American Economic Review 2025, 115(6): 1857–1895
https://doi.org/10.1257/aer.20221189

Who Benefits from Online Gig Economy Platforms?†

By Christopher T. Stanton and Catherine Thomas*

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. Workers retain a significant share of the surplus because demand-side search frictions and worker differentiation reduce direct competition. 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. (*JEL* C83, F31, J22, J23, J31, J64, M51)

importance of the phenomenon

An estimated 160 million global workers are registered on online labor market platforms (Kässi and Lehdonvirta 2018; Kässi, Lehdonvirta, and Stephany 2021). As in other gig settings, there are many more workers than jobs at any point in time (Autor 2001; Prassl 2018; Fisher 2022). While numerous studies have explored the economics of alternative work arrangements (Katz and Krueger 2019; Collins et al. 2020; Mas and Pallais 2020), little is known about the overall surplus generated by online labor markets or the distribution of these benefits. Despite this, concerns about excessive competition have prompted policy proposals aiming to reclassify freelancers as employees, 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.

what we know and what we don't

Our analysis addresses three main questions: How much surplus do buyers and workers each receive from online labor markets? If workers receive surplus despite their relative abundance, what prevents competition from pushing workers to their

key questions

* Stanton: Harvard Business School, NBER, and CEPR (email: christopher.t.stanton@gmail.com). Thomas: London School of Economics, CEP, and CEPR (email: c.m.thomas@lse.ac.uk). Chinhui Juhn was the coeditor for this article. 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, Jesse Shapiro, Kathryn Shaw, Andrei Shleifer, and Nathan Seegert for helpful comments. We thank Shira Aronson and Cameron Greene for excellent research assistance. IRB approval was obtained from Harvard University (23–1,163).
† Go to https://doi.org/10.1257/aer.20221189 to visit the article page for additional materials and author disclosure statement(s).

reservation wages? 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 \$4.24 per hour worked. About 46 percent of this total goes to workers, a substantial share when compared to traditional employment arrangements (Card et al. 2018).

Second, both worker differentiation and limited buyer search allow workers to earn positive surplus even in the face of intense competition. Although workers are numerous, buyers view them as imperfect substitutes. For example, we estimate that a standard deviation change in an index of workers' characteristics has an impact on labor demand that is equivalent to a 38 percent 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. The reason is that job posting and hiring rates decline, offsetting 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 online labor market 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 affect 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 follows 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 on similar jobs that are posted by other buyers. A bidder who observes little competition for similar alternative jobs infers a greater likelihood of being hired for another online job, increasing the value of the outside option and reducing the incentive to submit a low bid on the focal job. On the other hand, when there are many applicants for similar jobs, the bidder anticipates heightened competition for alternative online positions. In addition, a worker choosing a bid on the focal job anticipates that other applicants will shift their bids in response to their own perceived likelihood of competition for other positions. We measure competition using the mean number of applications in the first 24 hours for jobs in the same category and week but posted by other buyers. To isolate idiosyncratic variation, we then use fixed effects to remove persistent variation at the job category level and aggregate shocks at the week level. This instrument shifts all workers' bids, not just those in countries with exchange rate variation, and allows us to identify buyers' price elasticity relative to the option of not hiring on the platform.

Our survey results confirm that workers consider both local exchange rates and the number of prior applicants when determining their bids. Several tests indicate that exclusion restriction violations are limited. For example, applicant résumé characteristics vary little with the instruments, suggesting the instruments primarily affect wage bids rather than the distribution of worker quality. The instruments are also unrelated to a proxy for worker application effort, and all of our estimates are similar when we control for worker application quantity.

Our results show there are gains from trade for both sides of the market. We estimate that buyers gain an average of \$2.27 of surplus per hour when they hire. Buyer surplus upon hiring is defined as the dollar amount that would yield 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 \$1.97 per hour. Workers' markups over their costs average 28 percent. Over our 30 months of data—from an era when online labor markets were smaller than today—total surplus to hired workers was about \$4.52 million. The total surplus to buyers was about \$5.71 million.[1] These 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 \$2.05 million dollars, if we assume that each application takes about 5 minutes and we net out a per application cost of \$0.59, which comes from our estimates of workers' hourly costs. 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 their search costs leaves buyers with \$1.97 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 bid variation maps to model-predicted markup variation. One clear prediction relates to how bids change with applicant order: Because buyers are observed to evaluate applications in the order they are received, a later applicant knows he will face more competition given that a buyer considers him. 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 3.6 percent lower as the sixtieth job applicant compared to when he is among the first 10 applicants.

lower bids for later applicants

[1] The platform collects 10 percent of all revenues, which, at average wage bids, yields platform revenues of \$1.14 per hour. We do not observe the platform's marginal costs of servicing transactions, and platform surplus is excluded from our surplus calculations.

Model-predicted markups fall by more than 2 percent between the first 10 applicants and the sixtieth. Second, in our 2023 survey, reported markups range from 23 percent to 38 percent, depending on the markup construct described in how we elicited responses. These estimates are broadly consistent with the 27 percent average markup from the model.

Finally, we use the estimates to assess whether traditional employment regulations can redistribute surplus toward the supply side of the market. We study the counterfactual introduction of an additional 10 percent 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 W-2 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 nonstandard work arrangements, complementing 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 et al. 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

[2] We do not find evidence that surplus in online labor markets arises because buyers of online labor have substituted away from regulated offline employment. In Supplemental Appendix D, difference-in-differences estimates show there is no change in online job posts or hiring in US 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 percent 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.

[3] Related work focusing on ride-sharing 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 et al. 2019; Katsnelson and Oberholzer-Gee 2021), and assessed offline spillover benefits (Gorback 2020).

[4] Online labor markets for microtasks, 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 et al. (2020).

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.

Finally, our results are relevant for the ongoing debate about how to regulate the gig economy. Legislative proposals, such as the PRO Act, which passed the United States House of Representatives before dying in the Senate, could extend traditional labor regulations to online labor platforms, potentially altering surplus distribution. 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 I describes our worker survey, introduces the administrative data, and provides summary statistics that guide our modeling choices. Section II presents the model and estimation strategy. Section III presents the main results and calculations of buyer and worker surplus. Section IV contains the counterfactual analysis. Section V concludes.

I. The Setting, Data, and Descriptive Statistics

A. *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 main 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 percent of workers' earnings during our main data sample. The typical hourly job requires 16 hours of work per week, and 61 percent of postings are for jobs expected to last less than three months. On average, hired workers bill about 71 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

[5] The closest work answering a regulatory question about online labor markets is Horton's (2025) study of a $3 wage floor. His experiment was carefully designed to measure hiring responses holding fixed preexisting 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 (2025) also shows that buyers exposed to the minimum wage treatment subsequently posted fewer jobs.

