

Roussille, Nina; Scuderi, Benjamin

Working Paper

Bidding for Talent: A Test of Conduct in a High-Wage Labor Market

IZA Discussion Papers, No. 16352

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Roussille, Nina; Scuderi, Benjamin (2023) : Bidding for Talent: A Test of Conduct in a High-Wage Labor Market, IZA Discussion Papers, No. 16352, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/279050>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 16352

**Bidding for Talent:
A Test of Conduct in a High-Wage Labor
Market**

Nina Roussille
Benjamin Scuderi

JULY 2023

DISCUSSION PAPER SERIES

IZA DP No. 16352

Bidding for Talent: A Test of Conduct in a High-Wage Labor Market

Nina Roussille

Massachusetts Institute of Technology and IZA

Benjamin Scuderi

University of Michigan

JULY 2023

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Bidding for Talent: A Test of Conduct in a High-Wage Labor Market*

We develop a procedure for adjudicating between models of firm wage-setting conduct. Using data on workers' choice sets and decisions over real jobs from a U.S. job search platform, we first estimate workers' rankings over firms' non-wage amenities. We document three key findings: 1) On average, workers are willing to accept 12.3% lower salaries for a 1-S.D. improvement in amenities. 2) Between-worker preference dispersion is equally large, indicating that preferences are not well-described by a single ranking. 3) High-paying firms have better amenities. Following the modern IO literature, we use these estimates to formulate a test of conduct based on exclusion restrictions. Oligopsonistic models incorporating strategic interactions between firms and tailoring of wage offers to workers' outside options are rejected in favor of simpler monopsonistic models featuring near-uniform markdowns. Misspecification has meaningful consequences: while our preferred model predicts average markdowns of 19.5%, others predict average markdowns as large as 26.6%.

JEL Classification: J31, J42, L21

Keywords: wage-setting conduct, markdowns, monopsony

Corresponding author:

Nina Roussille
Department of Economics
Massachusetts
Institute of Technology
50 Memorial Drive
Cambridge, MA 02142
USA
E-mail: nroussil@mit.edu

* We are indebted to our advisors Patrick Kline, Christopher Walters, and Matthew Backus for invaluable support and guidance. We thank Daron Acemoglu, Nikhil Agarwal, David Autor, David Berger, Sydnee Caldwell, David Card, Attila Lindner, Claudia Macaluso, Alan Manning, Alex Mas, Parag Pathak, Yotam Shem-Tov, Gabriel Ulyssea, Basit Zafar, 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, and seminar participants from several other institutions. 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 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. 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, a single mode of wage-setting conduct prevails in equilibrium: firms 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 whether firms internalize differences in preferences across workers. These assumptions then become key ingredients in the assessment of markdowns, 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.

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 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., 2023), and second, whether firms tailor wage offers to individual workers' outside options (Postel-Vinay and Robin, 2002; Flinn

and Mullins, 2021). Our analysis builds on the modern Industrial Organization (IO) 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 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. Yet, observing these choice sets is typically impossible outside the lab: matched employer-employee data, for instance, only records the realized transitions of workers between firms.¹ To overcome these data limitations, we leverage the unique matching process on a large, high-stakes online job board. 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. 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 similar in spirit to Sorkin (2018). In particular, the estimated amenity value of any firm depends on the estimated amenity

¹The issue of measuring worker choice sets is widely recognized, as noted by Sorkin (2018): “In the empirical context of [...] studying transitions between employers, it is hard to imagine a data set that perfectly captures worker choice sets.” When workers’ choice sets are not measured, researchers must make assumptions to proxy for the distribution of available options. However, erroneous inference of choice sets can introduce substantial bias (Barseghyan et al., 2021).

values of the firms it was revealed-preferred to: for a firm to be highly-ranked, that firm must be revealed-preferred by workers over other highly-ranked 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. As a result, 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. Modelling the correlation of types with observable characteristics also 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 quasi-randomly 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. First, 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. Second, we document substantial vertical differentiation of firms on the platform: the average

worker is willing to pay 12.3% of her ask salary to enjoy a one standard deviation (1-S.D.) increase in firm amenities.² However, horizontal variation is just as important—the average standard deviation in valuations of a particular firm’s amenities across workers is 14%. 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. Third, we find that augmenting differentials prevail: firms that pay well are also firms with better amenities, such that between-firm dispersion in amenities amplifies inequality. On average, a 1-S.D. increase in amenity values is associated with a 0.325-S.D. increase in the firm pay premium.

We use the labor supply estimates to implement our procedure for comparing models of firm behavior. To formulate the exclusion restriction we use for our test, we leverage knowledge of platform rules to construct an instrument that captures quasi-random fluctuations in potential competition over time and across sub-markets. 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 model, we find markdowns of 19.5% on average, while oligopsony alternatives imply average markdowns of 26.6%. We also find large differences between models in implied productivity dispersion across firms. Indeed, while firms with relatively better amenities are more productive under both oligopsony and monopsonistic competition, the slope of the relationship is nearly six times larger under the former than under the latter. This leads to large differences in implied productivity dispersion across firms: in the preferred model, firms with the best amenities ($+2\sigma$) are 3.4% more productive than firms with the worst amenities (-2σ). Under oligopsony, that difference is 10.6%. What drives these large differences? In all imperfect competition models, every firm internalizes an upward-sloping firm-specific labor supply curve. However, under oligopsony, firm-specific labor supply elasticities are increasing in the value of firms’ amenities. Under monopsonistic competition, however, firm-specific labor supply elasticities do not depend upon firms’ amenities.

We then turn to testing whether firms take advantage of the significant predictable

²Note that, as long as there is not perfect competition, the existence of vertical differentiation does not necessarily imply that there are compensating differentials in the labor market.

(horizontal) variation in firm-specific labor supply across workers when making hiring decisions, and find that they do not: firms are not type predictive. This is especially striking in the context of an online job board that ostensibly seeks to reduce information frictions in the search and matching process. This finding also has significant implications for the labor market: assuming type-predictive oligopsony would imply that the offers firms make made to the workers who value their amenities are marked down 3.0pp more than the offers they make to workers who value them the least.

This paper contributes to several strands of literature. First, our paper is related to a growing literature that employs tools from IO 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. \(2023\)](#), 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 IO literature that tests conduct in product markets ([Berry and Haile, 2014](#); [Backus et al., 2021](#); [Duarte et al., 2023](#)).

Next, our paper contributes to the literature on the estimation of non-wage amenities ([Rosen, 1986](#)). While recent papers have leveraged experimental settings to estimate the value of non-wage amenities ([Mas and Pallais, 2017](#); [Wiswall and Zafar, 2018](#)), we study the decisions of workers in a real-world, high-stakes environment. Our paper most relates to [Sorkin \(2018\)](#), who applies tools from numerical linear algebra to U.S. matched employer-employee data in order to estimate workers’ preferences for firms. [Taber and Vejlin \(2020\)](#) and [Lagos \(2021\)](#) also use data on equilibrium matches to infer amenity values from the workers flows across firms. By contrast, we observe the full set of options available to each worker on the platform, and therefore estimate non-wage amenity values of firms by aggregating candidates’ revealed preferences over these options.

Our paper also contributes to a broader literature exploring the nature of 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 extent of firms’ wage-setting power ([Azar et al., 2020](#); [Schubert et al., 2022](#); [Arnold, 2021](#)). These analyses echo the “Structure-Conduct-Performance” (SCP) paradigm ([Robin-](#)

son, 1933; Chamberlain, 1933; Bain, 1951), which posits that firm conduct is dictated by market structure (that is, a firm’s optimization problem is a deterministic function of the distribution of its competitors). But since wages and market structure are joint outcomes in models of labor markets, it is both conceptually and practically difficult to find instruments that affect wages only through their effect on market structure (Berry, 2021; Schmalensee, 1989). To avoid these endogeneity issues, we propose a method to characterize firms’ exercise of wage-setting power that does not rely on an assumed equivalence between (observed) market structure and conduct.

Online search is now the most widely used job-search method in the United States (Faberman and Kudlyak (2016)). Echoing the spectacular growth of these platforms, a recent strand of literature in labor and IO has focused on the nature of online competition. We contribute to this literature by proposing models of imperfect competition adapted to these settings, combining the characteristics of online auctions and terrestrial labor markets. The closest paper 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 job board 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.

2 Setting and Data

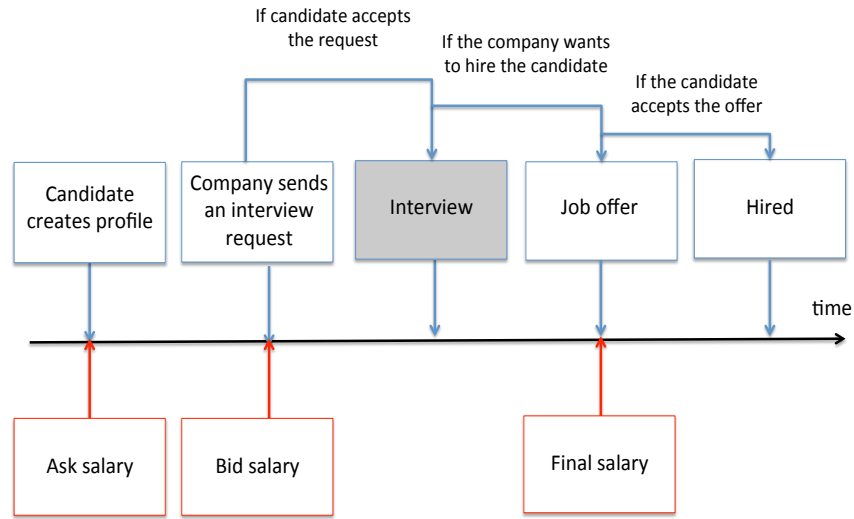
2.1 Market description

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 and 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 these limitations. First, wage bargaining on Hired.com is high-stakes: the average candidate on the the platform is a software engineer living in San Francisco with 11.3 years of experience looking for a full-time job with an expected salary of \$137,000. Candidates on Hired.com are highly qualified: 87.2% have at least a college degree (with 40.3% having additionally completed some form of graduate education), and 10.8% have

prior experience at a FAANG company. Importantly, most candidates are engaging in on-the-job search: 74.8% report being currently employed. 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 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 over it (whether or not they decided to interview with any given firm in the consideration set).

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



Note: This figure depicts the timeline of a recruitment on Hired.com. Salaries that are captured on the platform are denoted in red. The steps of the process, from profile creation to hiring, are colored blue. We do not have metadata from companies on the interview process.

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: the *ask salary*.³ Second, firms get access to candidate profiles that match standard requirements for the job they want to fill (i.e.,

³See Roussille (2022) for a more detailed description of a typical candidate profile. 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 this job, their

job title, experience, 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: the *bid salary*.⁴ 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 which *final salary*. It is important to note that the bid salary is non-binding, so the final salary can differ from it.

When modelling the recruitment process on Hired.com, we abstract from dynamic considerations for a number of reasons. Candidate profiles are only visible to firms for two weeks by default, and so candidates collect and consider bids over a short time frame. The median candidate receiving multiple bids collects those bids within a single week. Further, we find strong evidence that firms send most interview requests for the same job concurrently: the median time difference between sequential bids for the same job is about 13 minutes. Finally, firms do not observe the remaining time candidates have on the platform, and so cannot bid strategically over time.

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). Candidates in San Francisco represent 76% of all interview requests on the platform. Consequently, our analysis focuses on that subset of workers, which represents the largest homogeneous labor market on the platform. For this segment of the platform, 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 for inclusion in 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.

top skills, their education, their work history (i.e., firms they worked at), their contract preferences (remote or on-site, contract or full-time), as well as their search status, which describes whether the candidate is actively searching or simply exploring new opportunities. The ask salary is prominently featured on all profiles since it is a required field.

⁴See Roussille (2022) for details on the typical interview request messages sent by companies.

2.3 Stylized facts

Significant heterogeneity in bid acceptance. Panel (a) of Figure 2 plots the distribution of the share of each firm’s bids that are accepted. There are two important features of this distribution. First, perhaps because most candidates are currently employed, rejections are common: on average, candidates only accept 60.5% of the interview requests they receive. Second, there is significant heterogeneity across companies in the likelihood that a request is accepted: 10.2% of firms see less than 40% of their requests accepted, while 16.2% of firms see more than 75% of their requests accepted. 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, Panel (b) of Figure 2 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 10-15pp lower than when a bid is made exactly at the ask. However, the probability a bid is accepted when it is 10% more than the ask is only about 5pp 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.⁵ Panel (c) of Figure 2 plots the probability that a candidate selects “insufficient compensation” as the reason for rejecting a bid as a function of the ratio of bid to ask. The relationship between this probability and the bid to ask ratio is sharply kinked at bid=ask: the slope (and level) is almost exactly zero when bid>ask, and is strongly negative when bid<ask. In practice, this means that while virtually no bid is rejected due to “insufficient compensation” when the bid is above the ask, if the bid is, for instance, 20% below the ask, roughly 25% of interview rejections are due to “insufficient compensation”. We refer to this phenomenon as “kinked labor

⁵While this field is optional, 55% of candidates do fill it out.


supply” and formally allow for labor supply elasticities to vary above and below the kink in our model.⁶

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 \$19,697, which indicates that firms are willing to offer a wide range of salaries to candidates for the same vacancy. Indeed, only 2.6% of jobs offer the same bid salary to all candidates. Further, the bids firms make to candidates are highly individualized: 77.4% of bids are made *exactly* at the candidates’ ask. Panel (d) of Figure 2 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’ bid decisions as a fully-individualized process, allowing for systematic and idiosyncratic components of match-specific productivity.


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 not 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. Panel (e) of Figure 2 shows the relationship between the bid and final offer for the subset of candidates that receive one. Strikingly, this relationship is very linear,

⁶Leveraging 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. Both Panel (b) and (c) have a vertical grey dashed line at Bid=Ask. 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.

with a slope close to one. Additionally, about a third of all final offers are identical to the bid, and close to three-quarters 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 $i = 1, \dots, N$ post resume information x_i (which includes their ask a_i) before interacting with firms. Firms $j = 1, \dots, J$ have observable characteristics z_j . 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 Panel (e) of Figure 2 and described in Section 2.3.⁷ 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:

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

⁷Horton et al. (2021) also highlights the informative content of cheap talk about wages.

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 $\lim_{b \rightarrow a^-} \partial u(b, a) / \partial b > \lim_{b \rightarrow a^+} \partial u(b, a) / \partial b$. 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 $u(a, a) = 0$.⁸ 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 associates 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 : $\xi_{ij} \stackrel{iid}{\sim} F_\xi(\cdot)$, 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:

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

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.

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-

⁸For 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.

specific interview cost threshold c_j :

$$b_{ij}^* = \arg \max_b \pi_{ij}(b), \text{ and } B_{ij} = \mathbf{1} [\pi_{ij}(b_{ij}^*) \geq c_j]. \quad (2)$$

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:

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

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:

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

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:

$$\pi_{ij}(b) = \underbrace{\Pr_{ij}(D_{ij}^\circ(b) = 1)}_{\triangleq G_{ij}(b)} \times \underbrace{(\mathbb{E}_{ij}[\varepsilon_{ij}^\circ] - b)}_{\triangleq \varepsilon_{ij}}. \quad (3)$$

The first term, $G_{ij}(b)$, is j 's forecast of i 's labor supply decision, which we refer to as the firm's *beliefs* (or win probability).⁹ 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.

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 $G_{ij}(b)$:

$$b_{ij}(\varepsilon) = \begin{cases} \arg \max_b G_{ij}(b) \times (\varepsilon - b) & \text{if } \max_b G_{ij}(b) \times (\varepsilon - b) \geq c_j \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

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

$$G_{ij}(b) = \iint \Pr(u(b, a_i) + \Xi_{ij} = V_i^1 \mid V_i^1 = v, Q_i = q) \times dF_{V,Q}(v, q \mid \Omega_{ij}), \quad (5)$$

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

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:

$$F_{V,Q}(v, q \mid \Omega_{ij}) = \underbrace{F_{V|Q}(v \mid Q_i = q, \omega_{ij}^V)}_{=F_{V|Q}^\omega} \times \underbrace{F_Q(q \mid \omega_{ij}^Q)}_{=F_Q^\omega}. \quad (6)$$

⁹We 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.

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 $\Omega_{ij} = \{\omega_{ij}^V, \omega_{ij}^Q\}$. We specify two alternatives for each component—firms are either:
 - **Not Predictive**, with $\omega_{ij}^Q = \emptyset$ 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 $b_{ij}, \mathbf{A}_j \notin \omega_{ij}^V$ such that $\partial F_{V|Q}^\omega / \partial b = 0$; or **Oligopsonists**, with $b_{ij}, \mathbf{A}_j \in \omega_{ij}^V$ such that $\partial F_{V|Q}^\omega / \partial b > 0$.
- 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 Online Appendix (hereafter Appendix) C, 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 more than type-predictive: they

are 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). Type predictiveness has important labor market implications: in our setting, firms would make more offers and workers 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. \(2023\)](#), 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.¹⁰

Finally, our model of perfectly competitive firms serves as a 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 perfect competition, firms bid their valuations: $b_{ij}(\varepsilon) = \varepsilon$.

