_k = \frac{\rho_k L_k}{\sum_k \rho_k L_k}$.

discounted value of buyer surplus is

$$V_{kE} = \lambda_{kE} \times Surplus_{kE} \times Hours_{kE}/r \quad (14)$$

where λ_{kE} is the job-posting arrival rate for an experienced buyer of type k , $Surplus_{kE}$ is the average hourly surplus for an experienced buyer of type k multiplied by the average number of hours per opening, and r is the interest rate, which we set to 8.7% annually or 0.16% weekly. This says that the present value of the market for experienced buyers is equal to the arrival rate of postings per period multiplied by the expected surplus conditional on a posting, discounting the future surplus using rate r .

The present value for inexperienced buyers is similar but must account for the transition to becoming experienced. The value function for an inexperienced buyer is composed of a term that contains a transition to the experienced buyer value function upon hiring and a term that takes the inexperienced buyer value function forward upon a failure to hire. This can be written as:

$$V_{kI} = Pr(Hire|I, k) \times (Surplus_{kI}Hours_{kI} + \frac{1}{1+r}V_{kE}) + (1 - Pr(Hire|I, k))\frac{1}{1+r}\lambda_{kI}V_{kI},$$

which after rearranging gives:

$$V_{kI} = \frac{Pr(Hire|I, k) \times (Surplus_{kI} \times Hours_{kI} + \frac{1}{1+r}V_{kE})}{1 - \frac{1}{1+r}(1 - Pr(Hire|kI))\lambda_{kI}}. \quad (15)$$

The numerator contains the probability of hiring while inexperienced. Upon hiring, the buyer receives the surplus on a given job and the discounted continuation value of transitioning to becoming experienced, which is given in equation (14). The denominator accounts for the fact that buyers who do not hire return in the future based on the arrival rate for inexperienced employers, given by λ_{kI} in equation (4).

3.6.2 Provider Surplus

The surplus for providers who are hired is the difference between their wage bid and estimated cost, that is, their markup, which is given by the difference between equations (7) and (8). For applicants to a job, the surplus if hired is multiplied by the probability of being hired for the job, given in equation (3). These estimates measure the expected surplus from applying to the job, per hour and per job, and can be aggregated across all applicants $j \in J_o$ and multiplied by the number of hours to find the expected provider surplus at the job posting level. In an imperfectly competitive market, the amount of provider surplus upon hiring, the expected surplus from applying, and the expected provider surplus at the job posting level, can differ based on the the level of buyer experience in the market, as buyer segments may have different elasticities. We thus present estimates separately for inexperienced and experienced buyers.

Although providers are not tracked over time in the model, the present value of surplus to all providers can be measured by finding the discounted present value of expected surplus per job posting over all of the postings on the platform. This is given by:

$$\Sigma_{oI} \left(\frac{Hours_o}{1+r} \Sigma_{j \in J_o} p_{Ioj} \times (w_{Ioj} - c_{Ioj}) \right) + \Sigma_{oE} \left(\frac{Hours_o}{1+r} \Sigma_{j \in J_o} p_{Eoj} \times (w_{Eoj} - c_{Eoj}) \right), \quad (16)$$

where each summation is over all openings observed in the 30 months of data for buyers with different experience levels, I and E .

4 Results

4.1 Demand Parameters, Arrival Rates, and Type Shares

Table 4 presents the estimates of the key parameters governing buyer demand. Appendix Table A.2 shows the sensitivity of the estimated parameters when we add controls for provider applications. Because of the similarity of results, we focus on the estimates in Table 4. The estimated distribution of the three buyer types k is shown in panel A. There is an uneven mix. 80% of buyers are grouped together as Type 1, just over 17% are Type 3, and less than 3% are Type 2. Engagement with the platform varies significantly between these three types, as will be seen shortly.

Table 4 panel B shows some of the estimated applicant-specific choice parameters from equation (1) that govern buyers' indirect utility. The log bid coefficient is the estimated α_χ coefficient. This parameter is the disutility associated with higher wage bids, and the estimate is negative and significant, as expected. Column 4 shows how this disutility changes after the buyer has hiring experience. The coefficient is an additive interaction term with an indicator for buyer experience. These results do not have separate buyer type interactions with the α parameter, as these additions do little for model fit.

The other rows of panel B show how the constant term in equation (1) varies by buyer type and hiring experience. The term labeled constant is the value of μ , and can be interpreted as the baseline value for hiring on the platform relative to the outside option in the utility function. The coefficients shown under the Type 2 and Type 3 columns are additive interaction terms relative to the Type 1 baseline. The estimates under the row labeled Experienced Constant Shift indicate that the relative utility of hiring on the platform changes differently across buyer types as they gain experience. For Type 1 buyers, experience increases their average relative value for hiring on the platform by 2.6 on top of a baseline value while inexperienced of 2.1. Experience increases the average value for Types 2 and 3 by an additional 0.23 and 0.17, respectively. This variation across types implies that the hiring responsiveness to increased hiring costs in the counterfactuals will depend on the distribution of buyer types posting vacancies over time.

Estimates of other parameters in the vector $\beta_{k\chi}$ estimates are presented in Appendix Table A.1.²⁵

Panel C of Table 4 presents the estimates for the parameters in equation (4) that determine the frequency of job posting. Column 1 shows the estimate of δ_1 for inexperienced Type 1 buyers. Taking the exponent of the -3.674 coefficient gives the number of jobs posted on average in a month for these buyers, which is just below 0.03. The Type 2 and 3 interactions indicate that these buyers post jobs much more often. That is, they are much more likely to come back to the market, even in the absence of hiring. Adding these values to the term from column 1 in the exponent gives that Type 3s post 0.22 vacancies a month and Types 2s post 1.13 vacancies a month while inexperienced.

Panel C, column 4 gives the average increase in posting frequency from gaining hiring experience—the δ_2 estimate—which is positive and significant at 1.077. Including this term and exponentiating to deliver monthly average job posting frequency yields large increases in job postings for each type once they have prior hiring experience, to 0.07, 0.63, and 3.33 jobs per month on average for experienced Type 1, 3, and 2 buyers, respectively.

The second row of Panel C is the estimated coefficient for δ_3 in equation (4), allowing the job posting frequency to respond to variation in bids received on past postings. Recall from Section 3.2 that the bid variation used to estimate how the arrival rate depends on past bids utilizes deviations in log bids net of job category fixed effects and a time trend. The average bid deviation is zero and the standard deviation in the data is 0.04. The negative coefficient of -1.112 indicates that a one standard deviation increase in idiosyncratic bids would reduce posting rates by about 4.3%. Doubling bids approximately outweighs the increases in posting frequency that come from gaining hiring experience.

Overall, experience appears to be an important factor in determining platform use on the extensive margin, but buyer type heterogeneity also matters quite a lot. Eighty percent of buyers are relatively infrequent posters at baseline, whereas the remaining 20% are more likely to post and, consequently, are more likely to gain experience, magnifying posting differences across types. Panel D of Table 4 shows that, as a result, while Type 1 buyers are estimated to be 80 percent of the overall registered market users, they only account for 26% of the jobs posted by experienced buyers. Despite representing only 2.7% of unique buyers, Type 2 buyers post 28% of the jobs among experienced buyers. This heterogeneity in platform use suggests that small changes in surplus for the subset of very active buyers may have significant consequences for market volume.

²⁵Table A.1 shows the general pattern in comparing column 1 with columns 3 and 4 of the estimates is that for the applicant characteristics that appear either positively or negatively in Type 1’s hiring probability function, Type 2 and Type 3’s hiring decisions tend to respond less to variation in these characteristics. For example, Type 1 buyers are more likely to hire an applicant with a bachelors’ or higher degree, but these coefficients are smaller in absolute magnitude for other types of buyers. Column 2 shows that buyers’ preferences over applicant characteristics change with experience. Experienced buyers care less about applicants’ ability to speak English and applicant experience on the platform, as summarized by the number of times the applicant has been hired.

4.2 Wage Elasticities and Buyer Surplus

The demand parameters allow calculation of wage-bid elasticities both at the level of the applicant and the job posting. These are shown in Table 5 panel A (and Table A.3 when we include the provider application control). The elasticity of the probability that a particular applicant is hired with respect to the wage bid is $(1 - p_{\chi oj}) \alpha_{\chi}$, where the α_{χ} coefficients are given in Table 4. The first two rows in the panel present the applicant-level elasticities by buyer type and experience, all in levels rather than as additive interaction terms. The fourth column is the overall wage elasticity estimate, weighted by the share of buyers posting the jobs. For inexperienced buyers, the elasticity of hiring an applicant with respect to the wage bid is -4.5. This elasticity falls to -5.1 with hiring experience, reflecting the negative additional α_{χ} term for experienced buyers in Table 4 Panel B.

