Bidding for Talent:

Equilibrium Wage Dispersion on a High-Wage Online Job Board*

Nina Roussille Benjamin Scuderi

February 2023

Click here for the most recent version.

Abstract

This paper proposes a novel procedure for adjudicating between models of firm wage-setting conduct. Using granular data from a large online job search platform for the U.S. tech sector, we first estimate labor supply to differentiated firms without imposing restrictive assumptions on firm behavior. We then use those estimates to formulate a test of conduct based on exclusion restrictions. On average, workers are willing to pay 14% of their salary to enjoy a 1-s.d. improvement in non-wage amenities, with between-worker dispersion in preferences of a similar magnitude. Models incorporating strategic interactions between firms and tailoring of wage offers to workers' outside options are rejected in favor of simpler models featuring near-uniform markdowns. Misspecification has meaningful consequences: while our preferred model predicts average markdowns of 18%, more complicated models predict average markdowns of 26% (∼50% larger). Implied patterns of between- and within-firm productivity dispersion also differ markedly across models of conduct.

---

*Nina Roussille: Massachusetts Institute of Technology, nroussil@mit.edu. Benjamin Scuderi: Princeton University, bscuderi@princeton.edu. We are indebted to our advisors Patrick Kline, Christopher Walters, and Matthew Backus for invaluable support and guidance. We thank David Berger, Sydnee Caldwell, David Card, Claudia Macaluso, Alan Manning, Yotam Shem-Tov, Isaac Sorkin, Hadar Avivi, Kaveh Danesh, and Arlen Guarin for helpful comments and discussions. We also thank participants at the SITE Micro and Macro of Labor Markets conference, the 4th IDSC of IZA Workshop on Matching Workers and Jobs Online, and the UC Berkeley Labor Seminar. We gratefully acknowledge financial support from the Institute for Research on Labor and Employment, the UC Berkeley Opportunity Lab Labor Science Initiative, the Strandberg Family Graduate Fellowship Fund, and the Washington Center for Equitable Growth. The views expressed here are our own and do not necessarily reflect those of Hired.com. This manuscript was not subject to prior review by any party, except to verify that it does not contain any confidential data or information, as per the research contract signed at the outset of this project.

1 Introduction

Why do the wages earned by observably similar workers often vary widely across firms (Card et al., 2018)? A host of factors have been proposed to explain this dispersion: productivity (Abowd et al., 1999; Gibbons et al., 2005), non-wage amenities (Rosen, 1986; Mas and Pallais, 2017; Sorkin, 2018), and, increasingly, labor market power (Manning, 2005; Berger et al., 2022; Azar et al., 2022). Because most studies use data on equilibrium matches, they must rely on strong assumptions to decompose this dispersion in wages into the relative contributions of each of these factors.[1] A particularly important assumption is the form of firm wage-setting *conduct*: how firms determine which workers to hire, and how much to pay them.

When labor markets are perfectly competitive, only one form of wage-setting conduct can prevail in equilibrium: firms will equate the wage to the marginal revenue product of labor. But when labor markets are imperfectly competitive, a variety of forms of wage-setting conduct may prevail. Despite this, most studies estimate a *particular* model of imperfect competition and propose a reduced form test of that alternative relative to the perfect-competition null. In practice, this means that prior studies make untested assumptions about key aspects of firm wage-setting conduct, like whether firms interact strategically or the extent to which firms know workers' preferences. These assumptions then become key ingredients in the estimation of the size of markdowns and the analysis of welfare and efficiency.

Yet, different modes of conduct imply markedly different conclusions about the sources of wage dispersion and the extent to which firms exercise market power. For example, models in which firms interact strategically predict that larger firms should have larger markdowns. This would further imply that observed wage premiums at larger firms understate true differences in productivity across firm sizes. By contrast, models without strategic interactions need not imply differential markdowns by firm size, ceteris paribus (Boal and Ransom, 1997). Similarly, some modes of conduct give rise to compensating wage differentials between firms that offer different non-wage amenities, while others do not. Erroneous assumptions about the form of conduct can therefore lead to severely biased inferences about welfare and efficiency.

This paper provides direct evidence about the nature of firms' wage-setting conduct by developing a testing procedure to adjudicate between non-nested models of

[1] For instance, a form of random matching is often assumed: given a set of equilibrium wages, workers have no control over the vacancies they are matched to. An assumption of this kind is necessary when workers' choice sets are not measured. However, erroneous inference of these choice sets can introduce substantial bias (Barseghyan et al., 2021).

firms' behavior in the labor market. In particular, we focus on two sets of alternatives relevant to ongoing debates in the literature: first, whether firms compete strategically (Berger et al., 2022; Jarosch et al., 2021), and second, whether firms tailor wage offers to workers' outside options (Postel-Vinay and Robin, 2002; Flinn and Mullins, 2021). Our analysis builds on the modern Industrial Organization literature studying product markets (beginning with Bresnahan (1987) and recently reviewed by Gandhi and Nevo (2021)). At a high level, our strategy is a labor market analog of the marginal cost estimation procedure of Berry et al. (1995): given estimates of labor supply, applying an assumption about firm conduct immediately reveals implied equilibrium markdowns and therefore firms' willingness to pay for candidates. Consequently, in the first step of our analysis, we propose a novel technique for estimating the labor supply of workers to differentiated firms. Following Berry and Haile (2014) and Backus et al. (2021), we then use these model-implied markdowns to construct a test that can adjudicate between models of conduct.

To disentangle labor supply from labor demand without imposing restrictive assumptions on the underlying model of firm behavior, it is necessary to observe the choice sets of workers over jobs. This is typically not possible: matched employer-employee data, for instance, only records the realised transitions of workers between firms. To overcome these data limitations, we leverage the unique matching process on a large, high-stakes online job board. Specifically, on this platform, candidates do not directly apply to jobs—rather, firms looking to fill vacancies submit "bids" on candidates. A bid contains a description of the vacancy as well as an indication of how much the firm would be willing to pay the candidate (hereafter "the bid salary"). Candidates decide whether or not to interview with a firm based on its bid. This setting has several advantages. First, because candidates can only enter the recruitment process at firms that bid on them, we are able to measure the full set of options they choose from on the platform. Second, because we observe the candidates' decisions to accept or reject firms' bids, we can cleanly infer candidates' revealed preferences over firms. Last, our data on bids reveal detailed variation in firms' willingness to pay for candidates that extends beyond just those the firm ultimately hires.

Armed with these data, we turn to the analysis of worker preferences over firms. In a first step, we propose a novel method for estimating the non-wage amenity values candidates associate with firms. Our estimator ranks firms by aggregating the revealed preferences of candidates via a recursion is similar in spirit to Sorkin (2018). In particular, the estimated amenity value of any firm depends on the estimated amenity values of the firms it was revealed-preferred to: for a firm to rank highly, that firm must be revealed-preferred by workers over other highly-ranked firms. Importantly, our estimator flexibly models both the vertical differentiation (between-firm differences in amenity values common to all candidates) and horizontal differentiation (within-firm differences in amenity values across candidates) of firms. In contrast to existing estimates of amenity values, we neither assume that all candidates share the same (mean) ranking of amenities, nor that candidates' (mean) rankings are a deterministic function of their demographics. Instead, we describe candidates' preferences as a mixture over types, each with a unique mean ranking of firms, where the distribution of types can depend upon candidate characteristics. Modelling the correlation of types with observable characteristics allows us to test whether firms tailor offers to candidates based on the predictable component of their preferences in a later step.

Next, we propose a general blueprint for analyzing labor demand that allows us to adjudicate between many non-nested models of firm wage-setting conduct. Each model of conduct defines a unique mapping between labor supply and the marginal revenue product of labor (MRPL). Plugging in our first step estimates, we then invert these mappings to recover the match-specific markdowns implied by each alternative conduct assumption. In order to adapt models of conduct to our data, we analogize the behavior of firms on the platform to that of bidders in a large online auction marketplace: firms compete against each other by bidding for workers' talent. We draw upon insights from the empirical auction literature (e.g. Guerre et al., 2000; Backus and Lewis, 2020) to define an equilibrium concept, establish the identification of markdowns, and propose a method for estimating those markdowns. To test between the various models of conduct, we implement the Vuong non-nested model comparison test (Vuong, 1989; Rivers and Vuong, 2002). The logic of the Vuong test is simple: when comparing two alternative models, the one that is closer to the truth should "fit" better. Here, as in Backus et al. (2021) and Duarte et al. (2023), model "fit" is determined by an exclusion restriction: instruments that shift markdowns but that do not affect labor productivity should not be correlated with the model-implied MRPL recovered from our inversion. Instruments that generate differential shifts in markdowns across models can therefore be used to adjudicate between those models.

Our initial set of findings focuses on labor supply. Our preferred estimates of labor supply describe preferences as a mixture over three types of workers, therefore rejecting that preferences are well-described by a single ranking of firms. We document substantial vertical differentiation of firms on the platform: the average worker is willing to pay 14% of her desired salary to enjoy a standard deviation increase in firm amenities. However, horizontal variation is just as important—the average standard deviation in valuations of amenities across coworkers at the same firm is also 14%. Importantly, the existence of this large and predictable horizontal variation in preferences may lead to substantial firm market power in equilibrium. Indeed, if firms actively tailored wage offers to candidates on the basis of predictable horizontal variation in preferences, they could impose significant differential markdowns on the workers that most prefer them.

We use those estimates to implement our procedure for comparing models of firm behavior. As a baseline, we are able to resoundingly reject the perfect competition model against all possible imperfect competition alternatives. However, in every version of our test, models that assume firms ignore strategic interactions in wage setting significantly outperform models that incorporate strategic interactions. This finding has significant implications for our conclusions about the size of wage markdowns—under the preferred, monopsony model, we find markdowns of 18.2% on average, while the oligopsony model would have implied average markdowns of 25.8%. We also find large differences between models in implied productivity dispersion across firms: in the preferred model, the firms with the best amenities are 3% more productive than the firms with the worst amenities while under the oligopsony alternative, that difference is 8.5%. A simple variance decomposition exercise further highlights the contrasts between the two models. The monopsonistic competition assumption attributes almost none of the variation in bids to variation in markdowns: 91% of bid variation is driven by the systematic component of match productivity, while 9% is due to the idiosyncratic component. The oligopsony alternative, by contrast, apportions 10% of the variation in bids to variation in markdowns, 78% to variation in the systematic component, and 12% to variation in the idiosyncratic component.

We then turn to testing whether firms are type predictive and document that firms do not take advantage of the significant predictable variation in firm-specific labor supply when making hiring decisions. This is especially striking in the context of online labor markets that ostensibly seek to reduce information frictions in the search and matching process. This finding also has significant implications for the labor market: under oligopsony, had firms been type-predictive, the offers they would have made to the workers who value their amenities the most would have been marked down 2.6 percentage points more than the offers they makes to workers who value them the least. Finally, counterfactual simulations suggest that imperfect competition exacerbates gender gaps on the platform relative to a price-taking baseline. However, these exercises indicate that blinding employers to the gender of candidates would only lead to modest reductions in gender gaps.

This paper contributes to several strands of literature. First, our paper is related to a growing literature that employs tools from industrial organization to study the contribution of firms to labor market inequality in equilibrium. Card et al. (2018) and Lamadon et al. (2022) consider models in which firms are assumed to be monopsonistically competitive: firms internalize upward-sloping labor supply, but do not interact strategically. Berger et al. (2022) and Jarosch et al. (2021), on the other hand, write down models of non-atomistic firms that compete in local oligopolies. Our main contribution to this literature is to explicitly formulate a procedure for discriminating *between* these different modes of firm conduct, rather than assuming a single mode of conduct, building on the industrial organization literature that tests conduct in product markets (Bresnahan, 1989; Nevo, 2001; Berry and Haile, 2014; Backus et al., 2021; Duarte et al., 2023). We also focus on a single labor market in which it is likely that conduct of all firms is well-approximated by a single model, rather than applying our model to a national labor market. In this way, our study is related to a long tradition of single-industry studies in labor economics (Freeman, 1976; Staiger et al., 2010; Goldin and Katz, 2016).

Next, our paper contributes to the literature on the estimation of non-wage amenities and their role in wage dispersion (Rosen, 1986; Mas and Pallais, 2017; Wiswall and Zafar, 2018). In this literature, our paper most relates to Sorkin (2018), who applies tools from numerical linear algebra to matched employer-employee data in the U.S. and estimates search models that incorporate dispersion in non-wage amenities of firms. Taber and Vejlin (2020) and Lagos (2021) also use data on equilibrium matches to infer amenity values from the realized flows of workers across firms. By contrast, we observe the full set of options available to each worker on the platform, and therefore estimate amenity values by aggregating candidates' revealed preferences over these options without imposing restrictive assumptions on firm behavior.

Our paper also contributes to a broader literature exploring imperfect competition in labor markets (Boal and Ransom, 1997; Manning, 2005; Bhaskar et al., 2002). A number of recent studies have examined the relationship between measures of market structure—typically, concentration measures like the Herfindahl–Hirschman Index (HHI)—and wages across markets in order to gauge the importance of imperfect competition (Azar et al., 2020; Schubert et al., 2021; Arnold, 2021; Yeh et al., 2022). These analyses echo the "Structure-Conduct-Performance" (SCP) paradigm (Robinson, 1933; Chamberlain and Robinson, 1933; Bain, 1951), which posits that firm conduct is dictated by market structure. But since wages and market concentration are joint outcomes in models of labor markets, finding excludable instruments for market structure is challenging (Berry, 2021; Schmalensee, 1989). Our paper, in adopting an empirical strategy that sidesteps these endogeneity issues, provides complementary evidence on the extent of firms' exercise of labor market power.

Finally, our paper contributes to strands of the literature in labor and industrial organization on the nature of competition on online labor markets. We adapt models of imperfect competition to our setting, which combines the characteristics of online auction markets and terrestrial labor markets. The paper closest to ours in this literature is Azar et al. (2022), who gauge the potential market power of employers by estimating labor supply to individual firms on a large, online labor market using modern discrete choice methods. Our paper complements their analysis by further characterizing the nature of horizontal preference differentiation, and explicitly testing between models of firm conduct. Using experiments,Dube et al. (2020a) and Dube et al. (2020b) also demonstrate the importance of monopsony in online labor markets for task work, while a recent study by Horton et al. (2021) highlights the informative content of cheap talk about wages in online labor markets. We similarly find that cheap talk on Hired.com—in the form of firms' initial offers and workers' desired salaries—is an important signalling mechanism.

2 Setting and Data

2.1 Market description

A key limitation of the literature estimating labor supply over differentiated firms is that workers' choice sets are rarely observed, and almost never available in a high-stakes, real-world environment. Because of this, existing estimates of worker preferences are either computed in surveys and lab environments (e.g., Wiswall and Zafar (2018), Mas and Pallais (2017)), or reliant on strong assumptions applied to observational data. In survey or experimental settings, sample sizes and external validity can be limited. In observational settings, however, estimates may be confounded by differences in choice sets or erroneous inference of workers' options.

Two features of the recruitment process on Hired.com allow us to overcome this limitation. First, wage bargaining on Hired.com is high-stakes: as evidenced in Table 1, the typical candidate in our sample is a college-educated software engineer in San Francisco, currently employed and looking for a full-time job, with a salary of about \$137,000. Second, the recruitment process on Hired.com allows us to cleanly identify the choice set of candidates deciding which firms to interview with as well as the full set of observable characteristics of a candidate the firm accesses when deciding whether to

Table 1: Summary Statistics for Candidate characteristics

Variable (mean)	(1) All ($n = 43630$)	(2) Female (19%)	(3) Male (81%)
Salary			
Ask/Expectation	$137k	$126k	$140k
Education			
Has a BA+	0.872	0.913	0.862
Has an MA+	0.403	0.437	0.395
Has a CS degree	0.629	0.558	0.645
Attended an IvyPlus	0.154	0.185	0.147
Work History			
Years of experience	11.3	10.1	11.6
Software engineer	0.684	0.512	0.724
Worked at a FAANG	0.108	0.097	0.111
Employed	0.748	0.719	0.755

Note: This table reports summary statistics for candidates in the connected set (see Section 2.2).

send interview requests. Intuitively, this property of the data comes from the unique timeline of the recruitment on Hired.com: companies apply to candidates based on their profiles, and candidates decide whether or not to interview with companies based on the job descriptions and bid salaries they receive. Importantly, candidates have no way to directly view and apply to job postings without receiving an interview request. As a result, for each candidate on Hired.com, we know their consideration set (the set of all the firms that apply to them), and their choices (whether or not they decided to interview with any given firm in the consideration set).

Formally, the recruitment process can be divided into three sequential salary negotiation steps illustrated in Figure 1. First, candidates create a profile that contains standardized resume entries (education, past experience, etc.) as well as the salary that the candidate would prefer to make, which we label the *ask salary*. Figure A.1 is a screenshot of a typical candidate's profile, and Roussille (2021) further provides an exhaustive listing of profile fields.[2] Second, firms get access to candidate profiles that match standard requirements for the job they want to fill (i.e., job title, experience,

[2] In short, every profile includes the current and desired location(s) of the candidate, their desired job title (software engineering, web design, product management, etc.), their experience (in years) in this job, their top skills (mostly coding languages such as R or Python), their education (degree and institution), their work history (i.e., firms they worked at), their contract preferences (remote or on-site, contract or full-time, and visa requirements), as well as their search status, which describes whether the candidate is ready to interview and actively searching or simply exploring new opportunities. The ask salary is prominently featured on all profiles since it is a required field.

Figure 1: Timeline of the Recruitment Process on Hired.com

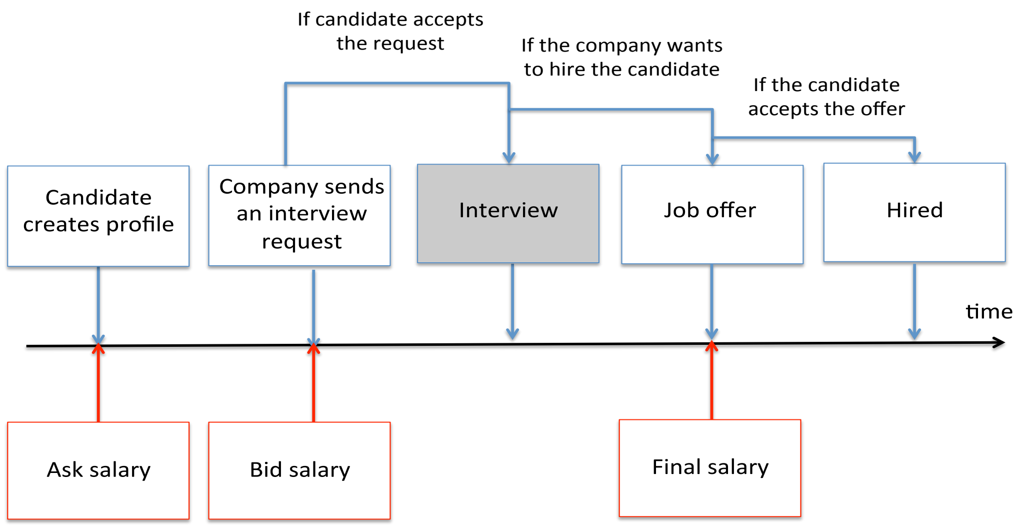

Note: This figure shows the timeline of a recruitment on Hired.com. In red are the different salaries that are captured on the platform. The blue boxes describe the steps of a recruitment, from profile creation to hiring. The grey shading for the interview stage indicates that we do not have meta data from companies on their interview process.

and location). To apply for an interview with a candidate, the company sends them a message—the interview request—that typically contains a basic description of the job as well as, crucially, the salary at which they would be willing to hire the candidate. We call this the *bid salary*.[3] Third, Hired.com records whether the candidate accepts or rejects the interview request. While interviews are conducted outside of the platform, Hired.com gathers information on whether the company makes a final offer of employment to the candidate and at what salary. We refer to this as the *final salary*. It is important to note that the bid salary is non-binding, so the final salary can differ from the bid. We also observe whether the candidate accepts the final salary offer, in which case the candidate is hired.

2.2 Sample restrictions: connected set

As is standard in the literature on firm fixed effects (Sorkin, 2018) , we only estimate amenity values for firms that are members of a connected set. To be a member of this set, a firm must have been both revealed-preferred to at least one member of the set, and have been revealed-dispreferred to at least one member of the set (the likelihood contribution of candidates with no choice variation is undefined). The candidate market is highly skewed towards tech workers in San Francisco, who represent 76% of

[3]Figure A.2 is a screenshot of a typical message sent to a candidate by a company. The bid salary is prominently featured in the subject line of the message and is required to be able to send the message. The equity field also exists but is optional.

all interview requests on the platform, and consequently our analysis focuses on those. For this segment of the market, 2,121 companies sent out 267,940 interview requests to 44,321 candidates, averaging 15.8 bids per job (median 5 bids) and 4.3 bids per candidate. 1,649 companies meet the requirements to qualify for the connected set. After making these restrictions, we retain 13,072 different jobs and 14,344 candidates, with 9.5 bids per job on average.

2.3 Stylized facts

We now document a number of empirical patterns in candidate and firm behavior that motivate the assumptions we make in our models of labor supply and demand.

Significant heterogeneity in bid acceptance. Figure 2a plots the distribution of the share of each firm's bids that are accepted. There are two important features of this distribution. First, firms are frequently rejected by candidates: on average, candidates only accept 60.5% of the interview requests they receive. Second, there is significant heterogeneity across companies in the likelihood that an interview request is accepted: while the mean share of bids accepted is 60.5%, 10.2% of firms see less than 40% of their interview requests accepted, while 16.2% of firms see more than 75% of their interview requests accepted. Candidates do not accept all the interview requests they receive; reflecting the fact that most candidates search from current employment. This motivates us to model their outside option as a key parameter in their interview decision (Section 3.1). Additionally, the wide variation in acceptance rates across firms is suggestive of significant vertical (between-firm) differentiation, which motivates our revealed-preference approach.

Reference-dependence of labor supply. Relevant to our modelling choices, Figure 2b plots the probability that an interview request is accepted as a function of the ratio of the bid salary to the ask salary. Perhaps unsurprisingly, higher bids are associated with a higher acceptance probability. But the slope of this relationship is steeper when bids are below the ask than when bids are above the ask: on average, the probability a bid is accepted when it is 10% less than the ask is roughly 15% lower than when a bid is made at the ask exactly. However, the probability a bid is accepted when it is 10% more than the ask is only about 5% higher than when the bid equals the ask. We take this pattern as suggestive evidence that candidates' labor supply is reference-dependent in their ask. Although it is not possible to definitively place a structural interpretation on these patterns without accounting for selection, we bolster this interpretation by using additional information that records the candidates' reason for rejecting a bid, which is available for a subset of the observations.[4] Figure 2c plots the probability that a candidate selects "insufficient compensation" as the reason for rejecting a bid as a function of the ratio of the bid salary to the ask. The relationship between this probability and the ratio of the bid to the ask is sharply kinked at bid=ask: the slope (and level) is almost exactly zero above bid=ask, and is strongly negative below bid=ask. We refer to this phenomenon in our model as "kinked labor supply" and formally model different slopes for the labor supply elasticities above and below the kink.[5]

Individualized pricing and the absence of wage posting. While wage posting is pervasive in many labor markets, it is not a feature of firm behavior in our setting. The average within-job standard deviation of bid salaries is roughly $23,041, which indicates that firms are willing to offer a wide range of salaries to candidates for the same vacancy. Indeed, only 1.4% of jobs offer the same bid salary to all candidates. Further, the bids firms make to candidates are highly individualized: 76.5% of bids are made *exactly* at the candidates' ask. Figure 2d synthesizes these two facts. It plots the relationship between the bid premium - the difference between bid and ask salaries - and the deviation of the ask from the average ask of candidates who receive bids for the same job. This figure illustrates the fact that there is a large heterogeneity of bid salaries for the same job, driven by the large underlying variation in the ask salaries of candidates who receive bids for that job. If firms were wage posting, they would offer every candidate the same bid salary, and the points would lie on the -45-degree red line. Empirically, we observe that the slope of the relationship is dramatically flatter than this "full compression" line: changes in the ask are almost entirely offset by changes in the bid. This indicates that, even for a given job, firms increase their bids almost one-for-one with the asks. We incorporate these patterns in our model of labor demand in two ways. First, firms internalize the reference-dependence of candidates' labor supply around the ask. This generates an incentive for firms to bunch at the kink, and rationalizes the large mass of offers made at ask. Second, we model firms' decisions to bid on each candidate as a fully-individualized process, allowing for systematic and idiosyncratic components of match-specific productivity.