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 best describes the true data-generating process.¹¹ To formulate our test, we first write ε_{ij} as a function of observ-

¹⁰Our 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).

¹¹The 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” ([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} .

ables and a mean-zero idiosyncratic component $\nu_{ij} \stackrel{iid}{\sim} F_\nu(\cdot)$ 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 $\tau(\gamma_j(\cdot, \cdot))$ 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 that $G_{ij}(b)$ is differentiable for all b . Then, under the true conduct assumption, all bids b_{ij} must satisfy the following first-order condition with equality:

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

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 , $\varepsilon_{ij}^m(\cdot)$. If m is misspecified, then this substitution introduces an additional error term:

$$\tau(\varepsilon_{ij}^m(b_{ij})) = \gamma(x_i, z_j) + \nu_{ij} + \zeta_{ij}^m.$$

The presence of specification error suggests an intuitive conclusion: if labor supply is determined in part by variables that are excluded from firms' valuations (ε_{ij}), then the further a model is from the truth, the higher the correlation between those excluded variables and the model's residuals. 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.

Following this logic, [Berry and Haile \(2014\)](#) establish that instruments that quasi-randomly shift demand but do not shift (are excluded from) the marginal cost function are necessary for both identification and conduct testing in product markets. [Backus et al. \(2021\)](#) implement a test of conduct that formalizes this logic: under true conduct assumptions, instruments that quasi-randomly shift markups but not marginal costs should not be correlated with recovered idiosyncratic cost shocks. Likewise, we need to find an instrument that quasi-randomly shifts labor supply but is excluded from firms' valuations, such that it is uncorrelated with true demand residuals ν_{ij} .¹²

¹²Our 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. An advantage of our granular data is that we can identify labor supply parameters by conditioning on the information available to firms when they bid, such that we do not have to rely on instruments for bids. Second, we analyze firms' initial *individualized* bids rather than uniform market prices. Our identification arguments therefore follow

To construct the instrument, we leverage quasi-random, high-frequency variation in *potential* labor market tightness generated by platform rules, both between and within granular sub-markets.¹³ Specifically, we take advantage of the fact that candidate profiles go live in batches and remain searchable for only two weeks.¹⁴ This generates large fluctuations in the number of candidates, relative to firms, that are live on the platform in a particular sub-market at any given time. Since the pool of candidates turns over every two weeks, variation in candidate quality between two-week periods is not endogenously determined by platform conditions—and so this variation 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 fewer active candidates there are per active firm, the more bids those candidates tend to receive. Our use of potential market tightness as an instrument mirrors 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)).

Formally, let v_{ow} denote the number of firms searching for experience and occupation o during two-week period w and let u_{ow} be the number of candidates with active profiles with experience and occupation o during two-week period w . The prevailing level of (inverse) potential market tightness when j bids on i is: $t_{ij} = u_{o_i w_{ij}} / v_{o_i w_{ij}}$. Our instrument exogeneity assumption can be formalised as:

Assumption 1. (Instrument Exogeneity) *Conditional on worker and firm observables x_i and z_j , the instrument t_{ij} (potential tightness) obeys:*

- a) **(Quasi-Random Assignment)** *Across ij pairs, the prevailing level of potential tightness is as-good-as randomly assigned, and*
- b) **(Exclusion Restriction)** *Potential tightness is not a determinant of the idiosyncratic component of labor demand,*

and so t_{ij} is (conditionally) independent of the idiosyncratic component of demand:

$$t_{ij} \perp\!\!\!\perp \nu_{ij} \mid x_i, z_j. \quad (7)$$

the empirical auction literature (Guerre et al., 2000; Backus and Lewis, 2020) by assuming that firms’ behavior must satisfy rational expectations rather than a market-clearing condition.

¹³We call our instrument potential tightness because it measures the relative number of firms that *may* bid on candidates during a two-week period, whether or not they actually decide to participate. We define the instrument within occupation and experience bins, since those categories are the primary search fields recruiters use when browsing candidates.

¹⁴Candidates 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.

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. That is, we consider each pair of models in turn and select the model that has the lowest correlation between the excluded variables and the model's residuals. To operationalize this test, we specify a scalar moment condition in the residuals of fitted models and excluded instruments, as in Backus et al. (2021). Because we estimate demand under each conduct assumption via maximum likelihood, our test is based on *generalized residuals* defined by the scores of the likelihood (Gourieroux et al., 1987).

Formally, 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 $\hat{\Psi}^m$ is the vector that sets:

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

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 of the individual likelihood contributions. We then compute the covariance between the generalized residuals of model m and the excluded instrument t_{ij} as our scalar moment/lack-of-fit measure:

$$Q_s^m = \left(\frac{1}{s} \sum_{ij: B_{ij}=1} h_{ij}^m(\hat{\Psi}^m) \cdot t_{ij} \right)^2, \quad (8)$$

where $s = |\{ij : B_{ij} = 1\}|$.¹⁵ 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 resid-

¹⁵ Q_s^m can also be motivated as a 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.

ual.¹⁶ Following Backus et al. (2021),¹⁷ we formulate a pairwise 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:

$$T_s^{m_1, m_2} = \frac{Q_s^{m_1} - Q_s^{m_2}}{\hat{\sigma}_s^{m_1, m_2} / \sqrt{s}} \xrightarrow{D} \mathcal{N}(0, 1), \quad (9)$$

where $\hat{\sigma}_s^{m_1, m_2}$ is an estimate of the population variance of $Q^{m_1} - Q^{m_2}$. We compute $\hat{\sigma}_s^{m_1, m_2} / \sqrt{s}$ 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.

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 :*

$$\Pr(\mathcal{B}_i \mid Q_i = q, x_i) = \Pr(\mathcal{B}_i \mid x_i). \quad (10)$$

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:

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

Assumptions analogous to Assumption 2 are implausible in administrative data, like linked employer-employee records, due to the various selection mechanisms at

¹⁶In Appendix G, we describe and implement an alternate testing procedure based on the Vuong (1989) likelihood ratio test. While our version of the Rivers and Vuong (2002) test isolates only the component of lack-of-fit directly correlated with the instrument, the alternate test combines all sources of residual variation and can be thought of as an omnibus version of our lack-of-fit measure.

¹⁷Backus 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 by selection and partial identification issues. While not pursued here, the formulation of optimal instruments is a promising avenue for future work.

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 $1, \dots, Q$. Denote the conditional probability of type membership by:*

$$\Pr(Q_i = q \mid x_i) \triangleq \alpha_q(x_i). \quad (11)$$

- 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 :*

$$\Pr(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \mathcal{B}_i, Q_i = q, x_i) = \Pr(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \mathcal{B}_i, Q_i = q) \triangleq \mathcal{P}_q(\mathcal{B}_i^1 \succ \mathcal{B}_i^0). \quad (12)$$

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:¹⁸

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

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

¹⁸Mixtures 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. Soufiani et al. (2013) also provide simulation evidence that estimation methods can correctly recover the true number of underlying types.

depend on candidate type, and write it as:

$$u_q(b, a) = (\theta_{q0} + \theta_{q1} \cdot \mathbf{1}[b < a]) \cdot [\log(b) - \log(a)] = \begin{cases} \theta_{q0} \cdot \log(b/a) & \text{if } b \geq a, \\ (\theta_{q0} + \theta_{q1}) \cdot \log(b/a) & \text{if } b < a, \end{cases}$$

and so $u_q(b, a)$ is continuous, but kinked, at $b = a$.¹⁹ We specify the distribution of types as a multinomial logit in x_i with parameter β :

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

Because Q_i has finite support, we write $A_j(Q_i) = \mathbf{Q}_i' \mathbf{A}_j$, where \mathbf{A}_j is a $Q \times 1$ vector of type-specific mean amenity values at firm j with q -th component A_{qj} , and \mathbf{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: $\xi_{ij} \stackrel{iid}{\sim} EV_1$.

Estimation: First Step. We estimate labor supply parameters via a two-step procedure. We first estimate type distribution parameters β and amenity values \mathbf{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:

$$\mathbf{Q}_i'(\mathbf{A}_j - \mathbf{A}_k) \geq \xi_{ik} - \xi_{ij}. \quad (13)$$

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 $u_q(b, a)$, 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. Plugging in estimates \hat{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.²⁰

To derive the probability of observing an arbitrary partial ordering of firms, it

¹⁹Note that we have defined $u(b, a)$ relative to the outside option: when $b = a$, $\log(b/a) = \log(1) = 0$. When making utility comparisons between candidates, we add back the monetary component associated with the outside option: $u_q(b, a) + \theta_{q0} \cdot \log(a)$.

²⁰Using a two-step procedure allows us to sidestep the need for instruments for bid salaries, if at the cost of the additional precision afforded by a one-step procedure that optimally combined multiple sources of variation. In addition, our strategy allows us to isolate “clean” comparisons without imposing additional assumptions necessary to justify instruments.

is useful to work with the re-parameterization $\rho_{qj} \propto \exp(A_{qj})$, with $\sum_{j=1}^J \rho_{qj} = 1$. Let $\sigma(\cdot) : \{1, \dots, J\} \rightarrow \{1, \dots, J\}$ denote a complete ranking of all J alternatives. A multinomial logit model of rankings (also known as “exploded logit”, or Plackett-Luce (Plackett, 1975; Luce, 1959)) yields the following likelihood:

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

One complication is that we only observe candidates’ partial orderings of firms, not their complete ranking. 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 D, we show that:

$$\mathcal{P}(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \boldsymbol{\rho}_q) = \int_0^1 \prod_{j \in \mathcal{B}_i^1} \left(1 - v^{\rho_{qj} / \sum_{k \in \mathcal{B}_i^0} \rho_{qk}}\right) dv. \quad (14)$$

This expression, and its derivatives, can be quickly and accurately approximated by numerical quadrature.²¹

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.²²

Estimation: Second Step. Next, we estimate the remaining labor supply elasticity and outside option parameters $\Theta = \{\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, \mathbf{A}_0\}$ via GMM using the full set of bids made by firms which can be ranked. We first construct model-implied probabilities

²¹Appendix D provides details on the generalized EM-algorithm we use to estimate $\boldsymbol{\beta}$ and $\boldsymbol{\rho}$.

²²While we do not present a formal proof of consistency here, parameter consistency and asymptotic normality of the MLE for similar models (pairwise comparisons with a single type) 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 (Simons and Yao, 1999).

of accepting an interview request as a function of Θ , plugging in $\hat{\beta}$ and $\hat{\rho}$ 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:

$$m(b_{ij}, x_i \mid \Theta) = \sum_{q=1}^Q \alpha_q(x_i \mid \hat{\beta}) \cdot \Lambda((\theta_{q0} + \theta_{q1} \cdot \mathbf{1}[b_{ij} < a_i]) \cdot \log(b_{ij}/a_i) + \hat{A}_{qj} - A_{q0}).$$

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

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

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

$$\hat{\Theta} = \arg \min_{\Theta} \widehat{m}(\Theta)' \mathbf{W} \widehat{m}(\Theta)$$

for a symmetric, positive-semidefinite weighting matrix \mathbf{W} .²³

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 $\Lambda_i = \log \left(\sum_{k: b_{ik} > 0} \exp(u_{Q_i}(b_{ik}, a_i) + Q'_i A_k) \right)$:

$$\Pr(V_{ij} = V_i^1 \mid \Lambda_i, b_{ij} = b) = \exp(u_{Q_i}(b, a_i) + Q'_i A_j) / \exp(\Lambda_i). \quad (15)$$

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

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

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 in [Guerre et al. \(2000\)](#) (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.

²³We estimate these parameters using a Continuously-Updated GMM (CUGMM) procedure that sets $\mathbf{W} = \mathbf{W}(\Theta)$. An efficient two-step GMM procedure yields extremely similar results.

Estimation. We first construct inclusive values Λ_i using our labor supply parameter estimates. We then use the empirical distribution of Λ_i to construct approximations to $G_{ij}(b)$ under each model of conduct. A given model of conduct is defined as a combination 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 vs. Oligopsony: Monopsonistically-competitive firms do not account for the contribution of their own bid to the inclusive value Λ_i —in other words, $\{b_{ij}, \mathbf{A}_j\} \notin \omega_{ij}^V$. Under this assumption, firms' beliefs are:

$$G_{ij}(b) = \sum_{q=1}^Q \alpha_q(\omega_{ij}^Q) \cdot \left(\exp(u_q(b, a_i) + A_{qj}) \times \mathbb{E}[\exp(-\Lambda_{iq}) \mid \omega_{ij}^V] \right). \quad (16)$$

Since firms are assumed to have rational expectations conditional on ω_{ij}^V , the quantity $\mathbb{E}[\exp(-\Lambda_{iq}) \mid \omega_{ij}^V]$ is identified and can be estimated by constructing the sample conditional expectation of $\exp(-\Lambda_{iq})$ given the variables contained in ω_{ij}^V (which include x_i , z_j , and market-level covariates).²⁴

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}^V is given by $\Lambda_{iq} \mid \omega_{ij}^V \sim \exp(u_q(b_{ij}, a_i) + A_{qj}) + \exp(\Lambda_{iq}^{-j})$, where $\Lambda_{iq}^{-j} = \log(\sum_{k \neq j: B_{ik}=1} \exp(u_q(b_{ik}, a_i) + A_{qk}))$ 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:

$$G_{ij}(b) = \sum_{q=1}^Q \alpha_q(\omega_{ij}^Q) \cdot \int \left(\frac{\exp(u_q(b, a_i) + A_{qj})}{\exp(u_q(b_{ij}, a_i) + A_{qj}) + \exp(\lambda)} \times dF_{\Lambda_q^{-j}}(\lambda \mid \omega_{ij}^V) \right). \quad (17)$$

Again, since firms' beliefs are assumed to be consistent, $F_{\Lambda_q^{-j}}(\lambda \mid \omega_{ij}^V)$ is identified and can be estimated by constructing the empirical distribution of leave-one-out inclusive values in the sample conditional on the variables in ω_{ij}^V . These estimates can then be used to construct a numerical approximation to the integral over the distribution of leave- j -out inclusive values.²⁵

²⁴When there are no differences in labor supply elasticities by preference type ($\theta_{q0} = \theta_0$ and $\theta_{q1} = \theta_1$ for all q), the beliefs of monopsonistically-competitive firms 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 $b_{ij} = a_i$, $\mu_{ij}^m = a_i / \varepsilon_{ij} \in [\frac{\theta_0}{1 + \theta_0}, \frac{\theta_0 + \theta_1}{1 + \theta_0 + \theta_1}]$.

²⁵Unlike monopsonistic competition, there is no simple closed-form expression for markdowns in the oligopsony case when labor supply elasticities do not vary by type.

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

We produce approximations to $G_{ij}(b)$ under all four combinations of these assumptions (i.e. $\{\text{Monopsonistic Competition, Oligopsony}\} \times \{\text{Type Predictive, Not Predictive}\}$).

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:

$$\varepsilon_{ij}^m(b) = b + \frac{G_{ij}^m(b)}{g_{ij}^m(b)} = \gamma_j^m(x_i, \nu_{ij}^m).^{26} \quad (18)$$

Crucially, given a choice of model m and labor supply parameters, the inverse bidding function is *known*: 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 E](#) establishes that bidding strategies $b_{ij}^m(\cdot)$ and option values $\pi_{ij}^{m*}(\cdot)$ are nevertheless continuous, monotonic functions in ε_{ij} .²⁷ Bids therefore partially identify valuations, motivating

²⁶Labor economists may be more familiar with the equivalent formulation of the firms' first-order condition in terms of a multiplicative markdown $\mu_{ij}^m(b)$ expressed as a function of the elasticity of labor supply to the firm: $\mu_{ij}^m(b) = \eta_{ij}^m(b)/(1 + \eta_{ij}^m(b))$, where $\eta_{ij}^m(b) = b \cdot g_{ij}^m(b)/G_{ij}^m(b)$.

²⁷This is due to the log-concavity of F_ξ and shape restrictions on $u(b, a)$. In particular, $b_{ij}^m(\cdot)$ is

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 $b_{ij} > 0$.

