

Who Benefits from Online Gig Economy Platforms?

Christopher T. Stanton Catherine Thomas
Harvard, NBER, and CEPR LSE, CEP, and CEPR

May 2024*

Abstract

Online labor platforms for short-term, remote work have many more job seekers than available jobs. Despite their relative abundance, workers capture a substantial share of the surplus from transactions. We draw this conclusion from demand estimates that imply workers' wages include significant markups over costs and a survey that validates our surplus estimates. Demand-side search frictions and differentiated characteristics prevent workers from competing away supply-side surplus. Finally, we show that applying traditional employment regulations to online gig economy platforms would lower job posting and hiring rates, reducing aggregate surplus for all market participants, including workers.

*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. We thank Shira Aronson and Cameron Greene for excellent research assistance.

1 Introduction

An estimated 160 million global workers are registered on online labor market platforms (Kässi and Lehdonvirta, 2018; Kässi, Lehdonvirta, and Stephany, 2021). Similar to other gig settings, there are many more workers than jobs at any point in time (Autor, 2001; Prassl, 2018; Fisher, 2022). A broad literature now seeks to understand the potential costs and benefits of alternative work arrangements (Katz and Krueger, 2019; Collins, Garin, Jackson, Koustas, and Payne, 2020; Mas and Pallais, 2020), but we know little about the surplus online labor markets create, or who captures the benefits. Despite this, concerns about excessive competition have prompted policy proposals to change employment classifications for freelancers, including those using online labor markets. In this paper, we contribute to this debate by using data from a large online market to quantify market surplus and understand how it is distributed among participants.

Our analysis addresses three main questions: (1) How much surplus accrues to workers and to buyers in online labor markets? (2) If workers receive surplus despite their relative abundance, what prevents competition from pushing workers to their reservation wages? (3) Would traditional labor market regulation increase worker surplus in these markets?

We start by briefly summarizing our findings: First, total surplus—roughly analogous to the sum of producer and consumer surplus—is approximately \$4.28 per hour worked. About 46% of this total goes to workers, a substantial share when compared to traditional employment arrangements (Card, Cardoso, Heining, and Kline, 2018).

Second, both worker differentiation and limited buyer search allow workers to earn positive surplus even in the face of intense competition. Buyers view workers as imperfect substitutes, such that the impact of a standard deviation change in workers’ characteristics on labor demand is equivalent to a 43% change in wages. However, worker differentiation alone does not ensure that workers capture meaningful surplus. A second force—buyers’ limited attention while hiring—also contributes. The average buyer considers only 18 workers before making a hiring decision, and early applicants are the most likely to be considered, limiting head-to-head competition between similar workers.

Third, policy counterfactuals demonstrate that employment regulations that raise buyers’ labor costs lead to lower surplus for both sides of the market. Job posting and hiring rate reductions offset any gains that could be reallocated to workers.

To arrive at these conclusions, we set up a structural model of supply and demand in an online labor market for short-term tasks. As is typical in these markets, wage setting is decentralized and workers submit job-specific wage bids. A survey we conducted in 2023 shows that workers tailor

bids to job openings. In the model, we allow workers to choose their bids strategically in response to buyers' hiring elasticities. We estimate the model with a sample of 170,000 competitive job openings spanning January 2008 to June 2010 from the administrative data of a large platform.

The model contains three key components. First, a buyer's demand on a job post—whom to hire, if anyone—is a function of applicants' characteristics, wage bids, and buyer attributes. These attributes are the buyer's privately known value of the demand parameters (buyer type), her past hiring experience, and her search costs, which determine the number of applicants she considers prior to hiring. Second, workers' hourly wage bids are based on markups over their costs. Markups are determined by the perceived demand elasticity they face on a job opening. Third, buyer job posting is determined by an arrival process that depends on her type and past hiring experience.

We use two instruments for workers' wage bids to recover credible estimates of demand elasticities. The first instrument exploits exchange rate fluctuations. For workers with currencies that float relative to the US dollar, a shifting exchange rate varies the benefit of working online for US dollars relative to local alternatives. This instrument generates bid variation between workers from different countries. The second instrument exploits supply-side competition for similar jobs posted by other buyers. Workers on the platform observe signals of competition, and workers submit lower bids when there are more applicants to similar jobs. This instrument shifts all workers' bids relative to the outside option of not hiring on the platform.

Survey responses confirm that workers pay attention to both local exchange rates and the number of prior applicants when setting bids. Several tests suggest that exclusion restriction violations are limited. For example, applicant resume characteristics vary little with respect to the instruments, suggesting the instruments primarily affect prices rather than the distribution of worker quality. The instruments are also unrelated to a proxy for worker application effort, and all our estimates are similar when we control for worker application quantity.

Our estimates show there are gains from trade for both sides of the market. We estimate that buyers gain an average of \$2.30 of surplus per hour when they hire. Buyer surplus upon hiring is defined as the dollar amount that would induce indifference between hiring the chosen worker at the wage paid and the outside option of not hiring on the platform. We estimate worker surplus assuming that wage bids are optimal responses to buyers' demand elasticities. Hired workers' have surplus of about \$1.98 per hour. Workers' markups over their costs average 27%. Over our 30 months of data—from an era when online labor markets were smaller than today—total surplus to hired workers was about \$4.56 million. The total surplus to buyers was about \$5.77 million.¹ These

¹The platform collects 10% of all revenues, which, at average wage bids, yields platform revenues of \$1 per hour. We do not observe the platform's marginal costs of servicing transactions, and platform surplus is excluded from our surplus calculations.

figures represent net gains rather than total revenues or wage bills because the surplus is relative to each party’s outside option or reservation utility/wage.

Moving beyond hires, aggregate worker surplus remains positive, at about \$2.06 million dollars, after netting out an average cost of each application of \$0.56. We estimate that buyers’ search costs average about \$1.21 per applicant considered, which we infer from the model as a function of how expected surplus changes when we alter each buyer’s consideration set. Subtracting buyer search costs leaves them with \$1.98 million in net surplus.

We validate the model in two ways. First, our model is primarily about the bid markups that come from the demand-side elasticities. We can assess model fit by examining whether bidding variation maps to model-predicted markup variation. One clear prediction relates to how bids change with applicant order: A later applicant knows he will face more competition in the event that a buyer considers him because buyers evaluate applications in the order they are received. The model, hence, prescribes that optimal markups fall with applicant order. In a regression of log bids on applicant order dummies and worker-by-week and job category fixed effects, we find that a worker submits bids that are about 3.5% lower as the 60th job applicant compared to when he is among the first 10 applicants. Model-predicted markups fall by about 1.8% between the first 10 applicants and the 60th. Second, in our 2023 survey, reported markups range from 23% to 38%, depending on the markup construct described in the elicitation. These estimates are broadly consistent with the 27% average markup from the model.

Finally, we use the estimates to assess whether traditional employment regulations can redistribute surplus towards the supply side of the market. We study the counterfactual introduction of an additional 10% tax paid by buyers when hiring workers on the platform. This policy illustrates how payroll taxes, such as the employer portion of Federal Insurance Contributions Act (FICA) contributions for W2 employees, would impact the market. When participants do not receive any benefit from the tax, there is a large reduction in both buyer and worker surplus. Hiring rates on posted jobs fall, with no increase in net wages to workers, and job posts fall significantly. Rebating the tax revenue to workers in a lump sum ameliorates the workers’ surplus loss, but the overall change remains negative because of the smaller market size. ²

Our work contributes to the growing literature on gig economy work arrangements. We use microdata to underscore how surplus varies in non-standard work arrangements, complementing

²We do not find evidence that surplus in online labor markets arises because buyers of online labor have substituted away from regulated offline employment. In Appendix D, difference-in-differences estimates show there is no change in online job posts or hiring in U.S. states that raised the local minimum wage during our sample period compared to control states. In addition, [Horton, Johari, and Kircher \(2021\)](#) reports that only 15% of surveyed platform buyers would have hired locally if the platform were not available. This suggests a limited substitution elasticity between online and offline work, in which case any lost online jobs would be net surplus destroying.

research on workers’ valuations for these alternatives (Mas and Pallais, 2017). We add to a literature that examines how platform design can impact information frictions (Pallais, 2014; Moreno and Terwiesch, 2014; Ghani, Kerr, and Stanton, 2014; Agrawal, Horton, Lacetera, and Lyons, 2015; Horton, Kerr, and Stanton, 2017; Kässi and Lehdonvirta, 2018; Cullen and Farronato, 2021) by showing that these frictions can limit competition and shift surplus.³

There are two potential external validity limitations in interpreting our results. First, not all online labor markets use hourly wages as the contract form. The most common alternative is fixed price procurement auctions where buyers set a maximum budget and workers bid a price to deliver work. To offer some evidence that our conclusions apply to these types of jobs, we surveyed workers about their perceived surplus on fixed price jobs. Survey responses show worker surplus is positive in these arrangements.⁴ Second, the data we use to estimate the model come from an era where online labor markets were relatively new. Improvements in platforms’ technology or matching procedures may have impacted surplus, particularly if search frictions have declined in importance. However, our survey of current users validates our markup estimates. Along with comparisons to other papers’ descriptions of how platforms have evolved, the survey responses provide confidence that the structural assumptions of our model capture meaningful and persistent features of gig economy contracting.

Our results are relevant for the ongoing debate about how to regulate the gig economy. Proposals like the PRO Act, which passed United States House of Representatives before dying in the Senate, would potentially apply traditional labor regulations to this part of the gig economy. Prior work has shown how labor demand responds on a range of margins in response to regulation (Clemens, Kahn, and Meer, 2018).⁵ We show that regulatory changes to online labor markets would reduce static surplus, but the main negative welfare impact would be due to missing jobs.

The paper proceeds as follows. Section 2 describes our worker survey, introduces the administrative data, and provides summary statistics that guide our modeling choices. Section 3 presents the model and estimation strategy. Section 4 presents the main results and calculations of buyer and worker surplus. Section 5 contains the counterfactual analysis. Section 6 concludes.

³Related work focusing on ridesharing platforms has quantified demand and supply (Hall, Horton, and Knoepfle, 2021), analyzed surplus from surge pricing (Castillo, 2020), measured the value of flexibility and drivers’ support for regulation (Chen, Rossi, Chevalier, and Oehlsen, 2019; Katsnelson and Oberholzer-Gee, 2021), and assessed offline spillover benefits (Gorback, 2020).

⁴Online labor markets for micro-tasks, like Mechanical Turk, use a contract form that is closer to the fixed price contract, but there are some meaningful differences. For further details, see Benson, Sojourner, and Umyarov (2020) and Dube, Jacobs, Naidu, and Suri (2020).

⁵The closest work answering a regulatory question about online labor markets is Horton’s (2022) study of a \$3 wage floor. His experiment was carefully designed to measure hiring responses holding fixed pre-existing job openings and applicants. Our work uses data on buyers’ decisions over time to show how their experiences affect future job postings. This allows us to estimate the dynamic response margins to policy counterfactuals. Horton (2022) also shows that buyers exposed to the minimum wage treatment subsequently posted fewer jobs.

2 The Setting, Data, and Descriptive Statistics

2.1 Empirical Setting

Our data come from an online platform where buyers contract with workers selling labor services. This platform, along with several leading others, facilitates search and matching, task management, and payments. All jobs are done remotely and work output is delivered electronically. To purchase online labor services, a buyer must first create an account on the platform. She then posts a job, which requires selecting the work category and its expected duration, gives the job a title, and describes the work to be done and the skills required. Once the posting is live, potential applicants learn about it by searching on the site or receiving emails about new jobs. Buyers also have the option of searching worker profiles directly and inviting applications.

Job postings tend to be for short-term, spot transactions. There are two primary types of contracts on the platform. Hourly contracts, where workers are paid hourly wages for time billed, are the most common job type, accounting for over 80% of workers' earnings during our main data sample. The typical hourly job requires 17 hours of work per week and 61% of postings last less than three months. On average, hired workers bill about 72 total hours. Buyers may alternatively choose fixed price contracts, where workers are paid a set amount only upon successful completion of the project. Fixed price jobs tend to be shorter in duration than hourly jobs. Our main analysis focuses on hourly contracts because fixed price jobs vary in length and complexity in unobserved ways. We later supplement our analysis of hourly contracts with survey evidence on workers' perceptions of the surplus they earn on fixed price contracts.

As shown in the examples in Figure 1, workers observe information about the buyer and the job before applying, including the date of posting and the number of workers that have already applied. Interested workers submit applications, and each applicant proposes (bids) the hourly wage at which they are willing to work. Workers are located worldwide but all bids are denominated in US dollars. Workers' profiles contain information about their skills, education, location, prior offline work experience, experience on the platform, and feedback scores from past work. Information from past jobs has been shown elsewhere to influence workers' attractiveness and future career prospects on the platform (Ghani, Kerr, and Stanton, 2014; Pallais, 2014; Agrawal, Lacetera, and Lyons, 2016; Stanton and Thomas, 2016; Barach and Horton, 2021).⁶

Before hiring, buyers may request interviews with any candidate, although an interview is not required. After hiring, buyers monitor work via software that the platform provides. The platform

⁶Figure A.1 presents some initial evidence that workers are differentiated, showing that on openings with a hire, buyers choose applicants in the lowest bid decile only around 20% of the time.

manages all payments for completed work and guarantees workers are paid for hours billed, which means that payment risk is unrelated to buyer reputation or experience. When a contract ends, buyers and workers leave feedback for one another.

We have two data samples: responses from a detailed survey of 113 workers conducted in 2023, and a much larger historical administrative dataset of applications and hiring processes from 2008 to 2010. Many core platform features, and the institutions governing contracting, remained largely unchanged between 2008 and 2023, and are shared by other large platforms (Kässi and Lehdonvirta, 2018). We later discuss some platform features that have evolved from the time of our administrative sample.

The survey was designed to supplement the main analyses using the administrative data and to assess the relevance of our findings in a different time period. In particular, the survey elicited information about workers' bidding decisions and perceptions of surplus on the platform. To the best of our knowledge, no other data source contains similar information. Workers' responses were matched to the information in their online profiles.

The administrative data were obtained from the platform and allow us to analyze buyer responses to worker bids and to follow buyer and worker careers on the platform. We observe every buyer's hiring choices among candidates for all their job postings. For each job posting, the data contain information about the entire applicant pool; which candidate(s), if any, are interviewed or hired; buyers' stated reasons for not hiring individual workers (available for a subset of openings where buyers opted to provide this information); and the feedback that the buyer and the hired worker leave for one another. These data cover the 30 months from January 2008 to June 2010, a period when the matching process between buyers and workers was decentralized and did not involve algorithmic recommendations. During this time, 67,566 potential buyers posted 170,556 hourly jobs and received more than 4.4 million applications from 192,628 unique workers.

Figure 2 documents a striking feature of the market: there are many more job seekers than openings in any given month. On average, around 16,600 unique workers applied to just over 5,600 hourly jobs per month, and only 1,100 workers were hired. The majority of applicants were not hired at all during a typical month, motivating concerns that worker competition would have driven down their surplus. Hires are scarce relative to job seekers in both technical and non-technical job categories, which make up approximately equal shares of all job posts.⁷

⁷The technical job categories are Web Development, Software Development, and Networking and Information Systems. The non-technical job categories, in order of frequency, are Administrative Support, Writing and Translation, Design and Multimedia, Sales and Marketing, Business Services, and Customer Service.

2.2 Survey Evidence

In Fall 2023, we posted jobs on the platform with the aim of surveying workers across a range of job categories and backgrounds who had previously been hired at least once. Appendix A contains details about the recruitment and survey instrument. The 113 workers who completed the survey came from 39 countries and had a median of six prior jobs, with a range from 1 to 144. When asked about their sources of income, 41% reported that “work through platforms is my exclusive source of income” and 35% reported “I also have a traditional job where I am an employee.” These figures are comparable to the responses in a 2012 survey of workers on a similar platform (Kuek, Paradi-Guilford, Fayomi, Imaizumi, Ipeirotis, Pina, and Singh, 2015), where 48% said that freelancing full time was their sole source of income and 38% had either a full time or part time job as well.⁸ This suggests that, on average, the way workers interact with these platforms now is much the same as it was in 2012.

We then asked survey respondents how often they varied their hourly wage bids during a typical week. The majority reported that they did so at least sometimes, with 45% of workers reporting that they always submitted different or tailored bids to different openings and a further 45% saying they sometimes did. Only 10% selected the option “I always submit the same bid,” which was the first option among those presented in the question.

The next set of questions probed why workers submit varying bids. Nearly all, 96% of workers, said the skills they would need to complete a job were somewhat or very important when deciding what bid to submit. When asked whether variation in their opportunity costs affected their bids, 89% said they considered other on-platform work opportunities and 70% said they considered off-platform worker opportunities when forming bids. To preview our identification strategies: 82% of respondents from countries that do not use the dollar as the local currency indicated that the value of the dollar relative to their local currency affects their bids, and 83% of all respondents said they varied their bids based on how many people applied to the job before them. Our model will allow workers’ bids to vary based on the extent of perceived competition for a position, as captured by the number of prior applicants, and shocks to the value of their local versus online opportunities.⁹

⁸In our data, 28% of respondents said “I also work as an independent contractor outside of platforms”. This response option doesn’t correspond well with prior surveys’ options. Other studies have attempted to capture online workers’ alternatives to platform work. For example, Gray and Suri (2019) document that online workers often spend time doing non-paid work, such as caring for dependents. They focus on microtask workers who are paid for immediate task completion, which is not a category of work on the platform we study. Kässi, Lehdonvirta, and Stephany (2021) estimate that microtask and contest-based platforms make up a relatively small share of the overall online labor gig economy, and workers on these platforms are outnumbered by more than six to one by those on online labor markets of the type we study.

⁹Gee (2019) finds that observing the number of prior applicants increases the probability a worker applies for a job on LinkedIn compared to when this information is hidden. Relative to Gee’s experiment, the number of applicants was always shown on this platform, allowing us to examine differences in the state of competition holding fixed that

2.3 Administrative Data

Table 1 presents descriptive statistics about the 67,556 buyers and 192,628 workers in the 2008 to 2010 administrative data sample. Buyers range from private individuals to employees or owners of small firms to the occasional individual who hires on behalf of a large enterprise. The first two columns show that most buyers are from the US, and 75% are located in G10 countries with high levels of per capita income. Only 6% of buyers come from India or the Philippines.¹⁰ Roughly half (56%) of buyers post their first job in a technical job category, a share that is similar for platforms today (Stephany, Kässi, Rani, and Lehdonvirta, 2021).