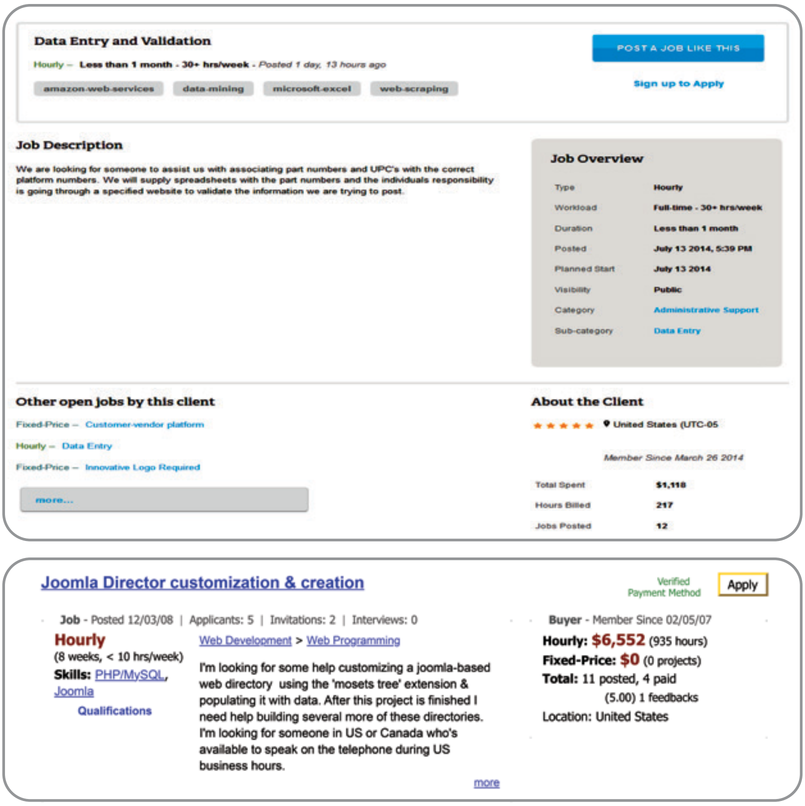

Figure 1. Two Examples of Job Postings

Notes: The first job post is a screenshot from the platform for a nontechnical job posting, taken in July 2014. The second job posting is for a technical job during our administrative data sample, taken from the WayBack Machine (with the name of the platform in top right corner hidden). The job was posted on December 3, 2008. Applicants can see the job has received five applications. Screenshot adapted with permission from the platform.

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 who 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]

[6]Supplemental Appendix Figure A1 presents some initial evidence that workers are differentiated, showing that on openings with a hire, buyers choose applicants in the lowest bid decile only slightly more than 10 percent of the time.

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 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 192,627 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 nontechnical job categories, which make up approximately equal shares of all job posts.[7]

B. *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. Sample inclusion required that

[7] The technical job categories are web development, software development, and networking and information systems. The nontechnical job categories, in order of frequency, are administrative support, writing and translation, design and multimedia, sales and marketing, business services, and customer service.

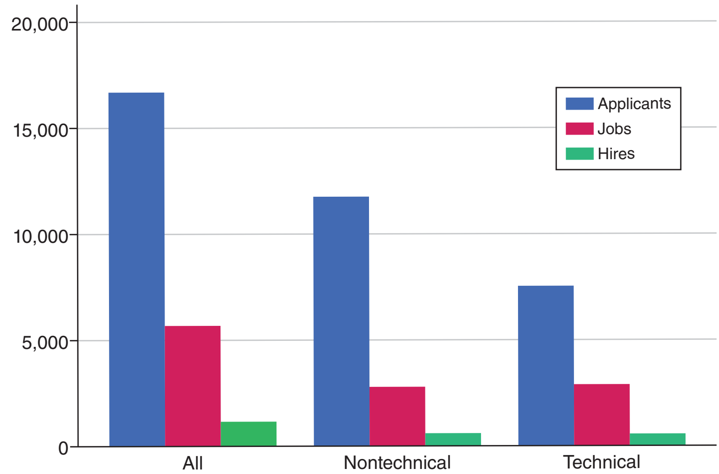

Figure 2. Job Posts, Applicants, and Hires, 2008–2010

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 of the administrative data. The left-hand columns are for all hourly jobs, the center columns are for nontechnical job openings, and the right-hand columns are for technical job openings. Because some applicants apply for both technical and nontechnical jobs in a month, the number of unique applicants in nontechnical and technical jobs sums to more than the total number of unique applicants.

workers had been hired at least once before. Supplemental 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 6 prior jobs, with a range from 1 to 144. When asked about their sources of income, 41 percent reported that "work through platforms is my exclusive source of income," and 35 percent 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 et al. 2015), where 48 percent said that freelancing full-time was their sole source of income and 38 percent 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 percent of workers saying that they always submitted different bids or tailored bids to different openings. A further 45 percent indicated they sometimes did. Only 10 percent selected the option, "I always submit the same bid," which was the first option among those presented in the question.

[8] In our data, 28 percent 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 nonpaid 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.

TABLE 1—DESCRIPTIVE STATISTICS ABOUT BUYERS AND WORKERS

	Buyers		Workers	
	Mean	SD	Mean	SD
In the United States	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	67,566		192,627	

Notes: The sample includes buyers and workers active on the platform from January 2008 to June 2010. Buyers and workers classified as being in other 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. Nontechnical jobs are those in administrative support, business services, customer service, design and multimedia, sales and marketing, and writing and translation.

The next set of questions explored why workers vary bids. Ninety-six percent 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 percent said they considered other on-platform work opportunities, and 70 percent said they considered off-platform worker opportunities when forming bids. To preview our identification strategies: 82 percent of respondents from countries that do not use the US dollar indicated that the value of the dollar relative to their local currency affects their bids; 83 percent 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]

C. *Administrative Data*

Table 1 presents descriptive statistics about the 67,566 buyers and 192,627 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 United States, and 75 percent are located in G10 countries with high levels of per capita income. Only 6 percent of buyers come from India or the Philippines.[10]

[9] Gee (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 (2019) experiment, the number of applicants was always shown on this platform, allowing us to examine differences in competition holding fixed that workers observe the number of prior applicants.