The next two rows of Table 5 panel A show the job-level wage elasticity, labeled “Vacancy Fill Elast.” These estimates are calculated as the elasticity of filling the job with respect to a percentage change in wage bids for all applicants. The vacancy fill elasticities are naturally smaller in magnitude than the applicant choice elasticities in the first two rows because they can be interpreted as the hiring response when all providers raise their bids rather than when only a single provider does so. The vacancy fill elasticities are around -3.5 for inexperienced and experienced buyers.

Panel B of Table 5 presents the expected surplus per hour worked on a job opening, as described in equation (13). Overall surplus per hour is about \$0.51 for inexperienced buyers. Type 1 buyers have the highest expected surplus per hour of paid work when inexperienced, followed by Type 3, and then Type 2. Each type sees a large increase in surplus after gaining experience, with an overall average of \$0.87 per hour of paid work. This increase in surplus is related to experienced buyers’ increased probability of hiring on each job.

Table 5 panel B next shows the expected value of the platform to buyers, as set out in equations (14) and (15). When inexperienced, the average value of the platform across buyers is \$412, but this average masks large heterogeneity by buyer type. Type 2 buyers are a very small share of all buyers, but they post jobs frequently. As a result of their intensive platform use, these buyers have a value for the platform that is over 14 times the value of the most frequent buyer type, Type 1. However, each buyer’s surplus increases by a large amount once they are experienced.²⁶ On average, experienced buyers have a lifetime value for the platform of about \$8,340.

Panel C displays details about job filling propensity and how the features of job posts differ by the posterior estimate of buyer type. As shown in the first two rows, Type 1 buyers are most likely to hire when inexperienced and experienced, while all types see a large increase in fill rate after gaining experience.

²⁶Note that the surplus for the inexperienced already includes the expected surplus arising from becoming experienced. Even so, there is a large jump upon making the first hire.

Since the expected surplus from not filling a job is zero relative to the outside option, dividing the expected surplus of posting a job (from panel B) by the probability of filling that job (from panel C) gives an estimate of the surplus conditional on filling a vacancy. Inexperienced buyers gain $\$3.18 = \frac{0.506}{0.159}$ when filling each vacancy, while the experienced gain $\$2.82 = \frac{0.865}{0.307}$. This yields an average buyer surplus of about \$3.00 per hour for a filled job.

The last two rows of Panel C show that different buyer types post jobs that require different skills. Type 1 buyers are the most likely to post technical jobs, which may be related to why their posts are relatively infrequent. Type 2 buyers are the least likely to post technical jobs, but they generate the most per-capita job posts.

4.3 Provider Costs, Markups, and Surplus

Table 6 panel A shows estimates of expected provider surplus per hour (see Table A.4 for the estimates controlling for provider applications). The expected surplus from applying to any job is \$0.014 for jobs posted by inexperienced buyers and \$0.018 for those posted by experienced buyers.²⁷ These expected surplus numbers appear low at first glance, but this should be expected because jobs often receive many applications and only 1 worker is typically hired. However, some workers are likely to receive more surplus than the average.

One determinant of provider surplus is the order of applications to a job. We include a flexible spline for the applicant’s arrival order in the X_{oj} matrix in equation (3), with knots that correspond to the pagination order that would require a buyer to click on a new page of applicant profiles. The elasticity estimates therefore control for an applicant’s arrival position, which is known to the applicant when making a bid. Early applicants are far more likely to be hired, and early applicants submit higher wage bids. Table 6 shows, accordingly, that expected provider surplus declines with arrival order. Providers among the first ten applicants obtain almost twice as much expected surplus as applicants 11 to 20. Figure 3 shows buyers rarely interview or hire late-arriving applicants. A buyer’s probability of choosing later applicants falls faster than the prospective benefits of competition from more applicants, allowing positive equilibrium markups.

Table 6 Panel B decomposes providers’ hourly bids into their hourly costs and markups based on equation (8) when applying to jobs posted by each buyer type and experience level. Estimated average hourly costs are between \$6.49 and \$7.12 per hour, and individual cost estimates correlate with characteristics that likely vary a provider’s opportunity cost of work, like the per-capita GDP in their country.

Estimated markups are fairly similar across buyer types, at 29% when applying to an inexperienced

²⁷Because the estimated semi-elasticities in Table 5 don’t vary significantly by buyer type, the provider’s optimal wage bid is insensitive to type. Differences in job surplus across types are included here for completeness, and arise mostly from difference in job characteristics across buyers.

buyer and around 24% when applying to an experienced buyer. Because there is little bargaining between offered wage bids and final wages, hired workers potentially earn significantly more than their outside option or reservation wage.²⁸

Auxiliary analysis of data from outside the model helps to validate that these markup estimates are reasonable. For example, the minimum wage bid for a provider in a narrow time interval is an upper bound on their reservation wage, although this bound is likely not sharp (reservation wages are likely lower) if all bids have some markup. Bids above the minimum observed wage either indicate surplus relative to the indifference point between work on the platform and an outside option or unobserved heterogeneity in the cost of working on a particular job. In raw data without accounting for heterogeneity in the cost of working, we find evidence suggesting substantial markups. For example, we consider a sample of providers with at least two applications to job postings in the same job category in the same month. The difference between the highest and lowest hourly bid in a month divided by the lowest bid, across producers and months, has a median value of 23%. This large difference in raw bids in the data suggests that an average markup of 24% to 29% is plausible.²⁹

Other patterns in the data suggest provider wage markups are significant even after controlling for the difficulty or cost of working on a job. For example, Figure 4 displays the coefficients of a regression of log bids on applicant order dummies, provider-by-week fixed effects, and measures of a job difficulty and expected feedback from the buyer. This simple focus on applicant order indicates that the same provider drops his or her bid by about 8% if applying to a job as applicant 100 rather than one of the first ten applicants. This pattern is consistent with providers tailoring markups based on the wage bid elasticity that is predictable over the applicant order distribution. These estimates display only a small portion of the total variation in bids, as they capture only predictable rather than idiosyncratic factors that lead providers to vary their markups.

Table 6 panel C shows the expected surplus for providers who are hired. These numbers are much larger than those in panel A. Being hired by an inexperienced buyer gives an expected surplus of \$2.18 per hour. Reflecting the lower optimal markups to experienced buyers, being hired by one generates a provider surplus of \$1.74 per hour. Jobs last on average 70 and 77 hours for inexperienced and experienced buyers, yielding total surplus from a hire of \$153 and and \$134.

²⁸In preparing Stanton and Thomas (2016), we investigated the extent that wages changed between first bids and contracted rates for those who are hired. We found very limited bargaining from initial offer to final wage.

²⁹This strategy is unlikely to be contaminated by time-varying reservation wages, as in Chen et al. (2019), as work on this platform is not instantaneous. To avoid issues with conflating costs that vary over time, we have also done these calculations while only looking at jobs of similar duration, yielding similar results. In the structural estimation, provider markups are calculated controlling for job category and expected job duration, which are included in the X_{oj} terms in equation (1).

5 Counterfactual Surplus Under Hiring Cost Increases

The estimated model permits evaluation of counterfactual policy environments. It is likely that any regulatory change aiming to shift the balance of power in favor of providers would increase buyers' hiring costs. Two types of changes are considered here. We first evaluate a 10% ad-valorem tax on wage bids in addition to the platform share of revenues. This mimics the FICA tax that is statutorily imposed on employers in traditional labor markets as well as compliance costs that may accompany such a policy. This first counterfactual does not affect the statutory incidence of the tax for providers because we focus on adjustment from the demand side. However, providers' bids are determined in equilibrium, and the incidence of the tax may be split between buyers and providers.³⁰ We then consider an hourly wage bid floor of \$7.00 per hour, which would affect the lowest-paid providers in the market. The buyer and provider responses to these changes are simulated to assess the impact on total surplus, and its distribution among buyers and providers.

In the counterfactual simulations, we assume that the alternative policy environment is in place for the entire data period studied, which allows us to assess how total surplus would evolve over the entire 30 months of data for the buyers and providers in the market. However, we note that the buyers who have prior experience at the start of the period also enter the counterfactuals as having prior experience, which means these counterfactuals, that potentially affect the arrival rates of jobs, are akin to studying a regime switch in market institutions.

In both counterfactuals, we find increases in buyer costs affect a variety of different margins that, when added together, have a significant impact on market surplus. In particular, in each case, providers re-optimize their bids, and wage bids to buyers increase. In response, buyers become less likely to hire on posted jobs, fewer buyers gain hiring experience and, as a consequence, fewer subsequent jobs are posted. This sequence reduces the opportunities for trade in tasks. The magnitude of these responses is informed by the counterfactual exercises.