The average buyer posts 2.52 jobs during the sample, while the standard deviation is nearly twice the mean, at 4.50 job posts. The variability in buyer engagement with the platform, which we allow for in our model, potentially reflects different types of buyers or different, idiosyncratic experiences when engaging with workers.

To the extent that buyers are heterogeneous, job applicants may observe signals of buyer type, but these signals are likely to be relatively coarse. For example, workers do not observe buyer identities or their line of business, but they do observe the country where a buyer is located and past platform activity, as illustrated in Figure 1.

The third and fourth columns of Table 1 present descriptive statistics about workers. Consistent with other studies, 42% are in India or the Philippines, while 31% are located in high-income countries. The average worker applies for 23.15 jobs over the 30 months of data, and there is a very large standard deviation. Despite this high average level of platform engagement, only 8% of all workers are hired over the time period studied, and the mean number of hires per worker is 0.18.

Table 2 presents descriptive statistics about the 170,556 job postings that form our main sample.¹¹ Column 1 shows that posts receive an average of 26 applications, the majority of which

workers observe the number of prior applicants.

¹⁰89% of the transactions in the market span international borders, with the worker typically located in the lower-wage country. See 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 between different countries. This market has evolved similarly to other prominent platforms (Stanton and Thomas, 2020). Key features of the time period we study remain in later periods. For example, the distribution of job categories and the geographic concentration of buyers and workers has remained stable over the last 15 years (Kuek et al., 2015; Kässi, Lehdonvirta, and Stephany, 2021). Our data show around 8% of workers were hired at least once, which is in line with the 10% share of workers who had positive earnings on large online labor markets in 2013 (Kuek et al. (2015)) and the 9.6% documented by Kässi, Lehdonvirta, and Stephany (2021) for 2020. Since the time of our administrative data, there have been some changes in the way platforms curate buyers and workers to increase the share of filled postings (e.g. using algorithmic recommendations). Some platforms now allow buyers to state a worker quality preference or impose a minimum hourly wage. These changes may have increased market efficiency overall, but they likely reduce the number of workers a buyer considers, which would strengthen the market power of those that are considered.

¹¹The sample includes job postings for hourly-billed work that are posted publicly and receive at least three applications. This filtering rule selects competitive openings, preventing us from picking up jobs that go to a worker that the buyer may have targeted from an offline source or a pre-existing online relationship (2,250 jobs). Such non-competitive openings involve buyers making hiring choices with different information about workers (Kahn (2013)). Any job that was declared to have been posted by mistake was also excluded. We drop spam job postings, which are

arrive within 24 hours of posting. The mean hourly wage bid is \$10.90, with a standard deviation of \$4.97.

In the model, we allow for differences in both demand and job posting frequency by buyer experience. To motivate this choice, columns 2 and 3 split the sample by buyer experience. Inexperienced buyers’ posts receive slightly fewer applications and receive slightly higher wage bids. Despite some minor differences, other rows in the table show that there is substantial overlap in applicants’ characteristics between the two levels of buyer experience.

Inexperienced buyers’ probability of filling a job is 15% compared to 26% for experienced buyers. The last row of Table 2 shows that experienced buyers post jobs more frequently than buyers who have yet to hire on the platform. Due to their increased job posting frequency, experienced buyers are responsible for around half of all postings on the platform despite the fact that only 27% of buyers are observed to hire at all. If hiring rates and/or job posting rates increase with experience (rather than arising due to buyer selection), then policies that alter whether buyers gain experience will potentially affect the size of the market.

2.4 Stylized Facts Motivating the Model

In this section, we tie together motivating data features and stylized facts that guide how we model buyers’ hiring choices and worker responses.

Buyer Consideration: Buyers are likely to consider only a subset of applicants rather than the entire applicant pool. An advantage of our setting is that the data allow us to measure the size of a buyer’s consideration set based on the actions they take. In other settings, consideration sets are often unobserved, necessitating computationally expensive procedures to simulate them (Abaluck and Adams-Prassl, 2021). In our setting, we say the buyer interacts with an applicant if the buyer messages them, hires them, or selects a reason for not choosing them. We can observe buyer-applicant interactions for 81% of the openings in the data.

Interactions and Applicant Order: At the time of our sample, the default when presenting applicants to buyers was to display them in the order in which they submitted the application. Appendix Table A.1 shows that ordering based on application timing strongly explains the likelihood of interaction (messaging, hiring, or selecting a reason for not choosing the worker).¹² The baseline

defined as those where the buyer sends over 60 interview requests. In addition, applications that buyers themselves flag as obvious spam are dropped. When we do not apply these filters, we find larger surplus estimates for both buyers and workers.

¹²The data record the calendar date when a buyer messages a worker, often for the purpose of requesting an interview. We do not observe whether an interview actually occurred because it would usually take place via an off-platform messaging or conferencing system. However, we do see buyers sending messages after only a subset of applications have arrived. This evidence suggests buyers’ search is targeted towards the subset of applicants who apply early, with a higher probability of “meeting” these applicants (Wright, Kircher, Julien, and Guerrieri, 2021).

probability of an interaction is 17.5%. Every ten positions an applicant falls reduces the probability of an interaction by 0.56 percentage points, or 3.2%. Application order is far more important than wage bid ranking for determining whether a buyer interacts with an applicant. Figure A.2 shows that this bias toward early applicants means applicant 45 receives an interview or hire request with less than 5% probability, a substantial reduction from the 20% probability for the first few applicants.¹³

Inferring Consideration Sets: Interacting with an applicant or extracting information from their profile appears costly for buyers because they do not interact with everyone. Part of their search costs likely entail the time costs of any hiring delay. Based on the strong ordering of the interaction measures in the data, we maintain that a consideration set consists of all applicants who applied prior to the last applicant the buyer interacts with explicitly. This concept is consistent with recent empirical work on costly browsing behavior in online settings.¹⁴ Figure 3 shows the distribution of consideration set sizes observed in the data.

Worker Responses: In the administrative data, 74% of worker-months and 62% of worker-week-job category cells have bid variation when workers submit multiple bids. For the median worker-week-job category cell, the maximum bid is 18.2 log points greater than the lowest bid. Workers who lack bid variation submit bids that are 5% higher on average than those who vary their bids. The survey suggests one source of bid variation over this narrow time window is that applicants set their bids taking into account how many workers have already applied for a job. In the model, we allow potential bidders to tailor their bids to the number of prior applicants because the probability they are considered by the buyer depends on their application order.

3 A Model of Online Labor Services Demand and Supply

This section presents our model of demand and supply for remote, online gig jobs. Buyers post hourly jobs with a frequency that depends on their type and past market experience. Workers observe open job posts and submit applications including tailored wage bids. Buyers choose how many candidates to consider for each job, inspect applications from the candidates they consider, and make a discrete choice about which, if any, applicant to hire.

¹³Since the time of our sample, many platforms have attempted to leverage algorithmic recommendations to reduce matching frictions. These recommendations typically supplement the pool of organic applicants that arrive according to the same processes that we document here (Horton, 2017).

¹⁴Consumers in other online markets are more likely to purchase products displayed first in their search results, explaining why advertising auction prices decline with slot position (Varian, 2007). Other papers that observe digital browsing behavior find that display order determines whether or not a product is considered. For example, Dinerstein, Einav, Levin, and Sundaresan (2018) assume a buyer considers all the listings seen up through the last one observed to be browsed in a search query on e-Bay. In a recent working paper, Yu (2023) studies ranked advertisements on Amazon, where the data show buyers search in order of rank.

We describe market participants’ decisions, starting with a buyer’s hiring decision on a posted job. We then focus on workers’ bidding strategies, and, finally, specify the process for how buyers post jobs. As we proceed, we set out how market participants may be heterogeneous with respect to preferences or their alternatives to online contracting. We then discuss our identification strategy and model estimation. Finally, we describe how the model estimates are used to calculate surplus for each side of the market. Further details about the model setup and estimation approach are given in Appendix B.

3.1 Setup and Timing

A job posting, or opening, denoted o , contains a work description and some information about the buyer, including their past experience hiring on the platform, which we denote χ . We assume that each buyer is one of $k \in K$, types, where type is not directly observed in a job posting but may be partially inferred from the buyer’s past experience or the job details.

Potential applicants observe the job posting some time after it opens and know that buyers will potentially hire a single candidate for the position. The platform automatically displays how many workers have already applied to the job (Figure 1). When applying, applicant j selects an hourly wage bid, w_{oj} , while taking into account her costs, c_{oj} .

The buyer considers a subset of applicants to the job at some time after the first applicant has applied. We denote the total set of applicants as \mathcal{J}_o and the subset of considered applicants as J_o . If she considers applicant j , the buyer observes w_{oj} , her work history and resume characteristics, denoted X_j , and an idiosyncratic preference shock, ε_{oj} . The buyer hires one of the considered applicants or chooses not to hire on the platform for the job. The hired worker completes the work and payments are made through the platform. Buyers then go on to post additional jobs, and the job arrival process can depend on past market conditions the buyer has encountered.

To accommodate buyers choosing from a subset of the applicant pool, we draw from the marketing and industrial organization literature on consideration sets (see the survey in [Honka, Hortacısu, and Wildenbeest \(2019\)](#)). Our general model does not specify the exact search process governing consideration set formation. We simply model buyer choices given an observed consideration set. However, when we estimate buyers’ search costs, we implicitly assume a fixed sample size search process where each buyer chooses the number of applicants to consider ex-ante.

The distribution of consideration set size by buyer type and experience is known to all market participants but applicants do not observe an individual buyer’s exact consideration set size or their type. We allow each applicant’s market power to depend on their application order and the

probability they will be considered, where workers may have signals about buyer type that provide inferences about consideration probability and buyer demand.

We make two more assumptions about the nature of bids and the availability of applicants to fit our setting. First, workers’ prices (wage bids) are assumed to be take-it-or-leave-it offers.¹⁵ Second, although workers apply for multiple jobs over a short time period (and the relative abundance of alternative jobs may affect their optimal bids), we assume that applicants are available when buyers are making a choice.¹⁶

3.2 Demand: Hiring Decisions on Posted Jobs

A buyer of type k with past online hiring experience χ considers a subset of all applicants to opening o , $J_o \subseteq \mathcal{J}_o$. She hires the considered applicant that yields the highest indirect utility per unit of wage, which is $\frac{\exp(X_j \beta_{k\chi} + \varepsilon_{oj})}{(w_{oj})^{\alpha_k}}$ for each $j \in J_o$, or she chooses the outside option of not hiring for the opening. X_j is a vector of worker j ’s observable characteristics that also contains a constant. $\beta_{k\chi}$ is a vector of buyer type- and experience-specific preferences for worker characteristics and for on-platform hiring. The term α_k measures the buyer’s type-specific wage disutility. The term ε_{oj} is an idiosyncratic utility shock. Hiring requires that the preferred considered applicant yields greater indirect utility than not hiring on the platform, denoted option 0.

After taking logarithms, the buyer’s choice, given J_o , can be expressed as an inequality such that applicant j is hired when

$$X_j \beta_{k\chi} - \alpha_k \log(w_{oj}) + \varepsilon_{oj} \geq X_l \beta_{k\chi} - \alpha_k \log(w_{ol}) + \varepsilon_{ol} \quad (1)$$

for all $l \in \{J_o, 0\}$. The right-hand-side of inequality (1) consists only of ε_{o0} when comparing applicant j to the outside option of not hiring. We assume ε_{oj} is a location- and scale-normalized type-1 extreme value distributed idiosyncratic shock, so the probability applicant $j \in J_o$ is hired takes the conditional logit form

$$p(j|J_o, k, \chi) = \frac{\exp(X_j \beta_{k\chi} - \alpha_k \log(w_{oj}))}{1 + \sum_{j \in J_o} \exp(X_j \beta_{k\chi} - \alpha_k \log(w_{oj}))}. \quad (2)$$

We now characterize the distribution of consideration sets J_o , which enters the probability that j is hired for job o accounting for all possible consideration sets. We let j index the application

¹⁵In preparing the data on the company’s servers, we investigated the extent to which wages change between first bids and contracted rates after hiring. We found very little change from the initial offer to the wage paid when beginning to work, but we did not download the raw data behind this analysis.

¹⁶Workers may be selective when choosing whether or not to apply, which is a possibility we address when discussing our identification strategy. Applying to multiple jobs also means they may receive several job offers within a short time interval, but only 0.6% of workers have two or more requests to interview or start a job on any one day. There are also only very few cases where the buyer or worker reports a scheduling conflict or lack of availability as a reason a worker was not hired. When buyers or workers explicitly report a realized conflict or lack of availability, we exclude the application from the buyer’s choice set.

order so that applicant 1 is the first arriving applicant and applicant $|\mathcal{J}_o|$ is the last applicant, where $|\cdot|$ denotes the number of elements in a set. Because buyers encounter applicants in order, applicant j is in the buyer's consideration set J_o only if this set is sufficiently large, that is, only if $j \leq |J_o| \leq |\mathcal{J}_o|$. Let $\eta_{k\chi}(|J_o|)$ denote the probability that the set of considered applicants is exactly $|J_o|$. The probability that applicant j is considered is then $\sum_{|J_o|=j}^{|\mathcal{J}_o|} \eta_{k\chi}(|J_o|)$.

The probability that applicant j is hired for job o by a type k buyer with experience χ is the consideration-weighted sum of the conditional hiring probabilities given in equation (2)

$$p(j|k, \chi) = \sum_{|J_o|=j}^{|\mathcal{J}_o|} \eta_{k\chi}(|J_o|) p(j|J_o, k, \chi). \quad (3)$$

When taking this to the data, we assume an exponential distribution for consideration set sizes, where the exact probability for a set of size $|J_o|$ is $\eta_{k\chi}(|J_o|) = \exp(-\lambda_{k\chi o}^{\text{CONSIDER}} |J_o|) - \exp(-\lambda_{k\chi o}^{\text{CONSIDER}} (|J_o| + 1))$. We let the parameter $\lambda_{k\chi o}^{\text{CONSIDER}}$ be a function of buyer type, experience, and whether the job is a technical or non-technical posting.

3.3 Supply: Workers' Optimal Wage Bids

If worker j is hired for job o , the buyer pays w_{oj} and the worker receives $\frac{w_{oj}}{(1+\tau)} = \exp(\log w_{oj} - \log(1 + \tau))$, which is her wage bid net of the ad-valorem platform fee, τ . When bidding, worker j takes into account the opportunity cost of her time if she is hired, c_{oj} , and her expectation of how the probability she is hired depends on w_{oj} . She chooses her bid to maximize

$$E(U_{oj}(w_{oj})) = \underbrace{E[\tilde{p}(j)]}_{\text{Pr(Hired)}} \times \underbrace{\exp(\log w_{oj} - \log(1 + \tau))}_{\text{Net Wage}} + (1 - E[\tilde{p}(j)]) \times \underbrace{c_{oj}}_{\text{Opp. Cost}}, \quad (4)$$

where $E[\tilde{p}(j)]$ is her best forecast of $p(j|k, \chi)$, the hiring probability in equation (3).¹⁷

The worker uses the information she has to form $E[\tilde{p}(j)]$. She knows buyer experience, χ , her own characteristics, X_j , and also her applicant order, j . She has imperfect information about buyer type k and the size of the buyer's consideration set. Specifically, she is aware of the overall buyer type distribution, where we denote the population probability that a buyer is of type k by ρ_k . She is also aware of the distribution of consideration set size conditional on type k and experience χ . A worker may conjecture the buyer's type from the information in a job post. We assume the conjecture is a noisy signal over each type, which we denote by $\tilde{\rho}_k$ for all k .

Imagine that the applicant knows buyer type with certainty: Because she knows her applica-

¹⁷Note that although we label the worker's costs as her opportunity cost of work, the first order condition is the same if costs of effort or of supplying labor relative to leisure reduce the wage benefit. We use the opportunity cost label to reflect the source of variation in wage bids that we use for identification, but our later surplus calculations should be interpreted to be net of workers' overall costs after landing a job. We later consider the costs of applying to jobs after recovering costs on the job.

tion will be the j 'th received, the relevant hiring function from equation (3) would be $E[p(j)] = \sum_{|J_o|=j}^{|J_o|} \eta_{k\chi}(|J_o|)p(j|J_o, k, \chi)$, where the summation is over all possible consideration sets that include her. However, if she does not know the buyer's type with certainty, she uses the signals $\tilde{\rho}_k$ together with her knowledge of the population buyer-type distribution to form an optimal bid. We assume she places some weight, denoted by the parameter b , on the signal she receives, so that her expected hiring function is

$$E[\tilde{p}(j)] = b \sum_k \tilde{\rho}_k \sum_{|J_o|=j}^{|J_o|} \eta_{k\chi}(|J_o|)p(j|J_o, k, \chi) + (1-b) \sum_k \rho_k \sum_{|J_o|=j}^{|J_o|} \eta_{k\chi}(|J_o|)p(j|J_o, k, \chi). \quad (5)$$

Substituting equation (5) into equation (4) gives worker j 's expected payoff function. Her optimal wage bid from the first-order condition when maximizing this expected payoff is the fixed point of

$$w_{oj}^* = c_{oj} (1 + \tau) \left(1 + E[\tilde{p}(j)] / \frac{\partial E[\tilde{p}(j)]}{\partial \log w_{oj}} \right)^{-1}, \quad (6)$$

where $\left(1 + E[\tilde{p}(j)] / \frac{\partial E[\tilde{p}(j)]}{\partial \log w_{oj}} \right)^{-1}$ is the optimal markup over costs that she includes in the wage bid. Worker j 's costs can be recovered from bid data and estimates of semi-elasticities as

$$c_{oj} = \frac{w_{oj}}{(1 + \tau)} \left(1 + E[\tilde{p}(j)] / \frac{\partial E[\tilde{p}(j)]}{\partial \log w_{oj}} \right). \quad (7)$$

Net wages less costs equal the surplus earned by the worker per hour if she is hired.