[10] Eighty-nine percent 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

Roughly half of buyers (56 percent) post their first job in a technical job category, a share that is similar for platforms today (Stephany et al. 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 percent are in India or the Philippines, while 31 percent 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 percent 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] The statistics relate to the mean characteristics of the applicant pools. Column 1 shows that posts receive an average of 26 applications, the majority of which arrive within 24 hours of posting. The mean hourly wage bid is \$11.43, with a standard deviation of \$6.54.

In the model, we allow for differences in both demand and job posting frequency by buyer experience. Columns 2 and 3 of Table 2 split the sample by buyer experience. Column 2 first shows that posts by inexperienced buyers, who have made zero prior hires, receive slightly fewer applications and also receive slightly higher wage bids. The applicants to these jobs have received slightly worse feedback, or no feedback, when they have been hired in the past. Applicants to inexperienced buyers have also been hired slightly fewer times before, although they received higher mean wages on these jobs. Despite these and other minor differences, however, the table shows there is substantial overlap in applicant characteristics for inexperienced versus experienced buyers.

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 percent of workers were hired at least once, which is in line with the 10 percent share of workers who had positive earnings on large online labor markets in 2013 (Kuek et al. 2015) and the 9.6 percent 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 raised market efficiency overall, but they likely reduce the number of workers a buyer considers, increasing the market power of those who are considered.

[11] The sample includes hourly billed, public job postings that do not entail repeat hires from prior online jobs and that receive more than three applications. The filtering rule based on a minimum number of applications selects competitive openings, preventing us from picking up jobs that go to a worker whom the buyer may have targeted from an offline source or a preexisting, unobserved online relationship (2,211 jobs). Such noncompetitive 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 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.

Table 2—Descriptive Statistics about Job Postings by Buyer Hiring Experience

	All postings	Postings by buyers with zero prior hires	Postings by buyers with 1+ prior hire
Number of applicants	26.14	23.52	28.74
	(32.85)	(26.67)	(37.83)
Number of applicants in first 24 hours	18.55	16.46	20.63
	(22.54)	(18.27)	(25.94)
Mean applicant wage bid	11.43	11.86	11.00
	(6.54)	(6.68)	(6.35)
Mean applicant share with good feedback	0.37	0.35	0.38
	(0.18)	(0.18)	(0.17)
Mean applicant share with no feedback	0.43	0.45	0.41
	(0.21)	(0.22)	(0.20)
Mean prior hires per applicant	6.60	6.10	7.09
	(4.94)	(4.63)	(5.19)
Mean applicant wage on last hire (if applicable)	10.18	10.57	9.79
	(5.39)	(5.34)	(5.40)
Mean applicant share in the United States	0.11	0.11	0.11
	(0.15)	(0.15)	(0.15)
Mean applicant share with a BA+ degree	0.35	0.35	0.35
	(0.15)	(0.15)	(0.15)
Probability of filling job	0.20	0.15	0.26
	(0.40)	(0.35)	(0.44)
Months between postings	1.15	1.53	0.89
	(2.58)	(3.24)	(1.97)
Number of posts	170,556	84,999	85,557

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 percent 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, autopopulated 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 bachelor's 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.

To motivate allowing demand and job posting frequency to vary by buyer experience, Table 2 then shows that experienced buyers' probability of filling a job is 26 percent compared to 15 percent for experienced buyers. Experienced buyers also post jobs more frequently, making them responsible for around half of all postings on the platform. Only 27 percent 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.

D. *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 percent of the openings in the data.

buyer-applicant interactions

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. Supplemental Appendix Table A.2 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 probability of an interaction is 17.5 percent. Every 10 positions an applicant falls reduces the probability of an interaction by 0.57 percentage points (column 3), or 3.3 percent. Application order is far more important than wage bid ranking for determining whether a buyer interacts with an applicant. Supplemental Appendix Figure A.2 shows that this bias toward early applicants means applicant 45 receives an interview or hire request with less than 5 percent probability, a substantial reduction from the more than 20 percent 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 are likely 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 percent of worker-months and 61 percent 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 18.6 log points greater than the lowest bid. In technical job categories, workers who lack bid variation submit bids that are 5 percent higher, on average, than those who vary their bids. The gap is even larger in nontechnical job categories. The survey

[12] The 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 have taken 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 toward the subset of applicants who apply early, with a higher probability of "meeting" these applicants (Wright et al. 2021).

[13] Since 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).

[14] Consumers 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 et al. (2018) assume a buyer considers all the listings seen up through the last one observed to be browsed in a search query on eBay. In a recent working paper, Yu (2023) studies ranked advertisements on Amazon, where the data show that buyers search in order of rank.

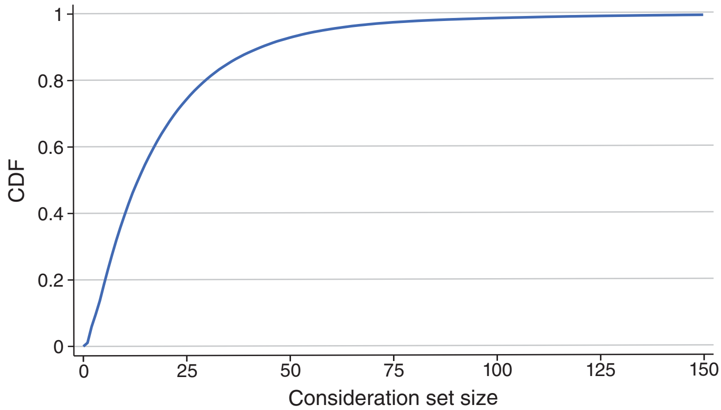

Figure 3. Cumulative Distribution of Consideration Set Sizes

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 the buyer messaged, hired, or indicated a reason for not hiring. Consideration set sizes are unobserved for about 19 percent of all openings.

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.

II. 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.

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 have heterogeneous preferences or 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 Supplemental Appendix B.

A. *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 set of all 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 résumé characteristics, denoted $\mathbf{X}_j$, 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, Hortaçsu, 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]

[15] In 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.

[16] Workers 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 0.66 percent of workers have more than 2 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.