Selection Correction and 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 E). 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 quantity. We denote these values by $\hat{\pi}_{ij}^{m*}$, and use them to construct a consistent estimate of each firm j 's interview cost threshold for each m by setting:

$$\hat{c}_j^m = \min_{i: B_{ij}=1} \hat{\pi}_{ij}^{m*} \xrightarrow{\text{a.s.}} c_j^m. \quad (19)$$

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

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: $\hat{\underline{\varepsilon}}_{ij}^m$ is the number that sets $\pi_{ij}^{m*}(\hat{\underline{\varepsilon}}_{ij}^m) = \hat{c}_j^m$. We use these lower bound estimates to construct the likelihood contribution of each bid:

$$\begin{aligned} \mathcal{L}_{ij}^m(\Psi^m) &= \Pr\left(\varepsilon_{ij} = \varepsilon_{ij}^m(b_{ij}) \mid \varepsilon_{ij} \geq \hat{\underline{\varepsilon}}_{ij}^m, \Psi^m\right)^{\mathbf{1}[b_{ij} \neq a_i]} \times \Pr\left(\varepsilon_{ij} \in [\varepsilon_{ij}^{m-}, \varepsilon_{ij}^{m+}] \mid \varepsilon_{ij} \geq \hat{\underline{\varepsilon}}_{ij}^m, \Psi^m\right)^{\mathbf{1}[b_{ij} = a_i]} \\ &= \left(\frac{f_\varepsilon(\varepsilon_{ij}^m(b_{ij}); \Psi^m)}{1 - F_\varepsilon(\hat{\underline{\varepsilon}}_{ij}^m; \Psi^m)}\right)^{\mathbf{1}[b_{ij} \neq a_i]} \times \left(\frac{F_\varepsilon(\varepsilon_{ij}^{m+}; \Psi^m) - F_\varepsilon(\max(\varepsilon_{ij}^{m-}, \hat{\underline{\varepsilon}}_{ij}^m); \Psi^m)}{1 - F_\varepsilon(\hat{\underline{\varepsilon}}_{ij}^m; \Psi^m)}\right)^{\mathbf{1}[b_{ij} = a_i]}, \quad (20) \end{aligned}$$

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, respectively, the model-implied upper and lower bounds on ε_{ij} when $b_{ij} = a_i$.²⁸

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} .

²⁸Our approach—concentrating c_j out of the likelihood by computing the minimum order

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

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

where both x_i and z_j include a constant (such that $z'_j \Gamma x_i$ includes a constant, and all main effects and interactions of x_i and z_j). 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.

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 under each method of clustering workers into types, we compute standard likelihood ratio statistics and compute the appropriate χ^2 p -value to test against the null hypothesis that the model with q types is equivalent to the model with $q - 1$ types. 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:

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

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

Table 1 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 numbered group of columns corresponds to the method used to assign 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

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).

Table 1: Candidate Preference Model Goodness-of-Fit

# Types (q)	(1)			(2)			(3)		
	Split on Gender			Split on Experience			Model-Based Clusters		
	Log L.	$p_{q>q-1}$	GOF	Log L.	$p_{q>q-1}$	GOF	Log L.	$p_{q>q-1}$	GOF
1	-47,207	-	0.677	-47,207	-	0.677	-47,207	-	0.677
2	-46,441	0.999	0.685	-46,287	0.015	0.687	-45,244	<0.001	0.744
3	-	-	-	-	-	-	-44,298	0.001	0.772
4	-	-	-	-	-	-	-43,507	0.987	0.798
<hr/>									
Number of:	Firms: 1,649			Candidates: 14,344			Comparisons: 235,827		

Note: This table reports maximized log likelihoods (Log L.), likelihood ratio test p -values ($p_{q>q-1}$), and goodness-of-fit (GOF) measures to adjudicate between labor supply models with different numbers of types. Each numbered group of columns represents a different way to cluster candidates into preference types. The GOF statistic is calculated as the fraction of pairwise comparisons correctly predicted by the model, $E[(\hat{A}_j(Q_i) > \hat{A}_k(Q_i)) \times (j \succ_i k)]$, and p -values are calculated against the null hypothesis that the model with q types is equivalent to the model with $q - 1$ types.

to define latent preference groupings.²⁹ 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 1, the one-type model increases GoF over that baseline to 0.677. This relatively large increase in explanatory power compared to the benchmark indicates significant vertical differentiation of firms.

Column 1 of Table 1 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 0.677 to 0.685), and the formal LR test fails to reject the null that the two-type and one-type models are equivalent ($p = 0.999$). This finding mirrors that of Sorkin (2017), who also finds that estimated average preference orderings of men and women 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 ($p = 0.015$), the GoF statistic increases by just 1pp. However, using the full set of observables to define types (Column 3) performs markedly better than the

²⁹In order to compare the results under separate groupings, we maintain the same sample of bids/comparisons in each column. However, not every firm in the overall connected set is accepted and rejected at least once by a candidate of each gender/experience level. When splitting by gender or experience categories, we therefore assign weights α_{iq} of 0.95 to each candidate's own-group and 0.05 to the other group, which maintains overlap.

gender- and experience-split models. With two types, the GoF statistic is 0.744, an almost 6pp larger increase 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 ($p \leq 0.001$). However, we are unable to reject the null hypothesis that the four-type alternative is equivalent to the three-type model ($p = 0.987$). We therefore adopt the three-type version as our baseline. Panel (a) of Figure A.1 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 (in expectation).

Plugging in the estimated rankings into our second-step GMM procedure yields the following labor supply elasticity parameter estimates:

$$u_q(b_{ij}, a_i) = \log(b/a_i) \times \begin{cases} 3.60 + 1.50 \cdot \mathbf{1}[b < a_i] & \text{if } Q_i = 1, \\ (0.21) & (0.25) \\ 3.95 + 1.62 \cdot \mathbf{1}[b < a_i] & \text{if } Q_i = 2, \\ (0.19) & (0.23) \\ 4.19 + 1.53 \cdot \mathbf{1}[b < a_i] & \text{if } Q_i = 3. \\ (0.18) & (0.22) \end{cases}$$

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 5.8.³⁰

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, we compute the probability that a firm was rejected for a job-related reason and regress these probabilities on firms' ordinal ranks under the one-type model (higher ranks are better). We estimate a strongly significant negative relationship, such that a one-percentile increase in estimated firm rank is associated with a -0.090 (0.014) decrease in the probability of rejection for job-related reasons. Figure A.2 plots this relationship. Finally, Figure A.3 leverages our access to firms' listed benefits

³⁰Note that, in contrast with other studies, our model allows for kinked labor supply and so we estimate elasticities of 5.1-5.7 below the kink, i.e. when $b < a_i$, and 3.6-4.2 above the kink.

on their Hired.com feature page³¹ to depict the relationship between listed benefits and estimated ranking for a sub-sample of firms for which these benefits could be collected. Panel (a) reports the distribution of the number of listed benefits in this sub-sample. Panel (b) depicts the strong positive correlation between the number of listed benefits and a firm’s estimated rank.

6.2 Significant Vertical and Horizontal Differentiation of Firms

Figure 3 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:

$$\left(\hat{\theta}_{q0} + \hat{\theta}_{q1} \times \mathbf{1}[\text{WTA}_{qj} < 1]\right) \times \log(\text{WTA}_{qj}) + \hat{A}_{qj} - \hat{A}_{q0} = 0.$$

where A_{q0} is the q -th component of the vector of mean outside option values. Panel (a) of Figure 3 plots the distribution of the mean WTA at each firm, averaging over the population probabilities of each type:

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

The average mean WTA is 0.985, indicating that candidates are willing to accept 1.5% less than their ask at the average firm. The standard deviation (S.D.) of mean WTA across firms is 0.123 (or 12.3% of the ask), which indicates 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) of Figure 3 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 S.D. of WTA is 0.140, suggesting that the scale of systematic horizontal differentiation is comparable to that of 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: substantial horizontal dif-

³¹Firms have profiles on Hired.com that candidates can consult and that contain a description of the firm’s mission as well as the benefits they offer (e.g. health insurance, vacations, remote work)

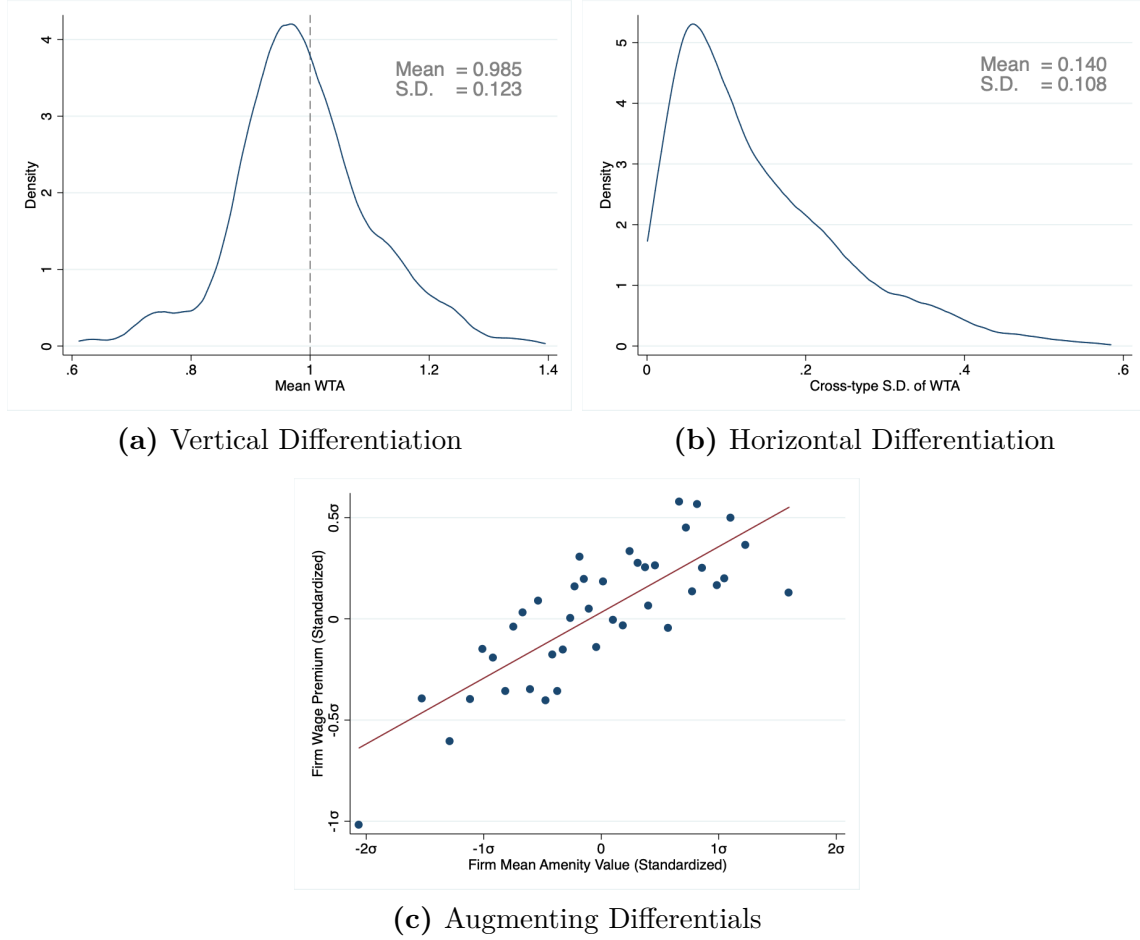
ferentiation implies that firms stand to gain significantly from accurately predicting which candidates are in which preference group, while substantial vertical differentiation implies that high-ranked firms, if acting strategically, can afford to mark down wages significantly. Given the significant scope for firms to set wages in response to preference heterogeneity, assessing firms’ true wage-setting conduct is crucial. Section 6.3 implements our formal test of conduct. Finally, in Panel (c) of Figure 3 we plot estimated firm pay premia—firm fixed effects from a regression of log bids on candidate characteristics interacted with market conditions—against mean firm amenity values. Our results suggest that augmenting differentials prevail: firms that pay well are also firms with better amenities, such that between-firm dispersion in amenities amplifies inequality. On average, a 1-S.D. increase in amenity values is associated with a 0.325 (0.030) S.D. increase in the firm pay premium.

What firm characteristics are associated with higher amenity values? To partially answer this question, we run regressions of (standardized) estimates of A_{qj} on firm covariates z_j in the sample for which those covariates are available. We report these estimates in Panel A of Table B.1.³² Even with the relatively coarse covariates available, some clear patterns are evident. In particular, our results 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 (e.g. startups) that may be more risky bets.

How are worker characteristics related to type membership? To assess this, we compute average posterior type probabilities for candidates with various observable characteristics (our discussion of the EM algorithm in Appendix D covers the construction of these probabilities). Panel B of Table B.1 reports these average posterior type probabilities. We find that women are 13.3pp more likely to belong to the risk-averse group and 9.4pp less likely to be in the risk-loving group than men. Candidates with above-median experience are 16.3pp more likely to be in the risk-loving group and 7.4pp less likely to be in the risk-averse group than those with below-median experience. While there is significant residual variation in preferences conditional on covariates, our estimates suggest that covariates are indeed predictive of preferences.

³²The covariates available to us represent only a small fraction of the full set of 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.

Figure 3: Firm Differentiation




Note: This figure illustrates the scale of vertical and horizontal differentiation of firms implied by our preferred model estimates. Willingness to Accept (WTA) is the fraction of a candidate's ask salary that the model implies a firm must offer to make her indifferent between accepting or rejecting an interview request, on average. Panel (a) plots the distribution of the mean WTA at each firm, averaging over the population probabilities of each type. The vertical grey dashed line indicates a WTA of 1, or Bid=Ask. Panel (b) illustrates the systematic component of horizontal differentiation, plotting the distribution of within-firm, cross-type standard-deviations of WTA. Panel (c) plots standardized firm pay premia (firm fixed effects from a regression of log bids on candidate characteristics and market conditions) against standardized firm amenity values.


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, we depict an illustrative example of the time-series variation in sub-market potential tightness that we use as our excluded instrument t_{ij} in Panel (a) of Figure 4. This figure plots the value of t_{ij} in the sub-market for software engineers with 2-4 years of prior experience over three years of our sample period, and is illustrative of the high-frequency idiosyncratic variation


Figure 4: Vuong Test



(a) Instrument Time-Series Variation



(b) First Stage



(c) Visualizing the Vuong Test

Note: Panel (a) depicts an example of the time-series variation in the excluded instrument t_{ij} for the sub-market of software engineers with 2-4 years of experience over three years of our sample period. Panel (b) is a binned scatterplot depicting the “first stage” relationship between model-implied inclusive values Λ_i and Λ_i^{-j} and t_{ij} , conditional z_j , x_i and two-week period dummies. Panel (c) plots the relationship between generalized residuals and the t_{ij} 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 misspecification.

in potential tightness captured by the instrument. Panel (b) of Figure 4 plots the “first stage” relationship between the model-implied inclusive values (Λ_i and Λ_i^{-j}) and t_{ij} , conditional on firm and candidate covariates and two-week period dummies. 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 Panel (a) of Figure 4: both full- and leave-one-out inclusive values are strongly negatively related to t_{ij} .

Table 2 reports the results of implementing our pairwise testing procedure on the

Table 2: Non-Nested Model Comparison Tests ([Rivers and Vuong, 2002](#))

Model	(1)	(2)	(3)	(4)
	Monopsonistic Comp.		Oligopsony	
	Not Predictive	Type Predictive	Not Predictive	Type Predictive
Perfect Competition	-64.94	-64.36	-55.89	-51.35
Monopsonistic, Not Predictive	–	4.00	4.00	10.57
Monopsonistic, Type Predictive		–	2.88	9.89
Oligopsony, Not Predictive			–	16.81
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.

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 models we tested. Among the remaining alternatives, the two monopsonistic competition models outperform the two oligopsony models, with the not-predictive monopsonistic competition alternative performing best. To assess model fit, in Panel (b) of Figure [A.1](#), we plot the relationship between observed bids and the systematic component of valuations $\gamma_j(x_i)$ in our preferred model. The two are strongly and positively correlated.

We visualize these results in Panel (c) of Figure [4](#), 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 misspecification. In the figure, the generalized residuals for the monopsonistic 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 [G](#), 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 monopsonistic competition model as our preferred model of conduct.