3.4 Demand: Job Post Arrival Rates

We model buyers' job posting decisions using an exponential arrival process that is a function of buyer type, hiring experience, and the wage bids they have encountered on prior job posts.¹⁸ We allow the relationship between job posts and past bids to depend on job category-level average log wage bids at the time of a buyer's past job posts rather than the specific bids a buyer received. This guards against unobserved correlation between a buyer-specific factor, bids, and platform use. Specifically, we define $\widetilde{\log(w_{oj})} = \frac{1}{o-1} \sum_1^{o-1} (\overline{\log(w_{t,Category})} - t_{Category} - 1_{Category})$ as the average log wage bid at the job category level when jobs 1 through $o-1$ were posted, net of job category fixed effects and a time trend. A buyer who posts jobs when wages are higher than trend for a job category will, on average, be exposed to a higher value of $\widetilde{\log(w_{oj})}$.

The job posting arrival rate $\lambda_{k\chi}^{\text{ARRIVAL}}$ is strictly positive, so we model the process as

$$\log(\lambda_{k\chi}^{\text{ARRIVAL}}) = \delta_{1k} + \delta_2 1(\chi > 0) + \delta_3 \widetilde{\log(w_{oj})}, \quad (8)$$

where the vector δ_{1k} contains buyer type-specific constants, δ_2 shifts the mean arrival hazard for

¹⁸The dependence of posting frequency on experience and past observations of prices/wages is a reduced form way to capture, for example, how buyers resolve uncertainty about market features (Nosko and Tadelis, 2015).

experienced buyers, and δ_3 is experienced buyers’ posting frequency sensitivity to past wage bids.

3.5 Instruments and Identification of Hiring Probabilities

Job applications potentially contain unobserved characteristics that are correlated with workers’ wage bids or with the quality of the match between a worker and a buyer, motivating an instrumental variables strategy. The expression for wage bids in equation (6) suggests a strategy based on finding instruments that shock workers’ costs, specifically workers’ opportunity costs. As confirmed in the survey responses, a worker’s opportunity cost comes from working in her offline local labor market or on another online job. Our instruments are based on these different components of workers’ opportunity costs.

The first instrument provides bid variation between workers. It utilizes changes in the dollar-to-local-currency exchange rate for a worker’s country, varying the opportunity cost of offline work. We assume that workers’ offline wages are paid in their local currency, whereas they receive US dollars for their platform work. Frictions limiting exchange rate pass through to local wages mean offline opportunities adjust to exchange rates more slowly than online transactions. Applicants’ wage bids are predicted to increase when the local currency appreciates relative to the dollar (i.e. one dollar earned on the site provides fewer local currency units). To account for level differences across countries and heterogeneous secular trends, each country’s exchange rate series is first detrended and then normalized to have zero mean. Figure A.3 illustrates the time-series variation in mean residual log wage bids and exchange rates that underpin this instrument for India, the largest non-US worker source country. Bids and exchange rates move together.¹⁹

While the exchange rate instrument provides both relative price variation among applicants and price variation relative to the outside option for workers in countries with exchange rates that move relative to the US dollar, the second instrument shifts all bids relative to the outside option. It does so by capturing variation in the intensity of competition for online job opportunities that applicants face. We need this second instrument because applicants living in countries with dollar-pegged exchange rates do not have any cost variation from the exchange rate instrument when evaluated against the buyer’s outside option of not hiring.

We construct the competition instrument by exploiting the fact that workers can observe the count of applications to each open job posting at any point in time. We calculate it as the log of the average number of applications received in the first 24 hours after posting for all jobs in

¹⁹This potential source of variation was revealed in conversations with buyers. As part of the monitoring software the platform provides to buyers, they are able to view screenshots from hired workers’ screens. Buyers commented on the frequency with which exchange rate calculators appear in these screenshots.

the same job category and the same week, excluding jobs for the focal buyer.²⁰ This captures that workers observe the extent of competition across open jobs, and tailor their bids upward or downward depending on the extent of relative competition from other applicants.

We use [Petrin and Train’s \(2010\)](#) control function approach to account for unobservables in bids that are not explained by the two instruments, denoted Z_{oj} , or worker characteristics, X_j . This entails taking the residuals from the first stage regression

$$\log(w_{oj}) = Z_j\gamma_1 + X_j\gamma_2 + \nu_{oj}, \quad (9)$$

and including them as controls in the choice model. As described above, the matrix X_j includes all of the non-wage characteristics describing the opening and applicant in the hiring probability given in equation (2).

Table 3 shows that both instruments have a substantive and precisely estimated effect on applicants’ bids. Column 1 confirms that log bids increase when the local currency appreciates, and log bids decline when the log number of applicants to similar jobs increases.²¹

Some concerns about our instruments are that they change (i) the quantity or mix of workers applying or (ii) worker application quality or effort in unobservable ways conditional on applying. We investigate these potential exclusion restriction violations with three suggestive tests. All three indicate that exclusion restriction violations are unlikely to affect the conclusions from our estimates. We also display other diagnostics about what worker behavior changes with the instruments in Appendix B.1.

In the first test, we assess how the first stage parameters (and, later, how the choice model estimates) change when we control for the number of applications an individual worker submits in a given month. For example, an appreciation of the local exchange rate may lead some workers to apply to fewer jobs. Table 3, column 2 shows the results when including the number of applications by the worker as a control in equation (9). If a non-random subset of workers changes their applications in response to the instruments, including this control will alter the coefficient on the main instruments. The parameter estimate on the exchange rate instrument goes from 0.078 to 0.069 and is within the confidence interval of the original estimate. The competition instrument parameter goes from -0.067 to -0.063 and also remains within the original confidence interval. The relatively small magnitude of changes in the estimates suggests that our inferences are unlikely

²⁰To mitigate the potential for correlated shocks to violate the exclusion restriction, we residualize the raw instrument to net out time fixed effects and job category fixed effects. Time fixed effects remove market-wide shocks across all job categories. Job category fixed effects account for competition differences across types of work. The residuals after removing these fixed effects capture competition differences across job categories that would be known immediately prior to a candidate’s application, holding fixed the average competition in the category at the same time and the average competition across all job categories.

²¹The F-statistic clustering at the job posting level is 84.5. The F-statistic clustering at the worker level is 48.2.

to be driven by changes in the composition of applicants.²² We later show that surplus estimates are qualitatively similar with or without controls for workers' applications.

Our second test examines how applicants' characteristics change with the instruments. Column 3 of Table 3 shows that when we omit all worker resume characteristics from the regression, the first stage coefficient estimate goes from .078 to .082 for the exchange rate instrument while the competition instrument estimate goes from -.067 to -.085. The change in the exchange rate instrument estimate is within the original confidence interval, while the movement in the competition instrument estimate is not. To understand if the mix of workers changes with the instruments, in Table A.2 we assess how applicant characteristics differ between top and bottom instrument terciles. Worker characteristics vary little with these large changes in the instruments. The only significant difference we find between the lowest and highest exchange rate tercile is a reduction in the good English skills indicator, from 0.910 to 0.893. The only difference for the competition instrument is for the agency affiliate indicator, which increases from 0.322 to 0.328.²³ Although statistically significant, these gaps are substantively small. There are no significant differences for other worker characteristics, including log wage rates on past hires, past numbers of jobs, education, or the presence of prior feedback. These minor differences offer reassurance that worker selection is likely vary little with the instruments.

A second concern is that the instruments may change workers' effort on a given application, altering how buyers perceive their quality. Although we do not have a direct measure of workers' effort, we do observe a proxy: the number of items workers upload files or examples to share in their resumes or with buyers directly. These items tend to include portfolios of past work, including writing samples, code, or graphics. Table A.3 shows that we detect no evidence that the instruments change how many items workers share.

Finally, we also require that buyers do not base their job posting decisions on exchange rate movements or on how many workers are competing for posts by other buyers. Horton (2021) confirms that there is little demand-side response to currency fluctuations. In addition, buyers do not see any statistics about competing jobs when they enter the platform. It is thus unlikely that buyers make strategic decisions to post based on perceived competition from other job openings.

²²In a related analysis, Horton (2021) finds that the collapse of the Russian ruble (outside of our sample period) drove large changes in the number of applications. In a larger panel from more countries, however, Brinatti, Cavallo, Cravino, and Drenik (2021) find larger rates of pass through of exchange rates to wages.

²³Agencies are groups of workers who share a common reputation score. See Stanton and Thomas (2016) for a discussion of the role of agencies on the platform.

3.6 Estimation and Inference

This section outlines how we form the likelihood and estimate the model. We discuss the separate sub-parts of the likelihood in turn: the hiring probability given a consideration set in equation (2); the size of the consideration set in equation (3); and the job arrival process in equation (8). We then discuss how we decompose wage bids into costs and markups as in equation (6). Estimates of markups include the weights that workers put on signals of buyer type, b ; we discuss how we estimate these weights at the end of the section.

3.6.1 Demand Side Estimation

We estimate the model by maximizing the likelihood of the observed sequences of buyer job postings and hiring decisions. The step-by-step estimation approach is as follows: First, we calculate the residuals from the first stage in equation (9) to form control functions that account for unobserved worker quality, denoted $CF_{oj} = \hat{v}_{oj}$. Second, we calculate choice probabilities conditional on buyer type, k , and the consideration set, J_o , as in equation (2), while including CF_{oj} ,

$$p(j|J_o, k, \chi) = \frac{\exp(X_j \beta_{k\chi} - \alpha_k \log(w_{oj}) + \psi_{k\chi} CF_{oj})}{1 + \sum_{j \in J_o} \exp(X_j \beta_{k\chi} - \alpha_k \log(w_{oj}) + \psi_{k\chi} CF_{oj})}. \quad (10)$$

We then form the joint likelihood for type k , yielding

$$\begin{aligned} L_k = & \prod_o \{p(j|J_o, k, \chi)^{(y_o=j)} \times \lambda_{k\chi o}^{\text{CONSIDER}} \exp(-\lambda_{k\chi o}^{\text{CONSIDER}} |J_o|)\}^{J_o \text{ Observed}} \\ & \times \{p(j|\hat{J}_o, k, \chi)^{(y_o=j)}\}^{J_o \text{ Unobserved}} \\ & \times (\lambda_{k\chi}^{\text{ARRIVAL}} \exp(-\lambda_{k\chi}^{\text{ARRIVAL}} t))^{o \in \{2, \dots, O-1\}} \\ & \times (\exp(-\lambda_{k\chi}^{\text{ARRIVAL}} T))^{o=O}. \end{aligned} \quad (11)$$

The term $p(j|J_o, k, \chi)^{(y_o=j)}$ is the conditional probability of choosing j on posting o given a consideration set J_o . The first line of equation (11) is the joint likelihood of choices and consideration set size when the consideration set is observed.²⁴ The second line evaluates the conditional choice probability given a draw of the consideration set size, \hat{J}_o , for postings where the consideration set is unobserved.²⁵ The third line captures the probability of waiting t months to post job o from the time when the previous job, $o - 1$, was posted.²⁶ The last line is the waiting time for the final job post, which accounts for right censoring of the next arrival time at the end of the data. The

²⁴We right-censor the size of the consideration set in the likelihood to 250, which is the top 0.2% of the data. This entails replacing $\lambda_{k\chi o}^{\text{CONSIDER}} \exp(-\lambda_{k\chi o}^{\text{CONSIDER}} |J_o|)$ with $\exp(-\lambda_{k\chi o}^{\text{CONSIDER}} \min(|J_o|, 250))$.

²⁵ \hat{J}_o is drawn from the exponential distribution given parameter guess $\lambda_{k\chi o}^{\text{CONSIDER}}$. We fix a set of uniform random draws, u , and draw the consideration set size given the parameters as $\frac{1}{\lambda_{k\chi o}^{\text{CONSIDER}}} \log(1 - u)$. We do not observe consideration sets for approximately 19% of the openings in the sample.

²⁶The first opening for every buyer, opening 1, does not have a lag prior to posting.

type-specific likelihood contribution is the product over the sequence of choice probabilities for the chosen alternative, consideration set size, and densities of posting times across job openings.

The overall likelihood is the weighted sum of the type-specific likelihoods, with type shares, ρ_k , that are parameters to be estimated (Train, 2009). We assume there are $K = 3$ distinct buyer types, which enables us to capture rich heterogeneity while keeping the model parsimonious.²⁷ This yields

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

The final likelihood has two different types of parameter heterogeneity: across buyer types and across buyer experience. To capture type heterogeneity, all features of the choice model, consideration set size, and job arrival rate are allowed to vary freely across types. To capture experience heterogeneity, we include an offset that shifts the type-specific resume characteristic preferences, consideration set size, and arrival rate parameters by a common amount after a buyer has hired once before on the site.

Maximizing the likelihood requires an iterative procedure because guesses of $\lambda_{k\chi o}^{\text{CONSIDER}}$ change the choice set discontinuously when the consideration set is unobserved and must be simulated. Appendix B.2 contains details about the iterative estimation algorithm. We compute standard errors for all estimates using a block-bootstrap procedure over buyers.

After estimating the parameters, we compute the posterior types for each buyer as

$$\hat{\rho}_{ik} = \frac{\rho_k L_k}{\sum_k \rho_k L_k}. \quad (13)$$

Because some types post more frequently than others, we compute market-level statistics by taking the mean of each posterior by job opening.

Examination of the posteriors and the likelihood provides intuition about the data features that identify the different buyer types and population type shares. In particular, the choice data contributes to identifying the population type shares when buyers take different actions in the face of similar choice sets. Consider a buyer who repeatedly hires when faced with low-quality applicants that submit high bids; this buyer is likely to be a type with a high valuation for the market, captured by a coefficient on price that is relatively small in absolute value. Another buyer who does not hire when faced with high-quality applicants and low bids is likely to be a type with a lower valuation. These patterns, combined with the joint distribution of posting frequency and consideration set size, contribute to identifying buyer type heterogeneity.

²⁷We stopped adding types at three, as one of the three types is estimated to be only around 4% of the total buyers in the data. The model is computationally expensive to estimate, and adding further buyer types would likely represent very niche segments while potentially leading to overfitting in sample.

3.6.2 Supply Side Estimation

Worker markups and costs are computed using data on wage bids and the demand-side estimates, as set out in Section 3.3, up to the weight that workers put on an individual buyer’s type, b . This weight appears in the worker’s estimate of her hiring probability, $E[\tilde{p}(j)]$, in equation (5).

To estimate b , we take logs of the bid equation in (6), which allows us to separate log markups from log costs. We assume that log costs are constant within a week by job category cell for each worker. Variation in a worker’s wage bids across openings within a cell then reflects differences in the worker’s perceptions of optimal markups due to buyer type heterogeneity. We use an estimator that searches over b in equation (5) to find the value that minimizes the difference between the model-predicted markup variation within a week-job category cell and the wage variation observed in the data. To calculate the model-predicted markups for each possible b , we substitute $\hat{\rho}$ (the model’s estimate about the buyer’s type) for $\tilde{\rho}$ (the worker’s information about the buyer’s type) in the hiring probability. We let b vary by buyer experience, reflecting that workers may have different information once a buyer has hiring experience on the platform.

As a preview, we find that the weight workers put on private information is small, at 0.1 when buyers are inexperienced and 0.08 when buyers are experienced. The worker’s perceived hiring probability in equation (5) puts over 90% of the weight on the average buyer type rather than on any signals they have about individual buyers.

3.7 Buyer and Worker Surplus

3.7.1 Buyer Surplus

Average buyer surplus conditional on hiring is defined as the wage bid change that would make the buyer indifferent between the chosen worker and not hiring on the platform (choosing option 0). One can think of this surplus concept as a measure of the realized gains from trade through the market, as it is the difference between the utility from the chosen alternative and not using the platform. To perform this calculation we use the estimated parameters and simulate the unobservables that rationalize observed hiring decisions. Appendix B.3.1 provides details.

Buyers’ expected surplus on a job opening uses the ex-ante distribution of unobservables. Expected surplus is closely related to familiar notions of consumer surplus, where we integrate over quantities as prices increase from the observed wage vector w_0 . Expected surplus for each hour of work (which is the wage unit) on an opening is given by

$$E(\text{Surplus}/\text{Hour})_{k\chi} = \int_{w_0}^{\infty} (1 - p(0|J_o, k, \chi)) \times (w - w_0) dw. \quad (14)$$

The expression $1 - p(0|J_o, k, \chi)$ is the “inside” hiring rate (probability of any hire) and is a function of each considered worker’s wage bid. If inside hiring rates change little with wages (i.e., demand is relatively inelastic), buyers will enjoy substantial surplus because they would be willing to hire at wages much in excess of w_0 . If, instead, demand across all workers is relatively elastic, buyer surplus will be lower because the inside hiring probability will fall rapidly as the wage increases. Our model also allows us to distinguish between static and dynamic surplus. Appendix B.3.2 contains details about how we compute dynamic surplus given the estimated arrival rate of jobs.

Although our estimates do not take a stand on how buyers’ consideration sets are formed (the empirical distribution is all that is needed for workers to determine their markups), making a few additional assumptions allows us to calculate buyers’ search costs. The calculation in equation (14) fixes the size of a consideration set. If we assume that the buyer chooses an optimal consideration set size before starting to search, as in fixed sample-size search, the additional expected consumer surplus from adding a worker to the set should be lower than the buyer’s search cost (Honka, Hortaçsu, and Wildenbeest, 2019). Similarly, the reduction in surplus from subtracting a worker should be greater than the search cost. We later use this insight to estimate buyer search costs based on how consumer surplus changes when we alter the consideration set. We then net out search costs from buyer surplus.

3.7.2 Worker Surplus

The surplus per hour worked for hired workers is the difference between their observed wage bid (net of platform fees) and their cost estimates from equation (7). Realized aggregate surplus to workers is the sum of hours worked across jobs multiplied by the per-hour difference between net wages and costs on each job. The expected surplus for an applicant is the surplus per hour conditional on being hired, multiplied by the hiring probability in equation (3) and the expected number of hours conditional on hiring. We treat the surplus calculations as gross of application costs. We later net out an estimate of application costs to understand the extent to which the job-finding process reduces the estimated benefits from the market. To the extent that workers’ cost estimates include the hassle of applying or the sunk time costs of an interview, our on-the-job surplus calculation represents a lower bound for hired workers.