B. *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 who yields the highest indirect utility per unit of wage, which is $\frac{\exp(\mathbf{X}_j'\boldsymbol{\beta}_{k\chi} + \varepsilon_{oj})}{(w_{oj})^{\alpha_k}}$ for each $j \in J_o$, or she chooses the outside option of not hiring for the opening. $\mathbf{X}_j$ is a vector of worker j's observable characteristics that also contains a constant. $\boldsymbol{\beta}_{k\chi}$ 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

$$\mathbf{X}_j'\boldsymbol{\beta}_{k\chi} - \alpha_k \log(w_{oj}) + \varepsilon_{oj} \geq \mathbf{X}_l'\boldsymbol{\beta}_{k\chi} - \alpha_k \log(w_{ol}) + \varepsilon_{ol} \tag{1}$$

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

$$p(j \mid J_o, k, \chi) = \frac{\exp(\mathbf{X}_j'\boldsymbol{\beta}_{k\chi} - \alpha_k \log(w_{oj}))}{1 + \sum_{j \in J_o} \exp(\mathbf{X}_j'\boldsymbol{\beta}_{k\chi} - \alpha_k \log(w_{oj}))}. \tag{2}$$

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 order so that applicant 1 is the first 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),

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

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 $\eta_{k\chi}(|J_o|) = \exp(-\lambda_{k\chi o}^{Consider} |J_o|) - \exp(-\lambda_{k\chi o}^{Consider}(|J_o| + 1))$. We let the parameter $\lambda_{k\chi o}^{Consider}$ be a function of buyer type, experience, and whether the job is a technical or nontechnical posting.

C. *Supply: Workers' Optimal Wage Bids*

If worker j is hired for job o, the buyer pays w_{oj}, and the worker receives $w_{oj}/(1+\tau) = \exp(\log w_{oj} - \log(1+\tau))$, 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

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

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

The worker uses the information she has to form $E\left[\tilde{p}(j)\right]$. She knows buyer experience, χ, her own characteristics, $\mathbf{X}_j$, 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 application will be the j'th received, the relevant hiring function from equation (3) would be $E\left[p(j)\right] = \sum_{|J_o|=j}^{|\mathcal{J}_o|} \eta_{k\chi}(|J_o|)p(j \mid J_o, k, \chi)$, 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

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

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

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

[17] Note 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.

where $\left(1 + \frac{E[\tilde{p}(j)]}{\partial E[\tilde{p}(j)]/\partial \log w_{oj}}\right)^{-1}$ 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 semielasticities as

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

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

D. *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 $\widetilde{\log(w_{oj})} = \frac{1}{O-1}\sum_{o=1}^{O-1}\left(\overline{\log(w_{t,Category})} - t_{Category} - \mathbf{1}_{Category}\right)$ 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 $\lambda_{k\chi}^{Arrival}$ is strictly positive, so we model the process as

$$\log\left(\lambda_{k\chi}^{Arrival}\right) = \delta_{1k} + \delta_2 \mathbf{1}_{\{\chi>0\}} + \delta_3 \widetilde{\log(w_{oj})}, \tag{8}$$

where the vector δ_{1k} contains buyer type-specific constants, δ_2 shifts the mean arrival hazard for experienced buyers, and δ_3 is experienced buyers' posting frequency sensitivity to past wage bids.

E. *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 first instrument provides exogenous variation in the value of her offline local labor market outside option, and our second instrument varies the value from working on another online job.

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

[18] The 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).

opportunity cost of offline work.[19] 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., $1 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. Supplemental Appendix Figure A.3 illustrates the time series variation in mean residual log wage bids and exchange rates that underpins this instrument for India, the largest non-US worker source country. Bids and exchange rates move together.[20]

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 currencies that move relative to the US dollar, the second instrument shifts all bids relative to the buyer's outside option. It does so by capturing variation in the intensity of competition for workers' alternative online job opportunities. 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, including other jobs that are similar to the focal job. The instrument is the log of the average number of applications received in the first 24 hours after posting for all jobs in the same job category and the same week, excluding jobs for the focal buyer.[21] If workers see that similar jobs are attracting many applicants in a given week, they anticipate relatively intense competition for these other jobs, making it harder for them to find work on another online job. Workers' ability to observe the extent of competition across similar open jobs varies the value of the part of their outside option that comes from different online jobs, causing them to tailor their bids on the focal job upward or downward.

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 $\mathbf{Z}_{oj}$, or worker characteristics, $\mathbf{X}_j$. This entails taking the residuals from the first-stage regression,

$$\log(w_{oj}) = \mathbf{Z}_{oj}'\gamma_1 + \mathbf{X}_j'\gamma_2 + \nu_{oj}, \tag{9}$$

[19] The exchange rate series is compiled using data from IMF (2007–2010) and Datastream International (2007–2010) accessed in 2015–2016 .

[20] This 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.

[21] To 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 when a candidate applies, holding fixed the average competition level in the job category across time and the average competition level in other job categories at the same time.

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

	(1)	(2)	(3)
Exchange rate instrument	0.087	0.076	0.085
	(0.010)	(0.010)	(0.011)
Competition instrument	−0.067	−0.063	−0.086
	(0.009)	(0.009)	(0.009)
Worker applications / month		−0.001	
		(0.000)	
R^2	0.613	0.614	0.552
Observations	4,458,737	4,458,737	4,458,737
F clustered on posting	89.50	74.24	92.18
F clustered on worker	48.17	41.63	47.62
Excludes résumé characteristics	No	No	Yes

Notes: 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 résumé 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.

and including them as controls in the choice model. As described above, the matrix $\mathbf{X}_j$ includes all of the nonwage 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.[22]

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 how worker behavior changes with the instruments in Supplemental 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. Column 2 of Table 3 shows the results when including the number of applications by the worker as a control in equation (9). If a nonrandom subset of workers changes their

[22] The F-statistic clustering at the job posting level is 89.50. The F-statistic clustering at the worker level is 48.17.

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 0.087 to 0.076 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 to be driven by changes in the composition of applicants.[23] 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 résumé characteristics from the regression, the first-stage coefficient estimate goes from 0.087 to 0.085 for the exchange rate instrument, while the competition instrument estimate goes from -0.067 to -0.086. 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, Supplemental Appendix Table A.3 assesses 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 0.912 to 0.893. For the competition instrument, the only difference is for the agency affiliate indicator, which increases from 0.319 to 0.323.[24] 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 likely varies 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 or examples to share in their résumés or with buyers directly. These items tend to include portfolios of past work, including writing samples, code, or graphics. Supplemental Appendix Table A.4 shows no evidence that the instruments change how many items workers share in their applications.

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.

[23] In 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 et al. (2021) find larger rates of pass-through of exchange rates to wages.

[24] Agencies 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.

F. *Estimation and Inference*

This section outlines how we form the likelihood and estimate the model. We discuss the separate subparts 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.