[4] While this field is optional, 55% of candidates do fill it out.

[5] Leveraging a survey of 6,000 job seekers in New Jersey, Figure 3 in Hall and Mueller (2018) plots the job offer acceptance frequency as a function of the difference between the log hourly offered wage and the log hourly reservation wage. A clear kink is observed at offered = reservation.

Figure 2: Empirical Patterns in Bid and Ask Strategies

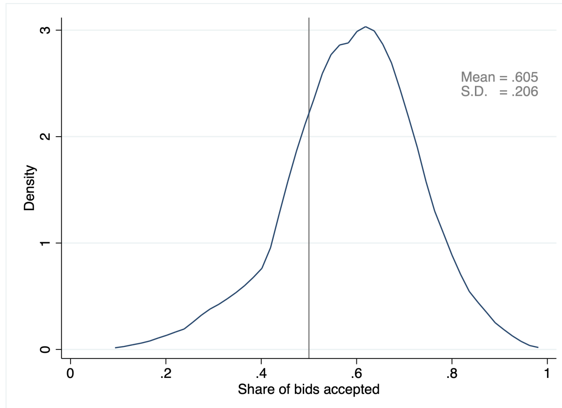

(a) Fraction of Interview Requests Accepted

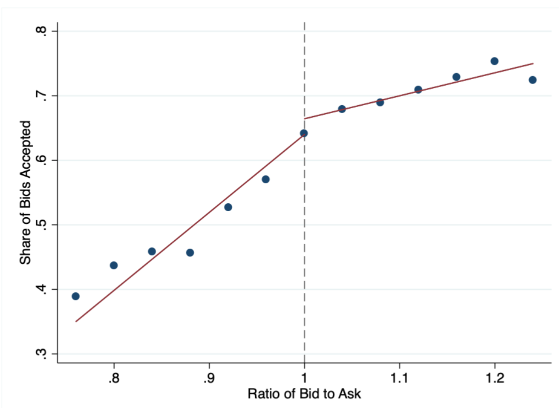

(b) Kink at Bid = Ask

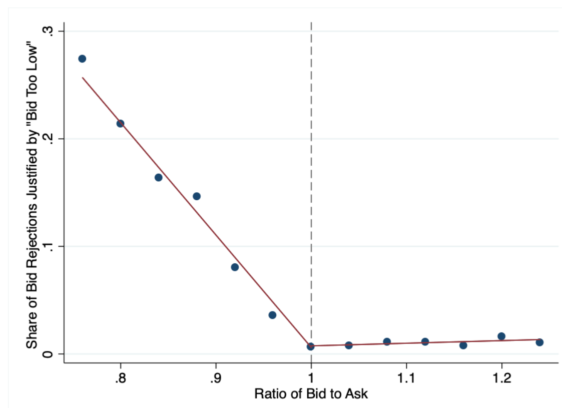

(c) Monetary Concerns Drive Rejections < Ask

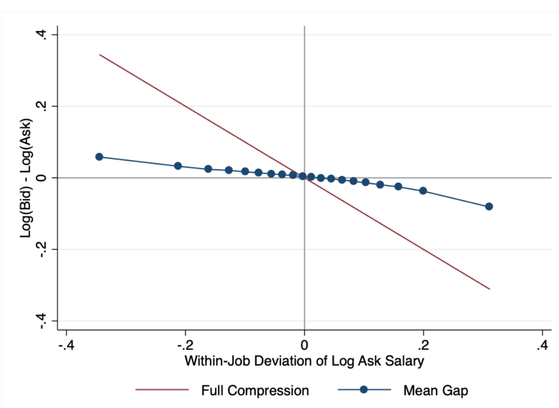

(d) Large Range of Bid Salaries for Same Job

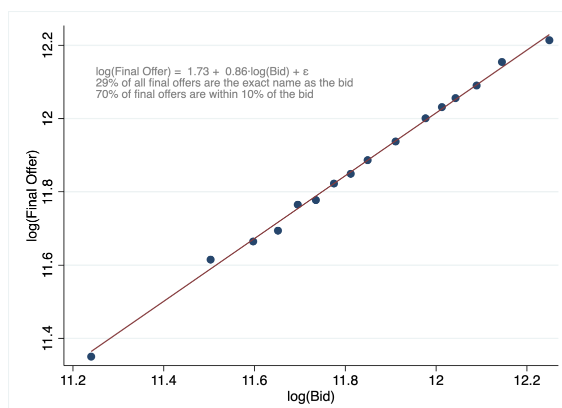

(e) Bids are Sticky in Expectation

Note: Panel (a) shows the distribution of the share of accepted interview requests for a given firm. Panel (b) plots the average probability that a candidate accepts an interview request against the ratio of the bid to ask salary. Panel (c) plots the average probability that a candidate accepts an interview request against the ratio of the bid to ask salary. Panel (d) plots the relationship between the premium—the difference between (log) bid and ask salary—and the within-job deviation of the (log) ask salary. Panel (e) plots the relationship between the bid and the final offer sent to candidates.

Bids are non-binding, but sticky. The bid salary is what firms declare they are willing to pay the candidate solely based on their profile, before any interaction with them. The final salary is offered to a candidate at the hiring stage. Given that companies are by no means contractually bound by their bids, final salaries may differ from bids. However, effectively, firms commit to making final offers that are close to the bids. Figure 2e shows the relationship between the bid and final offer for the subset of candidates that receive one. Strikingly, this relationship is very linear, with a slope close to one. Additionally, 31% of all final offers are identical to the bid and 72% of all final offers are within 10% of the bid. We correspondingly make the simplifying assumption that the expectation of the final salary is equal to the bid for both candidates and firms, such that we can estimate our model on the much richer data from the interview stage.

3 Defining Firm Wage-Setting Conduct

In order to particularize our definition of conduct—how firms determine which workers to hire and how much to pay them—to our setting, we first specify a general model of labor supply and demand on Hired.com. Candidates $i = 1, \ldots, N$ post resume information x_i (which includes their ask a_i) before interacting with firms. Firms $j = 1, \ldots, J$ have observable characteristics z_j. Both x_i and z_j include a constant. The outside option is denoted by $j = 0$. Firms browse active candidate profiles and decide, for each candidate, whether to send an interview request. We denote the bid salary of firm j on candidate i by b_{ij}, and let B_{ij} equal one if firm j sends an interview request to candidate i. After a candidate receives it, she decides whether to accept (and thereby move forward with the recruitment process) or to reject the request. After the interview, the firm can make a final offer of employment to the candidate.

Our analysis focuses on the interview stage of the recruitment process. In order to specify a tractable model of firm and candidate behavior at this initial stage, we make several simplifying assumptions about the final stages of the process. In particular, we assume that firms do not treat bids as cheap talk—rather, firms credibly expect to pay their bids, should they decide to make a final offer. In practice, this assumption is a fairly accurate description of firm behavior as documented in Figure 2e and described in Section 2.3. We also assume that candidates' choices at the interview request and final offer stages are governed by the same preferences. While our framework is consistent with certain forms of updating on the part of candidates after interviews take place, we remain agnostic about those mechanisms.

3.1 Labor Supply

We assume that the indirect utility candidate i associates with firm j at bid b_{ij} is:

$$V_{ij} = u(b_{ij}, a_i) + \Xi_{ij}, \text{ where } \quad \Xi_{ij} = A_j(Q_i) + \xi_{ij}, \tag{1}$$

where $u(b_{ij}, a_i)$ is the *monetary component* of utility and Ξ_{ij} is the *non-monetary component* of utility. Building on the stylized facts documented in Section 2.3, we first assume that labor supply is reference-dependent in the ask: $u(b, a)$ is continuous, strictly increasing, and twice continuously differentiable in its first argument, except at the point $b = a$, where $\lim_{b\to a^-} \partial u(b, a)/\partial b > \lim_{b\to a^+} \partial u(b, a)/\partial b$. We further assume that the ask serves as a sufficient statistic for the monetary component associated with the outside option, setting $b_{i0} = a_i$ and normalizing $u(a, a) = 0$.[6] The indirect utility associated with the outside option is therefore given by $V_{i0} = \Xi_{i0}$.

The non-monetary component of utility Ξ_{ij} can be further decomposed into the sum of a systematic *amenity value* $A_j(Q_i)$ and an idiosyncratic *taste shock* ξ_{ij}. The amenity value i associated with j is determined by i's *latent preference type* Q_i. Candidates i and ℓ with $Q_i = Q_\ell$ share a common mean valuation of amenities at all firms. Both preference types Q_i and taste shocks ξ_{ij} are private information: they are observed by workers, but not by firms. However, the distribution of types F_Q may depend on observables x_i: $F_{Q|X} \neq F_Q$. So, while Q_i is private information, it may be *partially revealed* to firms by x_i. By contrast, the ξ_{ij} are *iid* draws from a probability distribution that is independent of x_i: $\xi_{ij} \overset{iid}{\sim} F_\xi(\cdot)$, where $F_{\xi|x} = F_\xi$. We assume that F_ξ admits a continuous, log-concave density $f_\xi(\cdot)$ with support on the full real line.

Candidate i will accept firm j's interview request if and only if the utility associated with that request exceeds that of her outside option:

$$D_{ij} = B_{ij} \times \mathbf{1}[V_{ij} \geq V_{i0}]. \tag{2}$$

Candidates' final labor supply decision is given by choosing the final offer with the highest indirect utility. We assume the indirect utility i associates with a final offer from j is equal to V_{ij}, such that the same shocks that enter candidates' interview decisions also govern their final job choice. Because we focus on the ex-ante perspective of firms formulating bids, we view this as a simplifying abstraction.

[6]For the large fraction of workers on the platform engaging in on-the-job search, this assumption can easily be justified if asks are formulated as a function of current salary. Unemployed workers post lower asks even conditional on a rich set of covariates (the expected conditional gap is $8,366), suggesting that their asks reflect relatively worse outside options.

3.2 Labor Demand

For each candidate i it encounters, firm j formulates an optimal bid b_{ij}^* to maximize the *expected option value* of an interview request, given by the function $\pi_{ij}(b)$. Firms decide to bid on candidates if the maximized value of that function surpasses a firm-specific interview cost threshold c_j:

$$b_{ij}^* = \arg\max_b \ \pi_{ij}(b), \text{ and } B_{ij} = \mathbf{1}\left[\pi_{ij}(b_{ij}^*) \geq c_j\right]. \tag{3}$$

Realized bids are: $b_{ij} = B_{ij} \times b_{ij}^*$, where $b_{ij} = 0$ if $B_{ij} = 0$. The option value of an interview request from firm j to candidate i depends upon both i's labor supply decision and i's value to j. Encode i's *final* labor supply decision, given j's choice of bid b, via the potential outcome $D_{ij}^{\circ}(b)$, a binary random variable that equals one if i would accept j's final offer of employment given j's choice of bid salary $b_{ij} = b$. Denote the maximum utility of the offers available to i by V_i^1. Given our assumptions about candidate preferences, we have:

$$D_{ij}^{\circ}(b) = \mathbf{1}\Big[V_{ij} = V_i^1 \mid b_{ij} = b\Big]. \tag{4}$$

Denote the ex-post value firm j places on a match with candidate i as ε_{ij}°. Given these definitions, $\pi_{ij}(b)$ can be written as:

$$\pi_{ij}(b) = \mathbb{E}_{ij}\Big[D_{ij}^{\circ}(b_{ij}) \times (\varepsilon_{ij}^{\circ} - b_{ij}) \mid b_{ij} = b\Big], \tag{5}$$

where $\mathbb{E}_{ij}[\cdot]$ denotes an expectation taken over the *information set* of firm j when it evaluates candidate i, which we denote by Ω_{ij} (and which may include firm-, candidate-, and market-level variables). This objective function is nearly identical to that of a bidder in a standard first-price auction. In a first-price auction, a bidder's objective is to maximize her expected utility, where her bid affects both the net payoff should she win (here, $\varepsilon_{ij}^{\circ} - b$) and the probability that she wins the auction (here, $\mathbb{E}_{ij}[D_{ij}^{\circ}(b)]$). An "auction" on Hired.com differs from a standard first-price auction, however, because the firm that submits the highest monetary bid is not guaranteed to be the candidate's top-ranked choice.

We make two additional assumptions that simplify the form of $\pi_{ij}(b)$. Conditional on Ω_{ij}, we assume: 1) potential outcomes $D_{ij}^{\circ}(b)$ and ex-post match values ε_{ij}° are independent, and 2) ε_{ij}° is independent of the firm's bid b_{ij}. Since all firms must bid on candidates before the match value is revealed, the first assumption essentially establishes the sufficiency of the observables available to the firm for forecasting match values. It also rules out scenarios in which the event of winning the "auction" for candidate i reveals information about other firms' match values that is relevant to j's value (the "winner's curse"). The second assumption rules out behavioral effects of increasing bids on the value of a match (e.g. efficiency wages). Together, they imply:

$$\pi_{ij}(b) = \underbrace{\Pr_{ij}\Big(D^*_{ij}(b) = 1\Big)}_{\triangleq G_{ij}(b)} \times \Big(\underbrace{\mathbb{E}_{ij}[\varepsilon^*_{ij}]}_{\triangleq \varepsilon_{ij}} - b\Big). \tag{6}$$

The first term, $G_{ij}(b)$, is j's forecast of i's labor supply decision, which we refer to as firms' *beliefs* (or win probability).[7] The second term is the difference between j's forecast of i's ex-post match value (or *valuation*), ε_{ij}, and j's bid. Under certain assumptions ε_{ij}, coincides with i's true productivity at j (MRPL_{ij}).

3.3 Firm Conduct in Equilibrium

Before providing a precise definition of firm wage-setting conduct, we first define a notion of equilibrium. In Bayes-Nash equilibrium, players' actions are best responses given their beliefs, which are themselves consistent with equilibrium play. We explicitly define equilibrium such that beliefs are consistent *conditional on the information firms use to construct those beliefs*:

Definition 1 (Equilibrium). *Given information sets* $\{\Omega_{ij}\}_{i=1,j=1}^{N,J}$, *a pure strategy equilibrium is a set of tuples* $\{b_{ij}(\cdot), G_{ij}(\cdot)\}_{i=1,j=1}^{N,J}$ *satisfying:*

(Optimality) $b_{ij}(\varepsilon)$ *is* j*'s best response for valuation* ε *given beliefs* $G_{ij}(b)$*:*

$$b_{ij}(\varepsilon) = \begin{cases} \arg\max_b \, G_{ij}(b) \times (\varepsilon - b) & \text{if } \max_b \, G_{ij}(b) \times (\varepsilon - b) \geq c_j \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

(Consistency) *Conditional on* Ω_{ij}, *firm* j*'s beliefs* $G_{ij}(b)$ *obey:*

$$G_{ij}(b) = \iint \Pr\Big(u(b, a_i) + \Xi_{ij} = V_i^1 \mid V_i^1 = v, Q_i = q\Big) \times dF_{V,Q}\Big(v, q \mid \Omega_{ij}\Big), \tag{8}$$

where $F_{V,Q}(\cdot, \cdot \mid \Omega_{ij})$ *is the population joint CDF of* V_i^1, Q_i *conditional on* Ω_{ij}.

[7] We assume that firms' beliefs are stationary, such that firms behave as if they are in a steady state, as in Backus and Lewis (2020). We defer consideration of dynamics for future research.

To operationalize a notion of conduct in our setting, it is useful to partition each information set as $\Omega_{ij} = \{\omega_{ij}^V, \omega_{ij}^Q\}$, where ω_{ij}^V and ω_{ij}^Q encode the information j uses to forecast V_i^1 and Q_i, respectively. We write the joint CDF as:

$$F_{V,Q}\Big(v, q \mid \Omega_{ij}\Big) = \underbrace{F_{V|Q}\Big(v \mid Q_i = q, \omega_{ij}^V\Big)}_{=F^{\omega}_{V|Q}} \times \underbrace{F_Q\Big(q \mid \omega_{ij}^Q\Big)}_{=F^{\omega}_{Q}}. \tag{9}$$

We can now provide a definition of firm wage-setting conduct in our setting:

Definition 2 **(Conduct).** *Given the assumptions of Sections 3.1 and 3.2 and Definition 1, a model of firm wage-setting conduct is defined by specifying the form of firms' beliefs,* $G_{ij}(b)$*:*

- *When markets are* ***Imperfectly Competitive****, firms' beliefs are nondegenerate, and conduct is dictated by the contents of firms' information sets* $\Omega_{ij} = \{\omega_{ij}^{\Lambda}, \omega_{ij}^Q\}$*. We specify two alternatives for each component—firms are either:*
  - ***Not Predictive****, with* $\omega_{ij}^Q = \varnothing$ *such that* $F_Q^{\omega} = F_Q$*; or* ***Type Predictive****, with* $\omega_{ij}^Q = x_i$ *such that* $F_Q^{\omega} = F_{Q|X}$*; and either:*
  - ***Monopsonistically Competitive****, with* $b_{ij}, \boldsymbol{A}_j \notin \omega_{ij}^{\Lambda}$ *such that* $\partial F^{\omega}_{V|Q}/\partial b = 0$*; or* ***Oligopsonists****, with* $b_{ij}, \boldsymbol{A}_j \in \omega_{ij}^{\Lambda}$ *such that* $\partial F^{\omega}_{V|Q}/\partial b > 0$*.*
- *When markets are* ***Perfectly Competitive****, firms' beliefs are degenerate: every firm* j *believes that there exists a competitor with a valuation arbitrarily close to its own for each candidate* i*:* $G_{ij}(b) \propto \mathbf{1}[b \geq \varepsilon_{ij}]$*.*

Clearly, this notion of conduct does not encompass every interesting feature of firm behavior in wage setting. For instance, our model of firm behavior rules out common ownership effects, since firms are assumed to engage in separate maximization problems for each vacancy. However, our setting—one in which firms have the ability to offer fully individualized wages—is particularly well-suited for investigating how firms incorporate information about the distribution of preferences and competition into their recruitment decisions. In Appendix B, we illustrate with a simple model the implications of our conduct assumptions and how the conceptual framework of our study differs from those that relate measures of market structure to wages. Our approach can be extended to accommodate other conduct alternatives.

The first conduct assumption we test concerns ω_{ij}^Q, the information firms use to forecast types. This test is motivated by our assumption that observables may partially reveal candidates' preference types to firms. Whether firms do or do not use this information to offer different wages to candidates with identical productivity levels has been a matter of debate in the labor literature. For instance, Burdett and Mortensen (1998) assume that firms are not type-predictive, leading to efficiency losses that can be reduced by the introduction of a minimum wage. On the other hand, Postel-Vinay and Robin (2002) assume that firms are not type-predictive, but rather fully informed about the types of workers they meet, allowing them to engage in classic first-degree price discrimination. More recently, Flinn and Mullins (2021) analyze models in which firms differ in whether they commit to posted wages (akin to non-predictive conduct) or negotiate wages in response to outside offers (akin to type-predictive conduct). Similarly, whether firms use information on within-firm variation in price elasticities has been the subject of interest in the industrial organization literature on uniform pricing (DellaVigna and Gentzkow, 2019). In our setting, firms may make more offers and workers may capture a smaller share of match surplus when firms are type-predictive relative to when they are not.

The second conduct assumption we test concerns ω_{ij}^{V}, and the nature of interactions between vertically-differentiated firms. Under monopsonistic competition, firms are differentiated but view themselves as atomistic relative to the market: they ignore the effects of their behavior on the composite value of candidates' option sets. This assumption underlies a number of studies, including Card et al. (2018) and Lamadon et al. (2022). When firms are oligopsonists, on the other hand, they actively incorporate the effects of their behavior on the distribution of options available to each candidate into their wage-setting decisions. Models of oligopsony, as in Berger et al. (2022) and Jarosch et al. (2021), therefore feature *strategic interactions* between firms. Another distinction, as noted in Berger et al. (2022), is that, under monopsonistic competition, structural firm-specific labor supply elasticities are equal to reduced-form elasticities. In contrast, under oligospony, they depend upon both the firms' bid and the value of its amenities, in addition to competitor's bids and amenities.[8]

Finally, our model of perfectly competitive firms serves as a useful baseline against which we can compare more complicated models of conduct that incorporate additional sources of wage dispersion beyond differences in the marginal revenue product of labor. Under the perfect competition assumption, firms bid their valuations: $b_{ij}(\varepsilon) = \varepsilon$. Interview costs c_j are normalized to 0 without loss of generality.

[8] Our definition of oligopsonistic behavior encompasses multiple mechanisms that have been explored separately in prior work (for instance, our framework subsumes both size- and differentiation-based mechanisms by which oligopsonists generate markdowns).

4 A Test of Firm Wage-Setting Conduct

4.1 Setup: Testing via an Exclusion Restriction

Our objective is to determine which model of conduct is the best description of the true data-generating process.[9] To formulate our test, we first write ε_{ij} as a function of observables and a mean-zero idiosyncratic component $\nu_{ij} \overset{iid}{\sim} F_\nu(\cdot)$ that is unrelated to those observables by construction: $\varepsilon_{ij} = \gamma_j(x_i, \nu_{ij})$. We assume that there exists a transformation of that function $\tau(\cdot)$ such that the $\tau(\gamma_j(\cdot, \cdot))$ is additively separable in those components: $\tau(\varepsilon_{ij}) = \gamma(x_i, z_j) + \nu_{ij}$. The function $\gamma(x, z)$ encodes the systematic component of match values shared by candidates with $x_i = x$ at firms with $z_j = z$.

To illustrate the intuition of our testing procedure, assume for the moment that $G_{ij}(b)$ is differentiable for all b. Then, under the true conduct assumption, bids must satisfy the following first-order condition with equality:

$$\tau\Big(\varepsilon_{ij}(b_{ij})\Big) = \gamma(x_i, z_j) + \nu_{ij}, \tag{10}$$

where $\varepsilon_{ij}(b)$ is the *inverse bidding function* ($b = b_{ij}(\varepsilon_{ij}(b))$). This equation includes only one source of error: the idiosyncratic component of firms' valuations, ν_{ij}. Since the true model of conduct is unknown, in practice the true inverse bidding function $\varepsilon_{ij}(\cdot)$ is proxied by its counterpart under an assumed model of conduct m, $\varepsilon^m_{ij}(\cdot)$.[10] If m is misspecified, then this substitution introduces an additional error term:

$$\tau\Big(\varepsilon^m_{ij}(b_{ij})\Big) = \gamma(x_i, z_j) + \nu_{ij} + \zeta^m_{ij}. \tag{11}$$

The presence of misspecification error suggests an intuitive conclusion: if labor supply is determined in part by variables that are excluded from firms' valuations, then models that are further from the truth should yield residuals that are more strongly correlated with those excluded variables than those of models that are closer to the truth. In other words, if the true demand residuals (ν_{ij}) obey exclusion restrictions, then models can be compared by inspecting the degree to which their estimated residuals violate those restrictions (since ζ^m_{ij} need not obey those restrictions).

This is the basic logic of Berry and Haile (2014), who establish the necessity of

[9] The models we consider are *non-nested*: "Broadly speaking, two models (or hypotheses) are said to be 'non-nested' if neither can be obtained from the other by the imposition of appropriate parametric restrictions or as a limit of a suitable approximation; otherwise they are said to be 'nested"' (Pesaran, 1990). In our setting, models are non-nested as long as they generate distinct patterns of markdowns and selection that are not co-linear with the determinants of ε_{ij}.