4 Results

4.1 Demand Parameters, Arrival Rates, and Type Shares

Table 4 presents the parameter estimates related to buyer types, engagement with the platform, and demand elasticities. Most buyers (76%) are grouped together in one type, labeled Type 2, whose estimates are summarized in the second column of Table 4. Type 2 buyers' choice elasticity, defined as the elasticity of choosing an individual applicant in the consideration set with respect to that applicant's wage bid, is -4.31 . Their job fill elasticity is smaller, at -3.06 , because this elasticity assumes wage bids rise for all applicants, reducing substitution across considered applicants. Based on the parameter estimates on worker characteristics (see Appendix Table A.4), the standard deviation in worker characteristics scaled by the log wage coefficient, $\frac{X_j \beta_{kX}}{\alpha_k}$, is 0.38. Buyers appear to perceive substantial worker differentiation. Type 2 buyers consider an average of 15.4 applicants. Once experienced, they post 0.09 jobs per month. While they account for a large share of buyers in the sample, they post less often and are less likely to gain hiring experience than other types. At the job-opening level, they post 65% of the jobs by inexperienced buyers and only 38% of jobs by experienced buyers.

The first and third columns of Table 4 summarize the estimates for the two other buyer types, making up 4% and 20% of the buyers in the sample, respectively. The variation across columns in panels B to E of Table 4 confirms that buyer-type heterogeneity is an important factor in determining platform use. Type 1 is the most wage elastic, with an elasticity of -6.03 over applicants and -4.22 at the job level. This type considers 20.6 applicants on average and, once experienced, posts 4.26 jobs per month. As a result, they account for nearly one-fifth of all posts by experienced buyers. Type 3 buyers are the intermediate group that posts much more often than Type 2s but less often than Type 1s. Type 3s post 43% of all jobs by experienced buyers.

The fourth column of Table 4 shows the opening-weighted mean results. The wage bid elasticity for an applicant is -4.78 , the average consideration set size is 18.1, and the average number of jobs posted per month is 1.12. The final row in the last column of Panel D shows that the estimated sensitivity of job posting frequency to past wage bids received (δ_3 in equation (8)) is -2.05 . Buyers are less likely to use the platform if they have been exposed to high wage bids in the past. At the mean, a standard deviation increase in log wage bids of 0.036 reduces future job posting rates by about 7.6%.

4.2 Buyer Surplus

Table 5 presents results for buyer surplus. The first two columns display the mean and standard deviation of surplus estimates in our baseline specification. Realized buyer surplus per hour conditional on hiring is shown in the first row. The mean weighted by buyers' posterior types is \$2.30 per hour with a standard deviation of \$3.85. Panel B gives the expected surplus per hour when posting a job, as defined in equation (14). Because we do not condition on hiring, the mean is lower, at \$0.74, with a standard deviation of \$0.84. Panel C takes into account the expected hours of work per posted job and the expected arrival rate of future jobs to find the present discounted value of lifetime surplus. Future benefits are discounted at a 8.7% annual rate (see Appendix B.3.2 for calculation details). For the typical inexperienced buyer, the expected lifetime value of platform surplus is \$386. Lifetime surplus for buyers that hire at least once is much higher, at \$5,686. Experienced buyers post more frequently and hire more often, increasing their surplus. These estimates are all gross of buyers' search costs, which we return to in Section 4.4.

Columns 3 and 4 of Table 5 summarize buyer surplus when the estimated hiring probability function includes a control for the number of applications the worker submits that month. This robustness test helps to address the concern that workers select non-randomly in or out of applications in response to the instruments. Across all rows, the buyer surplus is slightly lower, but the magnitudes are similar.

The final two columns of Table 5 present estimates from a version of the model where buyers' consideration sets include all applicants. This is equivalent to allowing J_o in equation (2) to be the entire set \mathcal{J}_o . Two different forces are at play: Relative to the baseline specification, chosen workers come from a larger pool of applicants, meaning the unobservables required to rationalize hiring a given worker must be larger. This increases buyer surplus on a hire. However, because the considered applicant pool is assumed to be larger but the hiring rate does not change, buyers, on average, must value each worker less than in the model with limited consideration. These opposite forces mean the estimated buyer surplus per hour upon hiring rises to \$2.87, but expected surplus falls to \$0.49.

4.3 Worker Markups, Surplus, Model Validation, and Application Costs

4.3.1 Model estimates

Table 6 presents worker surplus estimates. Panel A focuses on hired workers. The mean worker surplus per hour is \$1.98, with a standard deviation of \$1.49. Hired workers' markups average 28% over costs, and hourly costs average \$6.97. Appendix B.4 shows that we get qualitatively

similar estimates of markups and surplus when we allow hired workers to rationally account for the possibility of wage growth after hiring. Panel B presents markup estimates of 27% for all applicants, regardless of whether they are hired, and costs of \$6.73. The \$5.00 estimated standard deviation in costs comes from the wide variation in offline opportunities for globally dispersed workers.

Figure 4 presents histograms of hired worker surplus for each of six large worker countries. Workers in the United States and United Kingdom have a larger mean value than those in other countries, consistent with markups multiplicatively magnifying these countries' relatively high offline opportunity costs. In contrast, the surplus in Bangladesh and the Philippines is more concentrated at low levels, reflecting both lower values of offline options and a job category mix weighted toward non-technical skills. Russian workers earn relatively high surplus per hour when hired because they tend to specialize in technical job categories.

Figure 5 shows how expected total surplus, which includes the expected hours required on an opening, varies by applicant order for the baseline model. Applicants in the first 10 positions have an expected surplus of \$2.21. Applicants 11 to 20 have a lower expected surplus, at \$1.18, while these figures are \$0.74 and \$0.55 for applicants 21 to 30 and 31 to 40, respectively. After applicant 40, the expected surplus is less than 8% of a worker's estimated hourly cost.

Returning to Table 6, columns 3 and 4 display estimates when we include the control for workers' total applications in a month. Mean estimates of surplus and markups are very similar. Finally, when we assume all workers are considered, in columns 5 and 6, estimates of workers' surplus and markups remain positive but are smaller than the baseline estimates. Mean surplus per hour on a hire is about \$1.53 and markups average 20%. Here buyers are more elastic to help rationalize the limited hiring rates in the data. Taken together, buyers' limited consideration appears to increase worker markups by about 7 percentage points (35%) relative to a simpler model that does not account for limited consideration. Still, in all models, workers benefit from the market.

4.3.2 Validation

The model predicts that markups and bids fall with applicant order for the same worker. This prediction allows us to validate the model-implied markups against the actual bids observed in the data. To do so, we regress log bids or model-implied markups on applicant order dummies and worker by job category by week fixed effects. This strategy removes worker-level costs over a short interval, isolating variation in how workers bid relative to model-implied markups. The red squares in Figure 6 plot the model-implied markups relative to those for the first 10 applicants. Markups decline with applicant order. Applicants 60 to 70 have model-implied markups that are around 2.2 percentage points (1.8%) lower than the earliest applicants. Driving these reductions is the fact

that later applicants know they will only be considered in the event that the buyer also considers many other applicants; that is, when there is more competition. For applicants who apply early, there is a chance that competition will be more limited, which raises their optimal markup.

Workers’ log bids in the data also decline with applicant order. Applicants 60 to 70 have log bids that are 3.5% lower than the first 10 applicants. Over half of the reduction in bids can be explained by changes in estimated markups with applicant order. Beyond the first 80 applicants, the raw bids decline at a faster rate than the model-implied markups, which likely arises from data sparsity; more than 95% of job openings have consideration sets that consist of fewer than 80 applicants, as shown in Figure 3.

4.3.3 Survey-Based Estimates of Surplus on Hourly and Fixed Price Contracts

We return to the survey data as another way to validate the model findings related to worker surplus. After asking respondents what factors led them to tailor their bids to different openings, we asked survey questions designed to elicit markups (the survey question wordings can be found in Appendix A.2). These questions asked respondents to provide the lowest rate at which they would have been willing to take their most recent job. We distinguish between markups at the time they accepted the job and after contract completion.²⁸ Differences between actual hourly wages and willingness to accept are summarized in Table 8. Column 1 reports the raw data for the respondents who answered the question, and column 2 reweights observations to be representative of a random sample of 1,485 worker profiles.²⁹

The first row of Table 8 reports that workers earned an hourly wage that was 24% higher than the lowest hourly wage they would have been willing to accept on the job ex-ante. In the reweighted data, their wage was 38% higher than the lowest acceptable hourly wage. That is, at the time they accepted their last job, surveyed workers anticipated earning positive markups over their willingness to accept. The second row shows that the markups they actually earned were lower, at 16% in the raw data and 15% in the reweighted data, suggesting that they incurred costs that were higher than anticipated on their last jobs. Nonetheless, respondents still reported a large markup and positive

²⁸The first of these questions was: “From the information that you had in the job posting on your last hourly job, what would have been the lowest hourly rate you would have been willing to accept to do the job? For example, if the job paid you \$10 US but you would have been willing to do the work for \$9.00 but not \$8.99, then \$9.00 is the lowest hourly rate you would have been willing to accept”. The question about ex-post willingness to accept is worded similarly and is given in Appendix A.2.

²⁹The probability of being in the surveyed sample is computed using a logistic regression of survey participation on the log of the number of prior jobs and the log of the hourly rate displayed on each profile. Responses are then weighted by the inverse of this probability, ensuring that responses from workers who are underrepresented in the survey are given more weight in the comparisons given in column 2. The random sample was hand collected from the site for workers with past hires in Administrative Support, Design, and Web Development job categories. Further details about the sampling procedure are given in the appendix.

surplus.

An alternative estimate of the markups included in workers' bids comes from another set of questions related to how the wages earned on their last jobs compared to their next best option off the platform. The average responses are given in the third row of Table 8. Markup estimates relative to off-platform options are again large, averaging 23% in the reweighted data.

Finally, to assess markups on fixed price contracts, we asked respondents to state the price they received, and the price they would have been willing to accept ex-post, on their last fixed price job. Respondents reported very high markups relative to their willingness to accept, averaging 62% in the reweighted data.

In sum, the survey responses suggest the majority of workers are strategic in setting their wage bids above their reservation wages, consistent with them tailoring bids to their perceived market power. The reservation wages elicited directly, via questions about willingness to accept the job, or indirectly, via questions about their next best alternative wage, are substantially lower than the hourly wages or fixed prices workers actually earned on recent jobs.

4.4 Total Surplus, Worker Application Costs, and Buyer Search Costs

Table 7 shows that the total surplus earned by workers who are ever hired during our sample has a mean of \$322 per worker and a standard deviation of \$863. Workers' aggregate gross surplus totals \$4.56 million. The second row shows that mean surplus per worker across all applicants, not just those who are hired, falls to \$23.70. From Table 1, 8% of the workers in our sample were ever hired. Because many workers search for jobs compared to those who find them, we must also account for application costs to characterize workers' total net surplus.

Aggregate worker surplus remains positive when we account for application costs. Our estimated hourly worker costs include the opportunity cost of time, allowing us to estimate application costs in dollars if we know the approximate time it takes to apply for a job. Based on our own experience applying to jobs and our assessments of the personalized messages in applications to jobs we have posted, we believe it is reasonable to assume that each application takes about five minutes. In dollar terms, this is about 8.3% of workers' hourly cost. We believe this estimate is conservative, as Figure 5 shows that 8.3% of workers' hourly costs corresponds to the ratio of expected surplus to hourly costs for the 40th applicant. The fact that many jobs receive more than 40 applications suggests our estimate of application costs likely overstates workers' search costs. Using 8.3% of each worker's estimated hourly costs and totaling the costs of all applications, we estimate that application costs for workers who are ever hired average \$95 per worker, totaling \$1.35 million in

the aggregate. Across all workers, the average is \$12.98 per worker with a total of \$2.5 million. Subtracting these application costs from gross surplus gives total net surplus of 3.21 million for hired workers and \$2.06 million for all applicants. Although only 8% of workers are hired, their surplus is large enough to offset not just their own application costs but those of the 92% of applicants who do not find jobs.

Several other points suggest that these calculations are reasonable even if the average worker takes more than 5 minutes to apply for a job. First, workers' costs come from model estimates that may entail portions of costs, like effort costs, that are not simply opportunity costs of time. If workers' opportunity costs are lower than the estimated total costs, then we overestimate application costs. Second, estimated costs increase by about 20% from the first to the fifth hire on the platform (see Figure A.4). This suggests many workers' alternatives are other online jobs. In the absence of the gains from trade enabled by the platform, these workers' opportunity costs, and hence application costs, would be lower.

Next, we turn to estimating buyer search costs. We use the insight that for a buyer's consideration set size to be optimal, the marginal benefit of adding the last applicant to the set must exceed the marginal search cost; adding an additional applicant beyond the observed set must have marginal benefits below the search cost. To estimate buyer search costs, we calculate the average change in consumer surplus from adding an unconsidered applicant. We then calculate the absolute value of the change in consumer surplus from deleting the last considered applicant. These two numbers provide an upper and lower bound for buyers' search costs. We average across these changes to get an estimate of buyers' search costs as the midpoint of the bounds.³⁰ We multiply the per-hour consumer surplus by the expected job size to get an average dollar value of buyer search costs of \$1.21 per considered applicant. Panel B of Table 7 shows that buyers who hire have search costs per opening that total about \$21.68 with a standard deviation of \$33.40. In total, search costs for those who hire total about \$0.76 million. When considering all openings, total search costs are \$3.79 million, or \$22.22 per opening. While search costs are meaningful, a comparison of columns 3 and 6 of Panel B shows that buyers' gains net of these costs total \$1.98 million.

³⁰The intuition for how to recover these search costs comes from a fixed sample size search process, but the approximation is likely to be a good one even if there are deviations from the fixed sample size model in practice. However, two pieces of evidence suggest that fixed sample size search likely fits the data well. First, the similarity between the search cost estimates for those that hire and those that do not suggests that consideration set sizes do not differ meaningfully based on whether a buyer hires. Second, buyer interview requests have date stamps that bunch together at early dates, whereas hiring happens with a lag. This suggests that it takes some time to setup interviews and make decisions out of a set of applicants that receive attention.

5 Policy Counterfactual

Our estimates of gig economy worker surplus are informative for debates about regulatory changes targeting worker wellbeing. Proposals have varied but would typically raise hiring costs by imposing taxes—like payroll taxes—or introducing minimum wages to bring independent contractors closer in line with W2 employees. We consider one such change here: an additional ad-valorem tax on wage bids that mimics the FICA tax that is statutorily imposed on employers in traditional labor markets. An additional motivation for this analysis is to consider whether governments could increase worker surplus by redistributing revenue raised from such a tax back to workers, in the form of benefits or other services.

We impose an additional counterfactual 10% tax on bids. Although the nominal incidence of the tax is on buyers, workers’ bids are determined in equilibrium, and the incidence will be split between buyers and workers. We assume that this tax is in place for the duration of the sample, which allows us to assess how surplus would evolve for buyers and workers in the market.³¹ We study two margins of adjustment: how surplus changes holding jobs fixed, and how the number of jobs changes as a result of lower buyer posting rates. Appendix C details the calculations.³²

We find that imposing the tax leads to changes on both the hiring and job posting margins that significantly reduce market surplus when added together. In particular, workers pass through the tax to buyers in the form of higher wage bids. In response, buyers become less likely to hire on posted jobs and fewer buyers gain hiring experience. As a consequence of both the higher wage bids and the reduced likelihood of gaining experience, buyers post fewer subsequent jobs. To illustrate how higher wage bids directly reduce job postings, we contrast the main counterfactual estimates with a second scenario that shuts down the direct effect of higher equilibrium wage bids on future job postings. Changes in surplus in this alternative come from differences in hiring decisions: changes to job posting rates arise only because fewer buyers transition to being experienced.

Table 9 summarizes changes in market outcomes under the tax. Column 1 presents changes in surplus in the main scenario. Panel A gives the outcomes for inexperienced buyers, and panel B for experienced buyers. Bids to inexperienced and experienced buyers increase by 8.8% and 8.9%, respectively. Hiring rates on posted jobs fall by 27% and 24%. Surplus to buyers on each posted job (labeled “Static Pct Change in Buyer Surplus”) falls by over 20% for both groups.

³¹The structure of the counterfactual 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.

³²Appendix C.2 presents a second policy change—a wage floor of \$7.00 per hour, which binds for the lowest-paid workers in the market.

The next rows in Panels A and B show that the dynamic implications magnify the losses to buyer surplus. Because of the reduction in hiring rates on posted jobs, fewer buyers gain hiring experience. While the number of jobs posted by inexperienced buyers increases slightly for this mechanical reason (by 3%), job postings by experienced buyers fall by 68%. The smaller number of buyers that do become experienced post fewer jobs because of the negative relationship between past bids and posting frequency. The reduction in the present value of experienced buyer surplus is much larger than their static loss, at 72%. Inexperienced buyers also see a 75% decrease in surplus because the value of gaining experience falls.

Column 2 of Table 9 presents the results when higher past bids do not affect the job arrival rate. The static results are the same as the base case in column 1. The number of jobs posted by inexperienced buyers increases relative to column 1 because there is a larger stock of inexperienced buyers who are undeterred from posting by high past bids. For experienced buyers, total job postings relative to the status quo remain negative because the stock of experienced buyers falls. Dynamic surplus reductions remain substantial for both groups of buyers, but breaking the link between price increases and future job posts partially ameliorates the decline in buyer surplus.

The reductions in buyer hiring and posting rates have large detrimental effects on worker surplus. Panel C shows the surplus changes for workers before any tax rebate. In the main scenario in column 1, static surplus from a given job falls by an average of 25% due to lower job fill rates. The small increase in hired workers' wages is offset by lower hiring rates. The present value of total worker surplus falls by 59% due to the reduction in hiring rates and, in particular, the reduction in the number of postings. Comparing columns 1 and 2 of Panel C shows that about half of the reduction in surplus to workers arises because of the decline in job postings.