A few technical details are required to understand how we calculate the counterfactual results. For each counterfactual scenario, our baseline results assume that the same applicants apply to jobs under the counterfactual (we later conduct robustness checks that relax this assumption). That is, providers' observable characteristics, other than their wage bids, are assumed unchanged. In the tax counterfactual, the additional tax is included as an additional wedge between the hourly wage providers receive and the price buyers pay. In the wage floor counterfactual, the floor is included in the providers' problem as a

³⁰We do not account for benefits that may accrue to providers under the policy, as many of the providers are coming from outside of the United States and would have difficulty accessing domestic programs financed by the FICA contribution. Instead the tax is assumed to provide equalized regulation to provide the same treatment between traditional employment and contract work.

constraint.³¹ In each case, providers select the optimal bid to maximize their payoffs in equation (5) given the composition of the applicant pool, the buyer’s semi-elasticity of demand to wage bids, and the additional 10% tax or wage constraint. Having posted a job, a buyer selects the option that maximizes her indirect utility out of the choice set $\{J_o, 0\}$ given the simulated wage bids, where the probability any provider $j \in J_o$ is chosen is given by equation (3). The wage bids observed on any posted job opening impact the rate at which each buyer posts subsequent jobs as set out in equation (4). When buyers post fewer jobs overall during the time period under the counterfactual, they are assumed to post the existing jobs in the same order, but with a delay when compared to the time of posting in the observed data. The number of jobs posted by each buyer is the number determined by equation (4) up to the cut off date when the simulated time period ends.

For each of the two main policy changes, we evaluate how altering several assumptions or parameter estimates influence our conclusions. To illustrate the role of buyer dynamics around job posting, we contrast the full estimates with those that impose the frequency of job posting is unrelated to the past wage bids a buyer has received. In this alternative, the buyer’s posting frequency, determined by equation (4), sets the term δ_3 to zero, breaking the link between past wage bids and job posting. Changes in surplus in this alternative result from differences in hiring decisions, and in changes to subsequent job posting rates, arising only from whether or not a buyer has past hiring experience on the platform. Hence, the differences between this alternative and the main scenarios reveal how much of the overall reduction in surplus to buyers and providers can be attributed to the job-posting-frequency elasticity with respect to historical wages.

We also illustrate the sensitivity of our findings to alternative assumptions about the composition of job applicants, which is especially relevant for the wage floor counterfactual. This adjustment accounts for the fact that the provider’s problem includes the implicit choice of whether or not to apply for a job. It is plausible that providers whose skills and qualifications lead them to submit bids below the relevant wage floor would view the probability of being hired to be so low in the counterfactual that the application costs exceed the expected benefits of applying. The new policy environment might also, however, attract new providers with the skills and qualifications that would make them competitive applicants at higher wage levels. This would allow for the possibility of “labor-labor” substitution Hamermesh (1986), where buyers would be willing to pay higher wages for more productive workers. This alternative is considered only for the wage floor counterfactual since the 10% tax applies across the board, for all providers, and is less likely to have a disproportionate effect on the participation decisions of providers currently bidding low wages.

To approximate the changes to the applicant pool under labor-labor substitution, applications with current bids below 90% of the counterfactual wage floor are removed from the applicant set and replaced

³¹For the hired providers to receive wages of at least \$7.00, the buyers are charged an hourly rate of at least \$7.70 since their costs also include the 10% of revenues charged by the platform.

with an equal number of new applications. The new applicants are assigned observable characteristics, X_{oj} , that are a random draw from the distribution of applicants whose original wage bids are above the wage floor.³² This change allows us to analyze static hiring and market dynamics under conditions where the market itself becomes more appealing for higher-quality applicants under the \$7.00 wage floor.

The counterfactual outcomes are summarized in Table 7. Column 1 presents the changes in surplus in the first counterfactual where buyers face a 10% tax. The applicant pool J_o remains the same as in the data and all providers re-optimize their bids. Bids to both inexperienced (experienced) buyers increase by 8.6%. Hiring rates fall by between 26- 27%. The surplus to buyers on a static opening, labeled “Static Pct. Change in Buyer Surplus” is calculated using equation (13). The bid increases and hiring rate reductions reduce buyer surplus created on posted jobs by 24% for inexperienced buyers and by 26% for experienced buyers.

The next several rows in Table 7 Panels A and B show the dynamic implications of the increases in buyer costs. Because of the reduction in hiring rates on posted jobs, fewer buyers gain hiring experience. More buyers mechanically remain inexperienced due to their reduced propensity to hire. Job postings for inexperienced buyers change little because there are more of them, but job posts fall 49.1% for experienced buyers because fewer inexperienced buyers become experienced. The percentage reductions in buyer surplus are calculated using equations (15) and (14). Because the number of jobs posted by experienced buyers falls so dramatically, the reduction in the present value of their surplus is larger than their static loss, at 56.8%. Inexperienced buyers also see a large decrease in surplus, of 68.5%, because lower hiring rates means fewer become experienced, as well as the fact that the continuation value of gaining experience is lower.

For providers, static surplus from a given job falls by an average of 28%. The tax does little to change their net take from a wage bid, in part because residual demand elasticity changes, so these reductions arise from reduced hiring rates. Because fewer jobs are posted and fewer jobs are filled, the present value of total provider surplus, calculated using equation (16), falls by 52.0%.

Column 2 presents results where the dependence of job posting frequency on past wage bids is set to zero for the 10% tax. The static results are the same as the base case. The impact of the δ_3 coefficient can be seen in the percentage reduction in the number of posted jobs. Setting the dependence on past bids to zero leads to an 11% reduction in jobs posted by experienced buyers. This 11% reduction in posts arises because fewer buyers become experienced. Despite job posts numbers falling by a fifth of the base

³²The random draw is over different candidates, so each application consists of a real resume from a different candidate. We do not assign a random draw of wages, but take the new candidates’ estimated opportunity costs of work as a function of observable characteristics and assume that bids are set as markups over costs. We continue to assume that providers’ first order conditions hold subject to the wage floor constraint, and applicants wage bids are computed to be optimal given buyers’ residual elasticity of demand. Buyers proceed by selecting the indirect utility-maximizing option as before. The effect of past wage bids on job posting rates is set to zero, otherwise the effect of past wage bids on future hiring would be confounded by past applicant quality changes under this alternative scenario.

reduction, surplus reductions are still substantial. Relative to the status quo baseline, surplus falls by 46% for inexperienced buyers and by 27% for experienced buyers. Provider surplus falls by 30%, largely because of the reduction in hiring rates.

Column 3 presents the second main counterfactual, where the applicant pool J_o remains the same as in the data and all providers re-optimize their bids in the presence of the \$7.00 wage floor. The hourly wage bids to inexperienced buyers increase by 34.0% and, for experienced buyers, by 46.8%. The larger increase for experienced buyers reflects the fact that they receive lower wage bids at baseline. Hiring rates on posted jobs fall in response to higher wage bids, by 31.1% and 35.4% respectively. The reduction in hiring rates may seem large when under 40% of applicants would be bound by the floor, but the reason becomes clear after considering which buyers are most affected by each counterfactual. Under the wage floor, some openings, like those in non-technical jobs, see extremely large wage bid increases, while jobs that require greater skills will see very small changes. A large share of the hiring reduction is arises from employers posting non-technical jobs.

The static buyer surplus on posted jobs falls by less than half as much under the wage floor counterfactual in column 3 as in the 10% tax scenario in column 1. This is because surplus estimates are averaged over all buyers, and the higher-skill jobs, which generate more surplus are relatively unaffected by the wage floor. However, taking the dynamic effects into account, the present value of inexperienced buyer surplus continues to fall by a large amount, 69%, and the surplus for the experienced buyers falls by 55%. Both reductions arise because the job posting frequency falls for experienced buyers, reducing the continuation value of being experienced, which is reflected in the surplus for both the experienced and inexperienced. The final rows of Table 7 column 3 shows that the present value of the surplus to providers falls by 41% under this counterfactual.

Column 4 contains the case where job posting rates do not depend on past wage bids ($\delta_3 = 0$). Surplus amounts remain negative for all parties, but the reductions in the present value of provider surplus are much smaller, at around 8% relative to the baseline, because job posting frequency remains elevated.

The final column of Table 7 presents the alternative scenario where the applicant pool is adjusted to mirror the characteristics of providers who typically submit bids above the \$7.00 floor. That is, it presents buyers with “better” applicant pools. The job posting rate here remains insensitive to past bids received, as in column 4. Because these applicants no longer face competition from lower-wage providers, the optimal bids in this case increase more than in the main scenario. However, hiring rates fall by less than in column 4 because buyers perceive applicants to be better quality.

The reduction in the number of jobs posted in column 5 is slightly smaller in magnitude than in column 4 because higher hiring rates mean more buyers transition to being experienced. The present value of buyer surplus reductions are hence also smaller. The increased hiring and experienced buyer job-posting rates, relative to column 4, translate into a smaller reduction in the present value of provider surplus in this

alternative. However, the change in the composition of applicants following the introduction of the wage floor does not reverse the finding that providers are made worse off. The present value of provider surplus falls by 7.6% even though this alternative scenario removes several of the features that served to reduce the size of the market in the main scenarios.