[10] Keeping in mind that under assumption m, we may treat $\varepsilon^m_{ij}(b_{ij})$ as data.

instruments that shift demand (here, labor supply) but that are excluded from the marginal cost function (here, firms' valuations or productivity) for both identification and testing of conduct in the product market setting with data only on market shares. Following this logic, Backus et al. (2021) implement a test of conduct that formalizes the above logic: under true conduct assumptions, instruments that affect markups (markdowns) but do not affect marginal costs (valuations) should not be correlated with recovered idiosyncratic cost shocks (demand residuals ν_{ij}). Our setting differs in two key ways from that of Berry and Haile (2014). First, we use micro data on individual choices, rather than market shares. Berry and Haile (2020) consider identification of differentiated products demand using micro data, demonstrating significantly reduced requirements for instruments. Section 5.1 illustrates how our ability to condition on the information available to firms when they bid allows us to identify labor supply parameters in our data without reliance on instruments for prices (bids). Second, we analyze *individualized* bids rather than uniform market prices. Bids are made before any negotiation has taken place and without direct knowledge of the competition, and so need not satisfy a strict market clearing condition (rather, we have assumed that firms' behavior must satisfy a conditional form of rational expectations). Our identification arguments therefore follow the empirical auction literature (Guerre et al., 2000; Backus and Lewis, 2020).[11]

To implement our testing procedure, we must find an instrument that is excluded from the determinants of firms' valuations, but that nevertheless shifts firms' behavior under each model of conduct differentially. We use an exogenous component of variation in market tightness generated by platform rules as our instrument, taking advantage of the fact that candidate profiles are only searchable for two weeks after they go live.[12] We define tightness within occupation and experience bins, since those categories are the primary search fields recruiters use when browsing candidates. Variation in tightness between two-week periods is driven primarily by the exogenous inflow of new candidates onto the platform. Since the pool of candidates turns over completely every two weeks, variation in candidate quality between two-week periods is not endogenously determined by platform conditions—and so this variation in tightness should not be related to firms' valuations (conditional on x_i and z_j). However, this variation should affect firms' expectations about the competition for i: the

[11] In Section 5.3 we generalize the differentiable case exposed above to our setting, where $G_{ij}(b)$ is not differentiable at the point $b = a_i$.

[12] Candidates can follow up with interview requests they received after their profiles are no longer live, but can only collect those requests during the two week period. Candidates may appeal to administrators to extend the time their profile is live, but in practice only a small fraction do so.

fewer active candidates there are per active firm, the more bids those candidates tend to receive. Our use of market tightness as an instrument mirrors the arguments of papers studying auctions with entry, which use exogenous variation in the potential number of entrants across auctions for identification (e.g. Gentry and Li (2014)). Let v_{ow} and u_{ow} denote the number of firms searching for experience and occupation o during two-week period w and the number of candidates with active profiles with experience and occupation o during two-week period w, respectively. The prevailing level of market tightness when j bids on i is: $t_{ij} = u_{o_i w_{ij}} / v_{o_i w_{ij}}$.[13] Our assumption is:

Assumption 1. (Instrument Exogeneity) *Labor market tightness is independent of idiosyncratic determinants of labor demand:*

$$t_{ij} \perp\!\!\!\perp \nu_{ij} \mid x_i, z_j. \tag{12}$$

Firms' information sets ω_{ij}^V include t_{ij} (as well as $u_{o_i w_{ij}}$ and $v_{o_i w_{ij}}$) in addition to x_i and z_j. Variation in tightness thereby drives variation in predicted markdowns that is independent of the determinants of firms' valuations.

4.2 The Rivers and Vuong (2002) Test

We implement the pairwise testing procedure of Rivers and Vuong (2002) to compare models of wage-setting conduct: we consider each pair of models in turn, asking whether one of those models is closer to the truth than the other. As in Backus et al. (2021), our test specifies a scalar moment condition in the residuals of fitted models and excluded instruments. Because we estimate demand under each conduct assumption via maximum likelihood, our test is based on *generalized residuals* (Gourieroux et al., 1987) defined by the scores of the likelihood. Let $s_{ij\ell}^m(\Psi) = \partial \mathcal{L}_{ij}^m(\Psi)/\partial \psi_\ell$ denote the ℓ-th component of the score vector for observation ij and model m, given parameters Ψ. The scores may be written as $s_{ij\ell}^m(\Psi) = h_{ij}^m(\Psi) \cdot \gamma_\ell(x_i, z_j)$, where $h_{ij}^m(\Psi)$ is the generalized residual and $\gamma_\ell(x_i, z_j) = \partial \gamma(x_i, z_j)/\partial \psi_\ell$. The maximum likelihood estimate $\widehat{\Psi}^m$ is the vector that sets:

$$\sum_{ij:B_{ij}=1} s_{ij\ell}^m\left(\widehat{\Psi}^m\right) = \sum_{ij:B_{ij}=1} h_{ij}^m\left(\widehat{\Psi}^m\right) \cdot \gamma_\ell(x_i, z_j) = 0 \;\; \forall\, \ell,$$

and so generalized residuals are constrained to be orthogonal to covariates. The generalized residuals for each model can be easily computed by taking the derivative

[13] This is technically the inverse of the usual definition of market tightness, v/u. We define tightness as u/v because we focus on exogenous variation in u.

of the individual likelihood contributions.

We form the generalized residuals for each model, and use them to compute the scalar moment/lack-of-fit measure:

$$Q_s^m = \left(\frac{1}{s} \sum_{ij:B_{ij}=1} h_{ij}^m \left(\widehat{\Psi}^m \right) \cdot t_{ij} \right)^2, \tag{13}$$

where $s = |\{ij : B_{ij} = 1\}|$. Q_s^m measures the covariance between the generalized residuals of model m and the excluded instrument t_{ij}. Under proper specification, the influence of the instrument on markdowns is completely summarized by the inverse bidding function, and so there should be zero correlation between the instrument and the generalized residual. The lack-of-fit measure Q_s^m can also be motivated as an unscaled version of the score test statistic for testing against the null hypothesis that the coefficient on t_{ij} in the labor demand equation is zero. In Appendix F, we describe and implement an alternate testing procedure based on the Vuong (1989) likelihood ratio test. This test can be thought of as an omnibus version of our lack-of-fit measure, while our version of the Rivers and Vuong (2002) test isolates only the component of lack-of-fit directly correlated with instrument variation.

Following Backus et al. (2021),[14] we formulate a pairwise test statistic for testing between models m_1 and m_2 as an appropriately-scaled difference between $Q_s^{m_1}$ and $Q_s^{m_2}$, which Rivers and Vuong (2002) show to be asymptotically normal:

$$T_s^{m_1,m_2} = \frac{Q_s^{m_1} - Q_s^{m_2}}{\widehat{\sigma}_s^{m_1,m_2} / \sqrt{s}} \xrightarrow{D} \mathcal{N}(0, 1). \tag{14}$$

$\widehat{\sigma}_s^{m_1,m_2}$ is an estimate of the population variance of $Q^{m_1} - Q^{m_2}$. We compute $\widehat{\sigma}_s^{m_1,m_2}/\sqrt{s}$ as the variance of $Q_s^{m_1} - Q_s^{m_2}$ across bootstrap replications. Given a significance level α with critical value c_α, we reject the null hypothesis that m_1 and m_2 are equivalent in favor of the alternative that m_1 is better than m_2 when $T_s^{m_1,m_2} < c_\alpha$, and vice versa if $T_s^{m_1,m_2} > c_\alpha$. If $|T_s^{m_1,m_2}| \leq c_\alpha$, the test cannot discriminate between the two models. The power of the test depends crucially on the ability of the instrument to predict differential markdowns and selection corrections (Duarte et al. (2023) consider weak instruments problems in conduct testing).

[14] Backus et al. (2021) formulate their moment-based test statistic by interacting residuals with an appropriate function of both the instrument and all other exogenous variables, and connect their choice of that function to the literature on optimal instruments (Chamberlain, 1987). In our setting, the formulation of such a function is complicated selection and partial identification issues. While we do not pursue it here, the formulation of optimal instruments is ripe ground for future work.

5 Identification and Estimation of Labor Supply and Demand

5.1 Labor Supply

Identification. Denote i's offer set by: $\mathcal{B}_i = \{b_{ij}, B_{ij}\}_{j=0}^{J}$. Our principal assumption for the identification of preferences from choice data is:

Assumption 2. (Conditional Independence) *Candidates' types Q_i are private information, so firms decide whether and how much to bid on the basis of x_i alone. In other words, i's offer set $\mathcal{B}_i$ is independent of her type Q_i conditional on her x_i:*

$$\Pr(\mathcal{B}_i \mid Q_i = q, x_i) = \Pr(\mathcal{B}_i \mid x_i). \tag{15}$$

A consequence of Assumption 2 is that the distribution of candidate types conditional on both $\mathcal{B}_i$ and x_i is equal to the distribution of types conditional on x_i alone:

$$\Pr(Q_i = q \mid \mathcal{B}_i, x_i) = \frac{\Pr(\mathcal{B}_i \mid Q_i = q, x_i)\Pr(Q_i = q \mid x_i)}{\Pr(\mathcal{B}_i \mid x_i)} = \Pr(Q_i = q \mid x_i). \tag{16}$$

Assumptions analogous to assumption 2 are implausible in administrative data, like linked employer-employee records, due to the various selection mechanisms at play in the formation of equilibrium matches. But in our setting, firms are required to make initial bids on the basis of candidate profiles alone—the same information available to us—before they have the chance to interact with candidates. Further, our data records not only the offers candidates accept, but also the ones they reject.

Next, denote i's sets of accepted and rejected bids by $\mathcal{B}_i^1 \subseteq \mathcal{B}_i$ and $\mathcal{B}_i^0 = \mathcal{B}_i \setminus \mathcal{B}_i^1$, respectively. The labor supply model of Section 3.1 implies that every option in $\mathcal{B}_i^1$ is revealed-preferred to every option in $\mathcal{B}_i^0$: $\min_{j\in\mathcal{B}_i^1} V_{ij} \geq \max_{k\in\mathcal{B}_i^0} V_{ik}$. We refer to this event as a *partial ordering* of i's offer set $\mathcal{B}_i$, which we denote by $\mathcal{B}_i^1 \succ \mathcal{B}_i^0$. We now formalize two additional assumptions about the structure of preferences:

Assumption 3. (Mixture Model) *The probability of observing any partial ordering is described by a finite mixture model over latent preference types:*

- **a) (Finite Support)** *The support of Q_i is restricted to the integers $1, \ldots, Q$. Denote the conditional probability of type membership by:*

$$\Pr(Q_i = q \mid x_i) \triangleq \alpha_q(x_i). \tag{17}$$

- **b) (Exclusion Restriction)** *Conditional on a candidate's latent type Q_i and $\mathcal{B}_i$, the probability of observing any partial ordering is independent of x_i:*

$$\Pr\left(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \mathcal{B}_i, Q_i = q, x_i\right) = \Pr\left(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \mathcal{B}_i, Q_i = q\right) \triangleq \mathcal{P}_q\left(\mathcal{B}_i^1 \succ \mathcal{B}_i^0\right). \tag{18}$$

Assumption 3a is a modelling choice about the form of unobserved heterogeneity in preferences over firms. Assumption 3b governs how preferences are related to individual characteristics: the variables in x_i shift the distribution of types, but provide no additional information about preferences conditional on those types. Note that Assumption 3b is an implication of the labor supply model specified in Section 3.1.

Combining Assumptions 2 and 3, the log-integrated likelihood of i's revealed partial ordering (given $\mathcal{B}_i$ and x_i) is:

$$\mathcal{L}(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \mathcal{B}_i, x_i) = \log\left(\sum_{q=1}^{Q} \alpha_q(x_i) \times \mathcal{P}_q\Big(\mathcal{B}_i^1 \succ \mathcal{B}_i^0\Big)\right).$$

Mixtures of random utility models (RUMs) of this form have been studied in both econometrics and computer science/machine learning. In particular, Soufiani et al. (2013) establish identifiability of a finite-mixture-of-types RUM for which the idiosyncratic error components follow a log-concave distribution, as assumed in our model. As in Sorkin (2018), identification is limited to firms that are members of the connected set defined in Section 2.2.

Parameterization. In order to estimate preferences, we first specify a parameterization of the labor supply model. We write the monetary component of utility as:

$$u(b,a) = \Big(\theta_0 + \theta_1 \cdot \mathbf{1}[b < a]\Big) \cdot \Big[\log(b) - \log(a)\Big] = \begin{cases} \theta_0 \cdot \log(b/a) & \text{if } \; b \geq a, \\ \Big(\theta_0 + \theta_1\Big) \cdot \log(b/a) & \text{if } \; b < a, \end{cases}$$

and so, $u(b,a)$ is continuous, but kinked, at $b = a$.[15] We specify the distribution of types as a multinomial logit in x_i with parameter $\boldsymbol{\beta}$:

$$\Pr(Q_{iq} = 1 \mid x_i) = \alpha_q(x_i \mid \boldsymbol{\beta}) = \frac{\exp(x_i'\beta_q)}{\sum_{q'=1}^{Q} \exp(x_i'\beta_{q'})}.$$

Because Q_i has finite support, we write $A_j(Q_i) = \boldsymbol{Q}_i'\boldsymbol{A}_j$, where $\boldsymbol{A}_j$ is a $Q \times 1$ vector of type-specific mean amenity values at firm j with q-th component A_{qj}, and $\boldsymbol{Q}_i$ is a $Q \times 1$ vector of type indicators with $Q_{iq} = 1$ if $Q_i = q$. Finally, we assume that the distribution of taste shocks is extreme value type 1: $\xi_{ij} \overset{iid}{\sim} EV_1$.

[15]Note that we have defined $u(b,a)$ relative to the outside option: when $b = a$, $\log(b/a) = \log(1) = 0$. To make comparisons of utility between candidates, we add back the monetary component associated with the outside option: $u(b,a) + \theta_0 \cdot \log(a)$.

Estimation: First Step. We estimate labor supply parameters via a two-step procedure. We first estimate type distribution parameters $\boldsymbol{\beta}$ and amenity values $\boldsymbol{A}_j$ via maximum likelihood. Our strategy is based on a simple observation: if i accepts an offer from j and rejects an offer from k when $b_{ij} = b_{ik}$, then by revealed preference:

$$\boldsymbol{Q}_i'(\boldsymbol{A}_j - \boldsymbol{A}_k) \geq \xi_{ik} - \xi_{ij}. \tag{19}$$

Candidates often have several offers at the same bid, most often equal to their ask or at round numbers. Therefore, we construct the connected set of firms using a subset of bids $S = \{b_{ij} \mid b_{ij} > 0 \text{ and } \exists\, k \neq j \text{ s.t. } b_{ik} = b_{ij}\}$. This subset contains more than half of all bids. Making this restriction allows us to non-parametrically difference out $u(b, a)$, thereby obviating the need for instruments for the wage: identification of the A_{qj} does not rely on comparisons of offers with wages that may differ endogenously. By plugging in estimates $\widehat{A}_{qj}$ in the second step, allows us to control for the key unobserved confound when we turn to the estimation of labor supply elasticities.

To derive the probability of observing an arbitrary partial ordering of firms, it is useful to work with the re-parameterization $\rho_{qj} \propto \exp(A_{qj})$, with $\sum_{j=1}^{J} \rho_{qj} = 1$. Let $\sigma(\cdot) : \{1, \ldots, J\} \to \{1, \ldots, J\}$ denote a linear order over (or complete ranking of) all J alternatives. A multinomial logit model of rankings, "exploded logit," or Plackett-Luce (Plackett, 1975; Luce, 1959) model yields the following likelihood:

$$\Pr(\sigma(\cdot) \mid \boldsymbol{\rho}_q) = \prod_{r=1}^{J} \frac{\rho_{q\sigma^{-1}(r)}}{\sum_{s=r}^{J} \rho_{q\sigma^{-1}(s)}}.$$

We only observe candidates' partial orderings of firms, however. Following Allison and Christakis (1994), we could compute the probability of observing any particular partial ordering by summing over all linear orders that are consistent with that partial ordering. Even with a small number of alternatives, however, this strategy is computationally intractable: the number of concordant linear orders grows exponentially in the number of alternatives. Simulation methods that sample linear orders (e.g. Liu et al., 2019) are likely to be slow, and introduce additional sources of noise.

We circumvent this issue by implementing a novel numerical approximation to the partial order likelihood that greatly reduces the computational burden of estimation. In Appendix C, we show that:

$$\mathcal{P}\Big(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \boldsymbol{\rho}_q\Big) = \int_0^1 \prod_{j \in \mathcal{B}_i^1} \left(1 - u^{\rho_{qj} / \sum_{k \in \mathcal{B}_i^0} \rho_{qk}}\right) du. \tag{20}$$

This expression, and its derivatives, can be quickly and accurately approximated by numerical quadrature. Appendix C also details how we use this approximation to estimate $\boldsymbol{\beta}$ and $\boldsymbol{\rho}$ via a generalized EM-algorithm.

As in Sorkin (2018) and Avery et al. (2013), the estimated rank of firm j depends not on j's raw acceptance probability, but rather on the composition of firms to which j was revealed preferred to. Sorkin (2018) summarizes this property as a recursion: highly-ranked firms are those that are revealed-preferred to other highly-ranked firms. Avery et al. (2013) note that producing rankings in this way is robust to the strategic manipulations of the units being ranked—a key property in our setting. While we do not present a formal proof of consistency here, parameter consistency of the MLE for similar models has been established under sequences in which the number of items to be ranked (here, the number of firms J) grows asymptotically, avoiding the usual incidental parameters problem (Neyman and Scott, 1948; Simons and Yao, 1999).[16]

Estimation: Second Step. Next, we estimate the remaining labor supply elasticity and outside option parameters $\Theta = \{\theta_0, \theta_1, \boldsymbol{A}_0\}$ via GMM. We first construct model-implied probabilities of accepting an interview request as function of Θ, plugging in $\widehat{\boldsymbol{\beta}}$ and $\widehat{\boldsymbol{\rho}}$ from the first step. Letting $\Lambda(x) = \frac{\exp(x)}{1+\exp(x)}$ denote the logistic CDF, the model-based estimate of $\Pr(D_{ij} = 1 \mid b_{ij}, x_i)$ given parameters Θ is:

$$m(b_{ij}, x_i \mid \Theta) = \sum_{q=1}^{Q} \alpha_q(x_i \mid \widehat{\boldsymbol{\beta}}) \cdot \Lambda\Big((\theta_0 + \theta_1 \cdot \mathbf{1}[b_{ij} < a_i]) \cdot \log(b_{ij}/a_i) + \widehat{A}_{qj} - A_{q0}\Big).$$

We compute the sample analogues of moment conditions of the form:

$$\mathbb{E}\Big[x_i \cdot (D_{ij} - m(b_{ij}, x_i \mid \Theta))\Big] = 0 \;\text{ and }\; \mathbb{E}\Big[z_j \cdot (D_{ij} - m(b_{ij}, x_i \mid \Theta))\Big] = 0,$$

stacking them in the vector $\widehat{m}(\Theta)$. Θ is estimated by minimizing the GMM criterion:

$$\widehat{\Theta} = \arg\min_{\Theta} \widehat{m}(\Theta)' \, \boldsymbol{W} \, \widehat{m}(\Theta)$$

for a symmetric, positive-semidefinite weighting matrix $\boldsymbol{W}$.[17]

[16] Simons and Yao (1999) established the consistency and asymptotic normality of the maximum likelihood estimator of the parameters of models of paired comparisons under asymptotics that hold fixed the number of comparisons available between each pair of choices, but let the number of choices tend to infinity. Han et al. (2022) generalizes this result, establishing uniform consistency of the MLE even for extremely sparse comparison matrices.

[17] In practice, we use an efficient two-step GMM procedure: we produce an initial estimate $\widehat{\Theta}^0$, setting $\boldsymbol{W}^0$ equal to the identity matrix. In the second step, we set $\boldsymbol{W}$ equal to the inverse of the

5.2 Constructing Firms' Beliefs

Identification. Definition 1 specified a general form for beliefs in equilibrium which depends upon the probability that a firm's bid ranks highest among all available options. Given our multinomial logit assumption, that probability depends on the *inclusive value* Λ_i, which takes the form $\Lambda_i = \log\left(\sum_{k:b_{ik}>0} \exp\left(u(b_{ik}, a_i) + Q_i' A_k\right)\right)$:

$$\Pr\left(V_{ij} = V_i^1 \mid \Lambda_i, b_{ij} = b\right) = \exp\left(u(b, a_i) + Q_i' A_j\right) \Big/ \exp\left(\Lambda_i\right). \tag{21}$$

Using this expression, we may re-write firms' beliefs as:

$$G_{ij}(b) = \sum_{q=1}^{Q} \alpha_q(\omega_{ij}^Q) \cdot \int \left[\exp\left(u(b, a_i) + A_{qj}\right) \Big/ \exp\left(\lambda\right)\right] dF_{\Lambda|Q}\left(\lambda \mid Q_i = q, \omega_{ij}^{\Lambda}\right). \tag{22}$$

In the classic first-price auction setting, $G_{ij}(b)$ is nonparametrically identified by the observed distribution of bids when bidders have rational expectations: because the seller accepts the highest bid, the empirical CDF of winning bids can be used as an estimate of $G_{ij}(b)$. This is the basic intuition of the approach of Guerre, Perrigne and Vuong (2000) (GPV). Our estimation strategy extends the logic of GPV to a setting where $G_{ij}(b)$ depends upon both the monetary and non-monetary components of the bid. Namely, given estimates of labor supply parameters, we can construct inclusive values for every candidate, which we then treat as data. The empirical distribution of Λ_i therefore identifies $G_{ij}(b)$, and can be used to construct estimates of firm beliefs under each conduct assumption.

Estimation. We construct approximations to $G_{ij}(b)$ under each model of conduct, which are combinations of assumptions about 1) firms' beliefs about the distribution of $\Lambda_{iq} = \Lambda_i \mid Q_i = q$ and 2) firms' beliefs about the distribution of preference types Q_i:

Monopsonistic Competition: Monopsonistically-competitive firms do not account for the contribution of their own bid in the inclusive value Λ_i—in other words, $\{b_{ij}, \boldsymbol{A}_j\} \notin \omega_{ij}^{\Lambda}$. Under this assumption, firms' beliefs are:

$$G_{ij}(b) = \sum_{q=1}^{Q} \alpha_q(\omega_{ij}^Q) \cdot \left(\exp\left(u(b, a_i) + A_{qj}\right) \times \mathbb{E}\left[\exp\left(-\Lambda_{iq}\right) \mid \omega_{ij}^{\Lambda}\right]\right). \tag{23}$$

Since firms are assumed to have rational expectations conditional on the information

covariance matrix of the moment conditions evaluated at $\widehat{\Theta}^0$.

ω_{ij}^{Λ}, the quantity $\mathbb{E}[\exp(-\Lambda_{iq}) \mid \omega_{ij}^{\Lambda}]$ can be approximated by regressing $\exp(-\Lambda_{iq})$ on a flexible function of the variables contained in ω_{ij}^{Λ} (which include x_i, z_j, and market-level covariates). The beliefs of monopsonistically-competitive firms $G_{ij}^{m}(b)$ are proportional to $(b/a_i)^{\theta_0+\theta_1 \mathbf{1}[b<a_i]}$, and markdowns are a constant fraction of the wage on either side of $b_{ij} = a_i$: $\frac{\theta_0}{1+\theta_0}$ when $b_{ij} > a_i$, and $\frac{\theta_0+\theta_1}{1+\theta_0+\theta_1}$ when $b_{ij} < a_i$. When $b_{ij} = a_i$, we have that $\mu_{ij}^{m} = a_i/\varepsilon_{ij} \in \left[\frac{\theta_0}{1+\theta_0}, \frac{\theta_0+\theta_1}{1+\theta_0+\theta_1}\right]$.