Panel C also shows the net change in worker surplus when the entire value of the tax collected is rebated to workers. In the static case with no changes in the number of job posts, the value of a 10% tax on all hires is sufficient to more than offset the reduced hiring rate, making workers 15% better off. However, the present value of net worker surplus remains negative after the rebate, falling by 37.5% when accounting for the reduction in job postings. Column 2 shows the impact on net worker surplus if past bid-related buyer dynamics are shut down. In this scenario, the present value of net worker surplus increases by about 12% because higher prices do not directly change the size of the market. To the extent that redistribution can raise workers' surplus, comparing columns 1 and 2 shows that any distortion to market size must be limited for these proposals to benefit workers. Our model estimates, however, suggest that buyers are sensitive to past prices, suggesting attempts to redistribute surplus to workers through tax-like policies would be unsuccessful.³³

³³When studying the impact of minimum wages in the Appendix, we find that changes in surplus are negative

Finally, regulatory policy may be motivated by the belief that access to global talent through platforms could erode traditional offline employment regulations or relationships. In Appendix D, we consider whether offline workers are hurt by the advent of online labor platforms. To do so, we study whether minimum wage changes across US states drive jobs online. While across-state designs have power issues, we fail to detect any evidence that states that impose higher minimum wages experience differential growth in online buyer demand. These results hold for relatively low-wage non-technical jobs and are consistent with survey evidence suggesting that platforms lead to new trade in tasks that otherwise would have been done by the hiring manager personally or not done at all (Ozimek and Stanton, 2022; Horton, Johari, and Kircher, 2021).

6 Conclusion

This paper estimates the surplus that online gig economy labor markets create for workers and buyers. Despite the relative abundance of applicants, workers earn significant surplus from the market due to perceived differentiation in their characteristics and buyers' limited search. Fears that intense competition among job applicants would limit surplus for workers using gig economy platforms do not appear to play out in this market.

Our model estimates suggest that the platform generates total gains from trade of around \$4.28 per hour worked online and that 46% of the surplus on a filled job goes to the hired worker. Applicants have local market power, enabling them to charge markups of more than 25% above their outside option when they bid for jobs. Even when taking into account the application costs of unsuccessful applicants, the net benefit to workers from the platform remains positive.

A recent survey of remote gig economy workers suggests the structural model setup is reasonable and provides evidence consistent with the models' main findings. Based on several different constructs, surveyed workers are found to perceive significant surplus from using online labor platforms on both hourly jobs and fixed price contracts.

Finally, we analyze the impact of counterfactual hiring taxes. Taxing buyers with the intention of redistributing surplus to workers would diminish value for each side of the market by reducing hiring rates and job posts.

for workers and buyers regardless of assumptions about how past prices change quantities. Adding "labor-labor" substitution and assuming buyer posting does not respond to past wage bids leads to small positive surplus changes.

References

- Abaluck, Jason and Abi Adams-Prassl. 2021. “What do consumers consider before they choose? Identification from asymmetric demand responses.” *The Quarterly Journal of Economics* 136 (3):1611–1663.
- 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.
- Autor, David H. 2001. “Wiring the Labor Market.” *The Journal of Economic Perspectives* 15 (1):25–40.
- Barach, Moshe A and John J Horton. 2021. “How do employers use compensation history? Evidence from a field experiment.” *Journal of Labor Economics* 39 (1):193–218.
- Benson, Alan, Aaron Sojourner, and Akhmed Umyarov. 2020. “Can reputation discipline the gig economy? Experimental evidence from an online labor market.” *Management Science* 66 (5):1802–1825.
- Bertrand, Marianne, Esther Dufo, and Sendhil Mullainathan. 2004. “How much should we trust differences-in-differences estimates?” *The Quarterly Journal of Economics* 119 (1):249–275.
- Brinatti, Agostina, Alberto Cavallo, Javier Cravino, and Andres Drenik. 2021. “The International Price of Remote Work.” Tech. rep., National Bureau of Economic Research.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. “Firms and labor market inequality: Evidence and some theory.” *Journal of Labor Economics* 36 (S1):S13–S70.
- 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. 2020. “Has the gig economy replaced traditional jobs over the last two decades? Evidence from tax returns.”
- Cullen, Zoë and Chiara Farronato. 2021. “Outsourcing tasks online: Matching supply and demand on peer-to-peer internet platforms.” *Management Science* 67 (7):3985–4003.
- Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan. 2018. “Consumer price search and platform design in internet commerce.” *American Economic Review* 108 (7):1820–1859.
- Dingel, Jonathan I and Brent Neiman. 2020. “How many jobs can be done at home?” *Journal of public economics* 189:104235.
- Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri. 2020. “Monopsony in online labor markets.” *American Economic Review: Insights* 2 (1):33–46.
- Fisher, Jack. 2022. “Worker Welfare in the Gig Economy.” *Working Paper* .
- Gee, Laura K. 2019. “The more you know: Information effects on job application rates in a large field experiment.” *Management Science* 65 (5):2077–2094.
- 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.” *Working Paper* .

- Gray, Mary L and Siddharth Suri. 2019. *Ghost Work: How to stop Silicon Valley from building a new global underclass*. Eamon Dolan Books.
- 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.
- Honka, Elisabeth, Ali Hortaçsu, and Matthijs Wildenbeest. 2019. “Empirical search and consideration sets.” In *Handbook of the Economics of Marketing*, vol. 1. Elsevier, 193–257.
- 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.
- . 2017. “The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment.” *Journal of Labor Economics* 35 (2):345–385.
- . 2021. “The Ruble Collapse in an Online Marketplace: Some Lessons for Market Designers.” Tech. rep., National Bureau of Economic Research.
- . 2022. “Price Floors and Employer Preferences: Evidence from a Minimum Wage Experiment.” Tech. rep.
- Horton, John J, Ramesh Johari, and Philipp Kircher. 2021. “Cheap Talk Messages for Market Design: Theory and Evidence from a Labor Market with Directed Search.” Tech. rep., National Bureau of Economic Research.
- Kahn, Lisa B. 2013. “Asymmetric Information Between Employers.” *American Economic Journal: Applied Economics* 5 (4):165–205.
- 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)* .
- Katsnelson, Lura and Felix Oberholzer-Gee. 2021. *Being the boss: Gig workers' value of exible work*. Harvard Business School.
- 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.
- Kuek, Siou Chew, Cecilia Paradi-Guilford, Toks Fayomi, Saori Imaizumi, Panos Ipeirotis, Patricia Pina, and Manpreet Singh. 2015. “The global opportunity in online outsourcing.” *World Bank, Washington, DC* .
- Mas, Alexandre and Amanda Pallais. 2017. “Valuing alternative work arrangements.” *American Economic Review* 107 (12):3722–59.
- . 2020. “Alternative work arrangements.” *Annual Review of Economics* 12:631–658.
- Moreno, Antonio and Christian Terwiesch. 2014. “Doing business with strangers: Reputation in online service marketplaces.” *Information Systems Research* 25 (4):865–886.
- 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.
- Ozimek, Adam and Christopher Stanton. 2022. “Remote work has opened the door to a new approach to hiring.” *Harvard Business Review* 11.

- 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.
- Prassl, Jeremias. 2018. *Humans as a service: The promise and perils of work in the gig economy*. Oxford University Press.
- 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.
- Stephany, Fabian, Otto Kässi, Uma Rani, and Vili Lehdonvirta. 2021. "Online Labour Index 2020: New ways to measure the worlds remote freelancing market." *Big Data & Society* 8 (2):20539517211043240.
- Train, Kenneth E. 2009. *Discrete Choice Methods with Simulation*. Cambridge University Press.
- Varian, Hal R. 2007. "Position auctions." *international Journal of Industrial Organization* 25 (6):1163–1178.
- Wright, Randall, Philipp Kircher, Benoît Julien, and Veronica Guerrieri. 2021. "Directed search and competitive search equilibrium: A guided tour." *Journal of Economic Literature* 59 (1):90–148.
- Yu, Chuan. 2023. "The Welfare Effects of Sponsored Product Advertising." *Working Paper* .

Figures

Figure 1: Two Examples of Job Postings

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

Joomla Director customization & creation

Job - Posted 12/03/08 | Applicants: 5 | Invitations: 2 | Interviews: 0

Hourly
(8 weeks, < 10 hrs/week)

Skills: [PHP/MySQL](#)
[Joomla](#)

[Qualifications](#)

[Web Development](#) > [Web Programming](#)

I'm looking for some help customizing a joomla-based web directory using the 'mosets tree' extension & populating it with data. After this project is finished I need help building several more of these directories. I'm looking for someone in US or Canada who's available to speak on the telephone during US business hours.

[more](#)

oDesk Verified
Payment Method

[Apply](#)

Buyer - Member Since 02/05/07

Hourly: \$6,552 (935 hours)

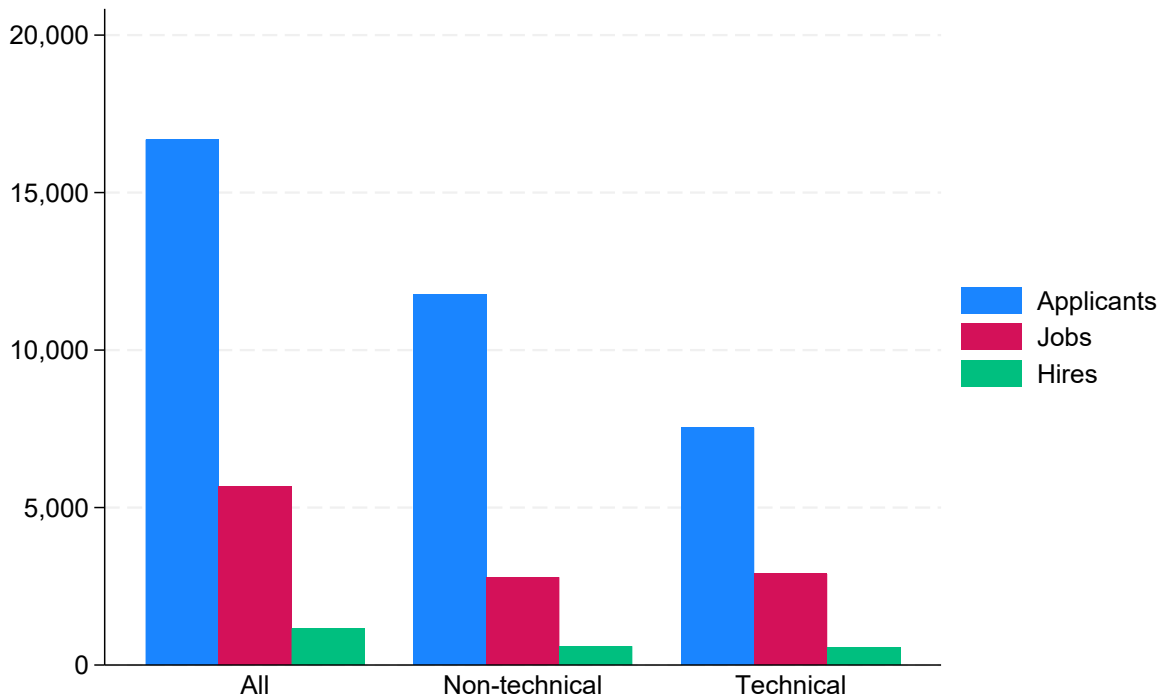
Fixed-Price: \$0 (0 projects)

Total: 11 posted, 4 paid
(5.00) 1 feedbacks

Location: United States

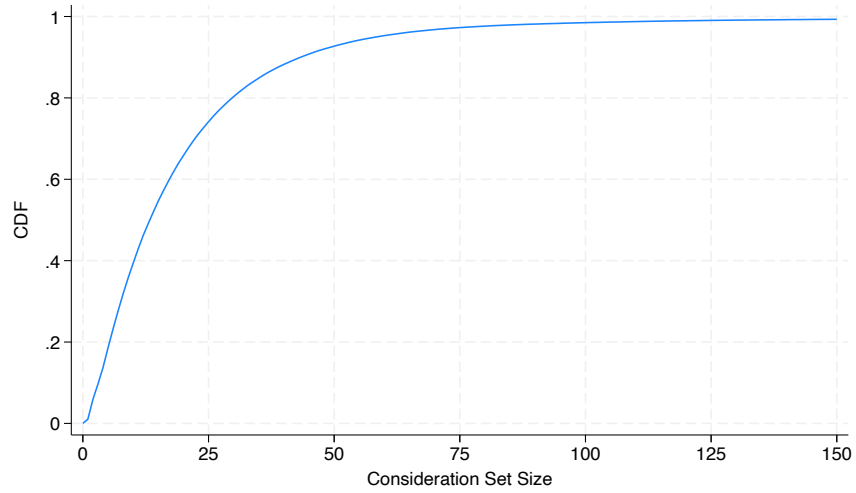
Notes: The first job post is a screenshot from the platform for a non-technical job posting, taken in July 2014. The second job posting is for a technical job during our administrative data sample, taken from the WayBackMachine. The job was posted on December 3, 2008. Applicants can see the job has received five applications.

Figure 2: Job Posts, Applicants, and Hires, 2008-2010




Notes: This figure shows the total count of workers submitting at least one application in a month, the number of new job postings in a month for hourly contracts, and the number of filled hourly job openings in a month, averaged across the 30 months in the administrative data. The left hand columns are for all hourly jobs, the center columns are for non-technical job openings, and the right hand columns are for technical job openings. Because some applicants apply for both technical and non-technical jobs in a month, the number of unique applicants in non-technical and technical jobs sums to more than the total number of unique applicants.

Figure 3: Cumulative Distribution of Consideration Set Sizes




Notes: This figure plots the empirical CDF of the size of buyers' consideration sets for openings where the consideration set is observed. The consideration set size is computed based on the last applicant that the buyer messaged, hired, or indicated a reason for not hiring. Consideration set sizes are unobserved for about 19% of all openings.

Figure 4: Distribution of Worker Surplus Across Countries




Notes: This figure plots the distribution of each workers' estimated surplus per hour for every hire involving a worker from six large countries.

Figure 5: Expected Worker Surplus as a Function of Applicant Order



Notes: This figure plots workers' average expected total surplus on a job application by applicant order. To compute expected total surplus, we take the hiring probability from the baseline model given the buyer's observed or simulated consideration set, multiply that probability by the worker's wage bid less cost, and multiply by expected hours. The x-axis is truncated at 120 applicants, as expected surplus is approximately zero for later applicants.

Figure 6: Worker Wage Bids and Applicant Order for Jobs in the Same Week



Notes: The first series in this figure plots the coefficients from a regression of log bids observed in the data on 10 applicant order categorical dummies (e.g. positions 11-20, 21-30, ..., 91-100, 100+) while controlling for worker week job category fixed effects. The excluded category is the first 10 applicants. Standard errors are clustered by worker. The second series plots model-implied markups relative to the first 10 applicants from a regression of markups on worker week job category fixed effects.

Tables

Table 1: Descriptive statistics about buyers and workers

	Buyers		Workers	
	Mean	Std. Dev.	Mean	Std. Dev.
In the USA	0.58	0.49	0.26	0.44
In another G10 country	0.17	0.38	0.05	0.22
In India or Philippines	0.06	0.24	0.42	0.49
First posting or application is in tech	0.56	0.50	0.28	0.45
Number of postings or applications	2.52	4.50	23.15	81.20
Hires or is hired at least once	0.27	0.44	0.08	0.27
Number of hires	0.52	1.40	0.18	0.95
Number of buyers or workers	67566		192628	

Notes: The sample includes buyers and workers active on the platform from January 2008 to June 2010. Buyers and workers 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 job postings by buyer hiring experience

	All postings	Zero prior hires	At least one prior hire
Number of applicants	26.14 (32.85)	23.52 (26.67)	28.74 (37.83)
Number of applicants in first 24 hours	18.55 (22.54)	16.46 (18.27)	20.63 (25.94)
Mean applicant wage bid	10.90 (4.97)	11.30 (4.93)	10.50 (4.98)
Mean share with good feedback	0.36 (0.18)	0.34 (0.18)	0.37 (0.17)
Mean share with no feedback	0.44 (0.21)	0.46 (0.22)	0.42 (0.20)
Mean number of prior hires	5.78 (4.43)	5.36 (4.15)	6.20 (4.66)
Mean last hire wage if applicable	10.16 (5.36)	10.55 (5.31)	9.77 (5.39)
Mean share in the U.S.	0.11 (0.15)	0.11 (0.15)	0.11 (0.15)
Mean share with a BA+ degree	0.35 (0.15)	0.35 (0.15)	0.35 (0.15)
Probability of filling job	0.20 (0.40)	0.15 (0.35)	0.26 (0.44)
Months between postings	1.15 (2.58)	1.53 (3.24)	0.89 (1.97)
Number of posts	170556	84999	85557

Notes: This table presents job posting-level averages (standard deviations) of posting patterns, application characteristics, and hiring rates. The mean wage bid is inclusive of a 10% ad-valorem platform fee. Statistics about applicant characteristics display averages by job posting. These characteristics are: the share of applicants with good feedback (defined as a feedback score of at least 4.5 out of 5, auto-populated in the profile), the share of applicants with no feedback, the average number of prior hires on the platform at the time of application, the average hourly rate on the last hire (for those with prior hires), the share of applicants in the United States, and the share with a bachelors or higher degree listed in their profile. Job filling is defined as hiring one of the applicants. The sample period is from January 2008 to June 2010. Standard deviations are in parentheses.

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

	(1)	(2)	(3)
Exchange Rate Instrument	0.078 (0.010)	0.069 (0.010)	0.082 (0.011)
Competition Instrument	-0.067 (0.008)	-0.063 (0.008)	-0.085 (0.009)
Worker Applications / Month		-0.001 (0.000)	
R-Squared	0.630	0.631	0.567
Observations	4458728	4458728	4458728
F Clustered on Posting	84.49	70.70	91.25
F Clustered on Worker	48.20	41.93	47.29
Excludes Resume Characteristics	No	No	Yes

Note: The exchange rate instrument is the log of the dollar to local currency exchange rate. The competition instrument is the log leave-buyer-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. Column 2 controls for the number of jobs a worker applies to in a given month. Column 3 omits worker resume characteristics. Standard errors clustered by job opening are in parentheses. Partial F statistics on the excluded instruments are also reported after clustering by job opening or by worker. Column 1 forms the basis of the main demand estimates reported in the paper. Columns 2 and 3 are the first stage estimates used for alternative models to assess sensitivity.