The most striking implication of the counterfactual analysis is the large reduction in the total number of posted jobs in each of the base cases. The impact of the buyer dynamics is further illustrated in Figure 5. Each panel plots the log of the observed number of job postings in the data over the 30 months. The left hand panel corresponds to the 10% tax scenario in column 1 of Table 7 and the right hand side panel corresponds to the \$7.00 wage floor scenario in column 3.

In each figure, the simulated baseline is the simulated log number of postings that would result from modeling the job posting of the buyers observed in the data under the parameters estimated and including the observed wage bids. The counterfactual is the log of the job postings simulated using the same buyers when the relevant counterfactual scenario is imposed from the start of the time period. The difference between the areas under the two dotted lines represents the reduction in job postings under each scenario. In both figures, these areas amount to reduction in job postings of at least 30%.

6 Online-Offline Wage Elasticities

The previous section showed that increased hiring costs would greatly reduce the number of jobs posted and the number of hires made on the platform. A full analysis of the welfare implications of these counterfactuals would incorporate changes in surplus in external labor markets, as off-platform workers may potentially serve as substitute sources of labor for the platform considered here. Such an analysis would require observing the outside option for potential and actual buyers. Unfortunately, the data do not contain this information, either for other online platforms or traditional labor market channels. However, to the extent that regulations would apply to other platforms, the substitution with offline hiring is most relevant for understanding welfare.

To offer insight on the extent to which online and offline sources of labor are viewed as substitutes for buyers of online labor services, we ask what happens to platform demand when there are exogenous changes in offline wages in buyers' local labor markets. The exogenous changes come in the form of increases in US State-specific minimum wages. In July 2009, which is just over half-way through the data sample, all but nine US States increased the hourly minimum wage, either because of the change in the Federal minimum wage at that time, or as a state-specific increase that coincided with the timing of the Federal minimum wage increase if the State minimum exceeded the Federal minimum.³³ Online job postings were

³³The minimum wage increase on July 24, 2009 is discussed here: https://www.epi.org/publication/mwig_fact_sheet/. The Federal Minimum wage increased by 10.6%, from \$6.55 to \$7.25. The nine states unaffected by the minimum wage increase

not subject to the same minimum wage regulations or to any increase in wages around this time. If online and offline labor are viewed as close substitutes, then the states affected by the offline minimum wage increase would have been likely to see an increase in online labor demand around July 2009 relative to demand from unaffected states.

We consider state-month-level numbers of jobs posted and the share of postings filled. We also divide jobs into technical and non-technical postings to isolate the non-technical postings that are most likely to be affected by minimum wage changes. The offline minimum wage increase might be more relevant for buyers with tasks that require fewer skills and qualifications and, on average in the data, tasks in non-technical jobs attract lower wage bids. Because the findings in previous sections highlight the differing job posting behavior of inexperienced and experienced buyers, any differences in response to offline wage increases can also be explored by further dividing jobs into those posted by each buyer group. A local increase in offline wages may encourage existing buyers of platform services to increase posting frequency, or might encourage more first-time job posters.

Figure 6, panels A and B, plot the log number of jobs posted on the platform in each month by buyers located in states that were and were not affected by the local minimum wage increase, relative to that group's mean number of postings prior to May 2008. In each figure, the top pair of lines are postings by experienced buyers, and the lower pair of lines are postings by inexperienced buyers. Panel A of the figure includes all job postings and panel B is all postings in non-technical tasks. The time of the minimum wage increase, in July 2009, is marked in the figures by vertical lines.

Comparing the series, there are no differential increases in the number of postings by buyers located in affected states at the time of the increase in the minimum wage. Regressions of the log counts of job postings in these figures on monthly indicators show no significant differences in trend for affected and unaffected states, and Wald tests reveal no evidence of structural breaks in trend for affected states around this time.

It is possible that the lack of online demand response is due to the fact that potential buyers are unaware of the online opportunity, but lack of awareness is not an issue for buyers who are already posting jobs in this labor market. To consider the behavior of existing buyers, we assess job fill rates after posting a job under different minimum wage regimes. We regress the probability of hiring on a posted job on month fixed effects, month fixed effects interacted with an indicator for whether the buyer is in a state with a minimum wage increase, and buyer fixed effects. Figure 7 panel A plots the difference in the average probability that a posting was filled when buyers are located in an affected state. When postings from affected states had a higher average fill rate than those from unaffected states in a given month, the corresponding point in

were California, Massachusetts, Michigan, Rhode Island, West Virginia, New Hampshire, Maryland, Iowa, and Hawaii. The Federal minimum wage had last increased to \$6.55 per hour on July 24, 2008. In this section of the paper, we consider online activity from mid-2008 onwards, and the data are relatively sparse from earlier months.

the figure is positive. The fill rate differences are generally very similar over time and there is no increase in the relative fill rate from affected states following the increase in offline minimum wages. Figure 7 panel B plots the difference in the hiring rate for non-technical tasks between jobs posted in affected states and unaffected states in each month. Again, here, the fill rate in affected states does not increase at or after the increase in local offline minimum wages.

Finally, we show whether being in a state affected by the minimum wage increase is related to the average number of job postings per month at the buyer level. Figure 7 panels C and D plot the interaction coefficients showing the differential number of jobs posted by buyers in affected states from a regression including month fixed effects as well as buyer fixed effects. There is no noticeable difference in this series around the time of the minimum wage increases in either figure.

In sum, the findings in Figures 6 and 7 suggest that increases in local offline wages induce very few buyers to look for labor on this platform, either on the extensive margin of bringing new buyers to the platform to post jobs for the first time, or on the intensive margins of increasing the job posting frequency or job-fill rate for those that were already using the platform. While it is not immediately clear that the cross-price elasticities are symmetric (so that increases in online hiring costs would have a limited impact on labor demand offline), other studies provide evidence for this symmetry. For example, Horton (2017b) discusses the prevalence of “multi-homing” on online platforms. Surveys of buyers on the large online platform studied there revealed that online and offline hiring are only very weak substitutes, and that multi-homing between platforms was relatively rare. In the experiment described there, the buyers who were unexpectedly and randomly subject to a minimum wage in the platform after posting jobs did not significantly decrease hiring rates but, instead, switched to higher-quality providers. The finding that platform demand responds very little to changes in local offline wages is also consistent with Horton, Kerr, and Stanton (2017), who find minimal cross price elasticities between US and foreign workers. Most of the applicants who would have been impacted by the imposition of a wage floor on the platform are from outside the US. However, the low degree of substitutability between US and foreign workers on the platform, and between platform and offline workers, suggests the changes in relative wages studied here would do little to increase labor demand in the US.³⁴

From the data observed in the present study describing the work to be done in each job, it is also clear that buyers continue to post task-based jobs, rather than hiring for long-term roles, as they gain hiring experience. That is, the nature of the labor demand tends to remain idiosyncratic and the work arrangements look quite different from those seen in traditional offline settings. The share of job postings expected to last for less than one week is 27% for inexperienced buyers and 28% for those with prior hiring experience. The mean number of hours worked on an experienced buyer’s filled job is similar for

³⁴While the static estimates in this paper are consistent with the conclusions of these studies, the dynamic implications of the buyer lifecycle online are not addressed elsewhere.

an inexperienced buyer's filled job, at 77 hours compared to 70. A regression of the total number of hours worked per hire on the number of prior hires and buyer fixed effects shows that there are no within-buyer changes in job length as they gain more experience. Although this evidence is more descriptive, it suggests that experienced buyers continue to take advantage of the flexibility offered by the online setting to post jobs on an as-needed basis, something that may be harder to do offline.

7 Conclusion

How much value is created by the ability to hire and supply work through online labor markets, and how is the surplus split between the demand and supply sides? This paper uses historical data from the early days of a large online labor market to estimate the surplus it creates for buyers and providers. The approach recovers estimates of buyers' surplus from a demand model of buyers' propensity to hire individual providers. The model also accounts for unobserved buyer heterogeneity in value for using the platform and history dependence in posting jobs that depends on past hiring. The global nature of the platform and idiosyncratic variation in competition provide instruments for the prices buyers face, allowing us to isolate demand elasticities.

Even though the potential providers outnumber the buyers on the platform by almost three to one, around two fifths of the surplus created when a job is filled goes to the provider. Providers are able to exert local market power in the applicant pool, and charge average markups of about 25% over their outside option. Given the abundance of providers, why are markups not competed away? Two sources of market power are differentiation among applicants and buyers' limited search, such that early applications to a job opening include higher markups in their wage bids. As a result, fears that intense competition among job applicants would limit surplus for the supply side do not appear to play out in practice.