Oligopsony: Oligopsonistic firms accurately account for the contribution of their bid on the inclusive value Λ_i. Under this assumption, the distribution of inclusive values conditional on ω_{ij}^{Λ} is given by $\Lambda_{iq} \mid \omega_{ij}^{\Lambda} \sim \exp(u(b_{ij}, a_i) + A_{qj}) + \exp(\Lambda_{iq}^{-j})$, where $\Lambda_{iq}^{-j} = \log(\sum_{k \neq j: B_{ik}=1} \exp(u(b_{ik}, a_i) + Q_i' A_k))$ denotes i's leave-j-out inclusive value. Denote the probability distribution of Λ_{iq}^{-j} by $F_{\Lambda_q^{-j}}$. Firms' beliefs are then:

$$G_{ij}(b) = \sum_{q=1}^{Q} \alpha_q(\omega_{ij}^{Q}) \cdot \int \left(\frac{\exp\Big(u(b, a_i) + A_{qj}\Big)}{\exp\Big(u(b_{ij}, a_i) + A_{qj}\Big) + \exp\Big(\lambda\Big)} \times dF_{\Lambda_q^{-j}}\Big(\lambda \mid \omega_{ij}^{\Lambda}\Big) \right). \quad (24)$$

Again, since firms' beliefs are assumed to be consistent, $F_{\Lambda_q^{-j}}\Big(\lambda \mid \omega_{ij}^{\Lambda}\Big)$ can be approximated by computing the distribution of leave-one-out inclusive values in the sample—for instance, by running a series of quantile regressions of Λ_{iq}^{-j} on a flexible function of the variables contained in ω_{ij}^{Λ}. These estimates can then be used to construct a numerical approximation to the integral over the distribution of leave-j-out inclusive values. Unlike monopsonistic competition, there is no simple closed-form expression for markdowns in the oligopsony case.

Type Predictive: Type-predictive firms use observed profile characteristics x_i to forecast candidate types ($\omega_{ij}^{Q} = x_i$). In this case, we approximate firms' beliefs using the estimated prior over types, $\alpha_q(\omega_{ij}^{Q}) = \alpha_q(x_i \mid \widehat{\boldsymbol{\beta}})$.

Not Predictive: Not-predictive alternative firms do not use observed profile characteristics x_i to forecast candidate types ($\omega_{ij}^{Q} = \varnothing$). In this case, we assume that firms weight type-specific win probabilities by the average probability of type membership, $\alpha_q(\omega_{ij}^{Q}) = \overline{\alpha}_q = \frac{1}{N} \sum_{i=1}^{N} \alpha_q(x_i \mid \widehat{\boldsymbol{\beta}})$.

We produce approximations to $G_{ij}(b)$ under all four combinations of these conduct assumptions, as well as the baseline perfect competition assumption.

5.3 Labor Demand

Identification: Let $G_{ij}^m(b)$ denote firms' beliefs under model m. It is useful to return to the case where $G_{ij}^m(b)$ is differentiable, with derivative $g_{ij}^m(b)$. As before, bids must satisfy the following first-order condition with equality in this case:

$$\varepsilon_{ij}^m(b) = b + \frac{G_{ij}^m(b)}{g_{ij}^m(b)} = \gamma_j^m(x_i, \nu_{ij}^m).\text{[18]} \tag{25}$$

Crucially, the inverse bidding function is *known* given choice of model m and labor supply parameters: in a Bayes-Nash Equilibrium, valuations are "revealed" by the bid. If the function $\varepsilon_{ij}^m(\cdot)$ is an injection, then a unique valuation $\varepsilon_{ij}^m = \varepsilon_{ij}^m(b_{ij})$ can be inferred for every bid b_{ij}. Conditional moment restrictions of the form $\mathbb{E}[\nu_{ij}^m \mid \Omega_{ij}] = 0$ can then be used to estimate $\gamma_j^m(x_i, \nu_{ij})$ (e.g. by regressing ε_{ij}^m on flexible functions of x_i and z_j). The parameters that govern $\gamma_j^m(\cdot,\cdot)$ are identified given sufficient variation in both ε_{ij}^m and covariates. This approach is taken by Backus et al. (2021) in their analysis of the common-ownership hypothesis. Our setting differs from this example in two important ways, both of which motivate our maximum likelihood framework.

First, $G_{ij}^m(b)$ is not differentiable at $b = a$ and so the first-order condition need not hold at that point. Appendix D establishes that bidding strategies $b_{ij}^m(\cdot)$ and option values $\pi_{ij}^{m*}(\cdot)$ are nevertheless continuous, monotonic functions in ε_{ij}, due to the log-concavity of F_ξ and shape restrictions on $u(b,a)$. In particular, $b_{ij}^m(\cdot)$ is strictly increasing in ε_{ij} outside an interval $[\varepsilon_{ij}^{m-}, \varepsilon_{ij}^{m+}]$, and is equal to a_i when ε_{ij} is inside that interval, while $\pi_{ij}^{m*}(\cdot)$ is strictly increasing over all ε_{ij}. Bids therefore partially identify valuations, motivating our use of a Tobit-style likelihood: $b_{ij} \neq a_i$ maps to a unique valuation, while $b_{ij} = a_i$ maps to an interval of possible valuations $[\varepsilon_{ij}^{m-}, \varepsilon_{ij}^{m+}]$. Second, selection is a key feature of our setting: firms only bid on candidates for whom $\pi_{ij}^{m*}(b_{ij}^m(\varepsilon_{ij})) \geq c_j$. The conditional moment restriction $\mathbb{E}[\nu_{ij}^m \mid \Omega_{ij}] = 0$ therefore cannot be used to estimate the labor demand parameters, since $\mathbb{E}[\nu_{ij}^m \mid \Omega_{ij}] > 0$ when $b_{ij} > 0$. While selection poses an estimation challenge, different conduct assumptions imply different patterns of selection, which increases the power of our test.

Estimation: We implement a selection correction using the fact that for each m, bids reveal not only ε_{ij}, but also the maximized value of firms' objective functions (See Appendix D). When $b_{ij} \neq a_i$, we construct the implied option value under model m, and when $b_{ij} = a_i$, we construct an upper bound on that qunatity. We denote these values by $\widehat{\pi}_{ij}^{m*}$, and use them to construct a consistent estimate of each firm j's interview cost threshold for each m by setting:

$$\widehat{c}_j^m = \min_{i:B_{ij}=1} \widehat{\pi}_{ij}^{m*} \overset{\text{a.s.}}{\to} c_j^m. \tag{26}$$

The consistency of our estimate of c_j necessarily depends upon the number of observations per firm growing without bound. See Appendix E for a proof of this result.[19]

Using this estimate, we can compute a lower bound on the valuation associated with each bid, which we use to implement a selection correction. Because $\pi_{ij}^{m*}(\cdot)$ is a strictly increasing function, there is a unique lower-bound valuation $\underline{\varepsilon}_{ij}^m$ at which firm j is indifferent between bidding and not bidding on candidate i. This lower bound controls the selection into bidding: employer j must draw a valuation of at least $\underline{\varepsilon}_{ij}^m$ to make a bid on candidate i, and so the distribution of valuations is censored from below by $\underline{\varepsilon}_{ij}^m$. We construct candidate-specific lower bounds by numerically inverting the option value function; $\widehat{\underline{\varepsilon}}_{ij}^m$ is the number that sets $\pi_{ij}^{m*}\left(\widehat{\underline{\varepsilon}}_{ij}^m\right) = \widehat{c}_j^m$. We use these lower bound estimates to construct the likelihood contribution of each bid:

$$\begin{aligned}\mathcal{L}_{ij}^m(\Psi^m) &= \Pr\left(\varepsilon_{ij} = \varepsilon_{ij}^m(b_{ij}) \mid \varepsilon_{ij} \geq \widehat{\underline{\varepsilon}}_{ij}^m,\ \Psi^m\right)^{\mathbf{1}[b_{ij}\neq a_i]} \times \Pr\left(\varepsilon_{ij} \in [\varepsilon_{ij}^{m-}, \varepsilon_{ij}^{m+}] \mid \varepsilon_{ij} \geq \widehat{\underline{\varepsilon}}_{ij}^m,\ \Psi^m\right)^{\mathbf{1}[b_{ij}=a_i]} \\ &= \left(\frac{f_\varepsilon\left(\varepsilon_{ij}^m(b_{ij}); \Psi^m\right)}{1 - F_\varepsilon\left(\widehat{\underline{\varepsilon}}_{ij}^m; \Psi^m\right)}\right)^{\mathbf{1}[b_{ij}\neq a_i]} \times \left(\frac{F_\varepsilon\left(\varepsilon_{ij}^{m+}; \Psi^m\right) - F_\varepsilon\left(\max(\varepsilon_{ij}^{m-}, \widehat{\underline{\varepsilon}}_{ij}^m); \Psi^m\right)}{1 - F_\varepsilon\left(\widehat{\underline{\varepsilon}}_{ij}^m; \Psi^m\right)}\right)^{\mathbf{1}[b_{ij}=a_i]},\end{aligned} \tag{27}$$

where Ψ^m denotes the parameters for model m, $f_\varepsilon(\cdot; \Psi^m)$ is the density of ε_{ij}, $F_\varepsilon(\cdot; \Psi^m)$ is the CDF of ε_{ij}, $\varepsilon_{ij}^m(\cdot)$ is the inverse bidding function for model m, and ε_{ij}^{m+} and ε_{ij}^{m-} are the model-implied upper and lower bounds on ε_{ij} when $b_{ij} = a_i$.[20]

Parameterization: We make the following assumptions about the functional forms of $\gamma_j(x_i, \nu_{ij})$ and the distribution of ν_{ij}:

$$\gamma_j(x_i, \nu_{ij}) = \exp\left(z_j'\Gamma x_i + \nu_{ij}\right),\ \ z_j'\Gamma x_i = \sum_k \sum_\ell \gamma_{k\ell} z_{jk} x_{i\ell}, \text{ and } \ \nu_{ij} \overset{iid}{\sim} N(0, \sigma_\nu).$$

where both x_i and z_j include a constant. For each model m, we estimate Γ^m and σ_ν^m by maximizing the log-likelihood of the full set of bids in the analysis sample.

[19] Our proof of the consistency of $\widehat{c}_j^m$ for each firm j (and model m) closely follows the proof of Lemma 1 (ii) of Donald and Paarsch (2002).

[20] Our approach—concentrating c_j out of the likelihood by computing the minimum order statistic—is similar to that of Donald and Paarsch (1993, 1996, 2002), who consider models in the classic procurement auction setting. Given m, the c_j are functions of only the labor supply parameters, which we treat as data. Because the c_j do not depend upon any of the labor demand parameters, our procedure yields a proper likelihood (unlike some of the cases they consider).

6 Results

6.1 Rejecting the Single Type Model of Labor Supply

We estimate several versions of the labor supply model in order to specify the number of latent preference types Q as well as how type membership is related to candidate observables. For each pair of models—a given number of types and a given set of observables used to define type membership—we calculate a standard likelihood ratio statistic and compute the appropriate χ^2 p-value. In addition to formal likelihood ratio (LR) statistics, we also compute a more directly-interpretable "goodness-of-fit" (GoF) statistic for each model. This statistic is simply the fraction of pairwise revealed-preference comparisons that are concordant with the estimated rankings:

$$\text{GoF} = N_{pw}^{-1} \sum_{i=1}^{N} \sum_{q=1}^{Q} \sum_{j \in \mathcal{B}_i^1} \sum_{k \in \mathcal{B}_i^0} \left(\alpha_q(x_i \mid \widehat{\boldsymbol{\beta}}) \cdot \mathbf{1}[\widehat{A}_{qj} \geq \widehat{A}_{qk}] \right),$$

where N_{pw} is the total number of pairwise comparisons implied by revealed preference.

Table 2 reports these goodness-of-fit statistics for several versions of our labor supply model. Each row corresponds to a number of types (from one to four) and each column corresponds to the observables leveraged to construct type membership. The first column allows men and women to have different rankings of firms, and the second column splits candidates between above- and below-median experience. The last column leverages all the observables we access for the candidates to define latent preference groupings. As benchmark, a model that assigned random numbers for each A_{qj} would in expectation yield a GoF statistic of 0.5. In contrast, as reported in the first row of Table 2, the one-type model increases goodness-of-fit over that baseline to 0.67.[21] This relatively large increase in explanatory power compared to the benchmark indicates significant vertical differentiation of firms.

Column 1 of Table 2, assigns women and men to distinct preference types. Doing so yields no additional explanatory power over the revealed preferences in the data relative to a one-type model: the GoF statistic increases imperceptibly (from 0.672 to 0.680), and the formal LR test fails to reject the null that the two-type and one-type models are equivalent ($p = 0.27$). This finding mirrors that of Sorkin (2017), who also finds that the implied preference orderings of men and women over firms

[21] The goodness of fit measure varies slightly across the three columns because the estimation samples are different. For instance, to be ranked in the model that splits types by gender, a firm needs to have been accepted once and rejected once by candidates of both genders while the connected set of firms in the model that splits types by experience is made of firms that have been accepted once and rejected once by candidates of each experience level.

Table 2: Candidate Preference Model Goodness-of-Fit

		(1) Split on Gender	(2) Split on Experience	(3) Model-Based Clusters
One Type	Log. L	-43,463	-45,184	-47,155
	GOF	0.672	0.673	0.677
Two Types	Log. L	-42,962	-44,535	-45,558
	GOF	0.680	0.684	0.744
	p(2,1)	0.271	<0.001	<0.001
Three Types	Log. L	-	-	-44,594
	GOF	-	-	0.779
	p(3,2)	-	-	<0.001
Four Types	Log. L	-	-	-43,857
	GOF	-	-	0.808
	p(4,3)	-	-	>0.999
Number of Firms		975	1,128	1,649
Number of Candidates		13,658	13,830	14,344
Number of Comparisons		209,934	222,935	235,827

Note: This table reports goodness-of-fit (GOF) measures and p-values to adjudicate between labor supply models with different numbers of types (rows). Each column represents a different way to split candidates into preference types. The GOF statistic is calculated as the fraction of pairwise comparisons correctly predicted by the model, $\mathbb{E}\big[\big(\widehat{A}_{qj} > \widehat{A}_{qk}\big) \times \big(j \succ_i k\big)\big]$, and p-values are calculated via the likelihood ratio. Each column corresponds to a different sample determined by (overlapping, if relevant) connected sets.

are extremely similar. Splitting by experience in Column 2 does only marginally better: while the LR test can reject the null that the two-type model is equivalent to the one-type model ($p < 0.001$), the GoF statistic only increases by 1.6 percentage points. However, using the full set of observables to define types (Column 3) performs markedly better than the gender- and experience-split models. With two types, the GoF statistic for the model-based clustering is 0.744, or 10.7 percentage points higher than for the gender or experience splits. Sequential LR tests between the one- and two-type models and two- and three-type models both reject the null that the more-complex models are equivalent to the simpler models ($p < 0.001$). However, we are unable to reject the null hypothesis that the four-type alternative is equivalent to the three-type model. We therefore adopt the three-type version as our baseline.[22]

Plugging in the estimated rankings into our second-step GMM procedure yields

[22] Figure A.3 provides additional evidence of the quality of the fit of the preferred 3-type model by plotting the relationship between the model-implied probabilities that a given bid will be accepted against the empirical acceptance probability. The Figure documents that the model-implied probabilities are extremely close to the actual acceptance probabilities.

the following labor supply elasticity parameter estimates:

$$u(b_{ij}, a_i) = \left[\underset{(0.33)}{4.05} + \underset{(0.28)}{1.58} \cdot \mathbf{1}[b < a_i]\right] \cdot \log(b/a_i).$$

These estimates are similar to others in the literature: for instance, Berger et al. (2022) report an estimate of 3.74, while Azar et al. (2020) report an estimate of 5.8.[23]

In order to validate the estimated rankings, we return to the reasons candidates provide when rejecting an interview request, described in Section 2.3. We now divide the list of reasons candidates choose from into two categories: personal reasons that should correspond to a low draw of ξ_{ij} and job-related reasons that should correspond to a low value of A_{qj}. If the model provides a good fit to the data, then we should find that candidates are more likely to reject highly-ranked firms for personal reasons than job-related reasons relative to lower-ranked firms. To test this hypothesis, Figure 3 plots the probability that a firm was rejected for a job-related reason as a function of firms' ordinal rankings (where lower ranks are better). This figure confirms that workers are significantly less likely to reject the most-preferred companies for job-related reasons than they are for lower-ranked companies.

6.2 Significant Vertical and Horizontal Differentiation of Firms

Figure 4 illustrates the scale of vertical and horizontal differentiation of firms implied by our preferred model estimates. To understand the importance of amenities relative to pay, we compute a willingness-to-accept statistic (WTA) for every firm. The statistic is equal to the fraction of a candidate's ask that the model implies a firm must offer to make that candidate indifferent between accepting or rejecting an interview request, on average. We compute WTA_{qj} as the number that solves:

$$\left(\widehat{\theta}_0 + \widehat{\theta}_1 \times \mathbf{1}[\text{WTA}_{qj} < 1]\right) \times \log(\text{WTA}_{qj}) + \widehat{A}_{qj} - \widehat{A}_{q0} = 0.$$

where A_{q0} is the q-th component of the vector of mean outside option values. Panel (a) of Figure 4 plots the distribution of the mean WTA at each firm, averaging over the population probabilities of each type:

$$\text{WTA}_j = \sum_{q=1}^{3} \overline{\alpha}_q \times \text{WTA}_{qj}.$$

[23] Note that, in contrast with other studies, our model allows for kinked labor supply and therefore our estimates of the parameter is 5.63 below the kink, i.e. when $b < w_i$, and 4.05 above the kink.

Figure 3: Interview Rejection Reasons as a Function of Firm Rankings

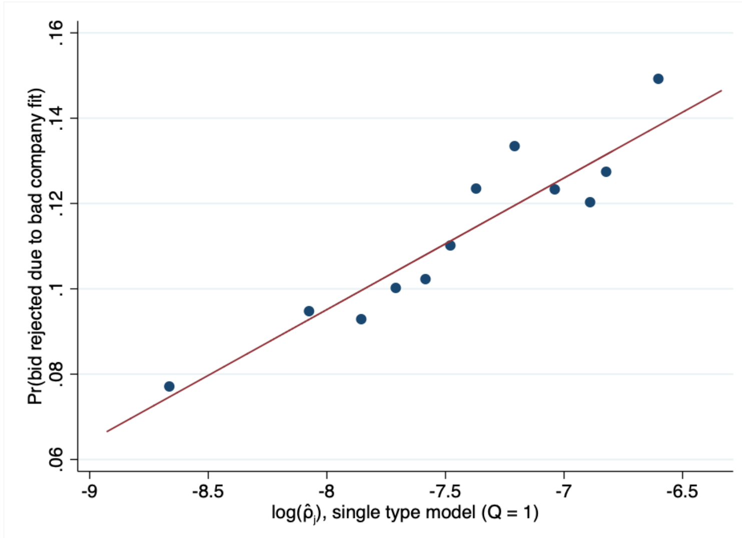

Note: This figure plots the probability that a firm was rejected for an amenity-related reason as a function of firms' ordinal rankings (where lower ranks are better). For a sub-sample (57%) of rejected bids, candidates opted to provide a justification. They can choose from justifications such as "company size", "insufficient compensation" or "company culture". The latter is the justification we label as "bad company fit". We plot the probability of rejection due to bad company fit against estimated rankings from the single-type model.

The average mean WTA is 0.99, indicating that candidates are willing to accept roughly 1% less than their ask at the average firm. The standard deviation (SD) of mean WTA across firms is 0.14, which suggests a large range of variability in the amenity values candidates attach to firms. Indeed, there is a nontrivial number of firms for which the average candidate would be willing to accept less than 80% of their ask, and an even larger number of firms for which candidates demand over 120% of their ask. Panel (b) illustrates the systematic component of horizontal differentiation. Here, we plot the within-firm standard-deviation of WTA_{qj} across preference types. The mean within-firm SD of WTA is 0.14, suggesting that horizontal differentiation is about as important as vertical differentiation. The implication of these estimates is that there is large scope for firms to exercise market power in the ways we have specified: due to significant horizontal differentiation, firms may stand to gain significantly from accurately predicting which candidates are in which preference groups, while the significant vertical differentiation implies that high-ranked firms, if acting strategically, can afford to mark down wages significantly. Given the significant scope for wage markdown based on preference heterogeneity, assessing whether firms are able to predict the types is crucial and Section 6.3 formally tests for it.

Figure 4: Differentiation between Firms

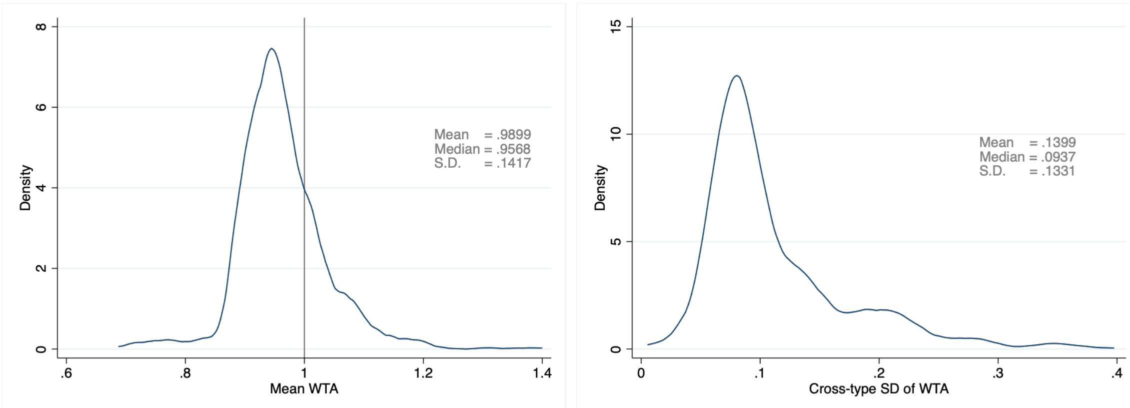

(a) Vertical Differentiation

(b) Horizontal Differentiation

Note: This figure illustrates the scale of vertical and horizontal differentiation of firms implied by our preferred model estimates. The Willingness to Accept (WTA) is equal to the fraction of a candidate's ask salary that the model implies a firm must offer to make that candidate indifferent between accepting or rejecting an interview request, on average. Panel (a) plots the distribution of the mean Willingness to Accept (WTA) at each firm, averaging over the population probabilities of each type. Panel (b) illustrates the systematic component of horizontal differentiation, plotting the distribution of the within-firm standard-deviation of (WTA) across preference types.

What firm characteristics are associated with higher amenity values? To partially answer this question, we report regressions of (standardized) estimates of A_{qj} on firm covariates z_j in the sample for which those covariates are available in Table 3. Here, larger values of A_{qj} correspond to better rankings. These covariates represent only a small fraction of the potential relevant characteristics candidates may consider when they choose among job offers. Importantly, the ("all-in") amenity values we estimate do not depend upon exhaustive knowledge of what candidates value. Even with the relatively coarse covariates available, some clear patterns are evident. In particular, the basic evidence in Table 3 suggests a loose classification of groups as "baseline" (group 2), "risk-averse" (group 3), and "risk-loving" (group 1). Relative to baseline, members of group 3 are more interested in working at larger, established firms for which there may be less employment risk, while members of group 1 are more interested in working at the smallest firms that may be more risky bets.

How do worker characteristics shift the probability of preference group membership? To answer this question, we compute the model-implied posterior probabilities of type membership for every candidate and correlate those probabilities with candidate characteristics (our discussion of the EM algorithm in Appendix C covers the construction of these probabilities). We find that women are 7 percentage points more likely to be in the risk-averse group and 7 percentage points less likely to be in the

Table 3: Which Firm Characteristics are Correlated with Amenity Values?