Table 4: Estimates of Buyer Types, Demand Elasticities, and Platform Engagement

	Type 1	Type 2	Type 3	Mean
Panel A: Buyer Types				
Buyer Type Share	0.04 (0.00)	0.76 (0.01)	0.20 (0.00)	
Panel B: Elasticities and Demand Features				
Job Fill Elasticity	-4.22 (0.77)	-3.06 (0.63)	-3.60 (0.59)	-3.40 (0.55)
Considered Applicant Elasticity	-6.03 (1.63)	-4.31 (1.07)	-5.01 (1.10)	-4.78 (1.00)
Std. Dev. of $X\beta$ (log wage units)	0.37 (0.02)	0.38 (0.02)	0.36 (0.01)	0.37 (0.03)
Panel C: Applicant Consideration				
Mean Applicants Considered	20.59 (0.48)	15.39 (0.17)	21.24 (0.25)	18.08 (0.07)
Panel D: Job Posting Frequency				
Mean Jobs Per Month (Experienced)	4.26 (0.44)	0.09 (0.00)	0.70 (0.04)	1.12 (0.05)
Sensitivity of Posts to Past Bids				-2.05 (0.23)
Panel E: Share of Job Posts by Type				
Share of Inexperienced Job Posts	0.08 (0.00)	0.65 (0.01)	0.27 (0.01)	
Share of Experienced Job Posts	0.18 (0.01)	0.38 (0.01)	0.43 (0.01)	

Notes: This table presents estimates of buyer types and behavior. The first three columns correspond to individual types, while the last column is a weighted average across types at the job posting or application level. Panel A displays latent buyer types in the sample. The first two rows of Panel B display estimated elasticities for filling a job and for hiring each individual applicant that the buyer considers. The job fill elasticity assumes that all wage bids increase, while the applicant hire elasticity considers a price change for each considered applicant. The final row of Panel B displays the log-wage scaled standard deviation of worker resume characteristics x the parameters on those characteristics. Panel C provides estimates of the mean number of applicants considered as implied by the estimated parameters of the exponential distribution. Because the consideration process is assumed to be exponential, the standard deviation for each type equals the mean, whereas the overall distribution is a mixture across types. Panel D provides estimates of average monthly job posting frequencies by buyer type. The row labeled Sensitivity of Posts to Past Bids is an estimate of the elasticity of job posting frequency with respect to the level of bids that a buyer has historically experienced. As detailed in the text, this is estimated using deviations in aggregate bids from a time trend by job category based on the timing when a buyer has posted prior jobs. Standard errors are estimated from 20 block-bootstrap iterations of the entire estimation procedure (drawing buyers with replacement).

Table 5: Estimates of Buyer Surplus

	Baseline		App Control		Full Consideration	
	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
Panel A: Surplus on Filled Jobs						
Surplus-Per-Hour	2.30 (1.23)	3.85 (4.04)	2.13 (0.94)	3.68 (3.07)	2.87 (1.89)	4.05 (62.07)
Panel B: Surplus on All Jobs						
E(Surplus-Per-Hour)	0.74 (0.27)	0.84 (0.31)	0.68 (0.24)	0.78 (0.28)	0.49 (0.18)	0.45 (0.21)
Panel C: Lifetime Value						
Lifetime Surplus (Inexper)	386 (138)	950 (412)	360 (129)	904 (411)	300 (161)	1155 (863)
Lifetime Surplus (Exper)	5686 (1959)	7307 (3855)	5395 (1997)	7065 (3927)	6499 (3953)	7991 (6280)

Notes: This table presents estimates of buyer surplus. The first two columns correspond to the specification in our main model. The next two columns add a control for worker applications per month. The final two columns assume all applicants are considered. Means and standard deviations are weighted by buyer posterior types. Panel A displays simulation estimates of surplus-per-hour for buyers who hire. Details of the simulation algorithm, including the accept-reject procedure used to rationalize hiring decisions, are included in Appendix B.3.1. Panel B displays expected surplus-per-hour for all job openings. Panel C displays estimates of lifetime surplus after accounting for expected hours of work on each job, expected arrival rates of future jobs, and hiring rates conditional on future jobs. Future benefits are discounted at an 8.7% annual rate. Standard errors are estimated from 20 block-bootstrap iterations of the entire estimation procedure (drawing buyers with replacement).

Table 6: Surplus and Markups for Applicants and Hired Workers

	Baseline		App Control		Full Consideration	
	Mean	St. Dev.	Mean	St. Dev.	Mean	St. Dev.
Panel A: Estimates for Hired Workers						
Surplus-Per-Hour	1.98 (0.57)	1.49 (0.44)	1.86 (0.50)	1.40 (0.38)	1.53 (0.44)	1.16 (0.41)
Markups	1.28 (0.12)	0.02 (0.03)	1.26 (0.10)	0.02 (0.02)	1.20 (0.09)	0.02 (0.08)
Costs	6.97 (0.55)	5.12 (0.43)	7.09 (0.48)	5.21 (0.38)	7.42 (0.44)	5.45 (0.30)
Panel B: Markups and Costs for All Applicants						
Markups	1.27 (0.12)	0.01 (0.03)	1.25 (0.10)	0.01 (0.02)	1.20 (0.10)	0.02 (0.10)
Costs	6.73 (0.52)	5.00 (0.41)	6.84 (0.46)	5.09 (0.35)	7.15 (0.44)	5.32 (0.30)

Notes: Panel A displays estimates of surplus per hour, markups, and costs for hired applicants. Surplus per hour is the difference between the to-worker wage and the estimated cost for the hired worker. Markups and costs are recovered using the details in the text for the models with consideration. Markups and costs estimates use a simple own-bid semi-elasticity in the model without consideration. Estimates of markups and costs use the non-linear least squares procedure to weight the semi-elasticities across different types, based on our estimate of a bidders' private information. This procedure is described in Section 3.6.2. Panel B displays estimated costs and markups over costs for all applicants. See Table 7 for aggregate surplus estimates accounting for hours. Standard errors come from 20 block-bootstrap iterations (drawing buyers with replacement).

Table 7: Total Surplus Accounting for Application and Search Costs

	Surplus			Search Costs		
	Mean	St. Dev.	Total (mil.)	Mean	St. Dev.	Total (mil.)
Panel A: Workers						
Workers with Any Hires	321.99 (120.77)	863.14 (370.13)	4.56 (1.32)	95.31 (8.85)	173.11 (15.55)	1.35 (0.10)
All Applicants	23.70 (7.86)	248.78 (98.81)	4.56 (1.32)	12.98 (1.19)	56.66 (4.88)	2.50 (0.20)
Panel B: Buyers						
Buyer Openings with Hires	165.51 (88.29)	277.38 (290.96)	5.77 (3.07)	21.68 (8.17)	33.40 (11.34)	0.76 (0.28)
All Openings	33.85 (18.06)	142.09 (135.53)	5.77 (3.07)	22.22 (8.14)	33.07 (11.56)	3.79 (1.39)

Notes: This table presents statistics on total surplus per worker and buyer gross of application and search costs (Columns 1-2) and gross surplus aggregated across all workers and buyers (Column 3). Columns 4-6 in Panel A provide estimates of application costs assuming they are proportional to 8.3 percent of worker's estimated cost per hour, implying that a job application has a time cost of about 5 minutes. The first row considers workers who ever land a job on the platform while the second row considers all workers. Total surplus and application costs are given in millions of dollars for the sample of jobs from January 2008 to June 2010. Panel B considers buyer total surplus and search costs. Total surplus is the surplus on a hire multiplied by the number of hours on a job. Search costs come from assuming a fixed sample size consideration set and calculating the average change in benefits when adding and subtracting a worker from the consideration set. We multiply these benefits from a change in search by the expected number of hours on a hire given the job category and expected duration of a posting, which we use as an estimate of the opportunity cost of not searching more intensely. We take that number and multiply by consideration set size.

Table 8: Estimates of Surplus from Survey Conducted in 2023

	Raw Statistics	Reweighted
Markup Relative to Willingness to Accept (WTA)		
Mean	1.24	1.38
Std. Dev.	(0.57)	(0.95)
N:	84	84
Markup Relative to Ex-Post WTA		
Mean	1.16	1.15
Std. Dev.	(0.50)	(0.50)
N:	84	84
Markup Relative to Outside Wage		
Mean	1.50	1.23
Std. Dev.	(1.27)	(1.00)
N:	46	46
Markup on Fixed-Price Contracts Relative to Ex-Post WTA		
Mean	1.87	1.62
Std. Dev.	(2.55)	(2.07)
N:	99	99

Notes: This table presents surplus estimates from a survey of workers on a leading online platform conducted in September and October of 2023. There were 113 responses collected from applicants to jobs posted on the platform. Sample inclusion required workers to have some prior experience, as questions about surplus were framed around their prior jobs on the platform. Respondents were paid \$6 USD upon survey completion. The reweighted column reports inverse-probability weighted statistics of survey participation relative to a random sample of 1,485 profiles with non-zero prior jobs collected from Administrative Support, Design, and Web Development job categories. The probability of being in the sample for the reweighting estimator is computed using a logistic regression of survey participation on the log of the number of prior jobs and the log of the hourly rate displayed on each profile. Markups are computed as the pay rate on the last job relative to either willingness to accept or the respondents' typical outside wage. Willingness to accept for hourly jobs is taken from the question: From the information that you had in the job posting on your last hourly job, what would have been the lowest hourly rate you would have been willing to accept to do the job? For example, if the job paid you \$10 but you would have been willing to do the work for \$9.00 but not \$8.99, then \$9.00 is the lowest hourly rate you would have been willing to accept. The ex-post willingness to accept was elicited immediately afterward using the question: From what you know about the job now that you've worked on it, what would have been the lowest hourly rate you would have been willing to accept to accept to take the job? (Note: The lowest hourly rate may be higher than the rate at which you agreed to work). Workers' outside wages were elicited with two questions, depending on whether they worked exclusively on the platform or not. For those who did not have off-platform jobs, we used the question: If you were working in your local labor market rather than online, what hourly wage (in US dollars) do you think you would be earning? If you would not be paid hourly, please convert your total pay into an hourly wage (in US dollars) by dividing your US-dollar equivalent compensation by the typical hours worked over a pay period. For workers with an outside job, we asked: When you work outside of platforms, what hourly wage (in US dollars) do you earn on average? If you are not paid hourly, please convert your total pay into an hourly wage (in US dollars) by dividing your US-dollar equivalent compensation by the typical hours worked over a pay period. We elicit markups on fixed-price contracts (contracts that pay for work delivered, rather than billed by the hour) by dividing the contract price by an ex-post willingness to accept elicited with the question: From the information that you have about the job now that you've worked on it, what would have been the lowest contract price you would have been willing to accept to take your last fixed price job? For example, if the contract price was \$22 but you would have been willing to do the work for \$20.00 but not \$19.99, then \$20.00 is the lowest total contract price you would have been willing to accept.

Table 9: Counterfactual Changes in Hiring Rates, Postings, and Surplus with a 10 Percent Payroll Tax

	Baseline	No Price Impact on Job Arrival
Panel A: Inexperienced Buyers		
Change in log Bids to Buyers	0.088 (0.006)	0.088 (0.006)
Static Pct Change in Hiring Rates	-0.266 (0.048)	-0.266 (0.048)
Static Pct Change in Buyer Surplus	-0.240 (0.056)	-0.240 (0.056)
Pct Change in Jobs Posted	0.034 (0.031)	0.119 (0.028)
Pct Change in P.V. of surplus	-0.749 (0.052)	-0.366 (0.106)
Panel B: Experienced Buyers		
Change in log Bids to Buyers	0.089 (0.007)	0.089 (0.007)
Static Pct Change in Hiring Rates	-0.236 (0.040)	-0.236 (0.040)
Static Pct Change in Surplus	-0.226 (0.051)	-0.226 (0.051)
Pct Change in Jobs Posted	-0.676 (0.048)	-0.143 (0.031)
Pct Change in P.V. of surplus	-0.721 (0.044)	-0.288 (0.061)
Panel C: Workers		
Change in log Bids to Workers	0.001 (0.007)	0.001 (0.007)
Static Pct Change in Surplus	-0.252 (0.051)	-0.252 (0.051)
Pct Change in Static Surplus with Tax Rebate	0.155 (0.014)	0.155 (0.014)
Pct Change in P.V. of Surplus	-0.594 (0.036)	-0.278 (0.055)
Pct Change in P.V. of Surplus with Tax Rebate	-0.375 (0.049)	0.117 (0.006)

Notes: Estimates of changes in log bids and percent changes in surplus (by buyer experience) under a 10 percent payroll tax counterfactual. 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 (14). Present value calculations are described in the appendix. The percent change in the number of jobs is computed based on opening arrival rates simulating forward wage bids and buyers' endogenous experience. Static worker surplus is the to-worker hourly wage less platform fees multiplied by hiring probabilities. The present value of worker surplus is calculated as to-worker hourly wages \times average hours \times hiring probabilities \times the number of jobs posted monthly. We discount future surplus to the start of the sample. Rows that rebate taxes allocate tax revenue back to workers using lump-sum transfers. The second column holds fixed wage bids under the current regime without accounting for how higher bids change the arrival of future jobs. In this alternative, the buyer's posting frequency, governed by equation (8), simply uses the status quo bids rather than new prices. Standard errors come from 20 block-bootstrap iterations (drawing buyers with replacement).

Appendix

The appendix contains: **A.** Details about the survey instrument and recruitment; **B.** Model and estimation details; **C.** Counterfactual calculation details and analysis of a counterfactual minimum wage; **D.** Analysis of buyer online-offline substitution; **E.** Appendix figures and tables.

A Survey Details

A.1 Survey protocol

The survey was conducted on one of the leading online labor platforms as of 2023. Respondents were recruited by posting a job with the title “[Job Category]: Take a survey about your work”, where [Job Category] was one of Administrative Support, Web Development, Design, or Web Research. The posting contained a description of the study, a consent form, and a link to the Qualtrics site hosting the survey. The consent form made clear that the survey responses would be linked to the worker’s public-facing profile data on the online labor market. Upon providing consent and completing the survey, workers became eligible for a payment of \$6 that was processed through the platform. Payment was offered to all survey respondents who successfully answered two attention check questions that were contained in the survey instrument. The platform automatically converted US dollar payments to local currencies for non-US workers. We estimated that the survey took around eight minutes to complete.³⁴

Workers were eligible to take the survey if they applied for the job and had been hired previously on the platform at least once. There were 40 workers hired in Web Research, 36 in Administrative Support, 21 in Design, and 13 in Web Development. Another six workers were hired under a job titled “Take a compensated research survey” that was used as a test prior to specifying a job category. 113 workers passed the attention checks included in the questions and completed the survey.

Because our recruitment was not random, we also collected a random sample of current workers’ profiles from the platform to compare respondents with other workers who had prior work history. We logged into the site as a buyer and browsed the workers in Administrative Support, Design and Creative, and Web/Mobile/Software Development, which are three of the largest categories of jobs on the platform. Workers’ profiles are paginated within each category, with ten profiles per page. We collected profile data for three random profiles on pages one and two of the paginated search results. From there, we drew random numbers between one and the maximum number of pages for

³⁴This workflow does not violate the platform’s terms of service.

each job category, and collected three random profiles on each page that we sampled. We continued until we had 500 profiles per job category. We were unable to collect data from 15 sampled profiles using this procedure, leaving us with 1,485 of the intended 1,500 profiles that we sought to sample across the three main job categories. The workers who completed the survey were then compared statistically against the random sample. Table 8 summarizes the raw answers given in the sample and also provides statistics that are re-weighted to match the random sample of 1,485 profiles of workers with non-zero prior jobs.³⁵

The survey itself began by collecting background demographics, such as the worker’s country, education, off-platform labor market experience, and experience in online labor markets. We then asked the worker’s hourly rate on their most recent engagement. We also asked about any off-platform earnings, including the actual (or perceived) hourly rate that workers would anticipate receiving if they worked in their local or non-platform labor market. The survey proceeded to a module on how workers set bids when applying for jobs. All respondents saw the same set of questions—the only variations entailed skipping inapplicable questions or randomizing question order. Respondents were asked what the minimum hourly rate would have been for them to accept their last online job before knowing anything about the difficulty of the work. We then followed up by asking whether, ex-post, the minimum rate they would accept would have been different. After a survey was completed, the respondent got a feedback score of five out of five for the completed job.

For the 113 respondents, the median hourly wage rate posted on their profile was \$10, and ranged from \$1 to \$200. The respondents were spread across 39 countries, with the four most frequent being Pakistan, the US, the Philippines, and India. 40% had been on the platform for six months or less, and 10% had been on the platform for longer than four years.

A.2 Survey questions related to markups

The survey questions designed to elicit worker markups on the platform were: “From the information that you had in the job posting on your last hourly job, what would have been the lowest hourly rate you would have been willing to accept to do the job? For example, if the job paid you \$10 but you would have been willing to do the work for \$9.00 but not \$8.99, then \$9.00 is the lowest hourly rate you would have been willing to accept”. The second was, “From what you know about the job now that you’ve worked on it, what would have been the lowest hourly rate you would have been willing to accept to take the job? (Note: The lowest hourly rate may be higher than the rate

³⁵The probability of being in the sample was computed using a logistic regression of survey participation on the log of the number of prior jobs and the log of the hourly rate displayed on each profile.

at which you agreed to work).”

To determine how platform wages relate to off-platform options, we asked a series of additional questions. Those who did not have off-platform jobs were asked: “If you were working in your local labor market rather than online, what hourly wage (in US dollars) do you think you would be earning? If you would not be paid hourly, please convert your total pay into an hourly wage (in US dollars) by dividing your US-dollar equivalent compensation by the typical hours worked over a pay period.” Workers with an outside job were asked: “When you work outside of platforms, what hourly wage (in US dollars) do you earn on average? If you are not paid hourly, please convert your total pay into an hourly wage (in US dollars) by dividing your US-dollar equivalent compensation by the typical hours worked over a pay period.”

To ask about markups on fixed price jobs, we used the question: “From the information that you have about the job now that you’ve worked on it, what would have been the lowest contract price you would have been willing to accept to take your last fixed price job? For example, if the contract price was \$22 but you would have been willing to do the work for \$20.00 but not \$19.99, then \$20.00 is the lowest total contract price you would have been willing to accept.”

