Estimating Labor Market Power*

JOSÉ AZAR, STEVEN BERRY, AND IOANA MARINESCU

June 13, 2019

PRELIMINARY AND INCOMPLETE – DO NOT CIRCULATE

Abstract

How much power do employers have to suppress wages below marginal productivity? It depends on the labor supply elasticity. Leveraging data on job applications from the large job board CareerBuilder.com, we estimate the wage impact on workers' choice among differentiated jobs. We use nested logit models where the wage is instrumented. In the largest markets, the market-level labor supply elasticity is about 2.3. The firm-level labor supply elasticity is about 5.7; this implies that workers produce about 18% more than their wage level, consistent with employers having market power even in the most competitive labor markets.

*Azar: IESE Business School, Avinguda Pearson, 21, 08034 Barcelona, Spain; Email, jazar@iese.edu. Berry: Department of Economics, Yale University, 37 Hillhouse Ave., New Haven, CT 06520 and National Bureau of Economics; Email, steven.berry@yale.edu. Marinescu: University of Pennsylvania School of Social Policy & Practice, and NBER, 3701 Locust Walk, Philadelphia PA, 19104-6214; Email, ioma@upenn.edu.

1 Introduction

There is a growing literature on the possibility of widespread monopsony or, more plausibly, oligopsony power in US labor markets. This runs against a long-standing tradition in labor economics of assuming perfect competition. One strand of literature emphasizes anti-competitive labor contracts, while another strand focuses on describing correlations between labor market concentration and wages. There is also a small ongoing literature attempting to directly estimate firm-level labor supply elasticities. Perfect competition would imply that these elasticities are infinite, but the existing literature often finds extremely low elasticities—for example, in the range of 0.1, which would imply an extremely high degree of firm labor market power.

There is also a literature suggesting the possible importance of differentiation in the desirability of jobs as one source of labor market power. For example, Card, Cardoso, Heining and Kline (2018) point out that a simple "product differentiation" model of jobs can reconcile many interesting facts about labor markets. In their conclusion, they suggest that labor economics should perhaps move in the direction of Industrial Organization (IO), studying supply and demand in specific labor markets, perhaps with an emphasis on differentiated jobs.

In this paper, we take a first step in the empirical implementation of that suggestion. We use a large online job posting and applications dataset from CareerBuilder.com to estimate a model of the demand for differentiated jobs. Jobs vary in geographic location within commuting zones, as well as by job characteristics that are known to the market participants. Some of these characteristics are observed by us, but some are not. These unobserved (to us) job characteristics could, for example, involve employment benefits and/or the employment-related reputation of the firm. The presence of unobserved job characteristics creates a classic endogeneity problem, similar to the endogeneity of price in the estimation of differentiated product models of supply and demand. In fact, we make direct use of standard empirical differentiated product models in our application to differentiated employment opportunities.

In particular, we model job application choice using a nested logit framework. We define a job market as a commuting zone by 6-digit SOC and allow workers to apply both with and across markets. We focus our estimation on the largest 115 markets in terms of the number of vacancies posted. We discuss alternative possible instruments for wages and obtain IV estimates of vacancy demand wage elasticities.

Drawing on a set of results from the literature, we translate our vacancy demand result into firm and market-level labor supply wage elasticities. Our IV estimates imply firm-level labor supply elasticities of about 5, which are much larger than some published estimates, but which are still consistent with substantial firm level labor market power. These elasticities allow us to compute the wage markdown (or "exploitation rate") $(\text{MRP}_L - w)/w$. We find that the markdown is about 18%. In addition, we can determine if a labor market is a valid antitrust market using a critical elasticity test, shedding light on plausible definitions of a labor market.

2 Data and estimation

2.1 Data

The data includes information on job seekers, job vacancies and job seekers' applications to these vacancies on the online job board CareerBuilder.com, the largest US employment website[1]. We drop job seekers who do not live in the US, who live in Alaska or Puerto Rico, and job seekers whose location is unknown. There are 397,902 unemployed job seekers with accounts active between April and June 2012 on CareerBuilder.com. The dataset contains 694,735 vacancies. This dataset was used by Marinescu and Rathelot (2018), who show that it is broadly representative of vacancies and unemployed job seekers in the US.

The dataset contains information on posted wages for 20% of job vacancies. When a range is provided, we take the middle of the range, as in Marinescu and Wolthoff (2016). It also contains the zip code location of job seekers and job vacancies.