6.4 Comparing Demand Estimates

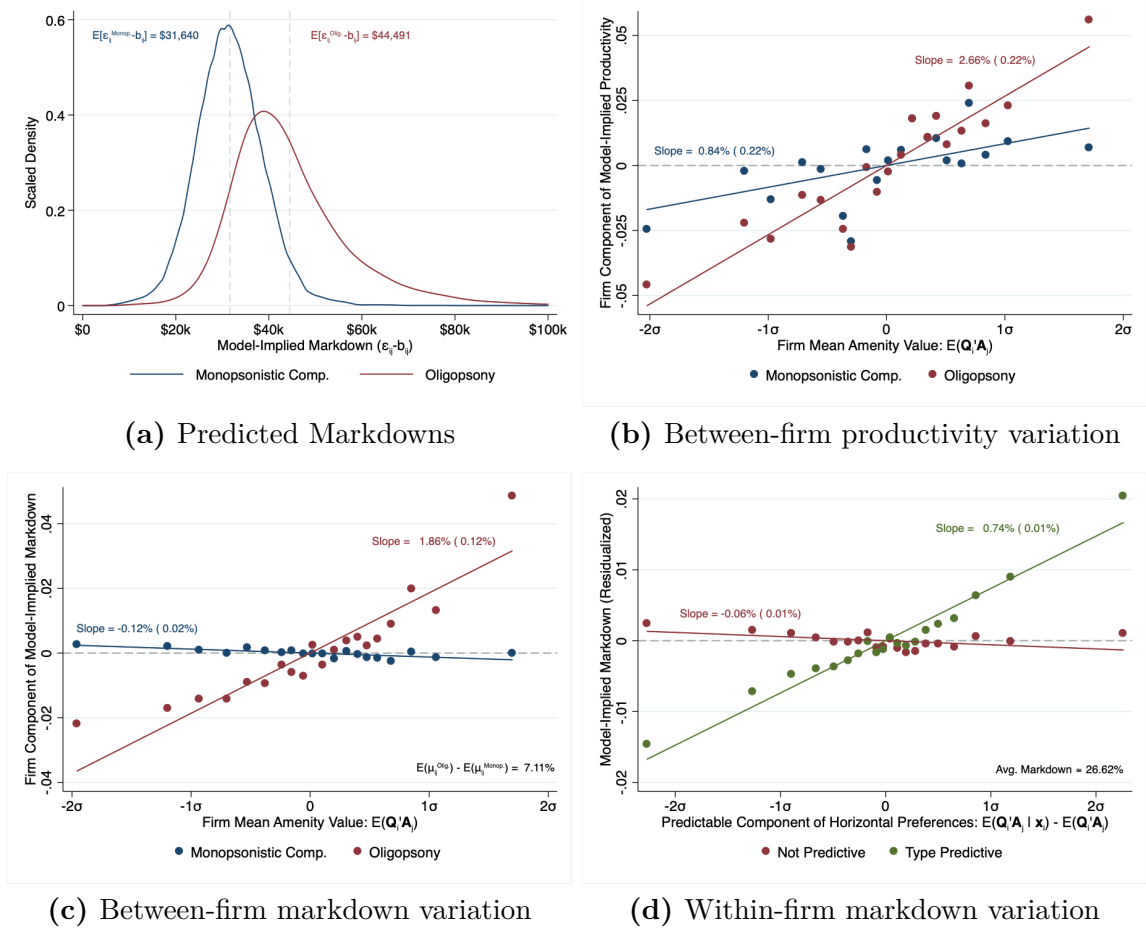
Our preferred model of conduct is the simplest of the four imperfect competition alternatives we specified. Under that model of conduct, there is little to no room for variation in markdowns between firms or differences in markdowns across candidates within firms. How much do the conclusions of the more 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, Panel (a) of Figure 5 plots the distributions of predicted markdowns in dollars under monopsonistic competition and oligopsony. We compute markdowns as the difference between the model-implied firm valuation and the observed bid: $\varepsilon_{ij}^m - b_{ij}$.³³ The two alternatives predict markedly different markdowns distributions. First, under the preferred, monopsonistic model, the average predicted markdown is \$31,640 (or 19.5% of productivity), with a standard deviation of \$6,976. In contrast, the oligopsony model predicts uniformly larger markdowns: the mean model-implied markdown under that assumption is \$44,491 (or 26.6% of productivity, roughly 36% larger than under monopsonistic competition). Second, the distribution of markdowns under oligopsony is significantly more variable, with a standard deviation of \$13,265. Third, under monopsonistic competition, the distribution of markdowns is relatively symmetric: its mean and median are separated by less than \$50, and its skewness is just 0.35. In contrast, the distribution of markdowns under oligopsony is highly skewed: its mean is \$2,405 larger than its median, and its skewness is 1.8. The two sets of markdowns are positively correlated but the correlation is far from one, at 0.25. The large contrasts highlighted by Panel (a) of Figure 5 illustrate the importance of understanding which form of conduct best describes firm behavior: different assumptions about the presence of strategic interactions lead to strikingly different conclusions about the size of markdowns.

Monopsonistic competition and oligopsony yield diverging implications not only for the marginal distribution of markdowns, but also for the joint distribution of

³³In 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: $(\varepsilon_{ij}^{m+} + \varepsilon_{ij}^{m-})/2 - b_{ij}$.

Figure 5: Contrasting labor market implications across models



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

markdowns and productivity across firms. Panel (b) of Figure 5 plots firm components of model-implied productivity against standardized mean firm amenity values. 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 over three times larger under oligopsony than under monopsonistic competition. This leads to large differences in implied productivity dispersion across firms: in the preferred model, firms with the best amenities ($+2\sigma$) are 3.4% more productive than firms with the worst amenities (-2σ). Under oligopsony, that difference is 10.6%.

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, the elasticity of which does not depend upon firms' amenities. Panel (c) of Figure 5 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 7.4pp more than firms with the worst amenities. Under monopsonistic competition, there is essentially no room for different firms to set different markdowns, and so the relationship is flat.

Next, we add another layer of complexity to wage setting: allowing firms to be type-predictive. Panel (d) of Figure 5 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 3.0pp more than the offers it makes to workers who value them the least.

The models also yield differing conclusions about labor demand and the sources of gender gaps (given the substantial gender ask gap (Roussille, 2022)). Under the preferred model, the average elasticity of productivity ε_{ij} with respect to the ask is 0.91, with small and statistically insignificant differences in firms' valuations of men and women (-0.44% (0.29%)). Under the oligopsony alternative, the average elasticity of ε_{ij} with respect to the ask is 0.80, with a larger and significant gender gap in productivity (-0.76% (0.27%)). Under the preferred model, 7.4% of the gender gap in ε_{ij} is accounted for by differences in firms' perceptions of productivity between men and women (conditional on ask), while differences in asks account for 92.6%. Under the oligopsony alternative, that share doubles to 14.4%. Appendix H presents further comparisons of estimated labor demand parameters between models. Among other things, we document that our labor demand estimates feature minimal complementary

between worker and firm covariates, suggesting that additive models of worker and firm effects (as in [Abowd et al. \(1999\)](#)) provide good approximations to log wages.

In a final exercise, we briefly consider implications of our findings for gender gaps in welfare. There exists a large gender gap in the number and average monetary value of bids received by men and women, which maps into a large average gap in welfare as measured by the inclusive values of candidates’ interview offer sets. These gaps are primarily driven by gender differences in the monetary value of bids received, but a nontrivial share of the gap can be attributed to the fact that women receive bids from firms with less attractive amenities than men. We conduct counterfactual simulations to quantify the impact of imperfect competition on welfare and gender gaps. Relative to a “price taking” baseline, we find that firms make significantly fewer offers with lower average wages under the preferred model. 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. Although a significant gender gap exists under price taking, relative gender gaps are larger under imperfect competition and increase further when firms are assumed to be type-predictive. Finally, we find that blinding employers to the gender of candidates generates only a modest reduction in gender gaps under the preferred model of conduct, but the size and direction of the predicted effect of blinding varies meaningfully across assumed conduct scenarios. Appendix [I](#) presents these decompositions and counterfactual exercises in greater detail

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., 2023](#)), and second, whether firms tailor wage offers to workers’ outside options ([Postel-Vinay and Robin, 2002](#); [Jäger et al., 2023](#); [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 19.5% 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 26.6% 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.

Granular online search data, such as job-seekers’ clicks, application behavior, and employment outcomes, are becoming increasingly available to researchers. Recent wage transparency laws also make the salary negotiation process more explicit on online platforms. This paper provides a blueprint for how to leverage these novel data to test for firm wage-setting conduct in the labor market.

References

- Abowd, J., Kramarz, F., and Margolis, D. (1999). “High Wage Workers and High Wage Firms.” *Econometrica*, 67(2), 251–333.
- Allison, P. D., and Christakis, N. A. (1994). “Logit Models for Sets of Ranked Items.” *Sociological Methodology*, 24, 199–228.
- Arnold, D. (2021). “Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes.” Working paper, UC San Diego.
- Avery, C. N., Glickman, M. E., Hoxby, C. M., and Metrick, A. (2013). “A Revealed Preference Ranking of U.S. Colleges and Universities.” *The Quarterly Journal of Economics*, 128(1), 425–467.
- Azar, J., Berry, S. T., and Marinescu, I. (2022). “Estimating Labor Market Power.” NBER Working Paper 30365.
- Azar, J., Marinescu, I., Steinbaum, M., and Taska, B. (2020). “Concentration in U.S. Labor Markets: Evidence from Online Vacancy Data.” *Labour Economics*, 66, 101886.
- Backus, M., Conlon, C., and Sinkinson, M. (2021). “Common Ownership and Competition in the Ready-to-Eat Cereal Industry.” NBER Working Paper 28350.
- Backus, M., and Lewis, G. (2020). “Dynamic Demand Estimation in Auction Markets.” NBER Working Paper 22375.
- Bain, J. S. (1951). “Relation of Profit Rate to Industry Concentration: American Manufacturing, 1936–1940.” *The Quarterly Journal of Economics*, 65(3), 293–324.

- Barseghyan, L., Coughlin, M., Molinari, F., and Teitelbaum, J. C. (2021). “Heterogeneous Choice Sets and Preferences.” *Econometrica*, 89(5), 2015–2048.
- Berger, D., Herkenhoff, K., and Mongey, S. (2022). “Labor Market Power.” *American Economic Review*, 112(4), 1147–1193.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). “Automobile Prices in Market Equilibrium.” *Econometrica*, 63(4), 841–890.
- Berry, S. T. (2021). “Market Structure and Competition, Redux.” Keynote, FTC Micro Conference.
- Berry, S. T., and Haile, P. A. (2014). “Identification in Differentiated Products Markets Using Market Level Data.” *Econometrica*, 82(5), 1749–1797.
- Bhaskar, V., Manning, A., and To, T. (2002). “Oligopsony and Monopsonistic Competition in Labor Markets.” *Journal of Economic Perspectives*, 16(2), 155–174.
- Boal, W. M., and Ransom, M. (1997). “Monopsony in the Labor Market.” *Journal of Economic Literature*, 35(1), 86–112.
- Bresnahan, T. F. (1987). “Competition and Collusion in the American Automobile Industry: The 1955 Price War.” *Journal of Industrial Economics*, 35(4), 457–482.
- Burdett, K., and Mortensen, D. T. (1998). “Wage Differentials, Employer Size, and Unemployment.” *International Economic Review*, 39(2), 257–273.
- Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). “Firms and Labor Market Inequality: Evidence and Some Theory.” *Journal of Labor Economics*, 36(1), 13–70.
- Chamberlain, E. (1933). *The Theory of Monopolistic Competition*. Harvard University Press.
- Chamberlain, G. (1987). “Asymptotic Efficiency in Estimation with Conditional Moment Restrictions.” *Journal of Econometrics*, 34(3), 305–334.
- Donald, S., and Paarsch, H. (1993). “Piecewise Pseudo-Maximum Likelihood Estimation in Empirical Models of Auctions.” *International Economic Review*, 34(1), 121–148.
- Donald, S., and Paarsch, H. (1996). “Identification, Estimation, and Testing in Parametric Empirical Models of Auctions within the Independent Private Values Paradigm.” *Econometric Theory*, 12(3), 517–567.
- Donald, S., and Paarsch, H. (2002). “Superconsistent Estimation and Inference in Structural Econometric Models using Extreme Order Statistics.” *Journal of Econometrics*, 109(2), 305–340.
- Duarte, M., Magnolfi, L., Sølvesten, M., and Sullivan, C. (2023). “Testing Firm Conduct.” *arXiv:2301.06720*.

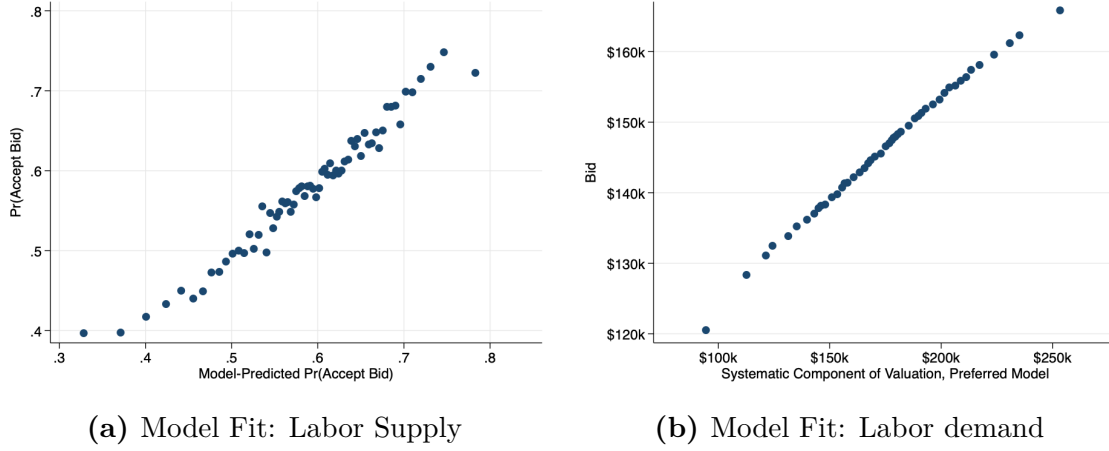
- Faberman, J., and Kudlyak, M. (2016). “What Does Online Job Search Tell Us about the Labor Market?” *Economic Perspectives*, 40(1).
- Flinn, C., and Mullins, J. (2021). “Firms’ Choices of Wage-Setting Protocols.” Working paper.
- Gandhi, A., and Nevo, A. (2021). “Empirical Models of Demand and Supply in Differentiated Products Industries.” In A. Hortascu, K. Ho, and A. Lizzeri (Eds.), *Handbook of Industrial Organization*, vol. 4, Elsevier, 1 edn.
- Gentry, M., and Li, T. (2014). “Identification in Auctions with Selective Entry.” *Econometrica*, 82(1), 315–344.
- Gibbons, R., Katz, L., Lemieux, T., and Parent, D. (2005). “Comparative Advantage, Learning, and Sectoral Wage Determination.” *Journal of Labor Economics*, 23(4), 681–724.
- Gourieroux, C., Monfort, A., Renault, E., and Trognon, A. (1987). “Generalised Residuals.” *Journal of Econometrics*, 34(1-2), 5–32.
- Guerre, E., Perrigne, I., and Vuong, Q. (2000). “Optimal Nonparametric Estimation of First-Price Auctions.” *Econometrica*, 68(3), 525–574.
- Hall, R., and Mueller, A. (2018). “Wage Dispersion and Search Behavior: The Importance of Nonwage Job Values.” *Journal of Political Economy*, 126(4), 1594–1637.
- Horton, J. J., Johari, R., and Kircher, P. (2021). “Cheap Talk Messages for Market Design: Theory and Evidence from a Labor Market with Directed Search.” NBER Working Paper 29445.
- Jarosch, G., Nimczik, J. S., and Sorkin, I. (2023). “Granular Search, Market Structure, and Wages.” Working paper.
- Jäger, S., Roth, C., Roussille, N., and Schoefer, B. (2023). “Worker Beliefs about Outside Options.” Working paper.
- Lagos, L. (2021). “Labor Market Institutions and the Composition of Firm Compensation: Evidence from Brazilian Collective Bargaining.” Working paper.
- Lamadon, T., Mogstad, M., and Setzler, B. (2022). “Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market.” *American Economic Review*, 112(1), 169–212.
- Liu, A., Zhao, Z., Liao, C., Lu, P., and Xia, L. (2019). “Learning Plackett-Luce Mixtures from Partial Preferences.” *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 4328–4335.
- Luce, R. D. (1959). *Individual Choice Behavior A Theoretical Analysis*. Wiley.
- Manning, A. (2005). *Monopsony in Motion*. Princeton University Press.