B Model and Estimation Details

B.1 Variation with the Instruments

This section provides additional detail about how worker behavior changes with the instruments. The main results in Tables 5 and 6 show that surplus estimates decline by about 7-8% and 6%, respectively when we control for workers’ total number of applications in a month. Workers appear to become slightly more selective about which jobs to apply to when the Dollar to Local exchange rate increases, which can be seen in Figure A.5. A 10% increase in the Dollar to Local rate is associated with a lower average applicant order, by about 1 position (3%). Because expected surplus is declining with applicant order, when the Dollar is worth less, workers appear less likely to apply to jobs as later applicants. However, this source of selection is unlikely to substantially impact our inference with respect to surplus. The marginal change in applicant position over the full range of the data is small. In addition, we condition on the observed choice set in the estimation problem, and Table 3 shows that the first stage relationship continues to remain strong when we condition on application quantity.

B.2 Maximizing the Likelihood

Our estimation algorithm proceeds in two steps. First, we hold fixed a guess of the parameters $\{\rho_k, \lambda_{k\chi o}^{\text{CONSIDER}}, \lambda_{k\chi}^{\text{ARRIVAL}}\}$ and estimate $\{\beta, \alpha, \psi\}$ for each buyer type conditional on this guess using the BHHH algorithm. Second, we fix $\{\beta, \alpha, \psi\}$ and update $\{\rho_k, \lambda_{k\chi o}^{\text{CONSIDER}}, \lambda_{k\chi}^{\text{ARRIVAL}}\}$ given the results in step 1 using a derivative-free Nelder-Mead algorithm. We iterate until convergence, after which we form the joint likelihood and use Matlab's `fminunc` function until the estimates converge in a local neighborhood of the starting values supplied by the iterative procedure.

B.3 Buyer Surplus

B.3.1 Buyer Surplus on a Hire

To compute buyer surplus conditional on hiring, we draw type 1 extreme value random variables for each applicant and the outside option. Using the choice parameters, we calculate the buyers' maximum utility including the draw for the unobservables. If the chosen alternative has the maximum utility, we accept the draw. Otherwise, we reject the draws and continue until the chosen alternative has the maximum utility. After we accept a draw, we compute surplus upon hiring as the difference in wage bids needed to equalize the utility of the chosen option with the outside option. Conceptually, we capture the amount the buyer would need to be compensated to make her indifferent between hiring her preferred applicant and not hiring on the platform. Denoting the draws for the hired worker and outside option as ε_{oj}^D and ε_{o0}^D , realized simulated surplus for a type k buyer is

$$RealizedSurplus_{k\chi} = \exp\left(\frac{X_j\beta_{k\chi} + \varepsilon_{oj}^D - \varepsilon_{o0}^D}{\alpha_k}\right) - w_{oj}. \quad (15)$$

Estimates of type-specific surplus per hour on an opening are weighted by the posterior distribution of buyer types, recovered by Bayes' rule in equation (13), to give a weighted average realized surplus of $\sum_k \hat{\rho}_{ik} RealizedSurplus_{k\chi}$.

B.3.2 Dynamic Surplus Estimates for Buyers

For experienced buyers, the present discounted value of surplus is

$$V_{kE} = \lambda_{kE} \times ExpectedSurplus_{kE} \times Hours_{kE}/r \quad (16)$$

where λ_{kE} is the job posting arrival rate for an experienced buyer of type k , $ExpectedSurplus_{kE}$ 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.³⁶ This says that

³⁶The choice of 8.7% is in line with the interest rate on a Small Business Administration loan during the sample, with the SBA being a plausible source of financing for most of the U.S. buyers in the sample.

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, while 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 contains the term $Pr(Hire|I, k)$, the probability an inexperienced buyer of type k hires and transitions to the experienced buyer value function. Failure to hire leaves the buyer with the inexperienced buyer value function carried forward in time. This can be written as $V_{kI} = Pr(Hire|I, k) \times (ExpectedSurplus_{kI} Hours_{kI} + \frac{1}{1+r} V_{kE}) + (1 - Pr(Hire|I, k)) \frac{1}{1+r} \lambda_{kI} V_{kI}$ which after rearranging gives:

$$V_{kI} = \frac{Pr(Hire|I, k) \times (ExpectedSurplus_{kI} \times Hours_{kI} + \frac{1}{1+r} V_{kE})}{1 - \frac{1}{1+r} (1 - Pr(Hire|I, k)) \lambda_{kI}}. \quad (17)$$

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 (16). The denominator accounts for the fact that buyers who do not hire return in the future based on the arrival rate for inexperienced buyers, given by λ_{kI} in equation (8).

B.4 Worker Bidding With Future Wage Growth in the Relationship

To assess whether anticipated future wage growth during the buyer-worker relationship might alter workers' markups on any one bid and, therefore, affect our surplus estimates, we generalize the model to account for this possibility. Suppose that with probability q the wage is w_{oj} for the entire relationship, and with probability $1 - q$ the average present value of the wage is $w_{oj} \times (1 + g)$, where g is some growth rate that reflects how wages evolve over the relationship. This yields a modified version of the objective function, equation (4), for the worker

$$\begin{aligned} E(U_{oj}(w_{oj})) &= \underbrace{E[\tilde{p}(j)]}_{Pr(Hired)} \times \{q \times \exp(\log w_{oj} - \log(1 + \tau)) + \\ &\quad (1 - q) \times \exp(\log w_{oj} + \log(1 + g) - \log(1 + \tau))\} + \\ &\quad (1 - E[\tilde{p}(j)]) \times c_{oj}. \end{aligned}$$

With this change, the optimal wage bid goes from

$$w_{oj}^* = c_{oj} (1 + \tau) \left(1 + E[\tilde{p}(j)] / \frac{\partial E[\tilde{p}(j)]}{\partial \log w_{oj}} \right)^{-1}$$

to

$$w_{oj}^* = c_{oj} (1 + \tau) \left[\left(1 + E[\tilde{p}(j)] / \frac{\partial E[\tilde{p}(j)]}{\partial \log w_{oj}} \right) \times (1 + g(1 - q)) \right]^{-1}, \quad (18)$$

which is smaller than the original bid because $1 + g(1 - q)$ is positive and is in the denominator.

If $g \times (1 - q)$ is large and workers correctly put weight on the prospect of wage growth over the relationship, this would suggest that our markup estimates on the initial wage bid are overstated. However, workers would still gain surplus over the course of the relationship, in expectation, and some of that surplus would come from expected wage growth.

Using billing data, we find wage growth since the beginning of the data for 6.6% of buyer-worker pairs, either because of wage changes on the current job or because of wage changes over multiple contracts. Conditional on a wage change, the mean (median) hours-weighted growth rate is 22.8% (13.6%). When we apply the adjustment factor of 1.015 in the denominator of the markup equation (for the probability of wage growth multiplied by the mean wage growth), markups over costs on the initial wage remain positive but get slightly smaller, moving from 27% to 25%.

C Counterfactuals

C.1 Calculation Details

When we impose the 10% tax, it is included as an additional wedge between the hourly wage workers receive and the price buyers pay. Our baseline results assume that the composition of workers does not change under the counterfactual tax.³⁷

Workers select the optimal bid to maximize their payoffs (equation (4)), given the composition of the applicant pool, the buyer’s semi-elasticity of demand to wage bids, and the additional tax. The size and composition of the buyer’s consideration set on each posting remains unchanged. A buyer selects the option that maximizes her indirect utility out of her consideration set $\{J_o, 0\}$ given the simulated wage bids, where the probability any worker $j \in J_o$ is chosen is given by equation (2). The wage bids observed on any posted job opening affect the rate at which each buyer posts subsequent jobs, as in equation (8). When buyers post fewer jobs 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 determined by equation (8) up to the cutoff date when the simulated time period ends. Buyers who have prior experience at the start of the period enter the counterfactuals with prior experience. Potential changes in the arrival rate of jobs due to buyers’ failure to gain experience will not impact those who start the sample with prior hiring experience.

³⁷In the analysis in Section C.2, which imposes a minimum wage counterfactual where worker sorting is likely to be more extreme, we conduct robustness checks that relax the assumption that applicant composition is static.

C.2 Imposing a Counterfactual Minimum Wage

This appendix considers the market surplus implications of an hourly wage floor of \$7.00 per hour. Our main counterfactual in Section 5, imposing a 10% tax, reduces surplus mainly through reducing the number of jobs posted. Horton (2022) provides experimental evidence that buyers respond to the imposition of a \$3.00 minimum wage in the platform by posting fewer jobs in the future. We relate our model to this finding by using our structural estimates to examine the impact of a minimum wage on market surplus. We choose a \$7.00 minimum wage to approximate the average hourly minimum wage at the time of the data in the United States, where the majority of the buyers in our data are located.

The wage floor of \$7.00 is included in the workers’ problem as a constraint, which directly binds for the lowest-paid workers in the market and has an indirect effect on the optimal wage bid of workers submitting higher wage bids who are no longer exposed to low-wage competition. We run the same base case as in the main counterfactual, and also the version that sets δ_3 to zero, breaking the link between past wage bids received and job posting.

For this counterfactual analysis, we also illustrate the sensitivity of our findings to alternative assumptions about the composition of job applicants. This adjustment accounts for the fact that the worker’s problem includes the implicit choice of whether or not to apply for a job. It is plausible that workers 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 attract new workers with skills and qualifications that would make them competitive applicants at higher wage levels. We hence want to 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 because the 10% tax counterfactual applies across the board, for all workers, and is less likely to have a disproportionate effect on the participation decisions of workers currently bidding low wages.³⁸ This change allows us to analyze static hiring and market dynamics under conditions where the market itself becomes more appealing for higher-

³⁸To 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 with an equal number of new applications. The new applicants are assigned observable characteristics, X_j , that are a random draw from the distribution of applicants whose original wage bids are above the wage floor. The 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 workers’ first order conditions hold subject to the wage floor constraint, and all 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.

quality applicants under the \$7.00 wage floor.

Table A.5 presents the results of the minimum wage counterfactual. Column 1 is the baseline specification, where job posting depends on past wage bids received and the applicant pool remains constant. Hourly wage bids increase by 34% and 47% to inexperienced and experienced buyers, respectively. The larger increase for experienced buyers reflects the fact that they receive lower wage bids in the data, so more of their applicants are bound by the counterfactual wage floor. Hiring rates on posted jobs fall in response to higher wage bids, by 30% and 34%, respectively. The reduction in hiring rates may seem large given that fewer than 40% of applicants are bound by the floor, but the reason becomes clear after considering which buyers are most affected. Under the wage floor, some openings, like those in non-technical jobs, see extremely large wage bid increases, while jobs that require greater skills experience little change. A large share of the hiring reduction arises for those posting non-technical jobs.

The static buyer surplus on posted jobs falls by less under the wage floor counterfactual than in the 10% tax scenario in the main text. This is because surplus estimates are averaged over all buyers, and the higher-skill jobs, which generate substantial surplus, are largely unaffected by the wage floor. However, the present value of inexperienced buyer surplus falls by 70%, and for experienced buyers falls by 64%. Both reductions arise because the job posting frequency falls for experienced buyers, reducing the future value of being experienced.

Columns 2 and 3 of Table A.5 show the additional results for the \$7.00 wage floor counterfactual. In column 2, job posting rates do not depend on past wage bids ($\delta_3 = 0$). The static changes are the same as in column 1, but column 2 shows a much smaller reduction in jobs posted by experienced buyers when posting does not depend on page wages. The reduction in buyer surplus in column 2 is also much reduced. Column 3 presents the scenario where the applicant pool is adjusted to mirror the characteristics of workers who typically submit bids above the \$7.00 floor. That is, it presents buyers with “better” applicants. The job posting rate here remains insensitive to past bids received, as in column 2. Because these applicants no longer face competition from lower-wage workers, the optimal bids increase more than in the main scenario. However, the number of jobs posted and hiring rates fall by slightly less than in column 2 because buyers perceive applicants to be better quality.

The final rows of Table A.5 show how worker surplus would be affected by the \$7.00 minimum wage counterfactual. The static percentage change in worker surplus in the baseline specification shown in column 1 is -2.6% , reflecting the balance of higher wages on filled jobs and the reduced share of posted jobs that are filled. However, the difference between columns 1 and 2 shows that the impact of the buyer posting dynamics has a large negative effect. Taking into account that buyers

post fewer jobs having received higher previous wage bids reduces the present value of workers' lifetime surplus by 45% when dynamics are present compared to only 7% when they are shut down. The minimum wage results in an increase in worker surplus only when buyer posting dynamics are shut down and the composition of applicants changes.³⁹

Overall, the main insight from the minimum wage counterfactual is similar to that of the 10% tax: Policies that increase hiring costs serve to shrink the size of the market and reduce total surplus going to the supply side.

D Online-Offline Substitution

The counterfactual 10% tax in Section 5 and minimum wage of \$7.00 in Appendix C.2 both lead to a large reduction in the number of jobs posted. However, the counterfactual findings leave open the possibility that buyers substitute away from online hiring to fill more jobs offline in response to the counterfactual policies that increase the costs of hiring on the platform. Any such substitution would generate surplus offline. While we do not have data on whether buyers would post more jobs in other labor markets under the counterfactuals, we provide context for the degree of substitutability between online and offline labor demand via an analysis of online buyers' response to an increase in the cost of offline hiring. We do this by asking what happens to labor demand on the platform when there are exogenous changes in offline wages in buyers' local labor markets. The exogenous change we explore comes from increases in US state-specific minimum wages.

In July 2009, all but ten US states increased their local hourly minimum wage, either because of the 10.6% increase in the federal minimum wage from \$6.55 to \$7.25 per hour that went into effect that month, or as a state-specific increase at the same time as the federal minimum wage increase.⁴⁰ Online job postings were not subject to the same minimum wage regulations or to any increase in wages around this time.

If buyers view online and offline labor as close substitutes, then we would expect a relative increase in the demand for online labor after July 2009 in affected states. We would also expect the demand increase i) to be concentrated in relatively low wage online job categories and ii) to come from those states with high shares of offline workers in low paid jobs that can be done at a distance via telework. For this reason, we focus on non-technical online jobs in this exercise, as

³⁹Unlike the main counterfactual of the 10% tax, with results shown in Table 9, there is no lump sum that could potentially be rebated to workers in this minimum wage counterfactual.

⁴⁰The minimum wage increase on July 24, 2009 is discussed here: https://www.epi.org/publication/mwig_fact_sheet/. The ten states unaffected by the minimum wage increase were California, Colorado, Hawaii, Iowa, Maryland, Massachusetts, Michigan, New Hampshire, Rhode Island, and West Virginia. The federal minimum wage had last increased to \$6.55 per hour on July 24, 2008.

the median hourly wage for online hires by US employers in non-technical categories in 2009 was \$4.44, compared to \$11.7 in technical job categories. Later, in regression analyses, we examine heterogeneous effects in states that had a relatively high share of employment in jobs within \$2.00 of the new minimum wage in teleworkable jobs.⁴¹

A first look at the raw data suggests minimal substitution between online and offline hiring. Figure A.6 plots the log number of hires at the bi-monthly level by buyers located in the group of affected and in the group of unaffected states relative to the group-specific normalized mean number of hires before May 2008. Panel A shows no differential increases in the number of hires in non-technical job categories by buyers located in affected states after the increase in the minimum wage. There is also no apparent increase in hires in technical categories in Panel B.⁴²

We test whether there are significant differences in job postings or hires in US states that increase the minimum wage using a dynamic event study around the July 2009 event. We start by finding the total number of posts and hires in each state in each month over the time period studied. Following Bertrand, Duflo, and Mullainathan (2004), we aggregate these observations into a pre- and a post-event period for each state, averaging over the six months prior to July 2009 and from July 2009, respectively. We estimate:

$$Y_{st} = A_s + B_t + \beta I_s T_t + \epsilon_{st}, \quad (19)$$

where A_s and B_t are state and aggregated time period fixed effects, respectively, I_s is an indicator for the state being affected by the minimum wage increase in July 2009, and T_t is an indicator for the post period. The coefficient β provides a test of whether the states with a minimum wage increase had significant difference in the level of the outcome Y_{st} in the six months after the event relative to the six months prior, compared to unaffected states. In a second specification, we also include interactions of the post period and having a high share of low paid telecommutable work, and a triple interaction of this term with the indicator for being in a state affected by the minimum wage increase. For each of postings and hires, we look at three different outcomes: the total level of each variable, only non-technical postings and hires, and then only postings and hires by buyers who had engaged with the platform prior to April of 2009, and so were aware of it before the minimum wage increase. Each specification has only 102 observations, one for each state and one for DC, in each of the aggregated six-month time periods.

⁴¹State-level variables are taken from the American Community Survey for 2007. Occupation teleworkability is taken from Dingel and Neiman (2020). The states with a higher-than-median share of teleworkable jobs in low wage administrative work in 2009 are Alabama, Arkansas, Florida, Iowa, Idaho, Kansas, Kentucky, Louisiana, Mississippi, Missouri, Montana, Nebraska, New Mexico, North Dakota, Oklahoma, South Carolina, South Dakota, Tennessee, Texas, Utah, West Virginia, and Wyoming.

⁴²The limited online demand response is not due to lack of awareness of the platform, as the patterns are similar for experienced and inexperienced buyers.

The difference-in-differences regression results from estimating equation (19) are shown in Table A.6. There are no statistically significant increases in total job postings or hires in states that increase the minimum wage. Column 1 shows that the number of posts and hires is actually lower in affected states in the post period, although the differences are small. If buyers viewed online labor as a substitute for local employment, we would instead expect to see positive coefficients on the post-period interaction. The interactions in column 2 show there is no significant differential increase in states where we would expect the greatest substitutability—those with high shares of workers near the minimum wage threshold in jobs that can be done remotely. The specifications in columns 3 and 4 find similar results when looking only at postings and hires in non-technical job categories. Columns 5 to 6 restrict the sample to buyers who were active on the platform at least once prior to April 2009. Although the point estimate turns positive in this subsample, it applies only to those states with a high share of low paid teleworkable jobs. The estimates are again very imprecise and we cannot reject the possibility that they are zero.