In order to be able to reliably estimate the value of each vacancy to job applicants, we restrict the sample to job vacancies with at least five applications. We define a labor market as a occupation at the 6-digit SOC level (for example, "Registered Nurses") and a commuting zone. We then restrict the data to labor markets with at least 100 posted vacancies each with at least five applications. We end up with 115 markets.

[1]Monster.com is the other leading job board and is comparable in size. Which of CareerBuilder or Monster is larger depends on the exact size metric used.

2.2 Econometric model: basic assumptions

We define a labor market as a commuting zone and occupation. Markets are indexed by m. We observe a labor market over several weeks t. Market m has N_m active users over our sample period. On any given week t, there is a set J_{mt} of active vacancies in market m (that is, in a given commuting zone and occupation combination). In any given week t, a worker decides to apply or not apply to market m. Choosing not to apply gives workers the value of the outside option, i.e. applying to another SOC-6 by CZ market, or not applying at all. Conditional on applying to market m, the worker chooses which vacancy to apply to.

The utility of individual i from applying to job vacancy $j \in J_{mt}$ in week t is:

$$u_{ijmt} = \delta_j + \gamma_m z_{ijm} + \theta_m \tilde{z}_{im} + \nu_{imt}(\lambda_m) + \lambda_m \epsilon_{ijmt}, \tag{1}$$

where z_{ijm} are market-specific regressors that potentially vary across users and jobs (specifically, the log of the distance between the user and a job vacancy in the market), $\tilde{z}_{im}$ are user-specific regressors that do not vary across jobs in a market (specifically, an indicator for whether the user is from the same CZ as the job or from a different CZ) and

$$\delta_j = \beta x_j - \alpha \log w_j + \xi_j \tag{2}$$

is McFadden's "alternative-specific constant" (or "mean utility") which captures the average level of utility of a job vacancy. w_j is the wage posted by the job vacancy. The observed characteristics of the job, x_j, include log employment of the firm posting the vacancy, squared log employment, and fixed effects for CZ×SOC. In some specifications, we also include job title fixed effects. Following Berry (1994), we allow for an unobserved component of the mean utility of a job, ξ_j. This captures unobserved job attributes that are not present in the data. These could include better or worse working conditions or benefits, but in the present context of job applications might also include perceived stringency of the job qualifications. It is natural to assume that the unobservable job attributes are correlated with the wage.

The term ϵ_{ijmt} is a random "match value" of worker to a particular job. The random utility of the outside option is $u_{i0mt} = \epsilon_{i0mt}$. The parameter λ_m is a market-specific nesting parameter, and $\nu_{imt}(\lambda_m)$ is

a "nested logit" random taste that is present in all job vacancies and differentiates the job vacancies in the market from the outside good.

We assume that the error term ϵ_{ijmt} is distributed extreme value, and that the error term $\nu_{imt}(\lambda_m)$ has the distribution that makes the distribution of the combined error terms

$$\nu_{imt}(\lambda_m) + \lambda_m \epsilon_{ijmt}$$

generalized extreme value (Cardell, 1998). This error structure generates a nested logit probability of application, where the two nests are: (i) an inside nest with all the jobs in market m in week t, which we call nest $g = 1$ and (ii) an outside nest which indicates that the user either didn't apply to any job that week, or applied to a job in another market, which we call nest $g = 0$ and has one option which we denote $j = 0$, and refer to as the outside option.

2.3 Bottom-level logit: choosing between jobs within a market

The "bottom-level" probability that user i applies to job j conditional to applying to a job in market m in week t is

$$s_{ijmt|g} = \frac{\exp\left[(\delta_j + \gamma_m z_{ijm})/\lambda_m\right]}{\sum_{k \in J_{mt}} \exp\left[(\delta_k + \gamma_m z_{ikm})/\lambda_m\right]}. \tag{3}$$

For each market, we estimate the bottom-level multinomial logit using the asclogit command in Stata. This model regresses the job seeker's choice (1 if applying to a particular vacancy – inlcuding the outisde option –, 0 otherwise) on her distance to the vacancy (z_{ijm}). The coefficient on z_{ijm} estimates γ_m/λ_m. We recover δ_j/λ_m. We also obtain the predicted probability of choosing each vacancy $s_{ijmt|g}$.

The inclusive parameter capturing user[2] i's utility of applying inside the market m is defined as:

$$I_{imt} = \log \sum_{k \in J_{mt}} \exp\left[(\delta_k + \gamma_m z_{ikm})/\lambda_m\right]$$

We use the linear prediction after the multinomial logit model to calculate the inclusive value.