- Mas, A., and Pallais, A. (2017). “Valuing Alternative Work Arrangements.” *American Economic Review*, 107(12), 3722–3759.
- Pesaran, M. H. (1990). “Non-Nested Hypotheses.” In J. Eatwell, M. Milgate, and P. Newman (Eds.), *Econometrics*, 167–173, London: Palgrave Macmillan UK.
- Plackett, R. L. (1975). “The Analysis of Permutations.” *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 24(2), 193–202.
- Postel-Vinay, F., and Robin, J.-M. (2002). “Equilibrium Wage Dispersion with Worker and Employer Heterogeneity.” *Econometrica*, 70(6), 2295–2350.
- Rivers, D., and Vuong, Q. (2002). “Model Selection Tests for Nonlinear Dynamic Models.” *The Econometrics Journal*, 5(1), 1–39.
- Robinson, J. (1933). *The Economics of Imperfect Competition*. Palgrave.
- Rosen, S. (1986). “The Theory of Equalizing Differences.” *Handbook of Labor Economics*, vol. 1, 641–692, Elsevier.
- Roussille, N. (2022). “The Central Role of the Ask Gap in Gender Pay Inequality.” Working paper, University of California, Berkeley.
- Schmalensee, R. (1989). “Inter-Industry Studies of Structure and Performance.” *Handbook of Industrial Organization*, vol. 2, 951–1009, Elsevier.
- Schubert, G., Stansbury, A., and Taska, B. (2022). “Employer Concentration and Outside Options.” Working paper.
- Simons, G., and Yao, Y.-C. (1999). “Asymptotics when the Number of Parameters Tends to Infinity in the Bradley-Terry Model for Paired Comparisons.” *The Annals of Statistics*, 27(3), 1041–1060.
- Sorkin, I. (2017). “The Role of Firms in Gender Earnings Inequality: Evidence from the United States.” *American Economic Review*, 107(5), 384–387.
- Sorkin, I. (2018). “Ranking Firms Using Revealed Preference.” *The Quarterly Journal of Economics*, 133(3), 1331–1393.
- Soufiani, H., Diao, H., Lai, Z., and Parkes, D. (2013). “Generalized Random Utility Models with Multiple Types.” *Advances in Neural Information Processing Systems*, 26.
- Taber, C., and Vejlín, R. (2020). “Estimation of a Roy/Search/Compensating Differential Model of the Labor Market.” *Econometrica*, 88(3), 1031–1069.
- Vuong, Q. H. (1989). “Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses.” *Econometrica*, 57(2), 307–333.
- Wiswall, M., and Zafar, B. (2018). “Preference for the Workplace, Investment in Human Capital, and Gender.” *The Quarterly Journal of Economics*, 133(1), 457–507.

Online Appendix


A Appendix Figures

Figure A.1: Assessing Model Fit



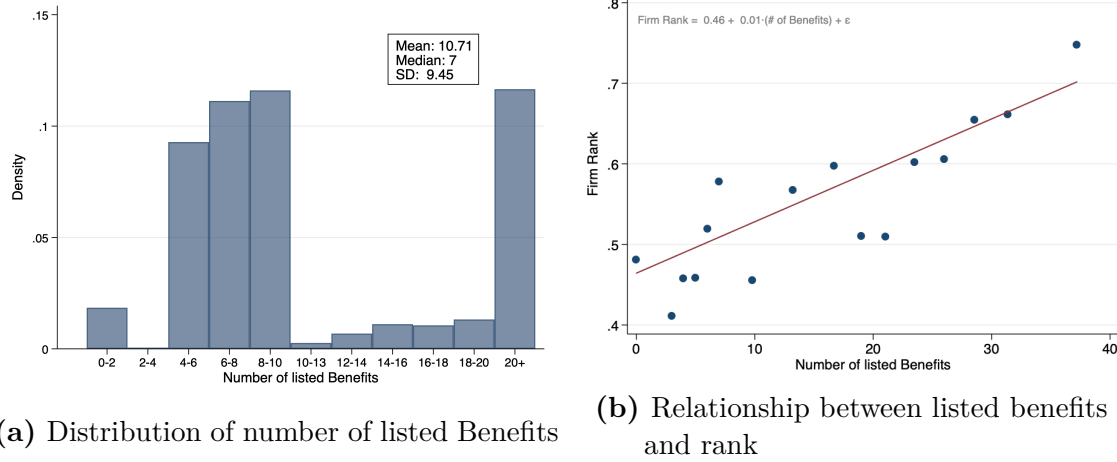
Note: Panel (a) plots the relationship between the empirical acceptance probability of a bid and the model-implied probabilities that the bid will be accepted. Panel (b) plots the relationship between observed bids and the systematic component of valuations $\exp(z_j' \Gamma x_i)$ in the preferred model, controlling for the ask salary. Unconditionally, the slope of the relationship between bids and the observed component of valuations is 0.83.

Figure A.2: Interview Rejection Reasons as a Function of Firm Rankings



Note: This figure plots the probability that a firm was rejected for a non-compensation-related reason as a function of firms' ordinal rankings (where higher ranks are better). For a sub-sample (57%) of rejected bids, candidates opted to provide a justification. They can choose from justifications such as "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.

Figure A.3: Benefits Listed by Firms



Note: This figure displays the distribution of benefits listed by firms in the subset of ranked firms for which information on benefits is available. 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 (S.D. 9.45), while the median is 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 1 centile.

B Appendix Table

Table B.1: Correlates of Amenity Values and Type Probabilities

	(1)	(2)	(3)	(4)
<i>Panel A: Correlates of Firm Amenity Values</i>				
	One-Type Model	Three-Type Model		
	\hat{A}_j	\hat{A}_{1j}	\hat{A}_{2j}	\hat{A}_{3j}
Year Founded	0.00153 (0.00394)	0.000951 (0.00165)	0.00846 (0.00436)	-0.00805* (0.00315)
15-50 Employees	0.161* (0.0742)	-0.237* (0.104)	0.0743 (0.0904)	0.153* (0.0763)
50-500 Employees	0.474*** (0.0743)	-0.320** (0.100)	0.218* (0.0875)	0.406*** (0.0738)
500+ Employees	1.144*** (0.118)	-0.373*** (0.103)	0.481*** (0.103)	0.743*** (0.0819)
Finance	-0.0610 (0.0902)	-0.0678 (0.0433)	0.0121 (0.0606)	-0.0490 (0.0509)
Tech	-0.188** (0.0635)	-0.0342 (0.0456)	-0.0716 (0.0500)	-0.0135 (0.0417)
Health	-0.102 (0.0953)	0.0133 (0.0637)	-0.0395 (0.0682)	0.0305 (0.0892)
adj. R^2	0.180	0.020	0.033	0.126
N	913	913	913	913

Panel B: Posterior Type Probabilities by Candidate Characteristics

	% of Sample	α_{i1}	α_{i2}	α_{i3}
All Candidates	100.0	0.290	0.315	0.395
Male	81.5	0.307	0.323	0.370
Female	18.5	0.213	0.284	0.503
Low Experience	50.0	0.208	0.360	0.432
High Experience	50.0	0.371	0.271	0.358
College or Less	61.8	0.300	0.378	0.323
Grad. Degree	38.2	0.274	0.215	0.511

Note: Panel A reports regressions of standardized estimates of firm amenity values by type, \hat{A}_{qj} , on firm characteristics z_j and a constant. The omitted category is 0-15 employees. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Panel B reports average posterior type probabilities conditional on a number of observable characteristics.

C 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 different 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 the road with location given by ξ , which is private information. These locations are uniformly distributed: $\xi \sim \text{Unif}[0, 1]$. The road has two sides (left and right) for a given location ξ . Workers' homes are on either on the left or right side, 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. The latter are a function of the workers' location along the road as well as whether the worker will have to cross the road to get to work. Worker utilities are given by:

$$u_{-1}^v(\xi) = w_{-1}^v - \beta(\xi + \alpha v); \quad u_{+1}^v(\xi) = w_{+1}^v - \beta(1 - (\xi + \alpha v)).$$

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

$$S_j^v(w_j^v; w_{-j}^v) = \frac{1}{2} + \frac{w_j^v - w_{-j}^v}{2\beta} + \alpha v j.$$

Labor demand is determined by profit maximization:

$$\pi_j(\mathbf{w}) = \frac{1}{2} \sum_{v=-1}^{+1} (\gamma_j - w^v) \times S_j^v(w^v; \hat{w}_{-j}^v),$$

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

$$w_j^v = \frac{1}{2}(\hat{w}_{-j}^v + \gamma_j - \beta) - \alpha\beta v j, \quad S_j^v(w_j^v; \hat{w}_{-j}^v) + S_{-j}^v(w_{-j}^v; \hat{w}_j^v) = 1.$$

We next define conduct as assumptions about the content of \hat{w}_{-j}^v and firms' use of v in wage setting. In the table below we map each conduct assumption with its

corresponding, distinct, equilibrium wage (and hence wage markdown):

Conduct	use v ?	Firm's \hat{w}_{-j}^v	Equilibrium Wage(s) w_j^v
Perfect Comp.	No	—	γ_j
Monopsonistic Not TP	No	\bar{w}	$\frac{3}{4}\gamma_j + \frac{1}{4}\gamma_{-j} - \beta$
Monopsonistic TP	Yes	\bar{w}^v	$\frac{3}{4}\gamma_j + \frac{1}{4}\gamma_{-j} - \beta(1 + \alpha v j)$
Oligopsony Not TP	No	w_{-j}	$\frac{2}{3}\gamma_j + \frac{1}{3}\gamma_{-j} - \beta$
Oligopsony TP	Yes	w_{-j}^v	$\frac{2}{3}\gamma_j + \frac{1}{3}\gamma_{-j} - \beta(1 + \frac{2}{3}\alpha v j)$

Note: TP stands for Type-Predictive

How can we adjudicate between these models? Each model, which we index by m , yields a wage equation of the form:

$$w_j^v = c_{\text{own}}^m \cdot \gamma_j + c_{\text{other}}^m \cdot \gamma_{-j} - c_j^{vm}.$$

where c_{own}^m and c_{other}^m are coefficients governing the pass-through of own-firm and other-firm productivity into wages, and where c_j^{vm} is a model-specific intercept. A traditional approach in labor economics is to estimate the vector of these coefficients $\hat{\mathbf{c}}$. 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_j^v 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 $\hat{\gamma}$, rather than $\hat{\mathbf{c}}$. A particular conduct assumption m , in combination with labor supply parameters estimated in a prior step, determines the coefficients \mathbf{c}^m . Rather than searching for instruments for productivity, find instruments for markdowns that are excluded from productivity. Then, regress $w_j^v + c_j^{vm}$ on c_{own}^m and c_{other}^m to recover $\hat{\gamma}_j^m$; for example, when firms do not use v in wage setting, we have:

$$\begin{bmatrix} \hat{\gamma}_{-1}^m \\ \hat{\gamma}_{+1}^m \end{bmatrix} = \begin{bmatrix} c_{\text{own}}^m & c_{\text{other}}^m \\ c_{\text{other}}^m & c_{\text{own}}^m \end{bmatrix}^{-1} \begin{bmatrix} w_{-1} + c_{-1}^m \\ w_{+1} + c_{+1}^m \end{bmatrix}$$

In order to adjudicate between different forms of conduct, we use the [Vuong \(1989\)](#) and [Rivers and Vuong \(2002\)](#) tests, which compare lack of fit between alternatives.

D 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 : if $F_\xi(x) = \exp(-\exp(-x))$ is the EV_1 CDF, then $\Pr\left(\max_{k \in \mathcal{B}_i^0} \log(\rho_{qk}) + \xi_{ik} < v\right) = F_\xi\left(v - \log\left(\sum_{k \in \mathcal{B}_i^0} \rho_{qk}\right)\right)$. 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:

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

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 first substitutes $u = F_\xi(v - \log(\sum_{k \in \mathcal{B}_i^0} \rho_{qk}))$, then substitutes $z = u^{1/\rho_{Riq}}$, where $\rho_{Riq} = \sum_{j \in \mathcal{B}_i^0} \rho_{qj}$, and $A_{Riq} = \log(\rho_{Riq})$. This expression, and its derivatives, can be quickly and accurately approximated by numerical quadrature.

We estimate $\boldsymbol{\beta}$ and $\boldsymbol{\rho}$ via a first-order EM algorithm (replacing full maximization in the M step with a single gradient ascent update). Applying successive minorizations yields parameter updates that monotonically increase the likelihood (Böhning and Lindsay, 1988; Wu and Lange, 2010). It is useful to define the shorthand: $f_i(\boldsymbol{\rho}_q) = \int_0^1 f_i(\boldsymbol{\rho}_q, z) dz = \mathcal{P}(\mathcal{B}_i^1 \succ \mathcal{B}_i^0 \mid \boldsymbol{\rho}_q)$, $f_{iq}^{(t)} = f_i(\boldsymbol{\rho}_q^{(t)})$, $g_{iq}(\boldsymbol{\beta}) = \alpha_q(x_i \mid \boldsymbol{\beta}) = \exp(x_i' \boldsymbol{\beta}_q) / \sum_{q'=1}^Q \exp(x_i' \boldsymbol{\beta}_{q'})$, $g_{iq}^{(t)} = g_{iq}(\boldsymbol{\beta}^{(t)})$. 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 log integrated likelihood by:

$$\mathcal{E}^{(t)}(\boldsymbol{\beta}, \boldsymbol{\rho}) = \sum_{q=1}^Q \alpha_{iq}^{(t)} \log\left(g_{iq}(\boldsymbol{\beta}) \cdot f_{iq}(\boldsymbol{\rho}_q)\right), \text{ where } \alpha_{iq}^{(t)} = \frac{g_{iq}^{(t)} \cdot f_{iq}^{(t)}}{\sum_{q'=1}^Q g_{iq'}^{(t)} \cdot f_{iq'}^{(t)}}.$$

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

We initialize our algorithm at 50 random starting values, and report the estimate that yields the highest likelihood. We now detail computation of gradient ascent steps.

Define $\mathcal{E}_g^{(t)}(\boldsymbol{\beta}) = \sum_{i=1}^N \sum_{q=1}^Q \alpha_{iq}^{(t)} \cdot \log(g_{iq}(\boldsymbol{\beta}))$, and $\mathcal{E}_{f_q}^{(t)}(\boldsymbol{\rho}_q) = \sum_{i=1}^N \alpha_{iq}^{(t)} \cdot \log(f_i(\boldsymbol{\rho}_q))$, such that: $\mathcal{E}^{(t)}(\boldsymbol{\beta}, \boldsymbol{\rho}) = \mathcal{E}_g^{(t)}(\boldsymbol{\beta}) + \sum_{q=1}^Q \mathcal{E}_{f_q}^{(t)}(\boldsymbol{\rho}_q)$. Since $\mathcal{E}^{(t)}$ is separable in $\boldsymbol{\beta}$ and $\boldsymbol{\rho}_q$, we consider each part separately.

The first component is $\mathcal{E}_g^{(t)}(\boldsymbol{\beta}) = \sum_{i=1}^N \sum_{q=1}^Q \alpha_{iq}^{(t)} \cdot (x'_i \beta_q - \log(\sum_{q'=1}^Q \exp(x'_i \beta_{q'})))$. Let $\boldsymbol{\alpha}_i^{(t)} = [\alpha_{i2}^{(t)} \dots \alpha_{iQ}^{(t)}]'$, $\mathbf{g}_i^{(t)} = [g_{i2}^{(t)} \dots g_{iQ}^{(t)}]'$, and $\mathcal{E}_g^{(t)} = \mathcal{E}_g^{(t)}(\boldsymbol{\beta}^{(t)})$. Then the gradient is given by: $\nabla \mathcal{E}_g^{(t)} = \sum_{i=1}^N (\boldsymbol{\alpha}_i^{(t)} - \mathbf{g}_i^{(t)}) \otimes \mathbf{x}_i$, and the Hessian is given by: $\nabla^2 \mathcal{E}_g^{(t)} = -\sum_{i=1}^N (\text{diag}(\mathbf{g}_i^{(t)}) - \mathbf{g}_i^{(t)} \mathbf{g}_i^{(t)'}) \otimes (\mathbf{x}_i \mathbf{x}_i')$. Our algorithm for $\boldsymbol{\beta}$ follows [Böhning \(1992\)](#). For any $Q-1 \times 1$ vector \mathbf{g} , where the elements of \mathbf{g} are nonnegative the sum of those elements is less than or equal to 1, we have: $\text{diag}(\mathbf{g}) - \mathbf{g}\mathbf{g}' \leq [\mathbf{I}_{Q-1} - Q^{-1} \mathbf{1}_{Q-1} \mathbf{1}_{Q-1}']$, where $\mathbf{A} \leq \mathbf{B}$ is the Loewner ordering: if $\mathbf{A} \leq \mathbf{B}$, then $\mathbf{B} - \mathbf{A}$ is positive semidefinite. Define the matrix $\mathbf{B}_0 = \frac{1}{2} [\mathbf{I}_{Q-1} - Q^{-1} \mathbf{1}_{Q-1} \mathbf{1}_{Q-1}'] \otimes (\mathbf{X}'\mathbf{X})$, where $\mathbf{X} = [\mathbf{x}_1 \dots \mathbf{x}_N]'$. It is straightforward to show that $\nabla^2 \mathcal{E}_g^{(t)} \geq -\mathbf{B}_0$. Now, consider the second-order Taylor approximation to $\mathcal{E}_g^{(t)}(\boldsymbol{\beta})$ at $\boldsymbol{\beta}^{(t)}$:

$$\begin{aligned} \mathcal{E}_g^{(t)}(\boldsymbol{\beta}) &\approx \mathcal{E}_g^{(t)}(\boldsymbol{\beta}^{(t)}) + (\boldsymbol{\beta} - \boldsymbol{\beta}^{(t)})' \nabla \mathcal{E}_g^{(t)} + (\boldsymbol{\beta} - \boldsymbol{\beta}^{(t)})' \nabla^2 \mathcal{E}_g^{(t)} (\boldsymbol{\beta} - \boldsymbol{\beta}^{(t)}) \\ &\geq \mathcal{E}_g^{(t)}(\boldsymbol{\beta}^{(t)}) + (\boldsymbol{\beta} - \boldsymbol{\beta}^{(t)})' \nabla \mathcal{E}_g^{(t)} - (\boldsymbol{\beta} - \boldsymbol{\beta}^{(t)})' \mathbf{B}_0 (\boldsymbol{\beta} - \boldsymbol{\beta}^{(t)}) = \tilde{\mathcal{E}}_g^{(t)}(\boldsymbol{\beta}) \end{aligned}$$

The second line is a quadratic lower bound approximation to $\mathcal{E}_g^{(t)}(\boldsymbol{\beta})$. We set:

$$\boldsymbol{\beta}^{(t+1)} = \arg \max_{\boldsymbol{\beta}} \tilde{\mathcal{E}}_g^{(t)}(\boldsymbol{\beta}) = \boldsymbol{\beta}^{(t)} + \mathbf{B}_0^{-1} \nabla \mathcal{E}_g^{(t)} = \boldsymbol{\beta}^{(t)} + \mathbf{B}_0^{-1} \left(\sum_{i=1}^N (\boldsymbol{\alpha}_i^{(t)} - \mathbf{g}_i^{(t)}) \otimes \mathbf{x}_i \right).$$

The matrix $\mathbf{B}_0^{-1} = 2 [\mathbf{I}_{Q-1} + \mathbf{1}_{Q-1} \mathbf{1}_{Q-1}'] \otimes (\mathbf{X}'\mathbf{X})^{-1}$ only needs to be computed once.