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 $CF_{oj} = \hat{\nu}_{oj}$. Second, we calculate choice probabilities conditional on buyer type, k, and the consideration set, J_o, as in equation (2), while including CF_{oj},

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

We then form the joint likelihood for type k, yielding

$$(11)\quad L_k = \prod_o \left[p\left(j\,|\,J_o,k,\chi\right)^{(y_o=j)} \times \lambda_{k\chi o}^{Consider} \exp\left(-\lambda_{k\chi o}^{Consider}|J_o|\right)\right]^{J_o Observed}$$
$$\times \left[p\left(j\,|\,\hat{J}_o,k,\chi\right)^{(y_o=j)}\right]^{J_o Unobserved}$$
$$\times \left[\lambda_{k\chi}^{Arrival}\exp\left(-\lambda_{k\chi}^{Arrival} t\right)\right]^{o\in\{2,\ldots,O-1\}}$$
$$\times \left[\exp\left(-\lambda_{k\chi}^{Arrival} T\right)\right]^{o=O}.$$

The term $p\left(j\,|\,J_o,k,\chi\right)^{(y_o=j)}$ 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.[25] 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.[26] The third line captures the probability of waiting t months to post job o from the time when

[25] We right-censor the size of the consideration set in the likelihood to 250, which is the top 0.2 percent of the data. This entails replacing $\lambda_{k\chi o}^{Consider}\exp\left(-\lambda_{k\chi o}^{Consider}|J_o|\right)$ with $\exp\left(-\lambda_{k\chi o}^{Consider}|J_o|\right)$.

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

the previous job, $o - 1$, was posted.[27] 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 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.[28] This yields

$$L = \sum_{k=1}^{3} \rho_k L_k. \tag{12}$$

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 résumé 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 $\lambda_{k\chi o}^{Consider}$ change the choice set discontinuously when the consideration set is unobserved and must be simulated. Supplemental 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

$$\hat{\rho}_{ik} = \frac{\rho_k L_k}{\sum_k \rho_k L_k}. \tag{13}$$

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 contribute 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 who 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.

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

[28] We stopped adding types at 3, as 1 of the 3 types is estimated to be only around 4 percent 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.

Supply Side Estimation.—Worker markups and costs are computed using data on wage bids and the demand-side estimates, as set out in Section IIC, 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, $E\left[\tilde{p}(j)\right]$, 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.

Under this model, the weight that workers put on private information is small, at 0.09 when buyers are inexperienced and 0.08 when buyers are experienced. The worker's perceived hiring probability in equation (5) puts over 90 percent of the weight on the average buyer type rather than on any signals they have about individual buyers.

G. *Buyer and Worker Surplus*

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. Supplemental Appendix B.3 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

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

The expression $1 - p\left(0\,|\,J_o,k,\chi\right)$ 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. Supplemental Appendix B.3 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 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.

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.

III. Results

A. *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 percent) 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 -4.38. Their job fill elasticity is smaller, at -3.24, because this elasticity assumes wage bids rise for all applicants, reducing substitution across considered applicants. Based on the parameter estimates on worker characteristics (see Supplemental Appendix Table A.5), the standard deviation in worker characteristics scaled by the log wage coefficient, $\mathbf{X}_j'\beta_{k\chi}/\alpha_k$, is 0.38. Buyers appear to perceive substantial worker differentiation. Type 2 buyers consider an average of 15.38 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 percent of the jobs by inexperienced buyers and only 38 percent of jobs by experienced buyers.

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	0.76	0.20	
	(0.00)	(0.01)	(0.00)	
Panel B. Elasticities and demand features				
Job fill elasticity	−4.53	−3.24	−3.80	−3.61
	(0.99)	(0.66)	(0.74)	(0.61)
Considered applicant elasticity	−6.12	−4.38	−5.04	−4.84
	(1.98)	(1.22)	(1.37)	(1.16)
SD of $\mathbf{X}'\beta$ (log wage units)	0.38	0.38	0.36	0.38
	(0.03)	(0.01)	(0.01)	(0.02)
Panel C. Applicant consideration				
Mean applicants considered	20.51	15.38	21.29	18.09
	(0.55)	(0.21)	(0.27)	(0.08)
Panel D. Job posting frequency				
Mean jobs per month (experienced)	4.23	0.09	0.70	1.11
	(0.40)	(0.00)	(0.03)	(0.05)
Sensitivity of posts to past bids				−2.02
				(0.42)
Panel E. Share of job posts by type				
Share of inexperienced job posts	0.08	0.65	0.27	
	(0.00)	(0.01)	(0.00)	
Share of experienced job posts	0.18	0.38	0.43	
	(0.01)	(0.01)	(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 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 résumé characteristics × 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 its 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).

The first and third columns of Table 4 summarize the estimates for the two other buyer types, making up 4 percent and 20 percent 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 −6.12 over applicants and −4.53 at the job level. This type considers 20.51 applicants on average and, once experienced, posts 4.23 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 percent of all jobs by experienced buyers.

Table 5—Estimates of Buyer Surplus

	Baseline		App. control		Full consideration	
	Mean	SD	Mean	SD	Mean	SD
Panel A. Surplus on filled jobs						
Surplus per hour	2.27	4.03	2.12	3.55	2.89	4.26
	(0.73)	(2.72)	(0.65)	(1.64)	(0.84)	(1.53)
Panel B. Surplus on all jobs						
E(Surplus per hour)	0.73	0.84	0.67	0.77	0.49	0.45
	(0.19)	(0.21)	(0.19)	(0.21)	(0.12)	(0.11)
Panel C. Lifetime value						
Lifetime surplus (inexper)	381	935	354	888	298	1,132
	(101)	(442)	(100)	(397)	(77)	(369)
Lifetime surplus (exper)	5,612	7,191	5,312	6,937	6,415	7,864
	(2,377)	(5,015)	(2,164)	(4,351)	(1,685)	(2,264)

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 Supplemental Appendix B.3. 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 percent annual rate. Standard errors are estimated from 20 block-bootstrap iterations of the entire estimation procedure (drawing buyers with replacement).

The fourth column of Table 4 shows the opening-weighted mean results. The wage bid elasticity for an applicant is -4.84, the average consideration set size is 18.09, and the average number of jobs posted per month is 1.11. 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 -2.02. Buyers are less likely to use the platform if they have been exposed to high wage bids in the past. The standard deviation of predicted average log wage bids is 0.037, so a standard deviation increase in past bids reduces future job posting rates by about 7.6 percent.

B. *Buyer Surplus*

Table 5 presents the 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 \$2.27 per hour, with a standard deviation of \$4.03. 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 \$0.73, 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 an 8.7 percent annual rate (see Supplemental Appendix B.3 for calculation details). For the typical inexperienced buyer, the expected lifetime value of platform surplus is \$381. Lifetime surplus for buyers that hire at least once is much higher, at \$5,612. 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 IIID.