	(1) $\widehat{A}_{1j}$	(2) $\widehat{A}_{2j}$	(3) $\widehat{A}_{3j}$
Year Founded	0.00521	0.00641	-0.00502
	(0.00374)	(0.00385)	(0.00358)
15-50 Employees	-0.0836	0.114	0.105
	(0.0881)	(0.0907)	(0.0843)
50-500 Employees	-0.0531	0.222**	0.337***
	(0.0829)	(0.0853)	(0.0793)
500+ Employees	-0.00169	0.287**	0.640***
	(0.0993)	(0.102)	(0.0950)
Finance	0.0153	0.0474	-0.105
	(0.0694)	(0.0715)	(0.0664)
Tech	-0.0179	-0.0312	-0.0594
	(0.0567)	(0.0584)	(0.0543)
Health	0.0174	0.117	-0.0778
	(0.0911)	(0.0938)	(0.0872)
adj. R^2	-0.004	0.009	0.085
N	913	913	913

Note: This table reports regressions of standardized estimates of firm amenity values, $\widehat{A}_{qj}$, on basic firm characteristics z_j. The omitted category for the number of employees is 0-15. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

risk-loving group, while candidates with above-median experience are 10 percentage points less likely to be in the risk-averse group and 9 percentage points more likely to be in the risk-loving group. While there is significant residual variation in preferences conditional on covariates, our preferred model estimates suggest that covariates are indeed predictive of preference type.

6.3 Testing Between Models of Conduct

We next describe the results of implementing our estimation and testing framework for labor demand. As a preliminary matter, Figure 5a plots the "first stage" relationship between the model-implied inclusive values (Λ_i and Λ_i^{-j}) and the instrumental variable (t_{ij}), conditional on firm and candidate covariates and two-week period dummies.

Figure 5: Vuong Test

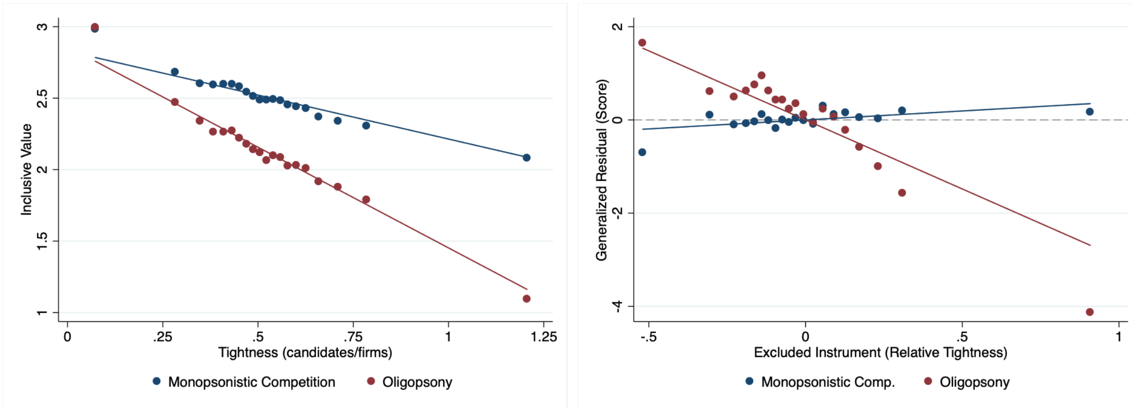

(a) First Stage

(b) Visualizing the Vuong Test

Note: Panel (a) plots the "first stage" relationship between the model-implied inclusive values Λ_i and Λ_i^{-j} and the instrumental variable t_{ij}, conditional on firm covariates z_j and candidate covariates x_i and two-week period dummies. Panel (b) plots the relationship between generalized residuals and the excluded instrument for the non-predictive monopsonistic competition and oligopsony models. Under proper specification, the correlation of the generalized residuals and the excluded instrument should be zero (the dashed line). The larger the deviation from zero, the greater the degree of mis-specification of the model.

Intuitively, the fewer candidates there are relative to firms (low t_{ij}), the more offers those candidates should receive, and the larger the inclusive values associated with their offer sets should be. This intuition is borne out in Figure 5a: both full- and leave-one-out inclusive values are strongly negatively related to labor market tightness.

Table 4 reports the results of implementing our pairwise testing procedure on the five models we estimated, using the moment-based versions of the Vuong test. In this table, positive values imply the row model is preferred to the column model. Under the null of model equivalence, the test statistics are asymptotically normal with mean zero and unit variance. The test statistics we report suggest that we can resoundingly reject the null hypothesis of model equivalence in most cases. The "Perfect Competition" model unambiguously performs the worst of all the models we tested. Among the remaining alternatives, the two monopsonistic competition models outperform the two oligopsony models, with the not-predictive monopsonostic competition alternative performing best.

We visualize these results in Figure 5b, which plots generalized residuals for two alternative models against the excluded instrument. Under proper specification, the generalized residuals should not be correlated with the instrument: the further a model's generalized residuals are from the x-axis, the greater the degree of mis-

Table 4: Non-Nested Model Comparison Tests (Rivers and Vuong, 2002)

Model	(1) Monopsonistic Comp.	(2)	(3) Oligopsony	(4)
	Not Predictive	Type Predictive	Not Predictive	Type Predictive
Perfect Competition	-54.84	-54.40	-39.92	-39.91
Monopsonistic, Not Predictive	–	7.83	3.98	2.69
Monopsonistic, Type Predictive		–	2.77	1.54
Oligopsony, Not Predictive			–	-3.67
Oligopsony, Type Predictive				–

Note: This table reports test statistics from the Rivers and Vuong (2002) non-nested model comparison procedure. Positive values imply the row model is preferred to the column model. Under the null of model equivalence, the test statistics are asymptotically normal with mean zero and unit variance.

specification. In the figure, the generalized residuals for the monopsonsitic competition alternative are closely aligned with the x-axis, while the generalized residuals for the oligopsony alternative are strongly negatively related to tightness.

Our tests therefore suggest that models of firm behavior in which firms both ignore strategic interactions in wage setting and do not tailor wage offers to candidates on the basis of predictable preference variation are closer approximations to firms' true bidding behavior on the platform than are models in which firms act strategically and tailor offers. In Appendix F, we report the testing results using the original Vuong (1989) likelihood comparison test, which yield qualitatively identical model comparisons. In the following analysis, we adopt the not-predictive monopsonostic competition model as our preferred model of conduct.

6.4 Comparing Demand Estimates

Our preferred model of conduct is the simplest of the four imperfect competition alternatives we specified. Under that model of conduct, there is little to no room for variation in markdowns between firms or differences in markdowns across candidates within firms. How much do the conclusions of the more complicated models of wage setting differ from those of the preferred model? To answer this question, we report comparisons between pairs of models of increasing complexity, adding one conduct assumption at a time. First, we compare the preferred model to the oligopsony model, maintaining the assumption that firms are not type-predictive. Then, we compare the not-predictive oligopsony model to its type-predictive version.

Assuming firms are not type predictive, Figure 6a plots the distributions of predicted markdowns in dollars under monopsonistic competition and oligopsony. We

Figure 6: Contrasting labor market implications across models

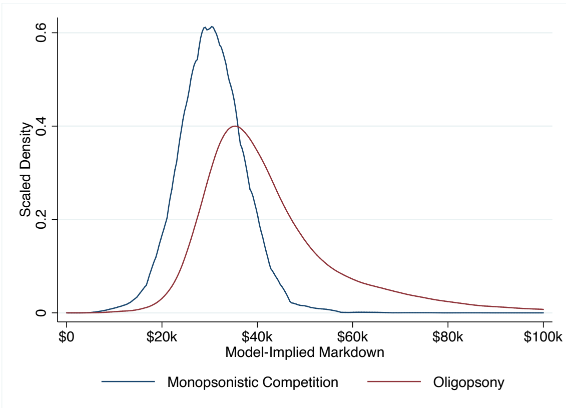

(a) Predicted Markdowns

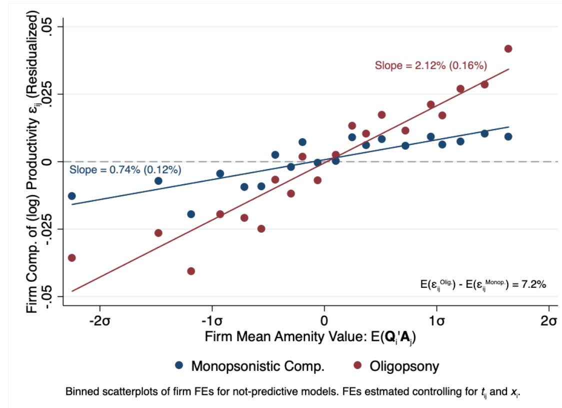

(b) Between-firm productivity variation

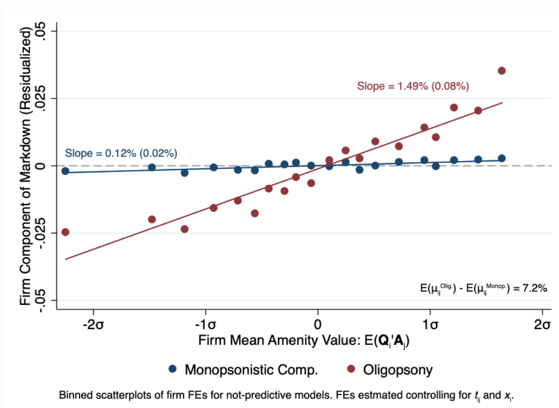

(c) Between-firm markdown variation

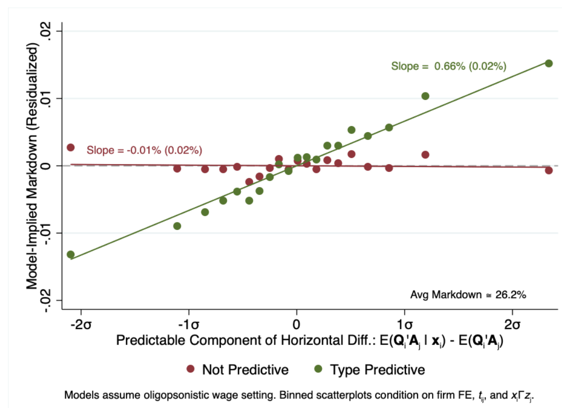

(d) Within-firm markdown variation

Note: Panel (a) plots the distribution of predicted markdowns under (not type-predictive) monopsonistic competition and oligopsony. Panel (b) plots the firm components of model-implied productivity for the preferred model and the not-predictive oligopsony against the standardized mean firm amenity value. Panel (c) plots de-meaned model-implied markdowns against mean firm amenity values, for the preferred model and the not-predictive oligopsony. Panel (d) plots de-meaned model-implied markdowns on the predictable component of horizontal preference variation, for the not-predictive and predictive oligopsony models.

compute markdowns as the difference between the model-implied firm valuation and the observed bid: $\varepsilon_{ij}^{m} - b_{ij}$.[24] The two alternatives predict markedly different distributions of markdowns. Under the preferred, monopsonistic model, the average predicted markdown is \$30,503, with a standard deviation of \$6,658. Further, the distribution of markdowns is relatively symmetric: the mean and median of the dis-

[24] In cases where the implied valuation is not point identified (the bid is equal to ask), we take the midpoint of the model-implied range of valuations: $(\varepsilon_{ij}^{m+} + \varepsilon_{ij}^{m-})/2 - b_{ij}$.

tribution are separated by less than \$300, and the skewness of the distribution of markdowns is just 0.35. By contrast, the oligopsony model predicts uniformly larger markdowns: the mean model-implied markdown under that assumption is \$43,385. Further, the distribution of markdowns under oligopsony is significantly more variable, with a standard deviation of \$16,357. Finally, the distribution of markdowns under oligopsony is highly skewed: the mean markdown is \$4,000 larger than the median markdown, and the skewness of the distribution is just over 2. The two sets of markdowns are positively correlated but the correlation is far from one, at 0.42. The large contrasts highlighted by Figure 6a illustrate the importance of understanding which form of conduct best describes firm behavior: different assumptions about the presence or absence of strategic interactions lead to strikingly different conclusions about the size of wage markdowns.

Monopsonistic competition and oligopsony not only yield different implications for the magnitude of aggregate markdowns, but also for firm-level variation in markdowns and productivity. In Figure 6b, we plot the firm components of model-implied productivity against the standardized mean firm amenity value. In both models, the relationship between amenities and productivity is positive: firms with relatively better amenities are more productive. But the slope of the relationship is nearly three times larger under oligopsony than under monopsonistic competition.[25] This leads to large differences in implied productivity dispersion across firms: in the preferred model, firms with the best amenities ($+2\sigma$) are 3% more productive than firms with the worst amenities (-2σ). Under the oligopsony alternative, that difference is 8.5%.

What drives the large differences between the conclusions implied by the two models? Under oligopsony, each firm internalizes a firm-specific labor supply elasticity that depends upon its amenities such that firms with better amenities should mark wages down more. Under monopsonistic competition, firms internalize upward-sloping firm-specific labor supply, but the elasticity of that supply does not depend upon firms' amenities. Figure 6c illustrates this empirically by reporting binned scatterplots of de-meaned model-implied markdowns against mean firm amenity values for the two models. Under oligopsony, firms with the best amenities mark down wages by 6 percentage points more than firms with the worst amenities. Under the monopsonistic competition assumption, there is essentially no room for different firms to set different markdowns, and so the relationship is essentially flat.

[25] Since wages and productivity move together in the monopsonistic competition model, our preferred estimates suggest the presence of augmenting differences (a positive (rather than negative) correlation between wages and amenities) similar to Lagos (2021).

Table 5: Variance Decomposition of Bids

	-log($\widehat{\mu}_{ij}$)	$z_j'\widehat{\Gamma}x_i$	$\widehat{\nu}_{ij}$
Panel A: Monopsonistic Competition			
Var($\cdot$)	0.003	0.897	0.097
Cov($\cdot$, log(b_{ij}))	-0.001	0.910	0.091
Panel B: Oligopsony			
Var($\cdot$)	0.133	0.680	0.219
Cov($\cdot$, log(b_{ij}))	0.101	0.777	0.122
Standard Deviation of log(b_{ij}) = 0.221.			

Note: This table shows the variance decomposition of log bids. Cells in the first row of each panel report the variance of each component, while cells in the second row of each panel report the covariance of each component with log bids. All cells are normalized by the variance of log bids. Results are for the not-predictive versions of each model.

Next, we add another layer of complexity to wage setting: allowing firms to be type-predictive. Figure 6d reports binned scatterplots of de-meaned model-implied markdowns on the predictable component of horizontal preference variation for the not-predictive and predictive oligopsony models. While the not-predictive model allows for systematic variation in markdowns between firms, it does not allow for systematic variation in markdowns within firms across candidates. This yields a flat relationship between markdowns and predictable horizontal preference variation. In contrast, the type-predictive alternative allows firms to optimally use the information about preferences revealed by observable candidate characteristics to mark down wages. Intuitively, the candidates who value a given firm's amenities relatively more will be offered lower wages. Our estimates imply that the wage offers a type-predictive firm makes to the workers who value its amenities the most are marked down 2.6pp more than the offers it makes to workers who value them the least.

To summarize the differences in labor market implications between alternative models, we report the results of a simple variance decomposition exercise in Table 5. For each estimated model, we use labor demand parameter estimates $\widehat{\Gamma}$ and model-implied match productivity $\widehat{\varepsilon}_{ij}$ to decompose bids into the sum of three components:

$$\log(b_{ij}) = -\underbrace{\log(\widehat{\mu}_{ij})}_{\text{markdown}} + \underbrace{z_j\widehat{\Gamma}x_i}_{\text{systematic comp.}} + \underbrace{\widehat{\nu}_{ij}}_{\text{idiosyncratic comp.}}, \tag{28}$$

where $\log(\widehat{\mu}_{ij}) = \log(b_{ij}) - \log(\widehat{\varepsilon}_{ij})$ and $\widehat{\nu}_{ij} = \widehat{\varepsilon}_{ij} - z_j\widehat{\Gamma}x_i$. We decompose the variation in log bids by computing the covariance of each term in the above equation with the log bid and dividing by the variance of log bids. Each term of this decomposition measures the relative contribution of each factor, and the terms sum to one. Table 5 reports this simple decomposition for both the preferred model and the not-predictive oligopsony model[26]. The monopsonistic competition assumption attributes almost none of the variation in bids to variation in markdowns: 91% of bid variation is driven by the systematic component of match productivity, while 9% is due to idiosyncratic match components. The oligopsony alternative, by contrast, apportions 10% of the variation in bids to variation in markdowns, 78% to variation in the systematic match component, and 12% to variation in the idiosyncratic match component.

In Appendix G, we present comparisons of estimated labor demand parameters between models. Under the preferred model, the average elasticity of ε_{ij} with respect to the ask is 0.89, with small and statistically insignificant differences in productivity between men and women (-0.47% (0.30%)). Under the oligopsony alternative, the average elasticity of ε_{ij} with respect to the ask is 0.74, with a larger and significant gender gap in productivity (-1.4% (0.44%)). These differences imply qualitatively different conclusions about the sources of gender gaps on the platform, given the substantial gender ask gap. Under the preferred model, just 7.5% of the gender gap in ε_{ij} is due to firms' perceived differences in productivity between men and women (conditional on ask), while the remaining 92.5% is driven by differences in asks. Under the oligopsony alternative, 22% of the gender gap in ε_{ij} is driven by perceived differences in productivity.

Finally, in Appendix H, we present a decomposition of gender gaps in welfare and counterfactual simulations of bidding behavior under various conduct assumptions to further explore gender gaps. There exists a large gender gap in the number and average amenity value of bids received by men and women, which maps into a large average gap in welfare as measured by the inclusive value of candidates' interview offer sets. These gaps are primarily driven by gender differences in the monetary value of bids received, but roughly one fifth of the gap in welfare can be attributed to the fact that women receive bids from firms with less attractive amenities than men. We then conduct counterfactual simulations to quantify the impact of imperfect competition on welfare and gender gaps. Relative to a perfect competition baseline, we find that firms make significantly more offers under the preferred model, but that the wages

[26] The results for the type-predictive versions of each model are nearly identical to their not-predictive counterparts, and so are not reported here

firms attach to those offers are lower. Relative to the preferred model, however, the average value of bids, the total number of bids, and welfare are significantly lower in simulated equilibria with strategic interactions. We also find that the form of conduct has important implications for gender gaps: relative to men, women receive significantly fewer bids when firms predict horizontal preference variation than when they do not. Imperfect competition exacerbates gender gaps relative to a price-taking baseline. Finally, we find that blinding employers to the gender of candidates generates only modest reductions in gender gaps.

7 Conclusion

This paper provides direct evidence about the nature of firms' wage-setting behavior by developing a testing procedure to adjudicate between many non-nested models of conduct in the labor market. In particular, we focus on two sets of alternatives relevant to ongoing debates in the labor literature: first, whether firms compete strategically (Berger et al., 2022; Jarosch et al., 2021), and second, whether firms tailor wage offers to workers' outside options (Postel-Vinay and Robin, 2002; Jäger et al., 2021; Flinn and Mullins, 2021). Applying our testing procedure, we find evidence against strategic interactions in wage setting as well as against the tailoring of offers to workers of different types. Although we study a specific labor market, these findings suggest that the relatively simple model of wage determination posited by Card et al. (2018) provides a reasonable approximation to firm wage-setting conduct in labor markets where many employers are competing for workers. Importantly, we find that incorrect conduct assumptions can lead to substantial biases: in our preferred model, wages are marked down by 18.2% on average, and markdowns do not vary systematically between firms or across workers at the same firm. Adopting alternate assumptions in which firms interact strategically in wage setting leads to average implied markdowns of 25.8% which vary substantially between firms. Further assuming that firms internalize predictable horizontal variation in preferences implies significant additional markdown heterogeneity across workers. Our results suggest that both of these patterns are inconsistent with the observed behavior of firms.

References

Abowd, John, Francis Kramarz, and David Margolis, "High Wage Workers and High Wage Firms," *Econometrica*, 1999, *67* (2), 251–333.

Allison, Paul D. and Nicholas A. Christakis, "Logit Models for Sets of Ranked Items," *Sociological Methodology*, 1994, *24*, 199–228.

Arnold, David, "Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes," Working Paper, UC San Diego 2021.

Avery, Christopher N., Mark E. Glickman, Caroline M. Hoxby, and Andrew Metrick, "A Revealed Preference Ranking of U.S. Colleges and Universities," *The Quarterly Journal of Economics*, 2013, *128* (1), 425–67.

Azar, José, Ioana Marinescu, Marshall Steinbaum, and Bledi Taska, "Concentration in US labor markets: Evidence from online vacancy data," *Labour Economics*, 2020, *66*, 101886.

_ , Steven T Berry, and Ioana Marinescu, "Estimating labor market power," NBER Working Paper 30365 August 2022.

Backus, Matthew and Gregory Lewis, "Dynamic Demand Estimation in Auction Markets," NBER Working Paper 22375 October 2020.

_ , Christopher Conlon, and Michael Sinkinson, "Common Ownership and Competition in the Ready-to-Eat Cereal Industry," NBER Working Paper 28350 January 2021.

Bagnoli, Mark and Ted Bergstrom, "Log-Concave Probability and Its Applications," *Economic Theory*, 2005, *26* (2), 445–69.

Bain, Joe S., "Relation of Profit Rate to Industry Concentration: American Manufacturing, 1936–1940," *The Quarterly Journal of Economics*, 1951, *65* (3), 293–324.

Barseghyan, Levon, Maura Coughlin, Francesca Molinari, and Joshua C. Teitelbaum, "Heterogeneous Choice Sets and Preferences," *Econometrica*, 2021, *89* (5), 2015–48.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor market power," *American Economic Review*, 2022, *112* (4), 1147–93.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–90.

Berry, Steven T., "Market Structure and Competition, Redux," Keynote, FTC Micro Conference 2021.

_ and Philip A. Haile, "Identification in Differentiated Products Markets Using Market Level Data," *Econometrica*, 2014, *82* (5), 1749–1797.

_ and _ , "Nonparametric Identification of Differentiated Products Demand Using Micro Data," NBER Working Paper 27704 August 2020.

Bhaskar, V., Alan Manning, and Ted To, "Oligopsony and Monopsonistic Competition in Labor Markets," *Journal of Economic Perspectives*, June 2002, *16* (2), 155–74.

Blinder, Alan, "Wage Discrimination: Reduced Form and Structural Estimates," *Journal of Human Resources*, 1973, *8* (4), 436–55.

Boal, William M. and Michael Ransom, "Monopsony in the Labor Market," *Journal of Economic Literature*, 1997, *35* (1), 86–112.

Bresnahan, Timothy F., "Competition and Collusion in the American Automobile Industry: The 1955 Price War," *Journal of Industrial Economics*, 1987, *35* (4), 457–82.

_ , "Chapter 17 Empirical studies of industries with market power," in "in," Vol. 2 of *Handbook of Industrial Organization*, Elsevier, 1989, pp. 1011–57.

Burdett, Kenneth and Dale T. Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 1998, *39* (2), 257–73.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36* (S1), 13–70.

_ , Jörg Heining, and Patrick Kline, "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, 05 2013, *128* (3), 967–1015.

Chamberlain, Edward and Joan Robinson, *The American Economic Review*, 1933, *23* (4), 683–85.

Chamberlain, Gary, "Asymptotic efficiency in estimation with conditional moment restrictions," *Journal of Econometrics*, 1987, *34* (3), 305–34.

DellaVigna, Stefano and Matthew Gentzkow, "Uniform Pricing in U.S. Retail Chains*," *The Quarterly Journal of Economics*, 06 2019, *134* (4), 2011–84.

Donald, Stephen and Harry Paarsch, "Piecewise Pseudo-Maximum Likelihood Estimation in Empirical Models of Auctions," *International Economic Review*, 1993, *34* (1), 121–48.

_ **and** _ , "Identification, Estimation, and Testing in Parametric Empirical Models of Auctions within the Independent Private Values Paradigm," *Econometric Theory*, 1996, *12* (3), 517–67.

_ **and** _ , "Superconsistent estimation and inference in structural econometric models using extreme order statistics," *Journal of Econometrics*, 2002, *109* (2), 305–40.

Duarte, Marco, Lorenzo Magnolfi, Mikkel Sølvsten, and Christopher Sullivan, "Testing firm conduct," *arXiv preprint arXiv:2301.06720*, 2023.

Dube, Arindrajit, Alan Manning, and Suresh Naidu, "Monopsony and Employer Mis-optimization Explain Why Wages Bunch at Round Numbers," NBER Working Paper 24991 January 2020.

_ , Jeff Jacobs, Suresh Naidu, and Siddharth Suri, "Monopsony in Online Labor Markets," *American Economic Review: Insights*, March 2020, *2* (1), 33–46.