The results in Table A.6 show no evidence that the increases in local offline wages in July 2009 induced buyers to look for labor on this platform. If cross-price elasticities are approximately symmetric, this evidence suggests that reducing online hiring would not result in additional offline hiring. Other papers make the case for this symmetry. For example, Horton (2022) discusses surveys of buyers and finds that online and offline hiring are only very weak substitutes. In his experiment, the buyers who were unexpectedly and randomly subject to a minimum wage after posting jobs on the platform did not significantly decrease hiring rates but, instead, switched to higher-quality workers.⁴³ 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 counterfactual changes in relative wages studied here would do little to increase offline hiring in the US.

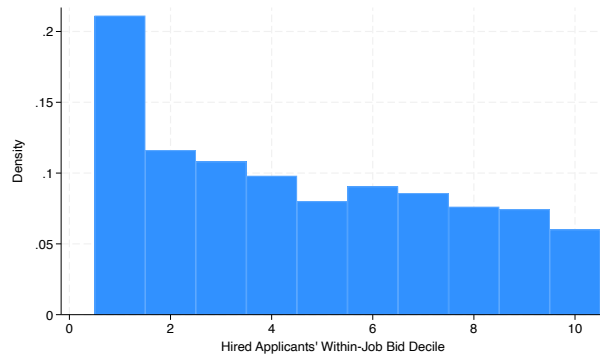
From the data, it is also clear that buyers post task-based jobs rather than hiring for long-term roles, even after they gain hiring experience. That is, the nature of labor demand tends to remain idiosyncratic and the work arrangements look quite different from those seen in traditional offline

⁴³In another paper, Horton, Johari, and Kircher (2021), buyers were asked what they would have done with their most recent project if the platform were not available. Only 15% of employers responded that they would have made a local hire. Online employers report that they are generally deciding among (a) getting the work done online, (b) doing the work themselves, and (c) not having the work done at all. The survey also found that 83% of employers said that they listed their last job opening only on the platform in question.

settings. 27% of jobs posted by inexperienced buyers are expected to last for less than one week; that number is nearly identical (28%) for those with prior hiring experience. 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 upon gaining experience. This evidence suggests that buyers take advantage of being able to post jobs online on an as-needed basis, something that is likely harder to do offline.

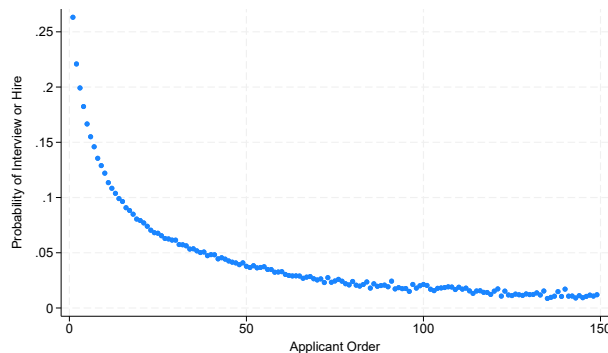
E Appendix Figures and Tables

Figure A.1: Bid Decile for Hired Worker



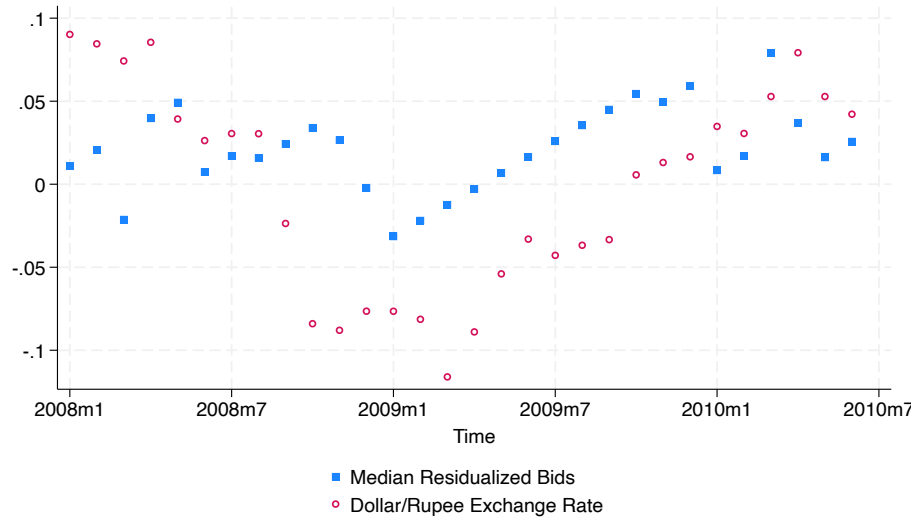
Notes: This figure shows the bid decile of the worker hired for a job when a hire is made, on hourly jobs that receive at least 10 applications.

Figure A.2: Interview or Hire Probabilities by Applicant Order



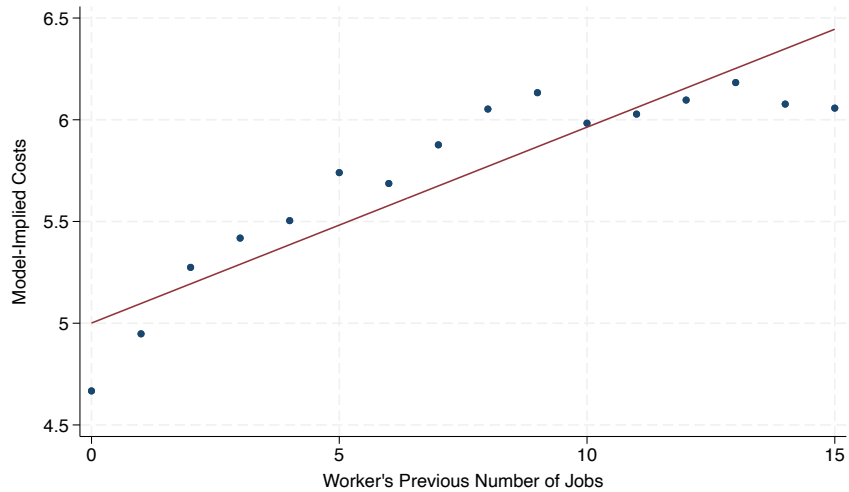
Notes: This figure plots the probability that a buyer either interviews or hires a worker as a function of their applicant order on the job posting.

Figure A.3: Median Residual Log Bids and Detrended Exchange Rates for India.



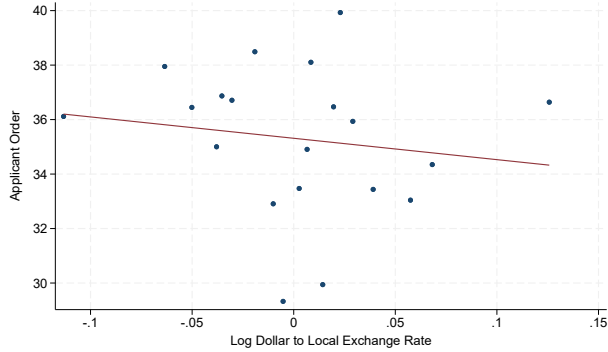
Notes: This figure plots median residual log wage bids from Indian applicants in each month against the log US Dollar to Indian Rupee exchange rate after removing a time trend and setting the series to have mean zero. Log wage bid residuals are net of job category fixed effects and a time trend.

Figure A.4: Workers' Estimated Costs Increase With Prior Jobs on the Platform



Notes: This figure plots the relationship between the number of past online hires and a worker's implied costs from equation (7). The sample contains workers who we can track from having zero to greater than 15 hires in the data. The upward sloping relationship suggests workers' opportunity costs of work on a given job opening are a function of online opportunities on other openings, as the probability of finding online work increases for those who have landed prior jobs.

Figure A.5: Average Applicant Position (Proxying for Expected Surplus) and the Exchange Rate Instrument




Notes: Relationship between average applicant order and the exchange rate instrument. The binscatter nets out country fixed effects and a time trend.

Table A.1: Tests of Sort Orderings and Applicant Interaction

	(1)	(2)	(3)
Applicant Order / 10	-0.0052 (0.0003)	-0.0055 (0.0003)	-0.0056 (0.0003)
Descending Bid Rank (1 is Highest) / 10	-0.0017 (0.0002)		0.0008 (0.0006)
Ascending Bid Rank (1 is Lowest) / 10		0.0018 (0.0002)	0.0025 (0.0006)
Mean DV	0.175	0.175	0.175
N	4379352	4379352	4379352
R^2	.036	.037	.037

Note: This table considers different rules for how buyers may potentially sort applicants when considering who to hire. We evaluate sorting based on wage bids (where a buyer may consider high or low bidders first) and on the sequential arrival of applicants. The dependent variable is an indicator that the buyer interacted with an application. Interaction is coded based on whether a buyer took an active action to evaluate an applicant that is recorded in the platform's database. In the database, we observe whether buyers message an applicant, hire, or select a reason for not hiring, but not all buyers are proactive about selecting reasons for not hiring. The right hand side variables are ranks based on different sorting rules. Because job openings with few applicants mechanically have ranks clustered around one, all specifications have fixed effects for job category - by - the number of applicants to the job. The sample contains only organic, worker-initiated applications. Standard errors are clustered by job opening.

Figure A.6: Log hires in states affected and unaffected by the July 2009 minimum wage increase



Notes: These series show the log number of hires from buyers in states affected and unaffected by a local minimum wage increase in July 2009. Each series is normalized relative to the mean level in the data in months prior to May 2008. Panels A and B split the sample for non-technical and technical jobs, respectively. The solid lines represent affected or treated states, while the dotted lines are for control states.

Table A.2: Tests of Applicant Differences For High and Low Values of the Instruments

	Exchange Rate IV			Competition IV		
	Low	High	P-Value	Low	High	P-Value
Number of Prior Jobs	5.532	5.833	0.140	5.269	5.296	0.671
No Prior Jobs	0.373	0.367	0.458	0.404	0.400	0.308
Log Rate Last Hire	1.136	1.160	0.191	1.068	1.077	0.130
Feedback Score if Non-Zero	4.475	4.458	0.116	4.473	4.474	0.760
Zero Feedback	0.449	0.443	0.505	0.481	0.475	0.202
BA+ Degree	0.369	0.364	0.427	0.350	0.351	0.962
Good English	0.910	0.893	0.000	0.901	0.901	0.892
Agency Affiliate	0.360	0.364	0.551	0.322	0.328	0.003

Notes: This table presents differences in applicant characteristics for low (bottom tercile) and high (top tercile) values of the exchange rate and competition instruments. For the exchange rate instrument, terciles are calculated across country-months. For the competition instruments, terciles are calculated at the job category -by- quarter level. Low values, high values, and the p-value come from the constant, the constant plus high value dummy, and the test statistic on the high-value dummy from a regression. For the exchange rate instrument, the regression includes provider country fixed effects and a country-specific time trend, with standard errors clustered b country. For the competition instrument, the regression contains job category -by- quarter fixed effects, with standard errors clustered by job category-quarter.

Table A.3: Regressions of the Number of New Uploaded Files or Work Samples on Instruments

	(1)	(2)
Exchange Rate	0.029 (0.030)	
Competition		0.002 (0.002)
R-Squared	0.002	0.003
Observations	680000	4428327

Note: Regressions of new items uploaded to workers portfolios or attached in applications on the two instruments. The sample is a worker-week panel for all weeks after a workers first through the 45th week of 2009 (when this particular dataset ends). The exchange rate instrument is the log of the dollar to local currency exchange rate. The competition instrument is the average of the instrument at the job category-week level and is merged into the worker-week panel based on the modal job category for each worker. Column 1 controls for country fixed effects and a time trend. Column 2 controls for modal job category fixed effects and a time trend. Standard errors clustered by worker. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A.4: Demand Parameter Estimates

	Type 1	Type 2	Type 3	Exper
Log Bid	-6.085 (1.641)	1.701 (1.488)	1.009 (1.660)	
Constant	5.307 (1.630)	-1.347 (1.474)	-0.809 (1.626)	
Constant x Experienced	5.242 (3.512)	-2.652 (3.247)	-1.467 (3.345)	
Log Wage on Last Job	1.126 (0.732)	-0.063 (0.105)	0.052 (0.121)	0.000 (0.000)
No Prior Jobs	2.070 (0.574)	-0.599 (0.509)	-0.363 (0.559)	-0.027 (0.025)
Number of Prior Jobs	2.993 (0.853)	-1.040 (0.737)	-0.790 (0.765)	-0.333 (0.135)
BA+ Degree	-0.066 (2.066)	0.185 (1.912)	0.161 (2.024)	-0.317 (0.722)
Agency Affiliate	-0.796 (2.077)	0.327 (1.901)	0.409 (2.017)	-0.163 (0.726)
Agency x No Prior Jobs	0.027 (0.004)	-0.004 (0.004)	-0.003 (0.004)	-0.007 (0.001)
Good English	0.054 (0.058)	0.043 (0.053)	0.004 (0.056)	0.020 (0.032)
Feedback	-0.239 (0.054)	-0.011 (0.050)	-0.073 (0.076)	0.025 (0.031)
Feedback Squared	0.714 (0.251)	-0.144 (0.212)	0.025 (0.275)	-0.023 (0.053)
Feedback Cubed	0.166 (0.114)	0.160 (0.103)	0.038 (0.110)	-0.231 (0.061)

Notes: This table presents selected parameters from the estimates of the buyer choice problem. The first column presents baseline parameters. Columns 2 and 3 report additive interactions for buyer types 2 and 3. Type-shares are reported in the text. The last column reports additive interactions for buyer experience. Each type, however, has a constant that is allowed to shift independently with experience, as displayed in the third row. Unreported parameters are on the control function, an indicator that the worker is experienced but has no feedback, an indicator for buyer initiation, a time trend and a separate trend for technical categories, job category dummies, a spline for applicant order, country dummies for the largest countries, and country-specific time trends. Standard errors are estimated from 20 block-bootstrap iterations of the entire estimation procedure (drawing buyers with replacement).

Table A.5: Counterfactual Changes in Hiring Rates, Postings, and Surplus with a \$7 Wage Floor

	Baseline	No Price Impact on Job Arrival	Labor-Labor Substitution
Panel A: Inexperienced Buyers			
Change in log Bids to Buyers	0.340 (0.004)	0.340 (0.004)	0.491 (0.016)
Static Pct Change in Hiring Rates	-0.303 (0.019)	-0.303 (0.019)	-0.233 (0.062)
Static Pct Change in Buyer Surplus	-0.101 (0.010)	-0.101 (0.010)	0.003 (0.114)
Pct Change in Jobs Posted	0.040 (0.015)	0.125 (0.009)	0.073 (0.038)
Pct Change in P.V. of surplus	-0.701 (0.034)	-0.307 (0.042)	0.031 (0.190)
Panel B: Experienced Buyers			
Change in log Bids to Buyers	0.468 (0.008)	0.468 (0.008)	0.655 (0.014)
Static Pct Change in Hiring Rates	-0.344 (0.020)	-0.344 (0.020)	-0.229 (0.023)
Static Pct Change in Surplus	-0.120 (0.011)	-0.120 (0.011)	0.081 (0.077)
Pct Change in Jobs Posted	-0.648 (0.036)	-0.151 (0.008)	-0.087 (0.043)
Pct Change in P.V. of surplus	-0.636 (0.036)	-0.150 (0.011)	0.153 (0.158)
Panel C: Workers			
Change in log Bids to Workers	0.411 (0.005)	0.411 (0.005)	0.582 (0.013)
Static Pct Change in Surplus	-0.026 (0.018)	-0.026 (0.018)	0.101 (0.043)
Pct Change in P.V. of Surplus	-0.450 (0.032)	-0.065 (0.023)	0.064 (0.049)

Notes: Estimates of changes in log bids and percent changes in surplus (by buyer experience) under a \$7 wage floor counterfactual. 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 (14). Present value calculations are described in the appendix. The percent change in the number of jobs is computed based on opening arrival rates simulating forward wage bids and buyers endogenous experience. Static worker surplus is the to-worker hourly wage less platform fees multiplied by hiring probabilities. The present value of worker surplus is calculated as to-worker hourly wages average hours hiring probabilities the number of jobs posted monthly. We discount future surplus to the start of the sample. The third column replaces applicants whose original bids were less than \$6.30.

Table A.6: Tests of Postings and Hires Changes in States that Raise the Minimum Wage

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Job Postings by State-Month						
Min Wage Increaser x Post	-12.621 (23.217)	-17.554 (28.675)	-5.254 (12.756)	-8.489 (15.675)	12.089 (17.084)	15.235 (21.036)
Post x Low Wage Tele		-44.125 (27.638)		-25.521 (14.913)		26.854 (20.928)
Increaser x Post x Low Wage Tele		36.145 (29.659)		21.689 (16.570)		-22.293 (21.701)
Mean DV	59.8	59.8	30.1	30.1	24.2	24.2
N	102	102	102	102	102	102
R^2	.964	.965	.945	.947	.919	.922
Panel B: Hires by State-Month						
Min Wage Increaser x Post	-0.327 (3.707)	-0.663 (4.647)	0.257 (2.102)	0.133 (2.658)	5.185 (5.628)	6.923 (6.891)
Post x Low Wage Tele		-6.667 (4.414)		-3.625 (2.450)		9.792 (6.887)
Increaser x Post x Low Wage Tele		4.621 (4.900)		2.392 (2.855)		-9.339 (7.098)
Mean DV	14.9	14.9	7.8	7.8	6.6	6.6
N	102	102	102	102	102	102
R^2	.98	.981	.972	.973	.886	.891
Non-tech jobs only:	No	No	Yes	Yes	No	No
Pre-event buyers only:	No	No	No	No	Yes	Yes

Note: The dependent variable in panel A is the average number of job postings by state and month in the six months prior to the minimum wage event. The dependent variable in panel B is the mean number of hires. Regressions in Columns 1, 3, and 5 are two-way fixed effects estimates with state and time fixed effects. Columns 2, 4, and 6 add interaction effects indicating that the state had an above-median share of low wage workers in teleworkable jobs prior to the minimum wage event date. This measure is computed from the 2007 ACS data, where low wage work is defined as earning less than \$9.25 per hour and teleworkable jobs are coded at the occupation level from Dingel and Neiman (2020). Column 3 and 4 restrict to postings in non-technical job categories. Columns 5 and 6 restrict to postings from buyers who were active prior to April of 2009 who were aware of the platform prior to the minimum wage changes. Standard errors are clustered by state.