The model jointly estimates buyer demand on posted jobs and job posting dynamics. These dynamics are particularly important when evaluating the welfare effects of counterfactual policies, like regulations, that may alter the nature of the market. Changes that increase the costs of hiring in the platform reduce the total surplus and negatively impact both buyers and providers. This mainly arises in the counterfactuals that impose regulations because buyers post far fewer jobs when it becomes more costly to hire. A 10% tax on market transactions, mimicking payroll taxes for W2 employees, would have reduced the present value of the platform by over 30% for the providers applying to jobs during the time period studied. Surplus would fall by over 20% for buyers on the platform. A \$7.00 hourly wage floor would have reduced the present value of the platform by 41% for incumbent providers, by at least 35% for inexperienced buyers, and by 6% to over 50% for experienced buyers depending on assumptions about how the composition of providers changes under the wage floor.

It is possible that the surplus lost from jobs not posted on the platform under the counterfactuals would

be offset by surplus gains in offline employment. Several pieces of evidence suggest that offsetting gains for offline work would be minimal. In particular, online labor demand did not respond to exogenous increases in the costs of offline hiring over the data period studied. The lack of responsiveness is consistent with buyers perceiving little substitutability between online and offline labor sources. In sum, increased hiring costs in online labor platforms appear likely to reduce the gains from trade in these settings by reducing market surplus and transaction volumes without increasing offline employment.

References

- Agrawal, Ajay, John Horton, Nicola Lacetera, and Elizabeth Lyons. 2015. "Digitization and the Contract Labor Market: A Research Agenda." In *Economic Analysis of the Digital Economy*. University of Chicago Press, 219–250.
- Agrawal, Ajay, Nicola Lacetera, and Elizabeth Lyons. 2016. "Does Standardized Information in Online Markets Disproportionately Benefit Job Applicants from Less Developed Countries?" *Journal of International Economics* 103:1–12.
- Barach, Moshe and John J Horton. 2017. "How Do Employers Use Compensation History?: Evidence from a Field Experiment." Tech. rep.
- Brinatti, Agostina, Alberto Cavallo, Javier Cravino, and Andres Drenik. 2021. "The International Price of Remote Work." Tech. rep., National Bureau of Economic Research.
- Castillo, Juan Camilo. 2020. "Who Benefits from Surge Pricing?" *Available at SSRN 3245533* .
- Chen, M Keith, Peter E Rossi, Judith A Chevalier, and Emily Oehlsen. 2019. "The value of flexible work: Evidence from uber drivers." *Journal of political economy* 127 (6):2735–2794.
- Clemens, Jeffrey, Lisa B Kahn, and Jonathan Meer. 2018. "The minimum wage, fringe benefits, and worker welfare." Tech. rep., National Bureau of Economic Research.
- Collins, Brett, Andrew Garin, Emilie Jackson, Dmitri Koustas, and Mark Payne. 2019. "Is gig work replacing traditional employment? Evidence from two decades of tax returns." *Unpublished paper, IRS SOI Joint Statistical Research Program* .
- . 2020. "Has the gig economy replaced traditional jobs over the last two decades? Evidence from tax returns."
- Datta, Nikhil, Giulia Giupponi, and Stephen Machin. 2018. "Zero hours contracts and labour market policy." *Economic Policy* .
- Ghani, Ejaz, William R Kerr, and Christopher Stanton. 2014. "Diasporas and Outsourcing: Evidence from oDesk and India." *Management Science* 60 (7):1677–1697.
- Gorback, Caitlin. 2020. "Your uber has arrived: Ridesharing and the redistribution of economic activity." Tech. rep., Wharton working paper.
- Hall, Jonathan V, John J Horton, and Daniel T Knoepfle. 2021. "Pricing in designed markets: The case of ride-sharing." Tech. rep., Working paper, Massachusetts Institute of Technology.
- Hamermesh, Daniel S. 1986. "The demand for labor in the long run." *Handbook of labor economics* 1:429–471.
- Horton, John, William R Kerr, and Christopher Stanton. 2017. "Digital labor markets and global talent flows." In *High-skilled migration to the United States and its economic consequences*. University of Chicago Press, 71–108.
- Horton, John J. 2010. "Online Labor Markets." In *International Workshop on Internet and Network Economics*. Springer, 515–522.
- . 2017a. "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment." *Journal of Labor Economics* 35 (2):345–385.
- . 2017b. "Price Floors and Employer Preferences: Evidence from a Minimum Wage Experiment." Tech. rep.
- . 2021. "The Ruble Collapse in an Online Marketplace: Some Lessons for Market Designers." Tech. rep., National Bureau of Economic Research.
- Jin, Li, Scott Duke Kominers, and Lila Shroff. 2021. "A Labor Movement for the Platform Economy." Tech. rep., Harvard Business Review.
- Kässi, Otto and Vili Lehdonvirta. 2018. "Online labour index: Measuring the online gig economy for policy and research." *Technological forecasting and social change* 137:241–248.

- Kässi, Otto, Vili Lehdonvirta, and Fabian Stephany. 2021. “How many online workers are there in the world? A data-driven assessment.” *A Data-Driven Assessment (March 23, 2021)* .
- Katz, Lawrence F and Alan B Krueger. 2019. “Understanding trends in alternative work arrangements in the United States.” *RSF: The Russell Sage Foundation Journal of the Social Sciences* 5 (5):132–146.
- Lyons, Elizabeth. 2017. “Team Production in International Labor Markets: Experimental Evidence from the Field.” *American Economic Journal: Applied Economics* 9 (3):70–104.
- Mas, Alexandre and Amanda Pallais. 2017. “Valuing alternative work arrangements.” *American Economic Review* 107 (12):3722–59.
- Nosko, Chris and Steven Tadelis. 2015. “The limits of reputation in platform markets: An empirical analysis and field experiment.” Tech. rep., National Bureau of Economic Research.
- Pallais, Amanda. 2014. “Inefficient Hiring in Entry-Level Labor Markets.” *American Economic Review* 104 (11):3565–99.
- Petrin, Amil and Kenneth Train. 2010. “A Control Function Approach to Endogeneity in Consumer Choice Models.” *Journal of Marketing Research* 47 (1):3–13.
- Riordan, Michael H. 1986. “Monopolistic Competition with Experience Goods.” *The Quarterly Journal of Economics* 101 (2):265–279.
- Stanton, Christopher T and Catherine Thomas. 2016. “Landing the First Job: The Value of Intermediaries in Online Hiring.” *The Review of Economic Studies* 83 (2):810–854.
- . 2020. “The Gig Economy Beyond Local Services and Transportation.” In *CESifo Forum*, vol. 21. München: ifo Institut-Leibniz-Institut für Wirtschaftsforschung an der , 21–26.
- Train, Kenneth E. 2009. *Discrete Choice Methods with Simulation*. Cambridge University Press.

Figures

Figure 1: A Job Posting.

Data Entry and Validation

Hourly – Less than 1 month - 30+ hrs/week - Posted 1 day, 13 hours ago

[amazon-web-services](#) [data-mining](#) [microsoft-excel](#) [web-scraping](#)

[POST A JOB LIKE THIS](#)

[Sign up to Apply](#)

Job Description

We are looking for someone to assist us with associating part numbers and UPC's with the correct platform numbers. We will supply spreadsheets with the part numbers and the individuals responsibility is going through a specified website to validate the information we are trying to post.

Job Overview

Type	Hourly
Workload	Full-time - 30+ hrs/week
Duration	Less than 1 month
Posted	July 13 2014, 5:39 PM
Planned Start	July 13 2014
Visibility	Public
Category	Administrative Support
Sub-category	Data Entry

Other open jobs by this client

Fixed-Price – [Customer-vendor platform](#)

Hourly – [Data Entry](#)

Fixed-Price – [Innovative Logo Required](#)

[more...](#)