Flinn, Christopher and Joseph Mullins, "Firms' Choices of Wage-Setting Protocols," Working Paper 2021.

Freeman, Richard B., "A Cobweb Model of the Supply and Starting Salary of New Engineers," *Industrial and Labor Relations Review*, 1976, *29* (2), 236–48.

Gandhi, Amit and Aviv Nevo, "Empirical Models of Demand and Supply in Differentiated Products Industries," in Ali Hortascu, Kate Ho, and Alessandro Lizzeri, eds., *Handbook of Industrial Organization*, 1 ed., Vol. 4, Elsevier, 2021.

Gentry, Matthew and Tong Li, "Identification in auctions with selective entry," *Econometrica*, 2014, *82* (1), 315–44.

Gibbons, Robert, Lawrence Katz, Thomas Lemieux, and Daniel Parent, "Comparative Advantage, Learning, and Sectoral Wage Determination," *Journal of Labor Economics*, 2005, *23* (4), 681–724.

Goldin, Claudia and Lawrence F. Katz, "A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation," *Journal of Labor Economics*, 2016, *34* (3), 705–46.

Gourieroux, Christian, Alain Monfort, Eric Renault, and Alain Trognon, "Generalised residuals," *Journal of Econometrics*, 1987, *34* (1-2), 5–32.

Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong, "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 2000, *68* (3), 525–74.

Guryan, Jonathan and Kerwin Kofi Charles, "Taste-based or Statistical Discrimination: The Economics of Discrimination Returns to its Roots," *The Economic Journal*, 2013, *123* (572), F417–F432.

Hall, Robert and Andreas Mueller, "Wage Dispersion and Search Behavior: The Importance of Nonwage Job Values," *Journal of Political Economy*, 2018, *126* (4), 1594–1637.

Han, Ruijian, Yiming Xu, and Kani Chen, "A general pairwise comparison model for extremely sparse networks," *Journal of the American Statistical Association*, 2022, pp. 1–11.

Horton, John J., Ramesh Johari, and Philipp Kircher, "Cheap Talk Messages for Market Design: Theory and Evidence from a Labor Market with Directed Search," NBER Working Paper 29445 November 2021.

Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin, "Granular Search, Market Structure, and Wages," NBER Working Paper 26239 September 2021.

Jäger, Simon, Chris Roth, Nina Roussille, and Benjamin Schoefer, "Worker Beliefs about Outside Options," *NBER Working Paper*, December 2021, (29623).

Lagos, Lorenzo, "Labor Market Institutions and the Composition of Firm Compensation: Evidence from Brazilian Collective Bargaining," Working Paper 2021.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler, "Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market," *American Economic Review*, 2022, (forthcoming).

Liu, Ao, Zhibing Zhao, Chao Liao, Pinyan Lu, and Lirong Xia, "Learning Plackett-Luce Mixtures from Partial Preferences," *Proceedings of the AAAI Conference on Artificial Intelligence*, Jul. 2019, *33* (01), 4328–35.

Luce, R. Duncan, *Individual Choice Behavior A Theoretical Analysis*, Wiley, 1959.

Manning, Alan, *Monopsony in Motion*, Princeton University Press, 2005.

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

Nevo, Aviv, "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 2001, *69* (2), 307–42.

Neyman, Jerzy and Elizabeth L. Scott, "Consistent Estimates Based on Partially Consistent Observations," *Econometrica*, 1948, *16* (1), 1–32.

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

Pesaran, M.H., "Non-nested Hypotheses," *Eatwell J., Milgate M., Newman P. (eds) Econometrics. The New Palgrave. Palgrave Macmillan, London.*, 1990.

Plackett, R. L., "The Analysis of Permutations," *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 1975, *24* (2), 193–202.

Postel-Vinay, Fabien and Jean-Marc Robin, "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity," *Econometrica*, 2002, *70* (6), 2295–2350.

Rivers, Douglas and Quang Vuong, "Model selection tests for nonlinear dynamic models," *The Econometrics Journal*, 2002, *5* (1), 1–39.

Robinson, Joan, *The Economics of Imperfect Competition*, Palgrave, 1933.

Rosen, Sherwin, "The Theory of Equalizing Differences," in "in," Vol. 1 of *Handbook of Labor Economics*, Elsevier, 1986, pp. 641–92.

Roussille, Nina, "The Central Role of the Ask Gap in Gender Pay Inequality," Working Paper, University of California, Berkeley 2021.

Schmalensee, Richard, "Inter-industry studies of structure and performance," in "in," Vol. 2 of *Handbook of Industrial Organization*, Elsevier, 1989, pp. 951–1009.

Schubert, Gregor, Anna Stansbury, and Bledi Taska, "Employer Concentration and Outside Options," Working Paper, University of Chicago 2021.

Simons, Gordon and Yi-Ching Yao, "Asymptotics when the number of parameters tends to infinity in the Bradley-Terry model for paired comparisons," *The Annals of Statistics*, 1999, *27* (3), 1041–60.

Sorkin, Isaac, "The Role of Firms in Gender Earnings Inequality: Evidence from the United States," *American Economic Review*, May 2017, *107* (5), 384–87.

_ , "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, 2018, *133* (3), 1331–93.

Soufiani, Hossein Azari, Hansheng Diao, Zhenyu Lai, and David C. Parkes, "Generalized Random Utility Models with Multiple Types," *Advances in Neural Information Processing Systems*, 2013.

Staiger, Douglas O., Joanne Spetz, and Ciaran S. Phibbs, "Is There Monopsony in the Labor Market? Evidence from a Natural Experiment," *Journal of Labor Economics*, 2010, *28* (2), 211–36.

Taber, Christopher and Rune Vejlin, "Estimation of a Roy/Search/ Compensating Differential Model of the Labor Market," *Econometrica*, 2020, *88* (3), 1031–69.

Vuong, Quang H., "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica*, 1989, *57* (2), 307–33.

Wiswall, Matthew and Basit Zafar, "Preference for the Workplace, Investment in Human Capital, and Gender," *The Quarterly Journal of Economics*, 2018, *133* (1), 457–507.

Yeh, Chen, Claudia Macaluso, and Brad Hershbein, "Monopsony in the US labor market," *American Economic Review*, 2022, *112* (7), 2099–2138.

A Appendix Figures

Figure A.1: Mandatory features of a candidate profile, at the time of the study

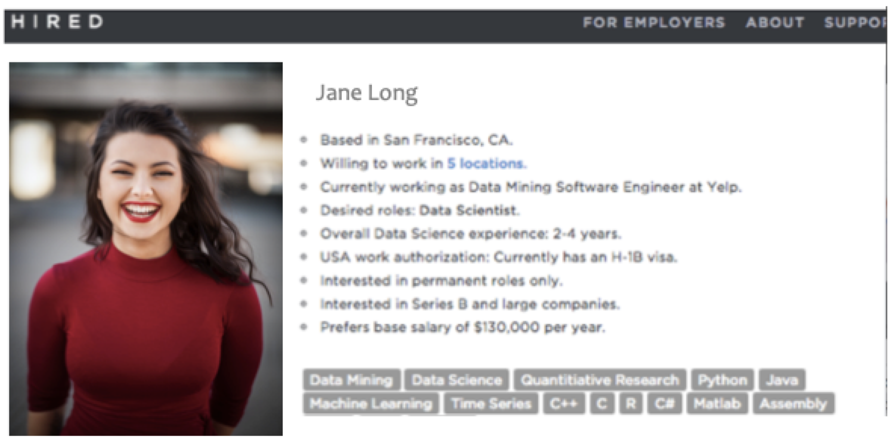

Figure A.2: Typical interview request message sent by a company to a candidate, at the time of the study

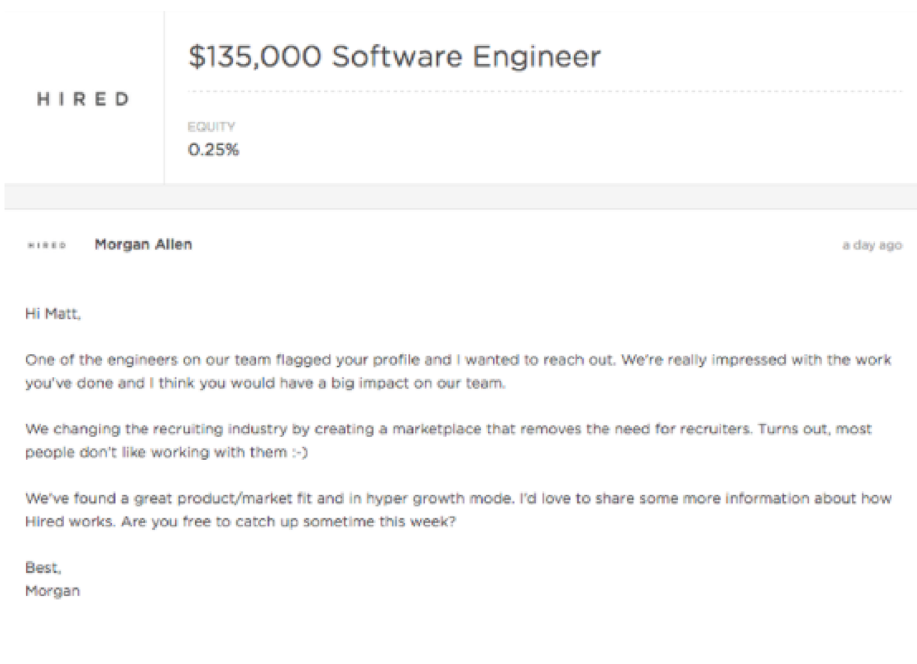

Figure A.3: Model Fit: Labor Supply

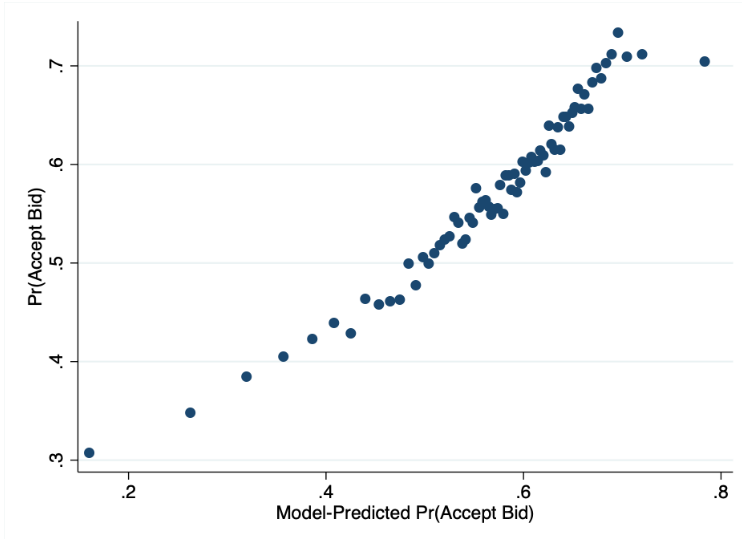

Note: This figure plots the relationship between the empirical acceptance probability of a bid and the model-implied probabilities that the bid will be accepted.

Figure A.4: Relationship between bids and systematic component of valuations, $\gamma_j(x_i)$

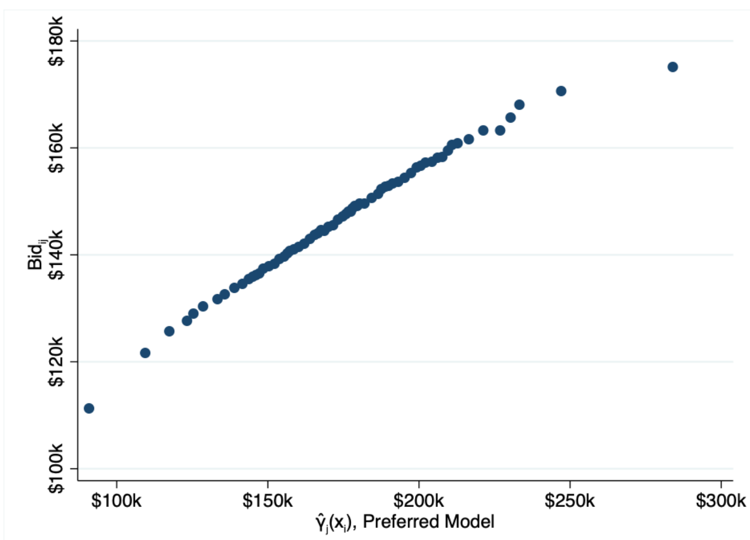

Note: This figure plots the relationship between observed bids and the systematic component of valuations $\exp(z_j'\Gamma x_i)$ in the preferred model, controlling for the ask salary. Unconditionally, the slope of the relationship between bids and the observed component of valuations is 0.83.

Figure A.5: Summary Statistics of Benefits listed by Firms

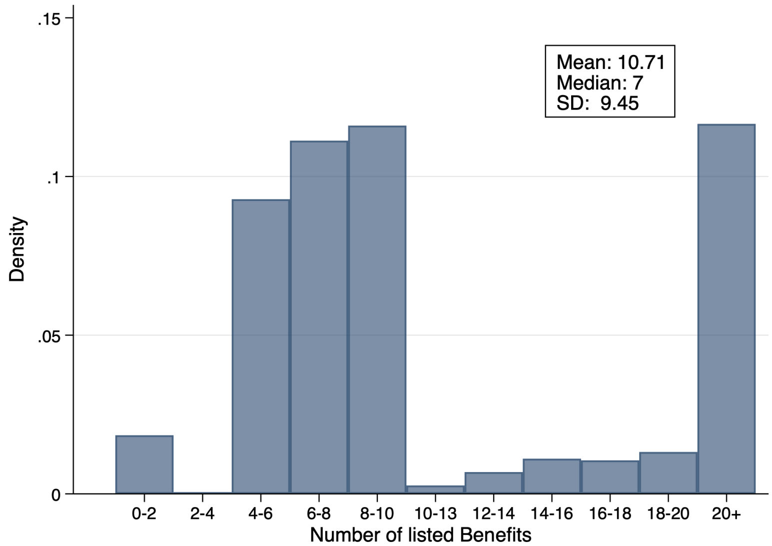

(a) Distribution of Number of listed Benefits

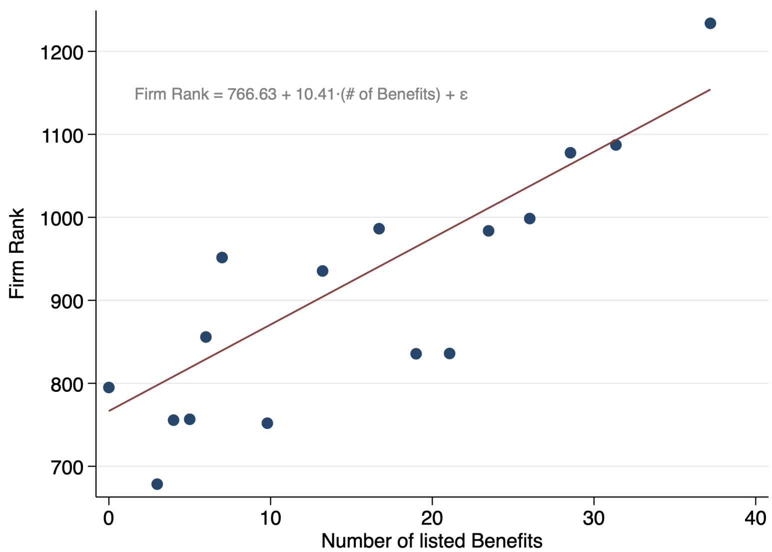

(b) Share of listed Benefits

Note: This figure displays the distribution of benefits listed by firms in the subset of ranked firms. Panel (a) plots the density of the number of listed benefits per firm. The bar "20+" includes numbers of listed benefits greater than 20 up to a maximum of 53. The mean number of benefits is 10.71 (SD 9.45), while the median lies at 7. Panel (b) illustrates the relationship between firm ranking and the number of listed benefits. On average an additional benefit increases the firm's ranking by 10.41.

B Illustration of conceptual framework

The following simple model, adapted from Bhaskar et al. (2002), can be used to illustrate the logic of our conduct testing procedure. In particular, the model illustrates the role of preference heterogeneity, the implications of conduct assumptions, and the basic logic of our estimation and testing framework. The basic message is that combinations of assumptions on competition and wage-setting flexibility deliver different wage equations, which can then be used to infer conduct.

In this model, there are two firms $j = -1, +1$. These firms are located on either end of a mile-long road, and have productivity $\text{MRPL}_j = \text{ARPL}_j = \gamma_j$. Workers' homes lie along road with location given be ξ, which is private information. These locations are uniformly distributed: $\xi \sim \text{Unif}[0, 1]$. Workers' homes are on either side of the road, recorded by v, which is public information observable to firms: $v \perp\!\!\!\perp \xi,\ v = \{-1, +1\}$ w.p. 1/2. Firms post wages (which may vary by v). Worker's preferences over firms depend upon the wage offered by each firm and commuting costs, which depend upon the workers' location along the road and whether the worker will have to make left turn across the road to get to work. Worker utilities are given by:

$$u^v_{-1}(\xi) = w^v_{-1} - \beta\Big(\xi + \alpha v\Big); \quad u^v_{+1}(\xi) = w^v_{+1} - \beta\Big(1 - (\xi + \alpha v)\Big).$$

Under these assumptions, type-v's labor supply to firm j is:

$$S^v_j(w^v_j; w^v_{-j}) = \frac{1}{2} + \frac{w^v_j - w^v_{-j}}{2\beta} + \alpha\, v\, j.$$

Labor demand is determined by profit maximization:

$$\pi_j(\boldsymbol{w}) = \frac{1}{2}\sum_{v=-1}^{+1} (\gamma_j - w^v) \times S^v_j(w^v; \widehat{w}^v_{-j}),$$

where the random variable $\widehat{w}^v_{-j}$ encodes j's knowledge of the competitive environment. Wages are determined by firms' first-order conditions and a market clearing constraint:

$$w^v_j = \frac{1}{2}(\widehat{w}^v_{-j} + \gamma_j - \beta) - \alpha\beta\, v\, j, \quad S^v_j(w^v_j; \widehat{w}^v_{-j}) + S^v_{-j}(w^v_{-j}; \widehat{w}^v_j) = 1.$$

We next define what we mean by firm conduct: in this setting, we define conduct as assumptions about the content of $\widehat{w}^v_{-j}$ and firms' use of v in wage setting. Applying each conduct assumption, we find that each conduct assumption implies a distinct markdown:

Conduct	**use v?**	**Firm's $\widehat{w}^v_{-j}$**	**Equilibrium Wage(s) w^v_j**
Perfect Comp.	No	—	γ_j
Monopsonistic	No	$\overline{w}$	$\frac{3}{4}\gamma_j + \frac{1}{4}\gamma_{-j} - \beta$
Monopsonistic	Yes	$\overline{w}^v$	$\frac{3}{4}\gamma_j + \frac{1}{4}\gamma_{-j} - \beta\big(1 + \alpha v\, j\big)$
Oligopsony	No	w_{-j}	$\frac{2}{3}\gamma_j + \frac{1}{3}\gamma_{-j} - \beta$
Oligopsony	Yes	w^v_{-j}	$\frac{2}{3}\gamma_j + \frac{1}{3}\gamma_{-j} - \beta\big(1 + \frac{2}{3}\alpha v\, j\big)$

Next, we consider estimation and model selection. Each model, which we index by m, yields a wage equation of the form:

$$w^v_j = c^m_{\text{own}} \cdot \gamma_j + c^m_{\text{other}} \cdot \gamma_{-j} - c^{vm}_j.$$

A traditional approach in labor economics is to estimate $\widehat{\boldsymbol{c}}$. To do so, one might first construct proxies for firm productivity γ_j and identify instruments that shift γ_j (and/or competitive environment). Then, one would regress w^v_j on γ_j, γ_{-j}, and concentration measures. To conduct inference, we might perform a simple Wald test on the parameter c_j, for instance: $H_0 : \;\; c_j \geq 1$, $H_a : \;\; c_j < 1$. Our approach (which follows the New Empirical Industrial Organization tradition) is to estimate $\widehat{\boldsymbol{\gamma}}$, rather than $\widehat{\boldsymbol{c}}$. A particular conduct assumption m, in combination with labor supply parameters estimated in a prior step, determines the coefficients $\boldsymbol{c}^m$. Rather than searching for instruments for productivity, find instruments for markdowns that are excluded from productivity. Then, regress $w^v_j + c^{vm}_j$ on c^m_{own} and c^m_{other} to recover $\widehat{\gamma}^m_j$; for example, when firms do not use v in wage setting, we have:

$$\begin{bmatrix} \widehat{\gamma}^m_{-1} \\ \widehat{\gamma}^m_{+1} \end{bmatrix} = \begin{bmatrix} c^m_{\text{own}} & c^m_{\text{other}} \\ c^m_{\text{other}} & c^m_{\text{own}} \end{bmatrix}^{-1} \begin{bmatrix} w_{-1} + c^m_{-1} \\ w_{+1} + c^m_{+1} \end{bmatrix}$$

Finally, in order to adjudicate between different forms of conduct, we use the Vuong (1989) and Rivers and Vuong (2002) tests, which compare model lack of fit between alternatives.

C EM algorithm details

Our strategy relies on the well known fact that the maximum of independent EV_1 random variables is also distributed EV_1:

$$\Pr\left(\max_{k\in\mathcal{B}_i^0} \log(\rho_{qk}) + \xi_{ik} < v\right) = F_\xi\left(v - \log\left(\textstyle\sum_{k\in\mathcal{B}_i^0}\rho_{qk}\right)\right),$$

where $F_\xi(x) = \exp(-\exp(-x))$ is the EV_1 CDF. Using this observation and a simple change of variables argument, we can re-write the probability of the partial ordering $\mathcal{B}_i^1 \succ \mathcal{B}_i^0$, conditional on preference parameters $\boldsymbol{\rho}_q$, as:

$$\begin{aligned}
\mathcal{P}\Big(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \boldsymbol{\rho}_q\Big) &= \Pr\left(\min_{j\in\mathcal{B}_i^1} \log(\rho_{qj}) + \xi_{ij} > \max_{k\in\mathcal{B}_i^0} \log(\rho_{qk}) + \xi_{ik} \mid \boldsymbol{\rho}_q\right)\\
&= \int_{-\infty}^{\infty} \prod_{j\in\mathcal{B}_i^1} \left(1 - F_\xi\left(v - \log(\rho_{qj})\right)\right) \times dF_\xi\left(v - \log\left(\textstyle\sum_{k\in\mathcal{B}_i^0}\rho_{qk}\right)\right)\\
&= \int_{-\infty}^{\infty} \prod_{j\in\mathcal{B}_i^1} \left(1 - F_\xi\left(v - \log\left(\textstyle\sum_{k\in\mathcal{B}_i^0}\rho_{qk}\right)\right)^{\rho_{qj}/\sum_{k\in\mathcal{B}_i^0}\rho_{qk}}\right) \times dF_\xi\left(v - \log\left(\textstyle\sum_{k\in\mathcal{B}_i^0}\rho_{qk}\right)\right)\\
&= \int_0^1 \prod_{j\in\mathcal{B}_i^1} \left(1 - u^{\rho_{qj}/\sum_{k\in\mathcal{B}_i^0}\rho_{qk}}\right) du.
\end{aligned}$$

The second line uses the independence of ξ_{ij} and the distribution of $\max_{k\in\mathcal{B}_i^0} \log(\rho_{qk}) + \xi_{ik}$, the third line uses the fact that $F_\xi(x - \log(a)) = F_\xi(x - \log(b))^{a/b}$, and the fourth line substitutes $u = F_\xi(v - \log(\sum_{k\in\mathcal{B}_i^0}\rho_{qk}))$. This expression, and its derivatives, can be quickly and accurately approximated by numerical quadrature.