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 nonrandomly 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 \$2.89, but expected surplus falls to \$0.49.

C. *Worker Markups*, *Surplus*, *Model Validation*, *and Application Costs*

Model Estimates.—Table 6 presents worker surplus estimates. Panel A focuses on hired workers. The mean worker surplus per hour is \$1.97, with a standard deviation of \$1.63. Hired workers' markups average 28 percent over costs, and hourly costs average \$7.08. Supplemental 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 percent for all applicants, regardless of whether they are hired, and costs of \$7.06. The \$9.08 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 nontechnical 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 \$2.18. Applicants 11 to 20 have a lower expected surplus, at \$1.15, while these figures are \$0.73 and \$0.55 for applicants 21 to 30 and 31 to 40, respectively. After applicant 40, the expected surplus is less than 5 percent 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 similar. Finally, when we assume all workers are considered, in columns 5 and 6, estimates of workers' surplus and markups remain positive but are

Table 6—Surplus and Markups for Applicants and Hired Workers

	Baseline		App. control		Full consideration	
	Mean	SD	Mean	SD	Mean	SD
Panel A. Estimates for hired workers						
Surplus per hour	1.97	1.63	1.85	1.52	1.53	1.27
	(0.43)	(0.36)	(0.42)	(0.35)	(0.38)	(0.32)
Markups	1.28	0.02	1.26	0.02	1.20	0.02
	(0.08)	(0.01)	(0.07)	(0.01)	(0.07)	(0.02)
Costs	7.08	5.71	7.20	5.82	7.52	6.08
	(0.41)	(0.37)	(0.40)	(0.35)	(0.37)	(0.31)
Panel B. Markups and costs for all applicants						
Markups	1.27	0.01	1.24	0.01	1.19	0.02
	(0.08)	(0.01)	(0.07)	(0.01)	(0.06)	(0.00)
Costs	7.06	9.08	7.18	9.24	7.49	9.65
	(0.41)	(0.54)	(0.40)	(0.52)	(0.37)	(0.48)

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 Baseline and App. control models that include consideration sets. Markups and costs estimates use a simple own-bid semielasticity in the model that assumes all applicants are fully considered. Estimates of markups and costs use the nonlinear least squares procedure to weight the semielasticities across different types, based on our estimate of a bidders' private information. This procedure is described in Section IIF. Panel B displays estimated costs and markups over costs for all applicants. See Table 8 for aggregate surplus estimates that take into account hours worked on each job. Standard errors come from 20 block-bootstrap iterations (drawing buyers with replacement).

smaller than the baseline estimates. Mean surplus per hour on a hire is about $1.53, and markups average 20 percent. 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 8 percentage points (40 percent) relative to a simpler model that does not account for limited consideration. Still, in all models, workers benefit from the market.

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 2.4 percentage points 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 61 to 70 have log bids that are 3.5 percent lower than the first 10 applicants. Over half of the reduction in bids can be explained by changes in estimated markups with applicant

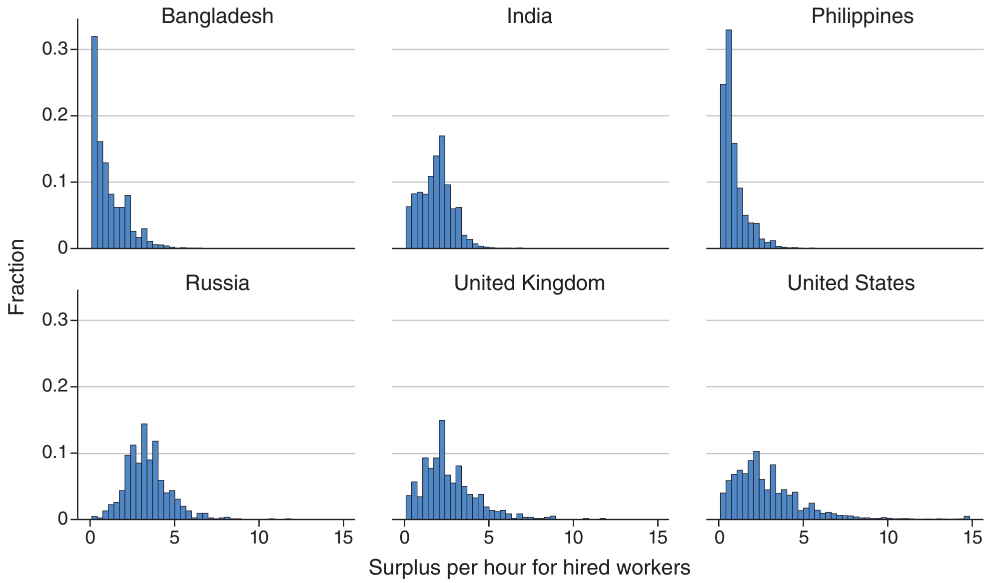

FIGURE 4. DISTRIBUTION OF WORKER SURPLUS ACROSS COUNTRIES

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

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 97 percent of job openings have consideration sets that consist of fewer than 80 applicants, as shown in Figure 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 Supplemental 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.[29] Differences between actual hourly wages and willingness to accept are summarized in Table 7. 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 1,488 worker profiles.[30]

[29] The 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 similarly worded and is given in Supplemental Appendix A.2.

[30] The 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

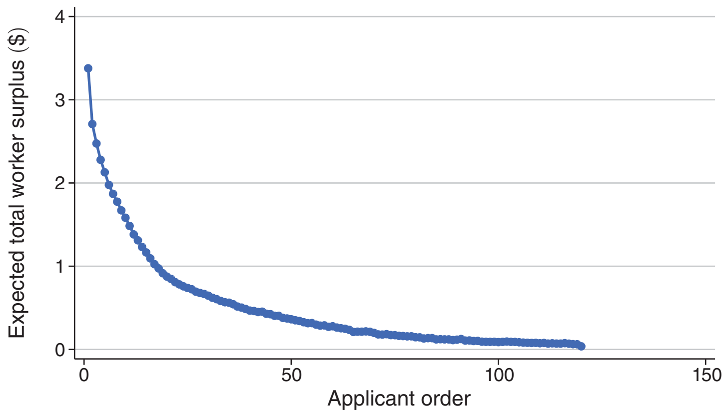

Figure 5. Expected Worker Surplus as a Function of Applicant Order

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 0 for later applicants.

The first row of Table 7 reports that workers earned an hourly wage that was 24 percent 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 percent 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 percent in the raw data and 15 percent 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 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 7. Markup estimates relative to off-platform options are again large, averaging 23 percent 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 percent 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

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 Supplemental Appendix A.1.