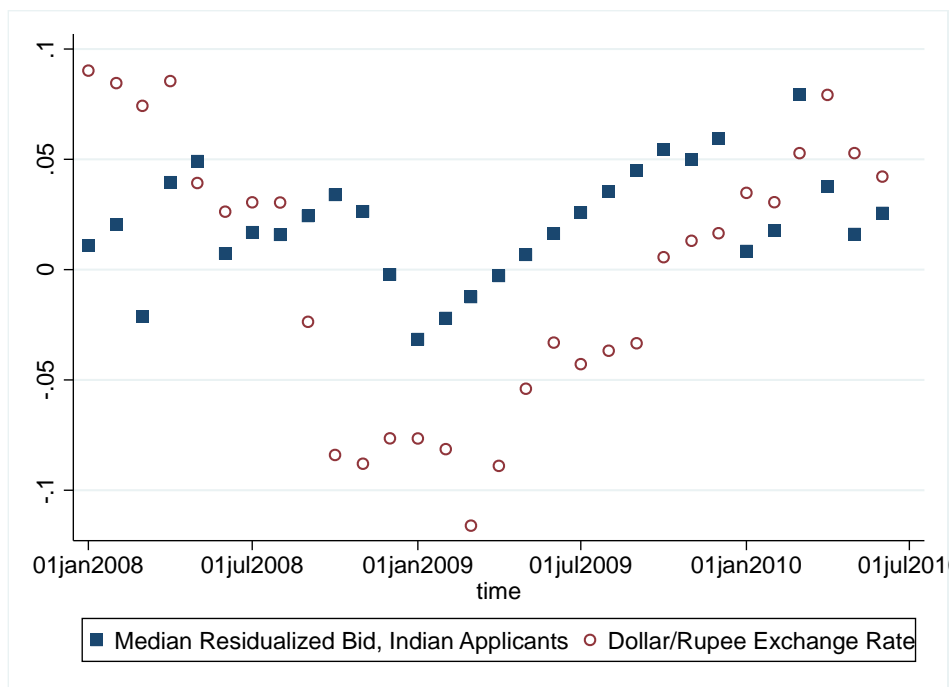
About the Client

★★★★★ United States (UTC-05)

Member Since March 26 2014

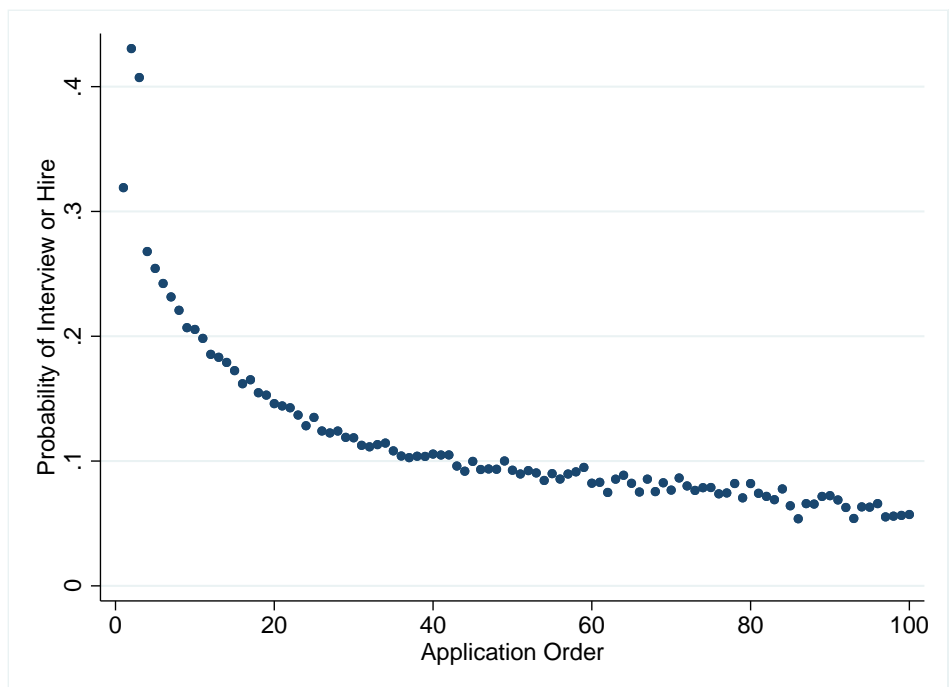
Total Spent	\$1,118
Hours Billed	217
Jobs Posted	12

Figure 2: Median Residual Log Bids and Detrended Exchange Rates for India.



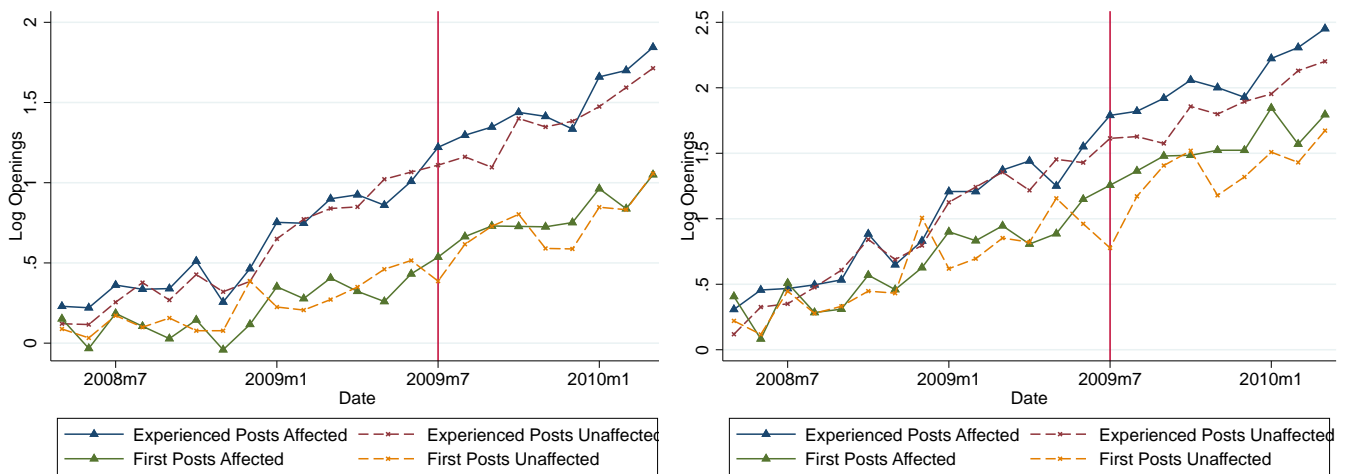
Notes: This figure plots median residual log wage bids from Indian applicants in each month against the log US Dollar to Indian Rupee exchange rate after removing a time trend. Log wage bid residuals are net of job category fixed effects.

Figure 3: Probability of A Provider Being Interviewed or Hired, by Applicant Order



Notes: This figure plots the probability that a buyer either interviews or hires a provider as a function of the order in which she applies.


Figure 6: Log Job Posts in States Affected and Unaffected by the July 2009 Minimum Wage Increase
 Panel A: All jobs
 Panel B: Non-technical jobs




Notes: The solid lines with triangle markers show the log number of job postings from buyers in states affected by an offline minimum wage increase in July 2009. Each series is normalized relative to the mean level in the data in months prior to May 2008. The dotted lines show the log number of job postings from buyers in unaffected states. In each panel, the top two lines are posts for experienced buyers and the lower two lines are posts for buyers posting their first job. Panel (a) shows all job postings; panel (b) shows non-technical job postings.

Figure 7: Difference in probability of hiring and posting for buyers in states affected by the minimum wage increase after July 2009


Panel A: Prob(Hire), All jobs




Panel B: Prob(Hire), Non-technical jobs



Panel C: Number of Posts, All jobs



Panel D: Number of Posts, Non-technical jobs.



Notes: These figures plot coefficients and 95% confidence intervals on the interaction of the affected state indicator and month fixed effects from a regression of the probability a job is filled (Panels A and B) or the number of postings per buyer per month (Panels C and D). The model includes month fixed effects and buyer fixed effects. Affected states are those that raise their minimum wage at the time the Federal minimum wage changed in July of 2009. Panels A and C show all job postings; panels B and D show non-technical job postings.

Tables

Table 1: Descriptive statistics about buyers

	Mean	Std. Dev.
In the USA	0.58	0.49
In another G10 Country	0.17	0.38
In India or Philippines	0.06	0.24
First Posting is Tech	0.57	0.50
Number of Jobs Posted	2.52	4.48
Number of Buyers	67292	

Notes: The sample includes buyers active on the platform from January 2008 to June 2010. Buyers classified as coming from another G10 countries include those in Belgium, Canada, France, Germany, Italy, Japan, the Netherlands, Sweden, Switzerland, and the United Kingdom. Technical jobs are those in Networking and Information Systems, Software Development, and Web Development. Non-technical jobs are those in Administrative Support, Business Services, Customer Service, Design and Multimedia, Sales and Marketing, and Writing and Translation.