We estimate the parameters of the the preference model via the EM algorithm. Specifically, we use a first-order (or "Generalized") EM (GEM) algorithm, in which we replace full maximization of the surrogate function in the M step with a single gradient ascent update. Our algorithm proceeds as follows:

- **Initialization:** provide an initial guess of parameter values $(\boldsymbol{\beta}^{(0)}, \boldsymbol{\rho}^{(0)})$.
- **E Step:** at iteration t, approximate the average log integrated likelihood at $\boldsymbol{\beta}^{(t)}, \boldsymbol{\rho}^{(t)}$ with the function:

$$\mathcal{E}(\boldsymbol{\beta}, \boldsymbol{\rho} \mid \boldsymbol{\beta}^{(t)}, \boldsymbol{\rho}^{(t)}) = \frac{1}{N}\sum_{i=1}^{N}\sum_{q=1}^{Q} \alpha_{iq}^{(t)} \log\left(\alpha_q(x_i \mid \boldsymbol{\beta}) \times \mathcal{P}\Big(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \boldsymbol{\rho}_q\Big)\right),$$

where the weights $\alpha_{iq}^{(t)}$ are given by:

$$\alpha_{iq}^{(t)} = \frac{\alpha_q(x_i \mid \boldsymbol{\beta}^{(t)}) \times \mathcal{P}\Big(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \boldsymbol{\rho}_q^{(t)}\Big)}{\sum_{r=1}^{Q} \alpha_q(x_i \mid \boldsymbol{\beta}^{(t)}) \times \mathcal{P}\Big(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \boldsymbol{\rho}_q^{(t)}\Big)}.$$

- **M Step:** Find $\boldsymbol{\beta}^{(t+1)}, \boldsymbol{\rho}^{(t+1)}$ by computing a single gradient ascent update (hence "first-order").

We initialize our algorithm at 50 random starting values, and report the estimate that yields the highest likelihood.

D Properties of bidding strategies

We leverage the log-concavity of $G_{ij}^m(\cdot)$ to establish several properties of bidding functions. A function f is log-concave if: $f(\lambda y + (1-\lambda)x) \geq f(y)^\lambda f(x)^{1-\lambda}\ \ \forall x, y \in \mathbb{R}, \lambda \in [0,1]$. Log-concavity of f implies that $F = \int_{-\infty}^{x} f(u)du$ and $1 - F = \overline{F}$ are also log-concave, that f/F is monotone decreasing, and that $f/\overline{F}$ is monotone increasing. A large number of common probability distributions admit log-concave densities, including the normal, logistic, extreme value, and Laplace distributions. Log-concave probability distributions are commonly used in models of search (Bagnoli and Bergstrom, 2005).

For clarity, we suppress dependence on m. Under each model m, we may generally write $G_{ij}(b) = \int \widetilde{G}_{ij}(b,\lambda)dH(\lambda)$, where either $\widetilde{G}_{ij}(b,\lambda) = \exp(u(b,a_i))/(\exp(u(b,a_i)) + \exp(\lambda))$ under oligopsony or $\widetilde{G}_{ij}(b,\lambda) = \exp(u(b,a_i) - \lambda)$ under monopsonistic competition. In the latter case, log concavity of $G_{ij}(b)$ follows directly from the fact that $u(b,a_i)$ is concave (by assumption), since $G_{ij}(b) = \exp(u(b,a_i)) \times \int \exp(-\lambda)dH(\lambda)$. Log concavity in the former case can also be shown via differentiation of $\log(G_{ij}(b))$.

Let the function $G_{ij}^+(b)$ (with derivative $g_{ij}^+(b)$) denote the right-hand side of the $G_{ij}(b)$ function, which replaces $\theta_0 + \theta_1 \cdot \mathbf{1}[b < w_i]$ with θ_0. We similarly let $G_{ij}^-(b)$ denote the left-hand side function, which replaces $\theta_0 + \theta_1 \cdot \mathbf{1}[b < w_i]$ with $\theta_0 + \theta_1$. Clearly, $G_{ij}(b) = \mathbf{1}[b \geq w_i] \cdot G_{ij}^+(b) + \mathbf{1}[b < w_i] \cdot G_{ij}^-(b)$. Under the assumption that both $G_{ij}^+(b)$ and $G_{ij}^-(b)$ are log-concave, we have that the functions $g_{ij}^+(b)/G_{ij}^+(b)$ and $g_{ij}^-(b)/G_{ij}^-(b)$ are both strictly decreasing functions of b. This implies that both the left-hand and right-hand inverse bidding functions, $\varepsilon_{ij}^-(b) = b + G_{ij}^-(b)/g_{ij}^-(b)$ and $\varepsilon_{ij}^+(b) = b + G_{ij}^+(b)/g_{ij}^+(b)$ are monotone increasing functions of the bid. This in turn implies that the left- and right-hand bidding functions, which we denote by $b_{ij}^-(\varepsilon_{ij})$ and $b_{ij}^+(\varepsilon_{ij})$ are also strictly increasing functions of ε_{ij}. We may also define the left- and right-hand indirect expected profit functions as $\pi_{ij}^{*s}(\varepsilon_{ij}) = G_{ij}^s(b_{ij}^s(\varepsilon_{ij}))^2/g_{ij}^s(b_{ij}^s(\varepsilon_{ij}))$ for $s \in \{-,+\}$, which are both strictly increasing functions of ε_{ij}. These results establish the monotonicity of firm strategies and payoffs in their unobserved valuations when firms bid on either side of the kink.

A necessary, but not sufficient, condition that the firm bids at the kink is that the derivative of the left-hand expected profit function is positive at the ask salary:

$$g_{ij}^-(w_i)(\varepsilon_{ij} - w_i) - G_{ij}^-(w_i) < 0.$$

We assume that $\varepsilon_{ij} > w_i$, since otherwise the firm would never choose to bid at ask. We additionally assume that both θ_0 and θ_1 are positive. Given these assumptions, we have that

$$g_{ij}^-(w_i)(\varepsilon_{ij} - w_i) - G_{ij}^-(w_i) < 0 \implies g_{ij}^+(w_i)(\varepsilon_{ij} - w_i) - G_{ij}^+(w_i) < 0,$$

since by construction $g_{ij}^+(w_i) < g_{ij}^-(w_i)$ and $G_{ij}^+(w_i) = G_{ij}^-(w_i)$. By the same logic, we can show:

$$g_{ij}^+(w_i)(\varepsilon_{ij} - w_i) - G_{ij}^+(w_i) > 0 \implies g_{ij}^-(w_i)(\varepsilon_{ij} - w_i) - G_{ij}^-(w_i) > 0.$$

These conditions guarantee that the firm's optimal choice of bid is unique, even incorporating the kink. Given these definitions, we can write the condition that firms bid at the kink as:

$$\varepsilon_{ij}^-(w_i) \leq \varepsilon_{ij} \leq \varepsilon_{ij}^+(w_i)$$

Therefore, we may write the firm's optimal bidding function as:

$$b_{ij}(\varepsilon_{ij}) = \begin{cases} b_{ij}^-(\varepsilon_{ij}) & \text{if } \varepsilon_{ij}^-(w_i) \geq \varepsilon_{ij} \\ w_i & \text{if } \varepsilon_{ij}^-(w_i) \leq \varepsilon_{ij} \leq \varepsilon_{ij}^+(w_i) \\ b_{ij}^+(\varepsilon_{ij}) & \text{if } \qquad\qquad \varepsilon_{ij} \geq \varepsilon_{ij}^+(w_i). \end{cases}$$

We have therefore shown that the firm's optimal strategy is a strictly increasing function of its valuation outside of the interval $[\varepsilon_{ij}^-(w_i), \varepsilon_{ij}^+(w_i)]$, and is flat within that region.

Next, we consider firms' participation decisions. The results established above imply that the firm's indirect expected profit function is a *strictly increasing* function of the firm's valuation:

$$\pi^*_{ij}(\varepsilon_{ij}) = \begin{cases} \pi^{*-}_{ij}(\varepsilon_{ij}) & \text{if } \ \varepsilon^-_{ij}(w_i) \geq \varepsilon_{ij} \\ G_{ij}(w_i)(\varepsilon_{ij} - w_i) & \text{if } \ \varepsilon^-_{ij}(w_i) \leq \varepsilon_{ij} \leq \varepsilon^+_{ij}(w_i) \\ \pi^{*+}_{ij}(\varepsilon_{ij}) & \text{if } \qquad\qquad \varepsilon_{ij} \geq \varepsilon^+_{ij}(w_i). \end{cases}$$

Firms participation decisions are therefore given by the condition:

$$B_{ij} = \mathbf{1}\left[\pi^*_{ij}(\varepsilon_{ij}) > c_j\right].$$

Since $\pi^*_{ij}(\varepsilon_{ij})$ is a strictly increasing function of the firm's valuation, an inverse indirect expected profit function exists and is also strictly increasing. Therefore, we may re-write the above equation as:

$$B_{ij} = \mathbf{1}\Big[\nu_{ij} > \pi^{*-1}_{ij}(c_j) - \gamma_j(x_i)\Big].$$

E Proof of the consistency of $\widehat{c}^m_j$

Our proof of the consistency of $\widehat{c}^m_j$ for each firm j (and model m) closely follows the proof of Lemma 1 (ii) of Donald and Paarsch (2002). For clarity, we omit j and m indices. Let n denote the total number of bids, with $n \to \infty$. A sufficient condition for establishing consistency is the existence of a vector of candidate characteristics $x \in \mathcal{X}$ (including ask salary a) occurring with positive probability such that there is a positive probability the firm optimally bids below ask for candidates with those characteristics: $\exists x \in \mathcal{X}$ such that $\Pr(a > b_i > 0 \cap x_i = x) > 0$. The vast majority of firms (92%) bid below ask at least once, which suggests that this assumption is reasonable. The vector x need not be the same for all firms. This assumption implies that the distribution of model-implied option value upper bounds $\widehat{\pi}_i$ is bounded below by c when $x_i = x$, and that $\Pr(\widehat{\pi}_i \in [c, c+\delta] \mid x_i = x) > 0$ for arbitrary $\delta > 0$. Let n_x denote the number of bids made to candidates with characteristics x and let $\widehat{c}^n_x$ denote the minimum implied $\widehat{\pi}$ among those bids (such that $\widehat{c}^n = \min_{x' \in \mathcal{X}} \widehat{c}^n_{x'}$). Our sampling assumptions imply $n_x \xrightarrow{\text{a.s.}} \infty$. For an arbitrary $\epsilon > 0$, note that $\Pr(|\widehat{\pi}_i - c| > \epsilon \mid x_i = x) = \Pr(\widehat{\pi}_i > c + \epsilon \mid x_i = x) = 1 - F_\pi(c + \epsilon \mid x_i = x) < 1$. Let $\overline{F}_{\pi|x}(a) = 1 - F_\pi(a \mid x_i = x)$. We then have that $\left(\overline{F}_{\pi|x}(c+\epsilon)\right)^{n_x} \xrightarrow{\text{a.s.}} 0$, and therefore $\Pr(|\widehat{c}^n_x - c| > \epsilon) = \Pr(\widehat{c}^n_x > c + \epsilon) = E\Big[\left(\overline{F}_{\pi|x}(c+\epsilon)\right)^{n_x}\Big]$. Since ϵ is arbitrary, $\widehat{c}^n_x \xrightarrow{\text{p}} c$, and since $\widehat{c}^n_x \geq \widehat{c}^n \geq c$, $\widehat{c}^n \xrightarrow{\text{p}} c$. Further, $\sup_{m>n} |\widehat{c}^m - c| = |\widehat{c}^n - c| \xrightarrow{\text{p}} 0$ since $\widehat{c}^n$ is non-increasing in n, and so $\widehat{c}^n \xrightarrow{\text{a.s.}} c$. □

F Alternative: The Vuong (1989) Likelihood Ratio Test

Because we estimate models by maximum likelihood, a natural option for our test of conduct is a straightforward application of the Vuong (1989) likelihood ratio test. The Vuong (1989) test is a pairwise, rather than ensemble, testing procedure: rather than explicitly identifying the "best" model among a set of alternatives, the test considers each pair of models in turn and asks whether one of those models is closer to the truth than the other. In the likelihood setting, the "better" of two models is the one with greatest goodness-of-fit, as measured by the maximized log-likelihoods.[27]

Let $s = |ij : B_{ij} = 1|$ denote the sample size. For a pair of models m_1 and m_2, denote the maximized sample log-likelihoods by $\mathcal{L}_s^{m_1}$ and $\mathcal{L}_s^{m_2}$, respectively, where:

$$\mathcal{L}_s^m = \max_{\Psi} \sum_{ij:B_{ij}=1} \log\left(\mathcal{L}_{ij}^m(\Psi)\right),$$

and Ψ^m denotes the arg max. The null hypothesis of our test is that m_1 and m_2 are equally close to the truth, or *equivalent*. In this case, the population expectation of the difference in log likelihoods is zero. There are two one-sided alternative hypotheses: that m_1 is closer to the truth than m_2, and vice versa. When m_1 is closer to the true data-generating process, the population expectation of the likelihood ratio $\mathbb{E}^0[\log(\mathcal{L}_{ij}^{m_1}(\Psi^{m_1})/\mathcal{L}_{ij}^{m_2}(\Psi^{m_2})]$ is greater than zero. Vuong (1989) shows that when m_1 and m_2 are non-nested, an appropriately-scaled version of the sample likelihood ratio is asymptotically normal under the null that the two models are equivalent:

$$Z_s^{m_1,m_2} = \frac{\mathcal{L}_s^{m_1} - \mathcal{L}_s^{m_2}}{\sqrt{s} \cdot \widehat{\omega}_s^{m_1,m_2}} \xrightarrow{D} \mathcal{N}(0,1),$$

where $\widehat{\omega}_s^{m_1,m_2}$ is the square root of a consistent estimate of the asymptotic variance of the likelihood ratio, $\omega_*^{2m_1,m_2}$. We set:

$$\widehat{\omega}_s^{m_1,m_2} = \left(\frac{1}{s} \sum_{ij:B_{ij}=1} \log\left(\frac{\mathcal{L}_{ij}^{m_1}(\Psi^{m_1})}{\mathcal{L}_{ij}^{m_2}(\Psi^{m_2})}\right)^2\right)^{1/2}.$$

We construct test statistics $Z_s^{m_1,m_2}$ for every pair of models we estimate. Given a significance level α with critical value c_α, we reject the null hypothesis that m_1 and m_2 are equivalent in favor of the alternative that m_1 is better than m_2 when $Z_s^{m_1,m_2} > c_\alpha$,

[27]The population expectation of the log-likelihood measures the distance, in terms of the Kullback-Liebler Information Criterion (KLIC), between the model and the true data generating process.

and vice versa if $Z_s^{m_1,m_2} < c_\alpha$. If $|Z_s^{m_1,m_2}| \leq c_\alpha$, the test cannot discriminate between the two models.

How does variation in the instrument increase the power of the test? The answer depends on the relevance of the instrument for predicting markdowns. Returning to the simplified example above, we may write the mis-specification error as:

$$\zeta_{ij}^m = \log\Big(\varepsilon_{ij}^m(b_{ij})\Big) - \log\Big(\varepsilon_{ij}(b_{ij})\Big).$$

To the extent that variation in tightness drives variation in markdowns under the true model, variation in tightness will also generate variation in ζ_{ij}^m if the assumed model m is mis-specified. This implies that relatively more mis-specified models will imply valuations that are more difficult to explain using observables than those that are closer to the truth. Table B.1 reports the results of implementing this testing procedure. The results are qualitatively extremely similar to the results of the moment-based testing procedure.

Table B.1: Non-Nested Model Comparison Tests (Vuong, 1989)

Model	(1) Monopsonistic Comp.	(2)	(3) Oligopsony	(4)
	Not Predictive	Type Predictive	Not Predictive	Type Predictive
Perfect Competition	-237.57	-237.67	-156.16	-154.34
Monopsonistic, Not Predictive	–	1.28	90.17	90.39
Monopsonistic, Type Predictive		–	88.45	89.81
Oligopsony, Not Predictive			–	6.88
Oligopsony, Type Predictive				–

Note: This table reports test statistics from the Vuong (1989) non-nested model comparison procedure. Positive values imply the row model is preferred to the column model. Under the null of model equivalence, the test statistics are asymptotically normal with mean zero and unit variance.

G Further model comparisons

As an additional comparison, we consider differences in estimated labor demand parameters $\widehat{\Gamma}$ between the preferred model and the not-predictive oligopsony alternative. Table B.2 reports the elasticity of the systematic component of labor demand with respect to the ask salary, along with implied semi-elasticities of the systematic component of labor demand with respect to a selection of binary covariates. All elasticies are evaluated at the (bid-weighted) mean values of firm characteristics. The estimated labor demand parameters represent the impacts of ceteris paribus changes in individual determinants of productivity. Since the ask salary co-varies strongly with the other observables, we report estimates of both the semi-elasticities of each binary covariate ℓ both holding the ask constant ($\widehat{\gamma}_\ell$) and adjusting for differences in the average ask salary. First, column 1 reports selected coefficients from a regression of the ask salary on all other included candidate characteristics. Women and unemployed candidates set lower asked salaries, while those with graduate degrees and FAANG experience set higher asked salaries. Columns 2 and 3 report results for the preferred model. Column 2 reports estimates of Γ. The ask salary is a powerful determinant of productivity: the estimated elasticity with respect to the ask salary is 0.89. The remaining semi-elasticities in column 2 are all relatively small and statistically insignificant. Column 3 reports semi-elasticities that do not condition on the ask salary, adding an adjustment to account for average differences in asks between groups. The significant differences between these columns suggest that the ask salary is a strong signal of quality. Columns 4 and 5 reproduce this analysis for the oligopsony alternative. The estimated elasticity with respect to the ask, 0.74, is significantly lower than in the preferred model, and the conditional semi-elasticities are much larger. We report the full set of labor demand parameter estimates for the preferred model in Tables B.3 and B.4 below.

How do our preferred estimates relate to models of additive worker and firm effects (Abowd et al., 1999)? Our model of productivity includes both firm-specific contributions (here captured by z_j), worker-specific contributions (captured by x_i), and the interactions of firm- and worker-specific covariates. Tables B.3 and B.4 provide evidence that interactions of worker and firm factors are statistically meaningful determinants of productivity. However, the interaction effects we estimate are generally small, which suggests that additive models might well-approximate productivity. To explore this, we regress bids, predicted ε_{ij}, and the predicted systematic component of productivity $\exp(z_j'\widehat{\Gamma}x_i)$ on all candidate and firm characteristics, without including interactions. Consistent with Card et al. (2013)'s informal assessment of the log-additivity of wages using mean residuals from Abowd et al. (1999) regressions, we find that the main effects of worker and firm characteristics separately explain the vast majority of variation in bids and productivity, as reflected in uniformly high (adjusted) R^2 values: 0.924 for bids, 0.905 for ε_{ij}, and 0.999 for $\exp(z_j'\widehat{\Gamma}x_i)$. In the context of the near-constant markdowns our preferred model implies, this further suggests that additive models of worker and firm effects provide good approximations to log wages.

Table B.2: Determinants of Match Productivity

	(1)	(2)	(3)	(4)	(5)
	$\mathbb{E}[\Delta\text{Ask}]$	Monopsonistic Comp.		Oligopsony	
	$\widehat{\beta}_\ell$	$\widehat{\gamma}_\ell$	$+\widehat{\beta}_\ell \cdot \widehat{\gamma}_{\text{ask}}$	$\widehat{\gamma}_\ell$	$+\widehat{\beta}_\ell \cdot \widehat{\gamma}_{\text{ask}}$
Ask Salary	–	0.8928	–	0.7437	–
		(0.0055)		(0.0081)	
Female	-0.0642	-0.0047	-0.0620	-0.0137	-0.0615
	(0.0013)	(0.0030)	(0.0058)	(0.0044)	(0.0075)
Unemployed	-0.0587	0.0022	-0.0502	-0.0034	-0.0471
	(0.0030)	(0.0067)	(0.0208)	(0.0099)	(0.0260)
Grad School	0.0275	0.0034	0.0279	0.0152	0.0357
	(0.0005)	(0.0025)	(0.0055)	(0.0038)	(0.0071)
FAANG	0.0541	-0.0046	0.0437	-0.0136	0.0266
	(0.0013)	(0.0033)	(0.0059)	(0.0050)	(0.0078)

Note: This table reports estimates of the elasticity of the systematic component of labor demand with respect to the ask salary and the semi-elasticities of that component with respect to a subset of binary covariates. Column (1) reports coefficients from a regression of all included candidate characteristics on the ask salary. Columns (2) and (3) report results for monopsonistic competition while Columns (4) and (5) report results for oligopsony (both models assume not-predictive conduct). Columns (2) and (4) report elasticities conditional on the ask salary while Columns (3) and (5) report unconditional versions. Robust standard errors are reported in parentheses.