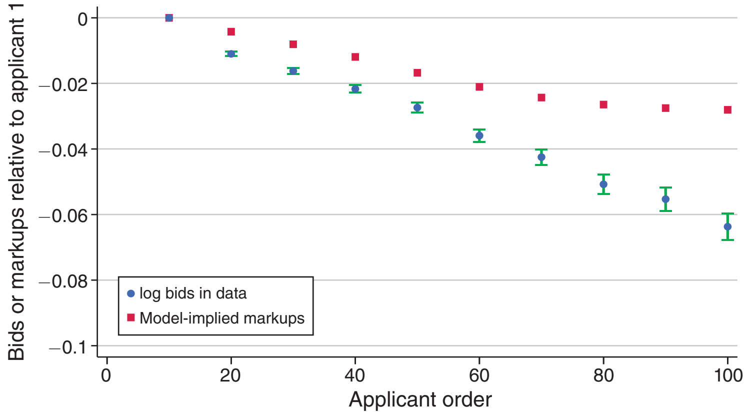

Figure 6. Worker Wage Bids and Applicant Order for Jobs in the Same Week

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.

next-best alternative wage, are substantially lower than the hourly wages or fixed prices workers actually earned on recent jobs.

D. *Total Surplus, Worker Application Costs, and Buyer Search Costs*

Table 8 shows that the total surplus earned by workers who are ever hired during our sample has a mean of $320 per worker and a standard deviation of $859. Workers' aggregate gross surplus totals $4.52 million. The second row shows that mean surplus per worker across all applicants, not just those who are hired, falls to $23.47. From Table 1, 8 percent 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 percent of workers' hourly cost. We believe this estimate is conservative, as Figure 5 shows that 8.3 percent of workers' hourly costs correspond to the ratio of expected surplus to hourly costs for the fortieth 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 percent 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 $97.12 per worker, totaling $1.37 million in the aggregate. Across all workers, the average is $13.61 per worker,

Table 7—Estimates of Surplus from Survey Conducted in 2023

	Raw statistics	Reweighted
Markup relative to willingness to accept (*WTA*)		
Mean	1.24	1.38
SD	(0.57)	(0.95)
Observations	84	84
Markup relative to ex post WTA		
Mean	1.16	1.15
SD	(0.50)	(0.50)
Observations	84	84
Markup relative to outside wage		
Mean	1.50	1.23
SD	(1.27)	(1.00)
Observations	46	46
Markup on fixed price contracts relative to ex post WTA		
Mean	1.87	1.62
SD	(2.55)	(2.07)
Observations	99	99

Notes: This table presents surplus estimates from a survey of workers on a leading online platform conducted between September and October 2023. There were 113 responses collected from a mix of direct invitations and 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 upon survey completion. The reweighted column reports inverse-probability-weighted statistics of survey participation relative to a random sample of 1,488 profiles with nonzero prior jobs collected from administrative support, design, and web development job categories. The probability of being in the 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. Markups are computed as the 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 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 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 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 work exclusively on platform or not. For those who do 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 US dollars 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 8—Total Surplus Accounting for Application and Search Costs

	Surplus			Search costs		
	Mean	SD	Total (mil.)	Mean	SD	Total (mil.)
Panel A. Workers						
Workers with any hires	319.97	859.11	4.52	97.12	175.90	1.37
	(93.85)	(309.72)	(1.02)	(6.30)	(10.91)	(0.07)
All applicants	23.47	247.21	4.52	13.61	57.97	2.62
	(6.04)	(81.73)	(1.02)	(0.88)	(3.55)	(0.15)
Panel B. Buyers						
Openings with hires	163.75	290.00	5.71	21.42	32.81	0.75
	(52.29)	(195.50)	(1.84)	(5.83)	(8.78)	(0.21)
All openings	33.49	146.83	5.71	21.92	32.63	3.74
	(10.70)	(90.15)	(1.84)	(5.88)	(8.70)	(1.01)

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.

with a total of \$2.62 million. Subtracting these application costs from gross surplus gives total net surplus of \$3.15 million for hired workers and \$1.90 million for all applicants. Although only 8 percent of workers are hired, their surplus is large enough to offset not just their own application costs but those of the 92 percent of applicants who do not find jobs.

Several other points suggest that these calculations are reasonable even if the average worker takes more than five 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 percent from the first to the fifth hire on the platform (see Supplemental Appendix 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.[31] 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 8 shows that buyers who hire have search costs per opening that total about \$21.42, with a standard deviation of \$32.81. In total, search costs for those who hire total about \$0.75 million. When considering all openings, total search costs are \$3.74 million, or \$21.92 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 \$1.97 million.

IV. Policy Counterfactual

Our estimates of gig economy worker surplus are informative for debates about regulatory changes targeting worker well-being. 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 W-2 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 percent 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.[32] 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. Supplemental Appendix C details the calculations.[33]

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

[31] The 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 set up interviews and make decisions out of a set of applicants who receive attention.

[32] The 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.

[33] Supplemental Appendix 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.

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 percent and 8.9 percent, respectively. Hiring rates on posted jobs fall by 27 percent and 24 percent. Surplus to buyers on each posted job (labeled "Static pct change in buyer surplus") falls by over 20 percent for both groups.

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 4 percent), job postings by experienced buyers fall by 67 percent. 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 percent. Inexperienced buyers also see a 75 percent 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 lessens 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 26 percent 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 percent 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 percent tax on all hires is sufficient to more than offset the reduced hiring rate, making workers 16 percent better-off. However, the present value of net worker surplus remains negative after the rebate, falling by 37.0 percent 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 percent

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	0.088
	(0.003)	(0.003)
Static pct change in hiring rates	−0.269	−0.269
	(0.054)	(0.054)
Static pct change in buyer surplus	−0.244	−0.244
	(0.063)	(0.063)
Pct change in jobs posted	0.036	0.120
	(0.036)	(0.035)
Pct change in P.V. of surplus	−0.748	−0.372
	(0.083)	(0.137)
Panel B. Experienced buyers		
Change in log bids to buyers	0.089	0.089
	(0.003)	(0.003)
Static pct change in hiring rates	−0.238	−0.238
	(0.049)	(0.049)
Static pct change in surplus	−0.230	−0.230
	(0.060)	(0.060)
Pct change in jobs posted	−0.672	−0.144
	(0.070)	(0.038)
Pct change in P.V. of surplus	−0.719	−0.292
	(0.069)	(0.080)
Panel C. Workers		
Change in log bids to workers	0.001	0.001
	(0.003)	(0.003)
Static pct change in surplus	−0.255	−0.255
	(0.054)	(0.054)
Pct change in surplus with tax rebate	0.155	0.155
	(0.011)	(0.011)
Pct change in P.V. of surplus	−0.593	−0.281
	(0.060)	(0.058)
Pct change in P.V. of surplus with tax rebate	−0.370	0.117
	(0.066)	(0.005)

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 Supplemental 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. Rows that rebate taxes allocate tax revenue back to workers using lump-sum rebates. The second column holds fixed wage bids without accounting under the current regime without accounting for how higher bids change the arrival of future jobs. Standard errors come from 20 block-bootstrap iterations (drawing buyers with replacement).