The second component is $\mathcal{E}_{f_q}^{(t)}(\boldsymbol{\rho}_q) = \sum_{i=1}^N \alpha_{iq}^{(t)} \cdot \log(f_i(\boldsymbol{\rho}_q))$. For now, we consider each term of the sum separately, and so we drop i and q subscripts. We have: $f(\boldsymbol{\rho}) = \int_0^1 f(\boldsymbol{\rho}, z) dz = \int_0^1 [\prod_{j \in \mathcal{B}^1} (1 - z^{\rho_j}) \cdot \rho^0 \cdot z^{\rho^0 - 1}] dz$. It is easy to show that this probability is invariant to positive scaling of the vector $\boldsymbol{\rho}$: for any $\alpha > 0$, $f(\alpha \boldsymbol{\rho}) = f(\boldsymbol{\rho})$. We set $\alpha = 1/\rho_R^{(t)}$ and re-write the expression for $f(\boldsymbol{\rho})$ as:

$$\frac{f(\boldsymbol{\rho})}{f(\boldsymbol{\rho}^{(t)})} = \frac{\int_0^1 f(\boldsymbol{\rho}/\rho_R^{(t)}, z) dz}{\int_0^1 f(\boldsymbol{\rho}^{(t)}/\rho_R^{(t)}, z) dz} = \int_0^1 \left(\frac{f(\boldsymbol{\rho}/\rho_R^{(t)}, z)}{f(\boldsymbol{\rho}^{(t)}/\rho_R^{(t)}, z)} \right) \cdot \underbrace{\left(\frac{f(\boldsymbol{\rho}^{(t)}/\rho_R^{(t)}, z)}{\int_0^1 f(\boldsymbol{\rho}^{(t)}/\rho_R^{(t)}, z') dz'} \right)}_{=\pi^{(t)}(z)} dz$$

Jensen's inequality implies: $\log(f(\boldsymbol{\rho})) - \log(f(\boldsymbol{\rho}^{(t)})) \geq \int_0^1 \log \left(\frac{f(\boldsymbol{\rho}/\rho_R^{(t)}, z)}{f(\boldsymbol{\rho}^{(t)}/\rho_R^{(t)}, z)} \right) \cdot \pi^{(t)}(z) dz$. Letting $H_\pi^{(t)} = -\int_0^1 \log(\pi^{(t)}(z)) \pi^{(t)}(z) dz \geq 0$, the above inequality can be rewritten as: $\log(f(\boldsymbol{\rho})) \geq \int_0^1 \log \left(f(\boldsymbol{\rho}/\rho_R^{(t)}, z) \right) \cdot \pi^{(t)}(z) dz + H_\pi^{(t)}$, which is an equality when

$\boldsymbol{\rho} = \boldsymbol{\rho}^{(t)}$, and is strict otherwise. We next analyze:

$$\log \left(f(\boldsymbol{\rho}/\rho_R^{(t)}, z) \right) = \sum_{j \in \mathcal{B}^1} \log \left(1 - z^{\rho_j/\rho_R^{(t)}} \right) + \log \left(\rho_R/\rho_R^{(t)} \right) + (\rho_R/\rho_R^{(t)} - 1) \log(z).$$

Note that: $\log \left(\rho_R/\rho_R^{(t)} \right) \geq \sum_{k \in \mathcal{B}^0} \log \left(\rho_k/\rho_k^{(t)} \right) \cdot \left(\rho_k^{(t)}/\rho_R^{(t)} \right)$, again by Jensen's inequality. Letting $H_\rho^{(t)} = -\sum_{k \in \mathcal{B}^0} \log \left(\rho_k^{(t)}/\rho_R^{(t)} \right) \cdot \left(\rho_k^{(t)}/\rho_R^{(t)} \right) \geq 0$, the above inequality can be rewritten as: $\log(\rho_R) \geq \sum_{k \in \mathcal{B}^0} \log(\rho_k) \cdot \left(\rho_k^{(t)}/\rho_R^{(t)} \right) + H_\rho^{(t)}$, where the inequality is again an equality when $\boldsymbol{\rho} = \boldsymbol{\rho}^{(t)}$, and is strict otherwise. Substituting this expression into the inequality above and lumping constant terms into the single term $H^{(t)}$ gives:

$$\log(f(\boldsymbol{\rho})) - H^{(t)} \geq \sum_{j \in \mathcal{B}^1} \log \left(1 - z^{\rho_j/\rho_R^{(t)}} \right) + \sum_{k \in \mathcal{B}^0} \frac{1}{\rho_R^{(t)}} \left[\log(\rho_k) \cdot \rho_k^{(t)} + \rho_k \cdot \int_0^1 \log(z) \pi^{(t)}(z) dz \right] = \tilde{f}^{(t)}(\boldsymbol{\rho}).$$

The function $\tilde{f}^{(t)}(\boldsymbol{\rho})$ is separable in the parameters $\boldsymbol{\rho}$, and so its Hessian is diagonal. To define the partial derivatives of $\tilde{f}^{(t)}(\boldsymbol{\rho})$, it will be useful to work with the following auxiliary functions: $h(z, x) = \log(z) \cdot \frac{z^x}{1-z^x}$, and $h^2(z, x) = \log^2(z) \cdot \frac{z^x}{(1-z^x)^2}$, and to define: $\tilde{\rho}_j^{(t)} = \rho_j^{(t)}/\rho_R^{(t)}$. We take derivatives with respect to $A_j = \log(\rho_j)$:

$$\begin{aligned} \nabla_j \tilde{f}^{(t)} &= \left. \frac{\partial \tilde{f}^{(t)}}{\partial A_j} \right|_{\boldsymbol{\rho}=\boldsymbol{\rho}^{(t)}} = \mathbf{1}[j \in \mathcal{B}^1] \left(-\tilde{\rho}_j^{(t)} \int_0^1 h(z, \tilde{\rho}_j^{(t)}) \pi^{(t)}(z) dz \right) + \mathbf{1}[j \in \mathcal{B}^0] \left(\tilde{\rho}_j^{(t)} + \tilde{\rho}_j^{(t)} \int_0^1 \log(z) \pi^{(t)}(z) dz \right) \\ \nabla_{jj}^2 \tilde{f}^{(t)} &= \left. \frac{\partial^2 \tilde{f}^{(t)}}{\partial A_j^2} \right|_{\boldsymbol{\rho}=\boldsymbol{\rho}^{(t)}} = \mathbf{1}[j \in \mathcal{B}^1] \left(-\tilde{\rho}_j^{(t)} \int_0^1 h(z, \tilde{\rho}_j^{(t)}) \pi^{(t)}(z) dz - \left(\tilde{\rho}_j^{(t)} \right)^2 \int_0^1 h^2(z, \tilde{\rho}_j^{(t)}) \pi^{(t)}(z) dz \right) + \mathbf{1}[j \in \mathcal{B}^0] \left(\tilde{\rho}_j^{(t)} \int_0^1 \log(z) \pi^{(t)}(z) dz \right) \end{aligned}$$

We construct a lower bound surrogate $\tilde{\mathcal{E}}_{fq}^{(t)}(\boldsymbol{\rho}_q)$ for the function $\mathcal{E}_{fq}^{(t)}(\boldsymbol{\rho}_q)$ by setting: $\tilde{\mathcal{E}}_{fq}^{(t)}(\boldsymbol{\rho}_q) = \sum_{i=1}^N \alpha_{iq}^{(t)} \tilde{f}_i(\boldsymbol{\rho}_q)$, $\nabla_j \tilde{\mathcal{E}}_{fq}^{(t)} = \sum_{i=1}^N \alpha_{iq}^{(t)} \nabla_j \tilde{f}_i^{(t)}$ and $\nabla_{jj}^2 \tilde{\mathcal{E}}_{fq}^{(t)} = \sum_{i=1}^N \alpha_{iq}^{(t)} \nabla_{jj}^2 \tilde{f}_i^{(t)}$, which are again defined with respect to $\mathbf{A}_q = \log(\boldsymbol{\rho}_q)$. Maximizing the second-order Taylor series approximation to $\tilde{\mathcal{E}}_{fq}^{(t)}(\boldsymbol{\rho}_q)$ yields the following Newton-Raphson step: $\mathbf{A}_q^{(t+1)} = \mathbf{A}_q^{(t)} - \left(\nabla^2 \tilde{\mathcal{E}}_{fq}^{(t)} \right)^{-1} \left(\nabla \tilde{\mathcal{E}}_{fq}^{(t)} \right)$. Because $\nabla^2 \tilde{\mathcal{E}}_{fq}^{(t)}$ is diagonal, this step takes a (relatively) simple form. When reintroducing iq subscripts, we have: $\tilde{\rho}_{ijq}^{(t)} = \rho_{jq}^{(t)}/\rho_{Riq}^{(t)}$ and $\pi_{iq}^{(t)}(z) = f_i(\tilde{\rho}_{ijq}^{(t)}, z) / \int_0^1 f_i(\tilde{\rho}_{ijq}^{(t)}, z') dz'$. It will again be helpful to define additional shorthand: $[h_{iq}^0]^{(t)} = -\int_0^1 \log(z) \pi_{iq}^{(t)}(z) dz$, $[h_{ijq}^1]^{(t)} = -\int_0^1 h(z, \tilde{\rho}_{ijq}^{(t)}) \pi_{iq}^{(t)}(z) dz$, and $[h_{ijq}^2]^{(t)} = \int_0^1 h^2(z, \tilde{\rho}_{ijq}^{(t)}) \pi_{iq}^{(t)}(z) dz$. The gradient ascent update for a single A_{qj} is:

$$A_{qj}^{(t+1)} = A_{qj}^{(t)} + \frac{\sum_{i=1}^N \alpha_{iq}^{(t)} \tilde{\rho}_{ijq}^{(t)} \left(\mathbf{1}(j \in \mathcal{B}^1) \cdot [h_{ijq}^1]^{(t)} + \mathbf{1}(j \in \mathcal{B}^0) \cdot \left(1 - [h_{iq}^0]^{(t)} \right) \right)}{\sum_{i=1}^N \alpha_{iq}^{(t)} \tilde{\rho}_{ijq}^{(t)} \left(\mathbf{1}(j \in \mathcal{B}^1) \cdot \left(\tilde{\rho}_{ijq}^{(t)} \cdot [h_{ijq}^2]^{(t)} - [h_{ijq}^1]^{(t)} \right) + \mathbf{1}(j \in \mathcal{B}^0) \cdot [h_{iq}^0]^{(t)} \right)}.$$

Because the scale of $\boldsymbol{\rho}_q$ (level of \mathbf{A}_q) is not identified, we renormalize the parameter vector at each step such that $\sum_{j=1}^J \rho_{qj} = 1$.

E Properties of bidding strategies

Log-concavity of $G_{ij}^m(\cdot)$ implies several properties of bidding functions. A function f is log-concave if: $f(\lambda y + (1 - \lambda)x) \geq f(y)^\lambda f(x)^{1-\lambda} \quad \forall x, y \in \mathbb{R}, \lambda \in [0, 1]$. Log-concavity of f implies that $F = \int_{-\infty}^x f(u)du$ and $1 - F = \bar{F}$ are also log-concave, that f/F is monotone decreasing, and that f/\bar{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).

Under each model, we may generally write $G_{ij}(b) = \int \tilde{G}_{ij}(b, \lambda) dH(\lambda)$, where either $\tilde{G}_{ij}(b, \lambda) = \exp(u(b, a_i)) / (\exp(u(b, a_i)) + \exp(\lambda))$ under oligopsony or $\tilde{G}_{ij}(b, \lambda) = \exp(u(b, a_i) - \lambda)$ 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 $\theta_0 + \theta_1 \cdot \mathbf{1}[b < a_i]$ with θ_0 . We similarly let $G_{ij}^-(b)$ denote the left-hand side function, which replaces $\theta_0 + \theta_1 \cdot \mathbf{1}[b < a_i]$ with $\theta_0 + \theta_1$. Clearly, $G_{ij}(b) = \mathbf{1}[b \geq a_i] \cdot G_{ij}^+(b) + \mathbf{1}[b < a_i] \cdot G_{ij}^-(b)$. 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 condition for the firm to bid at the kink is that the derivative of the left-hand expected profit function is positive at the ask salary and the derivative of the right-hand profit function is negative at the ask salary:

$$g_{ij}^-(a_i)(\varepsilon_{ij} - a_i) - G_{ij}^-(a_i) > 0 \quad \text{and} \quad g_{ij}^+(a_i)(\varepsilon_{ij} - a_i) - G_{ij}^+(a_i) < 0.$$

We assume that (1) $\varepsilon_{ij} > a_i$, (else the firm would never bid at ask) and (2) both

θ_0 and θ_1 are positive. Given these assumptions, we can write this condition as: $\varepsilon_{ij}^-(a_i) \leq \varepsilon_{ij} \leq \varepsilon_{ij}^+(a_i)$. To show that this implies a unique choice of bid (and is therefore both necessary and sufficient for establishing $b_{ij} = a_i$), consider the case where the derivative of the left-hand profit function is negative at a_i . This implies:

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

since by construction $g_{ij}^+(a_i) < g_{ij}^-(a_i)$ and $G_{ij}^+(a_i) = G_{ij}^-(a_i)$. By the same logic:

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

Therefore, if a firm finds it profitable to bid below (above) ask given its left-hand (right-hand) profit function, then it also finds it profitable to bid below (above) ask given its right-hand (left-hand) profit function. In other words, firms never face a situation in which they can increase expected profit relative to bidding at ask by bidding both slightly above or slightly below ask. These conditions guarantee that the firm's optimal choice of bid is unique, even incorporating the kink, and so we may write the firm's optimal bidding function as:

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

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

Next, we consider firms' participation decisions. Our results imply that the firm's indirect expected profit function is a *strictly increasing* function of the ε_{ij} :

$$\pi_{ij}^*(\varepsilon_{ij}) = \begin{cases} \pi_{ij}^{*-}(\varepsilon_{ij}) & \text{if } \varepsilon_{ij}^-(a_i) \geq \varepsilon_{ij} \\ G_{ij}(a_i)(\varepsilon_{ij} - a_i) & \text{if } \varepsilon_{ij}^-(a_i) \leq \varepsilon_{ij} \leq \varepsilon_{ij}^+(a_i) \\ \pi_{ij}^{*+}(\varepsilon_{ij}) & \text{if } \varepsilon_{ij} \geq \varepsilon_{ij}^+(a_i). \end{cases}$$