[2]In the rare cases where a user sends more than one application to a market in a week, each additional application is attributed to a newly created additional user.

2.4 Top-level logit: choosing whether to apply to a market

The "top-level" probability that user i applies in nest $g = 1$ to a job in market m in week t is

$$s_{igmt} = \frac{\exp\left(\theta_m \tilde{z}_{im} + \lambda_m I_{imt}\right)}{1 + \exp\left(\theta_m \tilde{z}_{im} + \lambda_m I_{imt}\right)}, \tag{4}$$

We estimate this equation for each market; this yields the coefficient on the inclusive value: the nesting parameter λ_m. We recover δ_j by multiplying the estimates of δ_j/λ_m from the bottom-level multinomial logit (eq. 3) by λ_m. We also recover the predicted top-level probability s_{igmt}.

The overall probability that user i applies to job j in market m in week t is:

$$s_{ijmt} = \frac{\exp\left[\left(\delta_j + \gamma_m z_{ijm}\right)/\lambda_m\right]}{\sum_{k \in J_{mt}} \exp\left[\left(\delta_k + \gamma_m z_{ikm}\right)/\lambda_m\right]} \times \frac{\exp\left(\theta_m \tilde{z}_{im} + \lambda_m I_{imt}\right)}{1 + \exp\left(\theta_m \tilde{z}_{im} + \lambda_m I_{imt}\right)}. \tag{5}$$

The expected share of applications to job vacancy j in market m in week t is simply the average of the probabilities across users:

$$s_{jmt} = \frac{1}{N_m} \sum_{i=1}^{N_m} s_{ijmt}. \tag{6}$$

The expected share of applications to any job vacancy inside the market in market m in week t is the average of the probabilities of applying inside the market across users:

$$s_{gmt} = \frac{1}{N_m} \sum_{i=1}^{N_m} s_{igmt}. \tag{7}$$

2.5 The impact of the wage on the utility of a job

We then run a regression of the job values δ_j on the log wage and the number of employees of the firm and the square of the number of employees of the firm (equation 2). We do this using either OLS or IV, and we control for job title fixed effects in some regressions.

Since the wage is endogenous, we need to instrument for it. We consider three types of instruments: "Hausman", "BLP", and "unemployment-based".

2.5.1 Instruments

The first set of instruments that we consider are a variant of the "Hausman instruments" that are often used in the industrial organization literature. The idea is to instrument for the price of a product by using its average price in other markets. For example, the price of Cheerios in Atlanta in a given week would be instrumented by the average price of Cheerios in that week across all US cities except Atlanta.

We adapt this idea to instrument a job vacancy's wage by using the average wage for job vacancies posted by the same firm in a given week in all other (CZ $\times$ SOC) markets. We also construct a variant of the instrument excluding from the calculation of the average any markets that are in the same CZ (but other SOCs) or the same SOC (but other CZs). This instrument is valid to the extent that the firm sets somewhat uniform wages for all employees rather than optimizing wages for each local labor market.

The second set of instruments that we consider are a variant of the "BLP instruments", which are also from the industrial organization literature. The idea is to instrument for the price of a product by using characteristics of other products in the market. If the model is correct, these are exogenous because they do not enter the formula for the utility of user i for product j. However, if product characteristics can respond relatively quickly to changes in demand, the exclusion restriction would be violated. This is an important drawback of these instruments in the context of labor market vacancies, as the vacancies change constantly. For this reason, we think that Hausman instruments may be preferable in this context, but provide the "BLP instrument" estimates for comparison.

Usually, the literature uses sums of other products' characteristics (plus, sometimes, their squares and interaction terms). We again consider two variants of instruments based on the job characteristics: first a simple version instrumenting only with number of vacancies in the market, and then a version instrumenting with number of vacancies in the market, the sum across other vacancies of log employment of the posting firm, and the sum across other vacancies of the square of log employment of the posting firm. In the latter case, since we have more than one instrument, we estimate the equation using Chamberlain (1987) optimal instruments (2-step GMM).

Finally, we use a third set of instruments based on the unemployment rates faced by the firm that is posting a given vacancy in its other markets. The idea is to go one step further than the Hausman instrument, by considering not the average wage of the firm in other markets, but the average of the

underlying determinants of the wage of the firm in other markets. We calculate CZ-level unemployment rates by aggregating Bureau of Labor Statistics data on county-level unemployment rates. We then take averages of the unemployment rates across vacancies posted by the same firm in other markets. As with the Hausman instruments we calculate two versions: one using the average unemployment rate for vacancies by the same firm in all other CZ $\times$ SOC markets, and one excluding markets that are in the same CZ or the same SOC from the average unemployment rate calculation.