Table 2: Descriptive statistics about postings, by number of prior buyer hires

	All Openings	0 prior hire	1 prior hire	2 prior hires	3+ prior hires
Months between postings	1.15 (2.58)	1.53 (3.24)	1.49 (2.93)	1.21 (2.27)	0.67 (1.47)
Number of applicants	26.28 (32.90)	23.66 (26.70)	25.76 (29.15)	27.50 (35.61)	30.00 (40.28)
Share of US applicants	0.11 (0.15)	0.11 (0.15)	0.11 (0.15)	0.11 (0.15)	0.11 (0.15)
Share of applicants with good feedback	0.34 (0.17)	0.33 (0.17)	0.35 (0.16)	0.35 (0.16)	0.35 (0.16)
Mean wage bid, USD	10.88 (4.94)	11.28 (4.89)	10.97 (4.85)	10.66 (4.83)	10.31 (4.99)
Share of wage bids below 7.00USD	0.34 (0.34)	0.32 (0.33)	0.33 (0.33)	0.35 (0.34)	0.38 (0.35)
Probability of making hire	0.21 (0.40)	0.15 (0.36)	0.28 (0.45)	0.28 (0.45)	0.25 (0.44)
Number of postings	169578	84486	16352	10032	58708

Notes: This table presents opening-level averages (standard deviations) of posting patterns, application characteristics, and hiring rates. Months between postings is calculated as the average time between job postings and is right-censored at the end of the sample period. Good feedback is defined as a feedback score greater than 4.5 out of 5. The mean wage bid is inclusive of a 10% ad-valorem platform fee. Sample period is from January 2008 to June 2010. Standard deviations in parentheses.

Table 3: First Stage Regressions of log Hourly Wage Bids on Instruments

	(1)	(2)	(3)	(4)
Exchange Rate Instrument	0.083*** (0.010)	0.074*** (0.010)	0.087*** (0.012)	0.077*** (0.012)
Competition Instrument	-0.070*** (0.008)	-0.068*** (0.008)	-0.069*** (0.010)	-0.068*** (0.010)
Exchange Rate x Exper Buyer			-0.007 (0.021)	-0.007 (0.020)
Competition x Exper Buyer			-0.002 (0.015)	-0.001 (0.015)
Provider Applications / Month		-0.001*** (0.000)		-0.001*** (0.000)
R-Squared	0.630	0.632	0.630	0.632
Observations	4456676	4456676	4456676	4456676
F Clustered on Posting	92.66	81.26	50.31	44.30
F Clustered on Provider	66.42	61.28	33.89	31.49

Notes: The exchange rate instrument is the log of the dollar to local currency exchange rate. The competition instrument is the log leave-own-opening out average number of applicants to openings in the same job category in that week to arrive within the first 24 hours after the initiation of a job posting. We then net out job category and week fixed effects via regression. Columns 3 and 4 add buyer experience interactions with the instruments. Columns 2 and 4 control for the number of jobs a provider applies to in a given month. Standard errors clustered by job opening are in parentheses. Partial F statistics on the excluded instruments are reported after clustering by job opening or by provider.

Table 4: Demand and Arrival Rate Parameter Estimates

	Type 1 (Baseline)	Type 2	Type 3	Experienced Additive Interactions
Panel A: Employer Types				
Type Share	0.805 (0.002)	0.026 (0.001)	0.169 (0.002)	
Panel B: Choice Model Parameters				
Log Bid	-4.497 (0.943)			-0.690 (1.156)
Constant	2.130 (1.640)	-1.374 (1.187)	-0.335 (0.936)	
Experienced Const. Shift	2.646 (2.085)	0.232 (0.055)	0.166 (0.034)	
Panel C: Arrival Rates				
Log Arrival Rate	-3.674 (0.008)	3.799 (0.007)	2.138 (0.007)	1.077 (0.007)
Past Log Bid Deviation	-1.112 (0.004)			
Panel D: Share of Job Posts by Buyer Type				
Inexperienced Buyer Posts	0.709	0.060	0.231	
Experienced Buyer Posts	0.261	0.276	0.463	

Notes: Estimates of the demand parameters, types probabilities, and arrival rates as described in the text. The parameters in columns 2-4 are additive interaction terms relative to the baseline of a Type 1, inexperienced buyer. The delta method is used to calculate standard errors for the type shares, as these parameters are estimated using a logit transformation to keep the probabilities in the unit interval. See the Appendix for detailed buyer-type-specific parameter estimates for the other X_{oj} characteristics.

Table 5: Buyer Demand Elasticities and Surplus

	Type 1	Type 2	Type 3	Overall Avg.
Panel A: Elasticities				
Applicant Choice Elast. (Inexp.)	-4.466 (1.282)	-4.483 (1.288)	-4.473 (1.284)	-4.469 (1.283)
Applicant Choice Elast. (Exper.)	-5.132 (1.702)	-5.153 (1.709)	-5.137 (1.704)	-5.140 (1.705)
Vacancy Fill Elast. (Inexp.)	-3.516 (0.768)	-3.778 (0.804)	-3.634 (0.780)	-3.536 (0.770)
Vacancy Fill Elast. (Exper.)	-3.446 (0.843)	-3.831 (0.893)	-3.554 (0.862)	-3.469 (0.846)
Panel B: Buyer Surplus				
Surplus-per-Hour (Inexp.)	0.538 (0.198)	0.222 (0.088)	0.396 (0.136)	0.506 (0.185)
Surplus-per-Hour (Exper.)	0.895 (0.391)	0.482 (0.235)	0.782 (0.352)	0.865 (0.380)
Value of Market (Inexper.)	142.2 (56.7)	2020.0 (656.3)	820.9 (310.8)	412.5 (148.3)
Value of Market (Exper.)	807.0 (332.8)	19411.4 (7271.1)	5979.8 (2290.5)	8337.6 (3024.1)
Panel C: Summary Statistics Weighted by Posterior Type				
Pr. Fill Job (Inexp.)	0.168 (0.003)	0.077 (0.005)	0.130 (0.004)	0.159 (0.002)
Pr. Fill Job (Exper.)	0.315 (0.006)	0.194 (0.008)	0.288 (0.006)	0.307 (0.005)
Share Technical Jobs (Inexp.)	0.583 (0.002)	0.424 (0.017)	0.489 (0.004)	0.552 (0.002)
Share Technical Jobs (Exper.)	0.553 (0.004)	0.400 (0.011)	0.468 (0.005)	0.471 (0.004)

Notes: Panel A displays estimated elasticities for hiring an individual applicants and elasticities for filling a vacancy (hiring any applicant). Panel B presents estimates of surplus per hour worked that accrues to buyers and total buyer value for the market. Total value for the market is computed as described in the text and accounts for the arrival rate of future openings. Estimates of overall surplus assume jobs last 70 (77) hours for inexperienced (experienced) buyers. Panel C displays how job fill rates and opening characteristics (technical or not) vary with buyer type. Standard errors come from 20 block-bootstrap iterations (drawing buyers with replacement).

Table 6: Surplus for Applicants and Hired Providers

Buyer Type:	Type 1	Type 2	Type 3	Overall Avg.
Panel A: Expected Surplus-Per-Hour for Applicants				
Applying to Inexp. Buyers	0.015 (0.005)	0.006 (0.002)	0.011 (0.003)	0.014 (0.005)
Applying to Exp. Buyers	0.022 (0.009)	0.009 (0.004)	0.016 (0.007)	0.018 (0.007)
Applicants 1-10	0.031 (0.007)	0.018 (0.007)	0.029 (0.009)	0.031 (0.008)
Applicants 11-20	0.016 (0.004)	0.010 (0.004)	0.015 (0.005)	0.016 (0.004)
Applicants 21-30	0.010 (0.002)	0.006 (0.002)	0.009 (0.003)	0.010 (0.003)
Applicants 31-40	0.008 (0.002)	0.005 (0.002)	0.007 (0.002)	0.007 (0.002)
Panel B: Applicant Costs and Markups				
Worker Costs (Inexp. Buyers)	7.355 (0.701)	6.263 (0.640)	6.683 (0.637)	7.116 (0.679)
Worker Costs (Exper. Buyers)	7.396 (0.709)	5.847 (0.551)	6.472 (0.607)	6.494 (0.616)
Markups (Inexp. Buyers)	0.289 (0.140)	0.287 (0.139)	0.288 (0.139)	0.288 (0.139)
Markups (Exper. Buyers)	0.242 (0.142)	0.241 (0.141)	0.242 (0.142)	0.242 (0.142)
Panel C: Surplus-Per-Hour for Hired Applicants				
Hired by Inexper. Buyers	2.274 (0.754)	1.663 (0.565)	1.914 (0.643)	2.176 (0.725)
Hired by Exper. Buyers	1.962 (0.780)	1.456 (0.612)	1.720 (0.680)	1.735 (0.694)

Notes: Panel A displays estimates of applicant expected surplus per hour, computed as applicants markups over costs multiplied by the individual hiring probability. The bottom rows show how this surplus varies over the based on the applicant arrival order. Panel B displays estimated applicant costs and the markups over costs that enter the surplus calculation. Panel C displays estimates of surplus per hour for hired applicants. Standard errors come from 20 block-bootstrap iterations (drawing buyers with replacement).