Table B.3: (Subset of) Labor Demand Parameter Estimates Γ: $\log(\varepsilon_{ij}) = z_j'\Gamma x_i + \nu_{ij}$

	(1) Constant	(2) log(Ask)	(3) Female	(4) Employed	(5) Grad School
Constant	11.9897***	0.7954***	-0.0079***	-0.0014	0.0094***
	(0.0523)	(0.0046)	(0.0025)	(0.0040)	(0.0021)
16-50 Employees	0.0305	0.0814***	0.0046	0.0006	-0.0022
	(0.0448)	(0.0039)	(0.0027)	(0.0044)	(0.0023)
51-500 Employees	0.0503	0.0832***	-0.0010	0.0037	-0.0069***
	(0.0510)	(0.0045)	(0.0025)	(0.0041)	(0.0022)
501+ Employees	0.0612	0.1073***	-0.0009	0.0011	-0.0090***
	(0.0516)	(0.0045)	(0.0026)	(0.0043)	(0.0022)
Finance	-0.0008	0.0156***	0.0055***	0.0024	0.0022
	(0.0526)	(0.0046)	(0.0016)	(0.0028)	(0.0013)
Tech	0.0052	0.0166***	0.0043***	-0.0028	-0.0001
	(0.0314)	(0.0027)	(0.0013)	(0.0023)	(0.0011)
Health	-0.0028	0.0011	0.0009	-0.0006	-0.0004
	(0.0462)	(0.0040)	(0.0022)	(0.0037)	(0.0017)
Std. Dev. of ν_{ij} ($\widehat{\sigma}_\nu$)	0.0743	(0.0001)	N = 181,927, Implied $R^2 = 0.888$		

Note: This table reports a subset of maximum likelihood parameter estimates from our preferred model. The parameters relate combinations of candidate and firm characteristics to the distribution of firms' valuations over each candidate (or, the ex-ante productivity of that candidate at that firm). The log of productivity/valuations is modelled as normally distributed, with mean $z_j'\Gamma x_i$ and variance σ_ν. Each cell reports the coefficient on the interaction of the variables specified in the corresponding row and column. Column variables are candidate characteristics (x_i), and row variables are firm characteristics (z_j). The second, third, and fourth rows correspond to dummies for firm size categories, such that the omitted category (subsumed into the constant, the first row of the table) corresponds to the smallest firms (between one and fifteen employees). The remaining three rows correspond to non-exclusive sector dummies. Column 1 reports the main effects of each firm characteristic. Column 2 reports the main effects and interactions for the log ask salary, where the log ask salary has been de-meaned. Columns 3-5 report coefficients on dummies recording whether the candidate is female, was employed, or has received at least a master's degree. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.4: Remaining Labor Supply Parameter Estimates: $\gamma_j(x_i) = z_j'\Gamma x_i$

	(1) Soft-Eng	(2) Experience	(3) (Experience)^2	(4) Unemployed	(5) Ivy Plus	(6) CS Degree	(7) FAANG	(8) Previous Jobs	(9) Fulltime	(10) Sponsorship	(11) Remote	(12) Java	(13) Python	(14) SQL	(15) C
Constant	0.0326***	0.0005	0.00002	0.0009	-0.0060*	0.0042	-0.0012	-0.0003	-0.0035	-0.0019	0.0029	-0.0002	0.0009	-0.0030	0.0077**
	(0.0029)	(0.0006)	(0.00002)	(0.0010)	(0.0023)	(0.0022)	(0.0028)	(0.0005)	(0.0022)	(0.0028)	(0.0020)	(0.0021)	(0.0020)	(0.0023)	(0.0026)
16-50 Employees	-0.0046	0.0007	-0.0000111	0.0003	-0.0035	-0.0028	-0.0017	-0.0008	0.0011	0.0150**	0.0032	-0.0006	-0.0004	0.0065*	-0.0136***
	(0.0031)	(0.0006)	(0.00002)	(0.0010)	(0.0025)	(0.0024)	(0.0030)	(0.0005)	(0.0024)	(0.0030)	(0.0022)	(0.0023)	(0.0022)	(0.0025)	(0.0029)
51-500 Employees	-0.0144***	0.0020***	-0.00005**	0.0002	0.0049*	-0.0031	-0.0020	-0.0011	0.0039	0.0058*	-0.0012	0.0042	-0.0018	0.0039	-0.0076**
	(0.0029)	(0.0006)	(0.0000176)	(0.0010)	(0.0024)	(0.0022)	(0.0028)	(0.0005)	(0.0022)	(0.0028)	(0.0021)	(0.0022)	(0.0020)	(0.0023)	(0.0027)
501+ Employees	-0.0167***	0.0016**	-0.00005**	-0.0006	0.0073**	-0.0020	0.0001	-0.0002	0.0034	0.0057*	-0.0020	0.0064**	-0.0029	0.0032	-0.0087**
	(0.0030)	(0.0006)	(0.00002)	(0.0010)	(0.0025)	(0.0023)	(0.0029)	(0.0005)	(0.0023)	(0.0028)	(0.0021)	(0.0022)	(0.0021)	(0.0024)	(0.0027)
Finance	0.0084***	-0.0006	0.00001	0.0006	-0.0077***	-0.0052***	-0.0047**	0.0003	-0.0023	0.0025	0.0003	-0.0063***	0.0011	0.0008	0.0012
	(0.0017)	(0.0004)	(0.00001)	(0.0006)	(0.0015)	(0.0013)	(0.0017)	(0.0003)	(0.0014)	(0.0016)	(0.0013)	(0.0013)	(0.0013)	(0.0014)	(0.0016)
Tech	0.0068***	-0.0005	0.00001	-0.0008	-0.0010	0.0016	-0.0022	-0.0003	-0.0028*	0.0004	0.0001	-0.0058***	0.0024*	0.0024	0.0021
	(0.0014)	(0.0003)	(0.00001)	(0.0005)	(0.0013)	(0.0011)	(0.0014)	(0.0003)	(0.0012)	(0.0013)	(0.0011)	(0.0011)	(0.0011)	(0.0012)	(0.0013)
Health	0.0074***	-0.0004	0.00001	-0.0007	-0.0027	-0.0049**	-0.0031	0.0004	0.0025	-0.0031	0.0027	0.0004	-0.0032	-0.0003	0.0013
	(0.0022)	(0.0005)	(0.00001)	(0.0008)	(0.0021)	(0.0018)	(0.0024)	(0.0004)	(0.0019)	(0.0021)	(0.0017)	(0.0018)	(0.0017)	(0.0019)	(0.0023)

Note: This table presents the remaining set of coefficients corresponding to Table B.3. The omitted category for the number of employees is "1-15 Employees". Every cell reports the coefficient on the interaction of the variables specified in the corresponding row and column. Column variables are candidate characteristics (x_i), and row variables are firm characteristics (z_j). In Column 1 interaction coefficients for software engineers are presented, Coefficients for years of experience in the candidates' field of occupation are shown Column 2 and squared in Column 3. Columns 4 - 7 display the coefficient for dummy variables of unemployment, whether the candidate received education in an Ivy+ school, has a degree in computer sciences and/or has worked at either Facebook, Amazon, Apple, Netflix or Google. In Column 8 the coefficients for the number of previous jobs are introduces, while Column 9 and 11 present values for whether candidates' wish to work full time, require VISA sponsorship for their work permit or want to work remotely. Lastly, Columns 12-15 display the coefficients for whether candidates are skilled in Java, Python, SQL or one C language respectively. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

H Welfare: decompositions and counterfactual simulations

H.1 Decomposing group differences in welfare

Given our estimates of amenity values and labor supply parameters, we may fully characterize the utility value candidates associate with the portfolios of bids they receive. Importantly, this allows us to ask whether observable differences in average bids between groups are reflective of underlying differences in welfare. We decompose mean differences in welfare using the Oaxaca-Blinder (OB) decomposition (Oaxaca, 1973; Blinder, 1973). The OB decomposition posits that variable Y_{ig} corresponding to individual i in group $g = 0, 1$ can be written:

$$Y_{ig} = X_{ig}'\beta_g + \epsilon_{ig},$$

where X_{ig} are covariates measured for all individuals and $\mathbb{E}(\epsilon_{ig}) = 0$. The average value of Y_{ig} in group g is therefore given by $\overline{Y}_g = \overline{X}_g'\beta_g$. We can decompose the difference in the average value of Y_{ig} between groups $g = 1$ and $g = 0$ as:

$$\overline{Y}_1 - \overline{Y}_0 = \overline{X}_1'\beta_1 - \overline{X}_0'\beta_0 = \underbrace{\left(\overline{X}_1 - \overline{X}_0\right)'\beta_0}_{\text{endowments}} + \underbrace{\overline{X}_0'\left(\beta_1 - \beta_0\right)}_{\text{coefficients/returns}} + \underbrace{\left(\overline{X}_1 - \overline{X}_0\right)'\left(\beta_1 - \beta_0\right)}_{\text{interactions}}.$$

The classic OB decomposition apportions the difference in the mean of a variable between two groups into components due to: 1) differences between those groups in *endowments*, or the distribution of relevant covariates; 2) differences between those groups in *coefficients* or *returns* associated with those covariates; and 3) the *interactions* between coefficient and endowment differences.[28] Roughly speaking, the greater the share of the mean difference the OB decomposition apportions to endowments relative to returns, the more we can conclude that a difference in means is driven by differences in characteristics between those groups, and not how those groups are treated conditional on those characteristics (i.e. differential returns). The OB decompositions we present should be interpreted as purely descriptive (Guryan and Charles, 2013). Importantly, we exclude the ask salary as an explanatory variable in our OB decompositions of welfare, because candidates formulate their asks as endogenous functions of their other characteristics (including gender). The endogeneity of the ask complicates the interpretation of decompositions that include the ask salary: if the asks are functions of gender, then gender differences in asks may not be

[28] Note that the OB decomposition is not unique. An equivalent "reverse" decomposition may be obtained by replacing β_0 with β_1 in the first term, $\overline{X}_0$ with $\overline{X}_1$ in the second term, and flipping the sign of the third term.

appropriately interpreted as reflecting differing endowments.[29]

We report decompositions of welfare-relevant quantities in Table B.5. The utility associated with each portfolio of bids depends both upon the number of bids received and the composition of those bids. In order to gauge the relative importance of quantity and quality, we compute the total number of bids received by each candidate, as well as the mean values of the components of utility associated with the bids each candidate received. We calculate the monetary component of utility for each bid as:

$$\overline{u}(b_{ij}, a_i) = \Big(4.05 + 1.58 \cdot \mathbf{1}[b_{ij} < a_i]\Big) \cdot \log(b_{ij}/a_i) + 4.05 \cdot \Big(\log(a_i) - \overline{\log(a_i)} \Big),$$

where we subtract the (grand) mean of the log of the ask salary ($\overline{\log(a_i)}$) without loss of generality, since the absolute level of utility is not identified. We also compute the mean amenity values associated with each bid, which we decompose into two parts: a common component of amenity valuations shared by all workers, and the worker-specific deviation from that common component: $A_{ij} = \overline{A}_j + \Delta A_{ij}$. The common component is the average candidates' amenity valuation: $\overline{A}_j = \sum_{q=1}^{Q} \overline{\alpha}_q \cdot \widehat{A}_{qj}$ (where $\overline{\alpha}_q$ is the population share of type q). The candidate-specific deviation is the difference between candidate i's amenity valuation and the average amenity valuation: $\Delta A_{ij} = \sum_{q=1}^{Q} \Big(\alpha_q\big(x_i \mid \widehat{\boldsymbol{\beta}}\big) - \overline{\alpha}_q \Big) \cdot \widehat{A}_{qj}$.

To understand how these differences map into welfare, we compute the (expected) inclusive value of every offer set:

$$\Lambda_i^* = \sum_{q=1}^{Q} \alpha_q\big(x_i \mid \widehat{\boldsymbol{\beta}}\big) \cdot \log\Big(\textstyle\sum_{j \in \mathcal{B}_i} \exp(\overline{u}(b_{ij}, a_i) + \widehat{A}_{qj}) \Big).$$

We decompose (expected) inclusive values into a monetary component and an amenity component. We compute the monetary component of the inclusive value by setting $\widehat{A}_{qj} = 0$ for all q and j:

$$\Lambda_i^b = \log\Big(\textstyle\sum_{j \in \mathcal{B}_i} \exp(\overline{u}(b_{ij}, a_i)) \Big).$$

We compute the amenity component of the inclusive value by setting $\overline{u}(b_{ij}, a_i) = 0$ for all i and j. We further decompose the amenity portion into a common component:

$$\bar{\Lambda}_i^A = \sum_{q=1}^{Q} \overline{\alpha}_q \cdot \log\Big(\textstyle\sum_{j \in \mathcal{B}_i} \exp(\widehat{A}_{qj}) \Big),$$

[29]Because we omit the ask salary from these decompositions, the effect of the ask salary will be apportioned between the endowments and coefficients components. Any differential patterns in the relationship between characteristics and asks will be reflected in the coefficients component, while mean differences in asks are reflected in the endowments component.

and a candidate-specific deviation:

$$\Delta\Lambda_i^A = \sum_{q=1}^{Q} \Big(\alpha_q\big(x_i \mid \widehat{\boldsymbol{\beta}}\big) - \overline{\alpha}_q\Big) \cdot \log\Big(\textstyle\sum_{j\in\mathcal{B}_i} \exp(\widehat{A}_{qj})\Big).$$

Because the inclusive value is a nonlinear function, the relative contributions of each component will not sum to one.

Table B.5 reports decompositions of mean gaps in these quantities by gender (here, the reference group corresponds to women, so positive differences correspond to larger values for men). Column 1 decomposes the gap in the number of bids received by men and women: on average, women receive fewer bids than men. However, slightly more than 100% of this raw gap is driven by differences in endowments: conditional on covariates, women and men receive nearly the same number of bids. Column 2 reports the decomposition of the mean gap in the monetary component of utility: the average monetary value of bids is significantly lower for women than for men. This result is driven by the fact that women ask for less (see Table 1), and therefore receive less, conditional on other characteristics—but as discussed above, the ask is an endogenous function of gender. Our decomposition, which excludes the ask as an explanatory variable, suggests that differences in characteristics between men and women can only explain about 1/3 of the raw gap in monetary values, with the rest explained by differential returns. Column 3 decomposes the mean difference in the common component of amenity values. Unconditionally, the bids men receive are from firms with better amenities than the bids women receive. Differences in the returns to characteristics, representing differential selection of firms into bidding by gender, explain 1/3 of this gap. In other words, even conditional on covariates, women receive bids from firms the average worker values relatively less than those that bid on men.

Column 4 decomposes differences in candidate-specific components of the amenity valuation. Here, we find a (small) reverse gap: women value the amenities associated with the bids they receive relatively more than the average worker would, and do so to a greater degree than men. What might be driving this pattern? Without knowing how firms behave, we cannot discriminate between possible explanations. One possibility is that the pattern is driven by differences in the degree of assortative matching of firms to male and female candidates—that is, firms' valuations over candidates might be more correlated with the preference of female candidates than male candidates. Another possibility is that firms are type-predictive and better at targeting offers to female candidates

Table B.5: Oaxaca-Blinder Decompositions of Gender Gaps

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Quantity	Composition			Inclusive Value			
	# Bids	$\overline{u}(b_{ij}, a_i)$	$\overline{A}_j$	ΔA_{ij}	Λ_i^b	$\overline{\Lambda}_i^A$	$\Delta\Lambda_i^A$	Λ_i^*
Mean Difference	0.507***	0.402***	0.021***	-0.008***	0.474***	0.085***	-0.013***	0.478***
$\overline{Y}_m - \overline{Y}_f$	(0.078)	(0.014)	(0.003)	(0.002)	(0.021)	(0.013)	(0.002)	(0.022)
Endowments	0.577***	0.151***	0.018***	0.025***	0.243***	0.111***	0.024***	0.287***
$\left(\overline{X}_m - \overline{X}_f\right)'\beta_m$	(0.045)	(0.009)	(0.002)	(0.002)	(0.013)	(0.008)	(0.002)	(0.015)
Coefficients	-0.083	0.242***	0.007*	-0.033***	0.215***	-0.026*	-0.037***	0.181***
$\overline{X}_f'\left(\beta_m - \beta_f\right)$	(0.074)	(0.013)	(0.004)	(0.002)	(0.020)	(0.013)	(0.002)	(0.021)
Interactions	0.012	0.010	-0.005*	-0.001	0.017	0.001	-0.000	0.010
$\left(\overline{X}_m - \overline{X}_f\right)'\left(\beta_m - \beta_f\right)$	(0.044)	(0.008)	(0.002)	(0.001)	(0.012)	(0.008)	(0.001)	(0.012)

Note: This table reports Oaxaca-Blinder decompositions of components of utility. Panel A reports decompositions by gender. Panel B reports decompositions by education. Column 1 decomposes the gap in the number of bids. Column 2 reports the decomposition of the mean gap in the monetary component of utility. Column 3 decomposes the mean difference in the common component of amenity values. Column 4 decomposes differences in candidate-specific components of the amenity valuation. Columns 5-8 report the decompositions of components of the inclusive value. The first row (Difference) reports the difference in means. The second row (Endowments) reports the component of the difference in means that can be attributed to differences in covariate values between the two groups. The third row (Coefficients) reports the component of the difference in means that can be attributed to differences in the returns to covariates between the two groups. The fourth row (Interactions) reports the component of the difference in means that cannot be attributed to differences in endowments or coefficients alone. The Endowments, Coefficients, and Interaction rows sum to the Difference row in every column, Robust standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

relative to male candidates, all else equal.[30] These qualitative patterns are reflected in the decompositions of components of inclusive values, reported in columns 5-8. Taken together, these results suggest that the large observed gender gap in bids is reflective of a large gender gap in welfare. Unconditionally, the gap in welfare between men and women is exacerbated by differences in the amenity values of the bids they receive. However, differences in covariates between men and women account for most of the unconditional gap.

H.2 Counterfactual scenarios of interest

To better understand the implications of imperfect competition for welfare, we use our supply and demand estimates to simulate bidding outcomes under all four conduct scenarios: {monopsonisitic competition, oligopsony}× {not predictive, type-predictive}. To

[30]Evidence from Section 6.3 that firms are in fact not type-predictive suggests the former explanation is more likely than the latter.

gauge the losses due to imperfect competition, we define a new form of conduct, which we term **price taking**. Under the price taking conduct alternative, firms have no discretion over the wages they offer. Instead, firms are constrained to offer a prevailing market wage, as if set by a Walrasian auctioneer. In our price-taking alternative, we set the equilibrium wage equal to the systematic component of firms' valuations, $b_{ij} = \exp(z_j' \Gamma x_i)$. Given this set of wages, the only decision firms have to make is whether to bid on each candidate. Because firms are price takers in this scenario, we assume that they view themselves as atomistic, as in monopsonistic competition.[31] In addition to these simulations, we also simulate the effects of a simple policy meant to reduce gender disparities in wages: blinding employers to candidates' gender. This counterfactual entails replacing gender-specific estimates of labor demand with cross-gender averages, and doing the same for estimates of labor supply.

H.3 Computing new counterfactual equilibria

In order to compute counterfactuals, we randomly select 500 firms and 500 candidates from the universe of firms and candidates in the analysis sample. For each firm-candidate pair, we compute the model-implied systematic component of firm valuations using our preferred estimates of labor demand parameters, $\exp(z_j' \widehat{\Gamma} x_i)$. Under a particular conduct assumption, equilibrium is determined by a set of beliefs over the distribution of the utility afforded by the best option in each candidates' offer set. The inclusive value is itself a sufficient statistic for the distribution of the maximum utility option for each candidate. At an equilibrium, firms' beliefs about inclusive values must be consistent with the true distribution of inclusive values generated by the bidding behavior of competing firms. We make the assumption that those beliefs depend only upon the expected value of the inclusive value to simplify our calculations here.

To compute new equilibria, we first conjecture an initial set of (expected) inclusive values Λ_i^1. We then iterate the following steps:

1. At iteration t, take *iid* draws from a normal distribution with mean zero and standard deviation $\widehat{\sigma}_\nu$ to produce a new set of idiosyncractic components of firms' valuations, ν_{ij}^t. Use these draws, plus the systematic components of valuations $z_j' \widehat{\Gamma} x_i$, to compute ε_{ij}^t.

2. Given ε_{ij}^t and Λ_i^t, compute b_{ij}^t as firm j's best response (under the assumed form of

[31] Because bids vary even conditional on our detailed controls, we automatically ruled out this form of price taking as a potential mode of conduct to describe firms' actual bidding behavior on the platform.

conduct). If there is no number b such that $G^m_{ij}(b)(\varepsilon_{ij} - b) \geq \widehat{c}_j$, then set $B^t_{ij} = 0$.

3. Given firms' best responses b^t_{ij} and B^t_{ij}, calculate the realized inclusive value for each candidate, $\Lambda^{t*}_i = \mathbb{E}[\log(\sum_{j:B^t_{ij}=1} \exp(u(b^t_{ij}, a_i) + A_{ij})]$. Compute the vector of expected inclusive values at the next iteration by taking a step $\alpha^t \in [0, 1]$ towards Λ^{t*}_i:
$$\Lambda^{t+1}_i = \alpha^t \Lambda^{t*}_i + (1 - \alpha^t)\Lambda^t_i.$$

We iterate this procedure until the distribution of inclusive values converges. We then use the equilibrium distribution of inclusive values to compute mean counterfactual outcomes by constructing the average across 50 simulations of firm bidding decisions.

H.4 Simulation Results

Table B.6 reports the results of our simulations. For each scenario, we compute the average bid, ratio of bid to ask, markdown, and number of bids received per candidate. We also compute the averages of (scaled) components of utility associated with each candidates' portfolio of bids. The absolute magnitudes of these components of utility do not have a direct interpretation, but relative differences across scenarios are meaningful.

The unconditional means of each of these variables across simulation repetitions are reported in Panel A of Table B.6. We first consider scenarios in which firms are assumed to be not predictive. Unsurprisingly, average bids are higher (\$169k vs \$145k), and markdowns are lower (10% vs 18%), in the price taking model (column 1) relative the the preferred monopsonistic competition model (column 2). Additionally, candidates receive markedly fewer bids (20 vs 43) under price taking than under monopsonistic competition, reflecting the increased labor costs under price taking. Even though they receive fewer bids under price taking, the increased monetary value of bids more than makes up for the substantial drop in the number offers: the average candidates' expected utility is higher under price taking than it is under monopsonistic competition. On the other hand, candidates fare far worse when firms act strategically (column 3): under oligopsony, candidates receive even fewer bids than when firms are price-takers (13.5), and the monetary value of those bids is even lower than under monopsonistic competition (\$139k). As a result, candidates' expected utilities are lowest under oligopsony. Interestingly, switching to modes of conduct in which firms are assumed to be type-predictive does little to change the unconditional means of each of the variables we summarize here (columns 4-6).

The lack of a difference between the type-predictive and not-predictive alternatives in unconditional mean outcomes obscures substantial differences in outcomes between

Table B.6: Counterfactual Simulations

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Unconditional Means						
	Not Predictive			Type-Predictive		
	PT	MC	OG	PT	MC	OG
Bid, b_{ij}	$169k	$145k	$139k	$169k	$145k	$139k
Ratio of Bid/Ask, b_{ij}/a_i	1.196	1.024	0.979	1.196	1.025	0.978
Markdown, $1 - b_{ij}/\varepsilon_{ij}$	0.099	0.182	0.182	0.099	0.183	0.183
# Bids Received/Candidate	20.1	43.2	13.5	19.6	42.0	13.2
Inclusive Value, Λ_i^*	0.930	0.886	0.822	0.932	0.888	0.822
Monetary Component, Λ_i^b	0.033	0.015	0.000	0.033	0.016	0.000
Common Amenity Comp., $\bar{\Lambda}_i^A$	0.282	0.357	0.315	0.281	0.355	0.314
Type-Specific Amenity Comp., $\Delta\Lambda_i^A$	0.002	0.004	0.004	0.005	0.008	0.007
Panel B: Differences, Women - Men						
	Not Predictive			Type-Predictive		
	PT	MC	OG	PT	MC	OG
# Bids Received/Candidate	-1.830	-3.793	-1.434	-2.411	-5.681	-2.529
Inclusive Value, Λ_i^*	-0.053	-0.069	-0.019	-0.056	-0.070	-0.019
Monetary Component, Λ_i^b	-0.026	-0.052	-0.016	-0.027	-0.051	-0.016
Common Amenity Comp., $\bar{\Lambda}_i^A$	-0.003	-0.005	-0.003	-0.004	-0.007	-0.004
Type-Specific Amenity Comp., $\Delta\Lambda_i^A$	0.005	0.010	0.013	0.003	0.010	0.011
Panel C: Differences, Women - Men, Gender Blind Firms						
	Not Predictive			Type-Predictive		
	PT	MC	OG	PT	MC	OG
# Bids Received/Candidate	-1.652	-3.749	-1.529	-2.776	-6.162	-2.549
Inclusive Value, Λ_i^*	-0.050	-0.066	-0.018	-0.053	-0.068	-0.019
Monetary Component, Λ_i^b	-0.025	-0.051	-0.016	-0.027	-0.050	-0.016
Common Amenity Comp., $\bar{\Lambda}_i^A$	-0.003	-0.005	-0.003	-0.002	-0.006	-0.002
Type-Specific Amenity Comp., $\Delta\Lambda_i^A$	0.004	0.011	0.013	0.005	0.009	0.011

Note: This table reports results of counterfactual simulations under various conduct assumptions. Columns labelled PT refer to the price-taking model of conduct, columns labelled MC refer to the monopsonistic competition model of conduct, and columns labelled OG refer to the oligopsony model of conduct. Each cell reports the average of the statistic over 50 simulation draws. In each simulation draw, we sample from the distribution of valuations for a set of 500 firms considering 500 workers (a single sample of workers and firms is used for all simulations). Panel A reports the unconditional means of various statistics. Panel B reports differences in means between women and men. Panel C reports differences in means between women and men for simulations in which firms are constrained to be gender blind.

men and women when firms are type-predictive relative to when they are not predictive. We report differences in mean outcomes across simulations between women and men in panel B of Table B.6. Across all conduct assumptions, women receive fewer bids than men (note, however, that this difference is not conditional on other characteristics). In absolute terms, the largest gender gaps in bids and welfare are predicted by the monopsonistic competition model, although these differences are partly driven by the fact that firms unconditionally make more bids under monopsonistic competition than they do under the other alternatives. Relative to the unconditional average, women receive 8-10% fewer offers when firms are not type predictive. The gap widens to 12-18% when firms are assumed to be type-predictive, and the oligopsony model predicts the largest relative gaps. Female candidates' expected utility also drops, although to only a relatively small degree. The upshot of these simulations is that firms have significant ability to exercise market power in ways that expand gender gaps, as first posited by Robinson (1933).

Can a simple policy that blinds employers to the gender of the candidates they consider narrow these gaps? Panel C reports differences between mean outcomes for men and women across simulation draws in which firms are constrained to no longer observe the candidate gender. The results from our simulations suggest that the efficacy of such a policy is relatively limited. Across all conduct possibilities, the policy is predicted to marginally increase the expected utility of female candidates relative to their male counterparts—across conduct scenarios, blinding employers to gender lowers the gender gap in expected utilities by 6-9.5%. Interestingly, while blinding not-predictive firms to gender modestly increases the number of offers women receive relative to men, the opposite is true when firms are type-predictive.