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. Firms' participation decisions are therefore given by the equivalent conditions:

$$B_{ij} = \mathbf{1} [\pi_{ij}^*(\varepsilon_{ij}) > c_j] \iff B_{ij} = \mathbf{1} [\nu_{ij} > \pi_{ij}^{*-1}(c_j) - \gamma_j(x_i)].$$

F Proof of the consistency of \hat{c}_j^m

Our proof of the consistency of \hat{c}_j^m 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 $n \rightarrow \infty$. 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 $\hat{\pi}_i$ is bounded below by c when $x_i = x$, and that $\Pr(\hat{\pi}_i \in [c, c + \delta] \mid x_i = x) > 0$ for arbitrary $\delta > 0$. Let n_x denote the number of bids made to candidates with characteristics x and let \hat{c}_x^n denote the minimum implied $\hat{\pi}$ among those bids (such that $\hat{c}^n = \min_{x' \in \mathcal{X}} \hat{c}_{x'}^n$). Our sampling assumptions imply $n_x \xrightarrow{\text{a.s.}} \infty$. For an arbitrary $\epsilon > 0$, note that $\Pr(|\hat{\pi}_i - c| > \epsilon \mid x_i = x) = \Pr(\hat{\pi}_i > c + \epsilon \mid x_i = x) = 1 - F_\pi(c + \epsilon \mid x_i = x) < 1$. Let $\bar{F}_{\pi|x}(a) = 1 - F_\pi(a \mid x_i = x)$. We then have that $\left(\bar{F}_{\pi|x}(c + \epsilon)\right)^{n_x} \xrightarrow{\text{a.s.}} 0$, and therefore $\Pr(|\hat{c}_x^n - c| > \epsilon) = \Pr(\hat{c}_x^n > c + \epsilon) = E\left[\left(\bar{F}_{\pi|x}(c + \epsilon)\right)^{n_x}\right]$. Since ϵ is arbitrary, $\hat{c}_x^n \xrightarrow{\text{P}} c$, and since $\hat{c}_x^n \geq \hat{c}^n \geq c$, $\hat{c}^n \xrightarrow{\text{P}} c$. Further, $\sup_{m > n} |\hat{c}^m - c| = |\hat{c}^n - c| \xrightarrow{\text{P}} 0$ since \hat{c}^n is non-increasing in n , and so $\hat{c}^n \xrightarrow{\text{a.s.}} c$. \square

G 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.³⁴

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:

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

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:

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

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

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

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 $Z_s^{m_1, m_2} > c_\alpha$,

³⁴The 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 misspecification error as:

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

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 misspecified. This implies that relatively more misspecified models will imply valuations that are more difficult to explain using observables than those that are closer to the truth. Table G.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 G.1: Non-Nested Model Comparison Tests (Vuong, 1989)

Model	(1)	(2)	(3)	(4)
	Monopsonistic Comp.		Oligopsony	
	Not Predictive	Type Predictive	Not Predictive	Type Predictive
Perfect Competition	-193.86	-192.57	-119.48	-117.93
Monopsonistic, Not Predictive	–	4.16	58.59	58.25
Monopsonistic, Type Predictive		–	54.64	58.77
Oligopsony, Not Predictive			–	3.96
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.

H Further model comparisons

We next consider differences in estimated labor demand parameters $\hat{\Gamma}$ between the preferred model and the (not-predictive) oligopsony alternative. Table H.1 reports estimated elasticities 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 elasticities 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 other observables, we report estimates of both the semi-elasticities of each binary covariate ℓ both holding the ask constant ($\hat{\gamma}_\ell$) and adjusting for differences in the average ask salary. 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 0.91. The remaining semi-elasticities in column 2 are all relatively small and statistically insignificant. Column 3 reports semi-elasticities adjusted to account for average differences in asks between groups. Columns 4 and 5 reproduce this analysis for the oligopsony alternative. The estimated elasticity with respect to the ask, 0.80, is significantly lower than in the preferred model, and the conditional semi-elasticities (Column 4) are much larger in magnitude and statistically significant in all but one case. The unconditional semi-elasticities under oligopsony (Column 5) are very similar to their counterparts under monopsonistic competition. In the preferred model, systematic differences in firms' average valuations between candidates of different groups (men vs women, lower- vs higher-educated) is in essence completely mediated by differences in the average asks of those groups. The oligopsony alternative apportions a nontrivial portion of the gaps in firms' average valuations between groups to autonomous differences that are independent of the ask (e.g. direct/taste-based discrimination).

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. Table H.2 reports the full set of labor demand parameter estimates for the preferred model. We find 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 $\exp(z'_j \hat{\Gamma} x_i)$ 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: 0.911 for bids, 0.920 for ε_{ij} , and 0.967 for $\exp(z'_j \hat{\Gamma} x_i)$. 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 H.1: Determinants of Match Productivity: Elasticities

	(1) $\mathbb{E}[\Delta \text{Ask}]$	(2) Monopsonistic Comp.	(3) Comp.	(4) Oligopsony	(5) Oligopsony
	$\hat{\beta}_\ell$	$\hat{\gamma}_\ell$	$+\hat{\beta}_\ell \cdot \hat{\gamma}_{\text{ask}}$	$\hat{\gamma}_\ell$	$+\hat{\beta}_\ell \cdot \hat{\gamma}_{\text{ask}}$
Ask Salary	—	0.9074 (0.0027)	—	0.7961 (0.0027)	—
Female	-0.0607 (0.0013)	-0.0044 (0.0029)	-0.0595 (0.0029)	-0.0076 (0.0027)	-0.0527 (0.0027)
Unemployed	-0.0568 (0.0030)	0.0022 (0.0063)	-0.0494 (0.0063)	-0.0026 (0.0044)	-0.0430 (0.0044)
Grad School	0.0253 (0.0010)	0.0033 (0.0025)	0.0262 (0.0025)	0.0113 (0.0024)	0.0234 (0.0024)
FAANG	0.0495 (0.0013)	-0.0024 (0.0033)	0.0425 (0.0033)	-0.0099 (0.0044)	0.0370 (0.0044)

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 H.2: Labor Demand Parameter Estimates $\hat{\Gamma}$ ($\log(\varepsilon_{ij}) = z_j' \Gamma x_i + \nu_{ij}$)

		(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Constant	Firm Size			Firm Industry		
Candidate Covariates			16-50	51-500	501+	Finance	Tech	Health
(1)	Constant	1.9374 (0.0183)	-0.6086 (0.0213)	-0.5183 (0.0259)	-0.7550 (0.0389)	-0.1447 (0.028)	-0.1788 (0.0269)	0.0174 (0.0408)
(2)	log(Ask)	0.8464 (0.0017)	0.0525 (0.0017)	0.0466 (0.0023)	0.0669 (0.0034)	0.0121 (0.0025)	0.0153 (0.0023)	-0.0021 (0.0034)
(3)	Female	-0.0057 (0.0024)	0.0036 (0.0026)	-0.0021 (0.0024)	-0.0025 (0.0025)	0.0040 (0.0015)	0.0035 (0.0012)	0.0004 (0.0021)
(4)	Software Eng.	0.0268 (0.0027)	-0.0037 (0.0029)	-0.0127 (0.0027)	-0.0156 (0.0028)	0.0068 (0.0016)	0.0054 (0.0013)	0.0064 (0.0021)
(5)	Experience	0.0001 (0.0006)	0.0008 (0.0006)	0.0016 (0.0006)	0.0015 (0.0006)	-0.0003 (0.0002)	-0.0003 (0.0003)	-0.0001 (0.0004)
(6)	Experience Sq.	0.0000 (0.0001)	0.0000 (0.0001)	0.0000 (0.0001)	0.0000 (0.0001)	0.0000 (0.0001)	0.0000 (0.0001)	0.0000 (0.0001)
(7)	Employed	0.0001 (0.0038)	0.0002 (0.0041)	0.0022 (0.0039)	0.0007 (0.0041)	0.0000 (0.0027)	-0.0033 (0.0022)	0.0019 (0.0035)
(8)	Time Unemp.	0.0012 (0.0009)	-0.0001 (0.001)	0.0001 (0.0009)	-0.0006 (0.001)	0.0000 (0.0006)	-0.0011 (0.0005)	-0.0004 (0.0008)
(9)	Attended Ivy+	-0.0009 (0.0023)	-0.0051 (0.0025)	-0.0020 (0.0024)	0.0003 (0.0025)	-0.0043 (0.0014)	-0.0008 (0.0012)	-0.0028 (0.002)
(10)	CS Degree	0.0069 (0.0021)	-0.0023 (0.0023)	-0.0031 (0.0022)	-0.0033 (0.0022)	-0.0039 (0.0013)	0.0014 (0.0011)	-0.0045 (0.0017)
(11)	Grad School	0.0080 (0.0021)	-0.0023 (0.0023)	-0.0053 (0.0021)	-0.0062 (0.0022)	0.0009 (0.0012)	-0.0001 (0.001)	-0.0011 (0.0016)
(12)	FAANG	0.0026 (0.0028)	-0.0017 (0.0029)	-0.0049 (0.0028)	-0.0046 (0.0029)	-0.0027 (0.0015)	-0.0007 (0.0013)	-0.0008 (0.0022)
(13)	No. Prior Jobs	-0.0008 (0.0005)	-0.0003 (0.0005)	-0.0009 (0.0005)	-0.0001 (0.0005)	0.0006 (0.0003)	0.0001 (0.0002)	0.0008 (0.0004)
(14)	Fulltime	-0.0042 (0.0021)	0.0017 (0.0023)	0.0029 (0.0022)	0.0029 (0.0023)	-0.0011 (0.0014)	-0.0022 (0.0011)	0.0032 (0.0018)
(15)	Sponsorship	-0.0029 (0.0027)	0.0146 (0.0029)	0.0072 (0.0027)	0.0084 (0.0027)	0.0012 (0.0015)	0.0002 (0.0012)	-0.0018 (0.002)
(16)	Remote	0.0008 (0.002)	0.0048 (0.0022)	0.0011 (0.002)	-0.0002 (0.0021)	0.0012 (0.0012)	0.0010 (0.001)	0.0029 (0.0016)
(17)	Java	0.0030 (0.0021)	-0.0007 (0.0023)	0.0036 (0.0021)	0.0046 (0.0022)	-0.0048 (0.0012)	-0.0044 (0.001)	0.0012 (0.0017)
(18)	Python	0.0028 (0.002)	-0.0007 (0.0021)	-0.0029 (0.002)	-0.0035 (0.002)	0.0015 (0.0012)	0.0024 (0.001)	-0.0028 (0.0016)
(19)	SQL	-0.0028 (0.0022)	0.0061 (0.0024)	0.0048 (0.0023)	0.0041 (0.0023)	0.0005 (0.0013)	0.0026 (0.0011)	0.0001 (0.0018)
(20)	C	0.0096 (0.0025)	-0.0136 (0.0028)	-0.0086 (0.0026)	-0.0093 (0.0027)	0.0001 (0.0015)	0.0009 (0.0013)	0.0011 (0.0022)
Std. Dev. of ν_{ij} ($\hat{\sigma}_\nu$)		0.0690	(0.0001)	—	$N = 182,550$	Implied $R^2 = 0.903$		

Note: This table reports maximum likelihood parameter estimates from our preferred labor demand model. The parameters relate combinations of candidate and firm characteristics to the distribution of firms' valuations. Each cell reports the coefficient on the interaction of the variables specified in the corresponding row and column. Row variables are candidate characteristics (x_i), and column variables are firm characteristics (z_j).

I Welfare: decompositions and counterfactual simulations

I.1 A Decomposition of (Expected) Inclusive Values

Given our estimates of amenity values and labor supply elasticities, it is possible to characterize the utility value candidates associate with the portfolios of bids they receive. This allows us to ask whether observable differences in average bids between groups are reflective of underlying differences in welfare. Recall that the utility candidate i of type q associates to firm j 's bid is:

$$V_{iqj} = u_q(b_{ij}, a_i) + A_{qj} + \xi_{ij}.$$

For the purposes of analyzing welfare, we add back a normalized outside option term to the monetary component utility function:

$$u_q(b_{ij}, a_i) = (\theta_{q0} + \theta_{q1} \cdot \mathbf{1}[b_{ij} < a_i]) \cdot \log(b_{ij}/a_i) + \theta_{q0} \cdot (\log(a_i) - \mathbb{E}[\log(a_i)]),$$

where $\mathbb{E}[\log(a_i)]$ is the average log ask across all candidates. We normalize candidates' outside options ($j = 0$) by setting $b_{i0} = a_i$ and $A_{q0} = 0$ (we therefore subtract A_{q0} from each A_{qj}). Let $\mu_{iqj} = \exp(u_q(b_{ij}, a_i))$ and recall that $\rho_{qj} = \exp(A_{qj})$. Then i 's type- q specific inclusive value Λ_{iq} can be written as:

$$\Lambda_{iq} = \mathbb{E} \left[\max_{j: b_{ij} > 0} V_{iqj} \right] = \log \left(\sum_{b_{ij} > 0} \exp(u_q(b_{ij}, a_i) + A_{qj}) \right) = \log \left(\sum_{b_{ij} > 0} \mu_{iqj} \cdot \rho_{qj} \right).$$

Next, define the following quantities:

$$\underbrace{N_i = \sum_{b_{ij} > 0} 1}_{\# \text{ Bids} + 1}, \quad \underbrace{\mu_{iq} = \frac{1}{N_i} \sum_{b_{ij} > 0} \mu_{iqj}}_{\text{Average Monetary Value}}, \quad \underbrace{\rho_{iq} = \frac{1}{N_i} \sum_{b_{ij} > 0} \rho_{qj}}_{\text{Average Amenity Value}}, \quad \underbrace{\gamma_{iq} = \frac{1}{N_i} \sum_{b_{ij} > 0} \frac{\mu_{iqj}}{\mu_{iq}} \cdot \frac{\rho_{qj}}{\rho_{iq}}}_{\text{Normalized Covariance}}.$$

Given these definitions, we may write:

$$\Lambda_{iq} = \log(N_i \cdot \mu_{iq} \cdot \rho_{iq} \cdot \gamma_{iq}) = \log(N_i) + \log(\mu_{iq}) + \log(\rho_{iq}) + \log(\gamma_{iq}).$$

Because types are not observed, we compute the expected inclusive value Λ_i by taking the average of Λ_{iq} 's over the conditional distribution of types given i 's observables. These probabilities are given by $\alpha_{iq} = \alpha_q(x_i \mid \hat{\beta})$, and unconditional type probabilities are denoted by $\bar{\alpha}_q$. We may then write the expected inclusive value as:

$\Lambda_i = \sum_{q=1}^Q \alpha_{iq} \Lambda_{iq}$. We decompose this value as follows:

$$\Lambda_i = \underbrace{\log(N_i)}_{\text{Scale Comp.}} + \underbrace{\sum_{q=1}^Q \bar{\alpha}_q \log(\mu_{iq})}_{\text{Monetary Comp.}} + \underbrace{\sum_{q=1}^Q \bar{\alpha}_q \log(\rho_{iq})}_{\text{Amenity Comp.}} + \underbrace{\sum_{q=1}^Q \bar{\alpha}_q \log(\gamma_{iq})}_{\text{Correlation Comp.}} + \underbrace{\sum_{q=1}^Q (\alpha_{iq} - \bar{\alpha}_q) \Lambda_{iq}}_{\text{Type-Specific Comp.}}.$$

This decomposition splits Λ_i into five components: 1) a *scale* component that increases in the number of bids i receives, 2) a *monetary* component that is a function only of i 's ask and the bid salaries (b_{ij}) i receives, 3) an *amenity* component that is a function only of the relative amenity values associated with the bids i receives, 4) a *correlation* component that captures the (cross-type average of the) direction of association between monetary and amenity values of bids i receives, and 5) a *type-specific* component that captures the difference between the expected valuation of i 's portfolio of bids with and without conditioning on i 's observables (note that the Monetary, Amenity, and Correlation components are all defined relative to the unconditional distribution of types). While γ_{iq} is not a standard covariance, $\text{sign}(\log(\gamma_{iq})) = \text{sign}(\text{Cov}_{iq}(\mu_{iqj}, \rho_{jq}))$ and is well-defined for positive random variables.