Table 7: Counterfactual Changes in Hiring Rates, Postings, and Surplus

	10% tax		\$7 Floor		
	No δ_3		No δ_3	Lab-Lab Sub.	
Panel A: Inexperienced Buyers					
Change in log Bids to Buyer	0.086 (0.000)	0.086 (0.000)	0.340 (0.004)	0.340 (0.004)	0.494 (0.035)
Static Pct Change in Hiring Rates	-0.267 (0.068)	-0.267 (0.068)	-0.311 (0.020)	-0.311 (0.020)	-0.298 (0.011)
Static Pct Change in Buyer Surplus	-0.240 (0.073)	-0.240 (0.073)	-0.100 (0.007)	-0.100 (0.007)	-0.057 (0.064)
Pct Change in Jobs Posted	0.005 (0.020)	0.079 (0.018)	-0.005 (0.013)	0.095 (0.007)	0.088 (0.006)
Pct Change in P.V. of Buyer surplus	-0.685 (0.048)	-0.464 (0.073)	-0.692 (0.022)	-0.409 (0.019)	-0.349 (0.053)
Panel B: Experienced Buyers					
Change in log Bids to Buyer	0.086 (0.000)	0.086 (0.000)	0.468 (0.008)	0.468 (0.008)	0.651 (0.012)
Static Pct Change in Hiring Rates	-0.255 (0.077)	-0.255 (0.077)	-0.354 (0.023)	-0.354 (0.023)	-0.339 (0.014)
Static Pct Change in Buyer Surplus	-0.255 (0.088)	-0.255 (0.088)	-0.119 (0.008)	-0.119 (0.008)	-0.066 (0.074)
Pct Change in Jobs Posted	-0.491 (0.037)	-0.109 (0.029)	-0.564 (0.033)	-0.131 (0.010)	-0.122 (0.008)
Pct Change in P.V. of Buyer surplus	-0.568 (0.068)	-0.266 (0.090)	-0.545 (0.036)	-0.128 (0.009)	-0.058 (0.107)
Panel C: Providers					
Change in log Bids to Provider	-0.001 (0.000)	-0.001 (0.000)	0.411 (0.005)	0.411 (0.005)	0.580 (0.015)
Static Pct Change in Provider Surplus	-0.278 (0.050)	-0.278 (0.050)	-0.045 (0.015)	-0.045 (0.015)	-0.042 (0.017)
Pct Change in P.V. of Provider surplus	-0.515 (0.040)	-0.301 (0.051)	-0.411 (0.042)	-0.079 (0.015)	-0.076 (0.015)

Notes: Estimates of changes in log bids and changes in surplus (by buyer experience) under different counterfactuals. The static percent changes in hiring rates and surplus are computed holding fixed the number of job openings. Surplus calculations for buyers come from equation 13. Present value calculations are described in the text. The percent change in the number of jobs is computed based on opening arrival rates simulating forward wage bids and buyers endogenous experience. Static provider surplus is the to-provider hourly wage less platform fees multiplied by hiring probabilities. The present value of provider surplus is calculated as to-provider hourly wages \times average hours \times hiring probabilities \times the number of jobs posted monthly and discounting future surplus to the start of the sample. The last column with Labor-Labor substitution sets $\delta_3 = 0$. Standard errors come from 20 block-bootstrap iterations (drawing buyers with replacement).

Appendix Figures and Tables

Figure A.1: Bid Decile of Hired Worker



This figure displays data for 26,067 job postings with at least 10 applications where the buyer ultimately hired an applicant. For each posting, the hourly wage bids received are sorted into deciles. The figure shows the decile that contains the bid of the hired provider. In 21% of the postings, the hired provider submitted a bid in the lowest decile of bids. However, in the remaining 79%, the buyer selects a provider whose bid is in a higher decile and 39% of hires are of providers who bid above the median bid received.

Table A.1 presents some of the estimates in $\beta_{k\chi}$ in equation (3) that are not reported in Table (4).

	Type 1 (Baseline)	Type 2	Type 3	Experienced
		Additive Interactions		
Log Rate Last Hire	1.516 (0.329)	-0.042 (0.021)	-0.029 (0.018)	0.231 (0.400)
No Prior Jobs	-0.076 (0.657)	0.539 (1.196)	0.332 (0.936)	-0.654 (0.881)
Number of Prior Jobs	0.025 (0.003)	-0.001 (0.001)	-0.001 (0.001)	-0.005 (0.003)
Prior Hires but no Feedback	0.914 (0.654)	0.576 (1.188)	0.171 (0.932)	-0.764 (0.877)
Feedback Score	-1.628 (0.625)	0.489 (1.136)	0.342 (0.886)	-0.732 (0.841)
Feedback Squared	0.692 (0.186)	-0.141 (0.342)	-0.144 (0.266)	0.190 (0.251)
Feedback Cubed	-0.069 (0.017)	0.013 (0.033)	0.017 (0.025)	-0.015 (0.023)
BA + Degree	0.103 (0.022)	-0.085 (0.039)	-0.038 (0.030)	0.022 (0.029)
Good English	0.354 (0.069)	-0.031 (0.088)	-0.061 (0.073)	-0.271 (0.093)

Notes: Estimates of additional demand parameters from the choice mode.

Table A.2: Demand and Arrival Rate Parameter Estimates Including Provider Applications per Month Control

	Type 1 (Baseline)	Type 2	Type 3	Experienced Additive Interactions
Panel A: Buyer Types				
Type Share	0.805 (0.010)	0.027 (0.002)	0.169 (0.009)	
Panel B: Choice Model Parameters				
Log Bid	-4.437 (0.083)			-0.850 (0.115)
Constant	1.939 (0.371)	-1.081 (0.137)	-0.373 (0.340)	
Experienced Const. Shift	0.950 (0.742)	0.237 (0.058)	0.163 (0.075)	
Panel C: Arrival Rates				
Log Arrival Rate	-3.674 (0.031)	3.800 (0.038)	2.137 (0.024)	1.076 (0.044)
Past Log Bid Deviation	-1.093 (0.155)			
Panel D: Share of Job Posts by Buyer Type				
Inexperienced Buyer Posts	0.708	0.061	0.231	
Experienced Buyer Posts	0.260	0.278	0.462	

Notes: This table replicates the Table A.2 estimates of the demand parameters, types probabilities, and arrival rates, with the inclusion of provider applications per month in the first stage regression and the X_{oj} characteristics.

Table A.3: Buyer Demand Elasticities and Surplus Including Provider Applications per Month Control

	Type 1	Type 2	Type 3	Overall Avg.
Panel A: Elasticities				
Applicant Choice Elast. (Inexp.)	-4.407	-4.423	-4.414	-4.410
Applicant Choice Elast. (Exper.)	-5.232	-5.253	-5.237	-5.240
Vacancy Fill Elast. (Inexp.)	-3.483	-3.745	-3.600	-3.503
Vacancy Fill Elast. (Exper.)	-3.497	-3.885	-3.604	-3.520
Panel B: Buyer Surplus				
Surplus-per-Hour (Inexp.)	0.546	0.224	0.404	0.514
Surplus-per-Hour (Exper.)	0.878	0.472	0.767	0.849
Value of Market (Inexper.)	139.1	1960.4	802.1	402.8
Value of Market (Exper.)	787.5	18935.5	5828.6	8155.5
Panel C: Summary Statistics Weighted by Posterior Type				
Pr. Fill Job (Inexp.)	0.168	0.076	0.131	0.159
Pr. Fill Job (Exper.)	0.315	0.194	0.288	0.308
Share Technical Jobs (Inexp.)	0.583	0.425	0.489	0.552
Share Technical Jobs (Exper.)	0.553	0.399	0.468	0.471

Notes: This table replicates Table 5 while including provider applications per month in the first stage regression and the X_{oj} characteristics.

Table A.4: Surplus for Applicants and Hired Providers Including Provider Applications per Month Control

Buyer Type:	Type 1	Type 2	Type 3	Overall Avg.
Panel A: Expected Surplus-Per-Hour for Applicants				
Applying to Inexp. Buyers	0.016	0.006	0.011	0.014
Applying to Exp. Buyers	0.022	0.009	0.016	0.017
Applicants 1-10	0.031	0.018	0.029	0.031
Applicants 11-20	0.016	0.010	0.015	0.016
Applicants 21-30	0.010	0.006	0.009	0.010
Applicants 31-40	0.007	0.005	0.007	0.007
Panel B: Applicant Costs and Markups				
Worker Costs (Inexp. Buyers)	7.327	6.245	6.655	7.089
Worker Costs (Exper. Buyers)	7.434	5.872	6.505	6.523
Markups (Inexp. Buyers)	0.294	0.292	0.293	0.293
Markups (Exper. Buyers)	0.237	0.235	0.236	0.236
Panel C: Surplus-Per-Hour for Hired Applicants				
Hired by Inexper. Buyers	2.306	1.684	1.941	2.206
Hired by Exper. Buyers	1.927	1.426	1.692	1.704

Notes: This table replicates Table 6 while including provider applications per month in the first stage regression and the X_{oj} characteristics. the X_{oj} characteristics.