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.[34]

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 Supplemental 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 the across-state design has 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 nontechnical 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).

V. 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 analysis indicates that the platform generates gains from trade of $4.24 per hour worked, with hired workers capturing 46 percent of the surplus on average. Applicants have local market power, enabling them to charge markups of more than 25 percent 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.

REFERENCES

Abaluck, Jason, and Abi Adams-Prassl. 2021. "What Do Consumers Consider before They Choose? Identification from Asymmetric Demand Responses." *Quarterly Journal of Economics* 136 (3): 1611–63.

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*, edited by Avi Goldfarb, Shane M. Greenstein, and Catherine E. Tucker, 219–50. University of Chicago Press.

[34] When studying the impact of minimum wages in Supplemental Appendix C.2, we find that changes in surplus are negative 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.

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." *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): 1802–25.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics* 119 (1): 249–75.

Brinatti, Agostina, Alberto Cavallo, Javier Cravino, and Andres Drenik. 2021. "The International Price of Remote Work." NBER Working Paper 29437.

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?" Unpublished.

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): 2735–94.

Clemens, Jeffrey, Lisa B. Kahn, and Jonathan Meer. 2018. "The Minimum Wage, Fringe Benefits, and Worker Welfare." NBER Working Paper 24635.

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." Unpublished.

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.

Datastream International . 2007–2010. *Exchange Rate Statistics in International Financial Statistics*. (accessed June 27, 2016).

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan. 2018. "Consumer Price Search and Platform Design in Internet Commerce." *American Economic Review* 108 (7): 1820–59.

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." Unpublished.

Gee, Laura K. 2019. "The More You Know: Information Effects on Job Application Rates in a Large Field Experiment." *Management Science* 65 (5): 2077–94.

Ghani, Ejaz, William R. Kerr, and Christopher Stanton. 2014. "Diasporas and Outsourcing: Evidence from oDesk and India." *Management Science* 60 (7): 1677–97.

Gorback, Caitlin. 2020. "Your Uber has Arrived: Ridesharing and the Redistribution of Economic Activity." Unpublished.

Gray, Mary L., and Siddharth Suri. 2019. *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Houghton Mifflin Harcourt.

Hall, Jonathan V., John J. Horton, and Daniel T. Knoepfle. 2021. "Pricing in Designed Markets: The Case of Ride-Sharing." Unpublished.

Hamermesh, Daniel S. 1986. "The Demand for Labor in the Long Run." In *Handbook of Labor Economics*, Vol. 1, edited by Orley C. Ashenfelter and Richard Layard, 429–71. Elsevier.

Honka, Elisabeth, Ali Hortaçsu, and Matthijs Wildenbeest. 2019. "Empirical Search and Consideration Sets." In *Handbook of the Economics of Mar- keting*, Vol. 1, edited by Jean-Pierre Dubé and Peter E. Rossi, 193–257. Elsevier.

Horton, John J. 2010. "Online Labor Markets." In *Internet Network Economics, 6th International Workshop, WINE 2020; Lecture Notes in Computer Science*, Vol. 6484, edited by Amin Saberi, 515–22. Springer.

Horton, John J. 2017. "The Effects of Algorithmic Labor Market Recommendations: Evidence from a Field Experiment." *Journal of Labor Economics* 35 (2): 345–85.

Horton, John J. 2021. "The Ruble Collapse in an Online Marketplace: Some Lessons for Market Designers." NBER Working Paper 28702.

Horton, John J. 2025. "Price Floors and Employer Preferences: Evidence from a Minimum Wage Experiment." *American Economic Review* 115 (1): 117–46.

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." NBER Working Paper 29445.

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*, edited by Gordon H. Hanson, William R. Kerr, and Sarah Turner, 71–108. University of Chicago Press.

IMF. 2007–2010. *Exchange Rate Archive* https://www.imf.org/external/np/fin/data/param_rms_mth.aspx.

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: 241–48.

Kässi, Otto, Vili Lehdonvirta, and Fabian Stephany. 2021. "How Many Online Workers Are There in the World? A Data-Driven Assessment." SSRN. http://dx.doi.org/10.2139/ssrn.3810843.

Katsnelson, Lura, and Felix Oberholzer-Gee. 2021. "Being the Boss: Gig Workers' Value of Flexible Work." Harvard Business School Working Paper 21-124.

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): 132–46.

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.

Mas, Alexandre, and Amanda Pallais. 2017. "Valuing Alternative Work Arrangements" *American Economic Review* 107 (12): 3722–59.

Mas, Alexandre, and Amanda Pallais. 2020. "Alternative Work Arrangements." *Annual Review of Economics* 12: 631–58.

Moreno, Antonio, and Christian Terwiesch. 2014. "Doing Business with Strangers: Reputation in Online Service Marketplaces." *Information Systems Research* 25 (4): 865–86.

Nosko, Chris, and Steven Tadelis. 2015. "The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment." NBER Working Paper 20830.

Ozimek, Adam, and Christopher Stanton. 2022. "Remote Work Has Opened the Door to a New Approach to Hiring." *Harvard Business Review*, March 11. https://hbr.org/2022/03/remote-work-has-opened-the-door-to-a-new-approach-to-hiring.

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." *Review of Economic Studies* 83 (2): 810–54.

Stanton, Christopher T., and Catherine Thomas. 2020. "The Gig Economy beyond Local Services and Transportation." *CESifo Forum* 21 (3): 21–26.

Stanton, Christopher T., and Catherine Thomas. 2025. *Data and Code for: "Who Benefits from Online Gig Economy Platforms?"* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E214341V1.

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 and Society* 8 (2): 20539517211043240.

Train, Kenneth E. 2009. *Discrete Choice Methods with Simulation*. 2nd ed. Cambridge University Press.

Varian, Hal R. 2007. "Position Auctions." *International Journal of Industrial Organization* 25 (6): 1163–78.

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." Unpublished.