I.2 Decomposing observed gender differences in welfare

We decompose mean differences in the components of inclusive values among the set of observed bids using the Oaxaca-Blinder (OB) decomposition (Oaxaca, 1973; Blinder, 1973). The OB decomposition posits that variable Y_{ig} corresponding to individual i in group $g = \{m, f\}$ can be written as $Y_{ig} = X'_{ig} \beta_g + \epsilon_{ig}$, 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 $\bar{Y}_g = \bar{X}'_g \beta_g$. Let $\Delta \bar{Y} = \bar{Y}_m - \bar{Y}_f$, $\Delta \bar{X} = \bar{X}_m - \bar{X}_f$, and $\Delta \beta = \beta_m - \beta_f$. The OB decomposition represents the difference $\Delta \bar{Y}$ as:

$$\Delta \bar{Y} = \bar{X}'_m \beta_m - \bar{X}'_f \beta_f = \underbrace{\Delta \bar{X}' \beta_f}_{\text{endowments}} + \underbrace{\bar{X}'_f \Delta \beta}_{\text{coefficients}} + \underbrace{\Delta \bar{X}' \Delta \beta}_{\text{interactions}}.$$

The classic OB decomposition apportions differences in the mean of a variable between two groups into components due to differences between those groups in: 1) *endowments* (the mean of X by group); 2) *coefficients* or *returns* associated with those covariates (β_g); and 3) the *interactions* between coefficient and endowment differences.³⁵ The OB decompositions we present should be interpreted as purely

³⁵OB decompositions are not unique: an equivalent “reverse” decomposition may be obtained by replacing f with m in the subscripts of the first two terms and flipping the sign of the third term.

descriptive. However, the size of the endowments component relative to the coefficients component can provide suggestive evidence about the sources of gender gaps. Roughly speaking, the larger the coefficients component relative to the endowment component, the stronger the suggestive evidence that that group differences are driven by differences in how those groups are treated conditional on characteristics. Importantly, we exclude the ask salary as an explanatory variable in our decompositions. The endogeneity of the ask salary complicates the interpretation of decompositions that include it as an explanatory variable: if the ask salary is a function of gender, then it may not be appropriate to interpret gender differences in asks as reflecting differing endowments.³⁶

We report decompositions of mean gaps in the number of bids received, log ask salary and the (expected) inclusive value and its five sub-components in Table I.1 (here, women are the reference group, and positive differences correspond to larger values for men). The first row decomposes the gap in the number of bids received by men and women: on average, women receive 0.248 fewer bids than men. The second row decomposes the ask gap. Two-fifths of the ask gap is driven by differences in endowments, while the remaining three-fifths is driven by differences in coefficients, suggesting that women set lower asks than men even when they have identical observables. The third row decomposes the significant gender gap in welfare as measured by the inclusive values associated with of candidates' offer sets. The decomposition apportions roughly 55% of this gap to differences in endowments, and 45% to differences in coefficients. While it is not possible to provide a causal interpretation of this decomposition, the substantial component associated with differences in coefficients is suggestive evidence of either differences in bargaining power or employer discrimination (or both). The remaining rows decomposes each of the five components of inclusive values. The Scale and Monetary components of inclusive values account for nearly the entire gap, although 2.3% of the gender gap in welfare is attributable to the fact that men receive bids from firms with better amenities than women do. 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. Gender differences in endowments account for the majority of the unconditional gaps.

³⁶Because we omit the ask salary from these decompositions, the effect of differences in 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.

Table I.1: Oaxaca-Blinder Decompositions of Gender Gaps

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Mean Difference		Endowments		Coefficients		Interactions	
	$\Delta \bar{Y}$	SE	$\Delta \bar{X}' \beta_m$	SE	$\bar{X}'_f \Delta \beta$	SE	$\Delta \bar{X}' \Delta \beta$	SE
Number of Bids	0.248	(0.040)	0.410	(0.029)	-0.143	(0.038)	-0.019	(0.028)
Log of Ask Salary	0.101	(0.003)	0.039	(0.002)	0.061	(0.003)	0.000	(0.002)
Inclusive Value =	0.443	(0.015)	0.243	(0.011)	0.205	(0.014)	-0.005	(0.010)
+ Scale Comp.	0.042	(0.007)	0.066	(0.005)	-0.027	(0.006)	0.002	(0.004)
+ Monetary Comp.	0.389	(0.011)	0.156	(0.008)	0.230	(0.010)	0.003	(0.007)
+ Amenity Comp.	0.010	(0.003)	0.018	(0.002)	-0.003	(0.003)	-0.005	(0.002)
+ Correlation Comp.	0.001	(0.000)	0.000	(0.000)	0.000	(0.001)	0.001	(0.000)
+ Type-Specific Comp.	0.001	(0.001)	0.003	(0.001)	0.004	(0.001)	-0.006	(0.001)

Note: This table reports Oaxaca-Blinder decompositions of gender gaps in components of utility. Each row corresponds to a particular quantity. Columns 1 and 2 report the mean differences for that quantity (1) and the standard error associated with that difference (2). Columns 3 and 4 report the Endowments component of the OB decomposition (3) and the standard error associated with that component (4). Columns 5 and 6 report the Coefficients component of the OB decomposition (5) and the standard error associated with that component (6). Finally, columns 7 and 8 report the Interactions component of the O decomposition (7) and the standard error associated with that component (8).

I.3 Counterfactual scenarios of interest

To better understand the welfare implications of imperfect competition, we use our supply and demand estimates to simulate bidding outcomes under all four conduct scenarios: $\{\text{monopsonistic competition, oligopsony}\} \times \{\text{not predictive, type-predictive}\}$. To gauge the losses due to imperfect competition, we define a new form of conduct, which we term **price taking**. Under this 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.³⁷ 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.

³⁷Because bids vary conditional on detailed controls, price-taking is automatically ruled out as a mode of conduct that can describe firms' actual bidding behavior on the platform.

I.4 Computing new counterfactual equilibria

In order to compute counterfactuals, we randomly select 500 candidates from the subset of candidates who are software engineers with 6-10 years of experience and 1,000 firms from the subset of firms who bid on such candidates (the 2-1 ratio of firms to candidates approximates the average level of labor market tightness for this submarket). For each firm-candidate pair, we compute the model-implied systematic component of firm valuations using our preferred estimates of labor demand parameters, $\exp(z'_j \hat{\Gamma} x_i)$. 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 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.

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

1. At iteration t , take *iid* draws from a normal distribution with mean zero and standard deviation $\hat{\sigma}_\nu$ to produce a new set of idiosyncratic components of firms' valuations, ν_{ij}^t . Use these draws, plus the systematic components of valuations $z'_j \hat{\Gamma} x_i$, to compute ε_{ij}^t .
2. Given ε_{ij}^t and Λ_{iq}^t , compute b_{ij}^t as firm j 's best response (under the assumed form of conduct m). If there is no number b such that $G_{ij}^m(b)(\varepsilon_{ij} - b) \geq \hat{c}_j$, then set $b_{ij}^t = 0$.
3. Given firms' best responses b_{ij}^t , calculate the realized inclusive values for each candidate, $\Lambda_{iq}^{t*} = \mathbb{E}[\log(\sum_{j: b_{ij}^t > 0} \exp(u_q(b_{ij}^t, a_i) + A_{qj}))]$. Compute the vector of expected inclusive values at the next iteration by taking a step $\alpha^t \in [0, 1]$ towards Λ_{iq}^{t*} :

$$\Lambda_{iq}^{t+1} = \alpha^t \Lambda_{iq}^{t*} + (1 - \alpha^t) \Lambda_{iq}^t.$$

We iterate this procedure until the distribution of inclusive values converges. We then compute mean counterfactual outcomes by averaging over firms' best responses given the equilibrium distribution of inclusive values across 50 draws of ν_{ij} .

I.5 Simulation Results

Table 1.2 reports the results of our simulations. For each scenario, we compute the number of bids received per candidate, the expected inclusive value of the candidate's portfolio of bids, and the five components of that expected inclusive value. We also compute the average monetary value of the bids candidates receive, the difference between those bids and candidates' asks (as a percent of the ask), and the markdown (as a percent of firms' valuations), conditional on having received at least one bid.

The unconditional means of each of these variables across simulation repetitions are reported in Panel A of Table 1.2. We first consider scenarios in which firms are assumed to be not predictive (columns 1-3). Unsurprisingly, average bids are higher (\$161k vs \$133k or \$130k) and markdowns are lower (12.33% vs 19.01% or 21.67%) in the price taking model (column 1) relative to the preferred monopsonistic competition model (column 2) or the oligopsony model (column 3). Additionally, candidates receive markedly more bids (19.33 vs 6.26 or 6.13) under price taking than under monopsonistic competition or oligopsony. These factors combine to make overall expected utility lower under monopsonistic competition or oligopsony than under price taking (with the caveat that absolute utility levels not possible to interpret). Strikingly, the simulations suggest that candidates' welfare losses relative to price taking are 44% larger under oligopsony than under monopsonistic competition. The lion's share of this difference is accounted for by a drop in the average amenity value of bids candidates receive under oligopsony relative to monopsonistic competition. While the story is broadly the same under type-predictive conduct (columns 4-6), there are some notable differences. First, the number of bids candidates receive and overall welfare is higher under type-predictive conduct, although markups are also slightly higher. These changes are more muted under oligopsony than under monopsonistic competition: the average candidate receives nearly one additional bid under type-predictive monopsonistic competition than under not-predictive monopsonistic competition, but just 0.1 additional bids under type-predictive oligopsony relative to not-predictive oligopsony. The average amenity value of candidates' bids drops for each of these conduct assumptions, but this drop is more than made up for by large increases in the type-specific component, suggesting that firms are able to target bids to the candidates who most strongly value their amenities. Interestingly welfare losses relative to price taking under type-predictive conduct are 9.7% lower under monopsonistic competition and 4.8% lower under oligopsony than under not-predictive conduct, suggesting that while increased targeting of bids can yield

additional market power to firms, that effect is more than counterbalanced by the increased value of the amenities candidates receive.

Panel B of Table [I.2](#) reports differences in these statistics by gender. Under all conduct scenarios, women receive fewer bids, lower bids, and higher markdowns than men. Although the absolute level of the difference in number and monetary value of bids is larger under price taking than under monopsonistic competition or oligopsony, the relative difference in these quantities is smaller: under not-predictive (type-predictive) conduct, women receive 6.9% (8.2%) fewer bids under price taking, but 7.7% (9.0%) fewer bids under monopsonistic competition and 13.2% (10.37%) fewer bids under oligopsony. Similarly, the relative difference between the bids men and women receive is roughly 9.6% under price taking, 9.8% under monopsonistic competition, and 10.5% under oligopsony (in both not-predictive and type-predictive scenarios). These gaps lead to substantial differences in welfare between women and men across all scenarios, and are larger under type-predictive conduct than under not-predictive conduct. The upshot of these results is that while firms' exercise of labor market power tends to lower welfare for all workers, it also tends to 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 of Table [I.2](#) reports differences between mean outcomes for men and women across simulation draws in which firms are constrained to no longer observe candidate gender. The results from our simulations suggest that the efficacy of such a policy is relatively limited. Under our preferred model of firm conduct (not predictive, monopsonistic competition), the gender gap in welfare declines by 11.0%. However, such a policy is predicted to increase the gap in welfare by 5.4% under not-predictive oligopsony conduct, and the predicted effect on welfare varies substantially across conduct scenarios. These policy simulations suggest that interventions to remove information will likely be less effective in closing gender gaps in labor market outcomes than interventions that nudge women to adopt bargaining positions closer to those of similar-qualified men (e.g. increase their ask salaries, as in [Roussille 2023](#)). Further, the variability in predicted policy effects across conduct scenarios further underscores the importance of testing assumptions around firm conduct for informing analysis of and policy for labor markets.

Table I.2: Counterfactual Simulations

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Unconditional Means</i>						
	Not Predictive			Type-Predictive		
	PT	MC	OG	PT	MC	OG
Number of Bids/Candidate	19.33	6.26	6.13	19.83	7.23	6.25
Bid Salary	\$161k	\$133k	\$130k	\$161k	\$133k	\$130k
Bid – Ask (as % of Ask)	20.24	-0.48	-2.81	20.27	-0.38	-3.09
Markdown (%)	12.33	19.01	21.67	12.28	18.92	21.79
Inclusive Value =	4.467	3.164	2.593	4.560	3.384	2.776
+ Scale Component	2.981	1.937	1.863	3.002	2.055	1.892
+ Monetary Component	0.706	0.005	-0.092	0.707	0.009	-0.104
+ Amenity Component	0.792	1.302	0.824	0.749	1.073	0.817
+ Correlation Component	0.003	-0.042	-0.049	0.005	-0.032	-0.060
+ Type-Specific Component	-0.016	-0.038	0.048	0.098	0.279	0.232
<i>Panel B: Differences, Women - Men</i>						
	Not Predictive			Type-Predictive		
	PT	MC	OG	PT	MC	OG
Number of Bids/Candidate	-1.330	-0.480	-0.809	-1.618	-0.654	-0.648
Bid Salary	-\$15.4k	-\$13.0k	-\$13.4k	-\$15.4k	-\$12.8k	-\$13.8k
Bid – Ask (as % of Ask)	0.72	0.46	-0.20	0.73	0.57	-0.22
Markdown (%)	0.05	0.25	0.72	0.07	0.15	0.90
Inclusive Value =	-0.457	-0.408	-0.483	-0.461	-0.464	-0.511
+ Scale Component	-0.067	-0.057	-0.159	-0.082	-0.063	-0.122
+ Monetary Component	-0.369	-0.335	-0.334	-0.369	-0.335	-0.346
+ Amenity Component	0.011	0.042	-0.014	-0.004	-0.071	-0.063
+ Correlation Component	-0.002	-0.024	-0.029	-0.002	-0.011	-0.012
+ Type-Specific Component	-0.031	-0.035	0.053	-0.004	0.017	0.031
<i>Panel C: Differences, Women - Men, Gender Blind Firms</i>						
	Not Predictive			Type-Predictive		
	PT	MC	OG	PT	MC	OG
Number of Bids/Candidate	-1.220	-0.353	-0.626	-1.341	-0.439	-0.684
Bid Salary	-\$14.7k	-\$12.7k	-\$12.9k	-\$14.7k	-\$12.6k	-\$12.7k
Bid – Ask (as % of Ask)	1.24	0.69	0.20	1.23	0.76	0.36
Markdown (as % of MRPL)	0.05	0.43	0.70	0.05	0.35	0.70
Inclusive Value =	-0.435	-0.363	-0.509	-0.440	-0.427	-0.463
+ Scale Component	-0.061	-0.038	-0.146	-0.066	-0.035	-0.128
+ Monetary Component	-0.352	-0.327	-0.324	-0.353	-0.327	-0.328
+ Amenity Component	0.011	0.048	-0.049	-0.001	-0.050	-0.026
+ Correlation Component	-0.002	-0.022	-0.020	-0.002	-0.012	-0.015
+ Type-Specific Component	-0.030	-0.024	0.030	-0.017	-0.003	0.034

Note: This table reports results of counterfactual simulations under various conduct assumptions. Each column corresponds to a combination of conduct assumptions (PT = price-taking, MC = monopsonistic competition, and OG = oligopsony). Each cell reports the average of a statistic over 50 simulation draws. Panel A reports the unconditional means, Panel B reports differences in means between women and men, and Panel C reports differences in means between women and men for simulations in which firms are constrained to be gender blind.

References

- Abowd, J., Kramarz, F., and Margolis, D. (1999). “High Wage Workers and High Wage Firms.” *Econometrica*, 67(2), 251–333.
- Bagnoli, M., and Bergstrom, T. (2005). “Log-Concave Probability and Its Applications.” *Economic Theory*, 26(2), 445–469.
- Bhaskar, V., Manning, A., and To, T. (2002). “Oligopsony and Monopsonistic Competition in Labor Markets.” *Journal of Economic Perspectives*, 16(2), 155–174.
- Blinder, A. (1973). “Wage Discrimination: Reduced Form and Structural Estimates.” *Journal of Human Resources*, 8(4), 436–455.
- Böhning, D. (1992). “Multinomial logistic regression algorithm.” *Annals of the Institute of Statistical Mathematics*, 44(1), 197–200.
- Böhning, D., and Lindsay, B. G. (1988). “Monotonicity of quadratic-approximation algorithms.” *Annals of the Institute of Statistical Mathematics*, 40(4), 641–663.
- Card, D., Heining, J., and Kline, P. (2013). “Workplace Heterogeneity and the Rise of West German Wage Inequality.” *The Quarterly Journal of Economics*, 128(3), 967–1015.
- Donald, S., and Paarsch, H. (2002). “Superconsistent Estimation and Inference in Structural Econometric Models using Extreme Order Statistics.” *Journal of Econometrics*, 109(2), 305–340.
- Oaxaca, R. (1973). “Male-Female Wage Differentials in Urban Labor Markets.” *International Economic Review*, 14(3), 693–709.
- Rivers, D., and Vuong, Q. (2002). “Model Selection Tests for Nonlinear Dynamic Models.” *The Econometrics Journal*, 5(1), 1–39.
- Robinson, J. (1933). *The Economics of Imperfect Competition*. Palgrave.
- Vuong, Q. H. (1989). “Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses.” *Econometrica*, 57(2), 307–333.
- Wu, T. T., and Lange, K. (2010). “The MM alternative to EM.” *Statistical Science*, 25(4).