2.6 Elasticity Formulas

We use the estimated nesting parameter λ_m, wage coefficient α, and the predicted probabilities of application such as S_{ijmt} to calculate the wage elasticity at the vacancy, market and firm level.

Vacancy-Level Elasticities The slope of the probability that user i applies to job vacancy j in week t with respect to the utility of job vacancy j is

$$\frac{\partial s_{ijmt}}{\partial \delta_j} = s_{ijmt}\left(\frac{1}{\lambda_m} - \frac{1-\lambda_m}{\lambda_m} s_{igmt} - s_{ijmt}\right). \tag{8}$$

The slope of the probability that user i applies to job vacancy j in week t with respect to the utility of job vacancy $k \neq j$ is

$$\frac{\partial s_{ijmt}}{\partial \delta_k} = -s_{ijmt}\left(\frac{1-\lambda_m}{\lambda_m} s_{igmt} + s_{ikmt}\right). \tag{9}$$

The slopes of the expected shares are the averages of the slopes across users:

$$\frac{\partial s_{jmt}}{\partial \delta_k} = \frac{1}{N_m}\sum_{i=1}^{N_m} \frac{\partial s_{ijmt}}{\partial \delta_k}. \tag{10}$$

The general formula for the elasticity of the share of job vacancy j in week t with respect to vacancy k's (where $k = j$ or $k \neq j$) wage is:

$$\frac{\partial \log s_{jmt}}{\partial \log w_k} = \frac{\alpha}{s_{jmt}} \cdot \frac{\partial s_{jmt}}{\partial \delta_k}. \tag{11}$$

The elasticity of the share of job vacancy j in week t with respect to vacancy j's own wage is

$$\frac{\partial \log s_{jmt}}{\partial \log w_k} = \frac{\alpha}{s_{jmt}} \cdot \frac{1}{N_m}\sum_{i=1}^{N_m} s_{ijmt}\left(\frac{1}{\lambda_m} - \frac{1-\lambda_m}{\lambda_m} s_{igmt} - s_{ijmt}\right) \tag{12}$$

The elasticity of the share of job vacancy j in week t with respect to vacancy $k, k \neq j$'s wage is

$$\frac{\partial \log s_{jmt}}{\partial \log w_k} = \frac{\alpha}{s_{jmt}} \cdot \frac{1}{N_m} \sum_{i=1}^{N_m} -s_{ijmt} \left(\frac{1-\lambda_m}{\lambda_m} s_{igmt} + s_{ikmt} \right) \tag{13}$$

Market-Level Elasticities The slope of the probability that user i applies to nest $g = 1$ in week t with respect to the inclusive value is

$$\frac{\partial s_{igmt}}{\partial I_{imt}} = \lambda_m s_{igmt}(1 - s_{igmt}). \tag{14}$$

The slope of the the inclusive value with respect to the log wage of job vacancy k is

$$\frac{\partial I_{imt}}{\partial \log w_k} = \frac{\alpha}{\lambda_m} s_{ikmt|g}. \tag{15}$$

Combining the last two equations, the slope of the probability that user i applies to nest $g = 1$ with respect to the log wage of job vacancy k is

$$\frac{\partial s_{igmt}}{\partial \log w_k} = \alpha s_{ikmt}(1 - s_{igmt}). \tag{16}$$

The total slope of user i's probability of applying to the inside nest with respect to the log wage of all vacancies inside the market is

$$\sum_{k \in J_{mt}} \frac{\partial s_{igmt}}{\partial \log w_k} = \alpha s_{igmt}(1 - s_{igmt}). \tag{17}$$

The total slope of the expected share of the inside nest with respect to the log wage of all vacancies inside the market is

$$\sum_{k \in J_{mt}} \frac{\partial s_{gmt}}{\partial \log w_k} = \alpha \frac{1}{N_m} \sum_{i=1}^{N_m} s_{igmt}(1 - s_{igmt}). \tag{18}$$

The total elasticity of the expected share with respect to the wages of all vacancies in the market is

$$\frac{1}{s_{gmt}} \sum_{k \in J_{mt}} \frac{\partial s_{gmt}}{\partial \log w_k} = \frac{\alpha}{s_{gmt}} \frac{1}{N_m} \sum_{i=1}^{N_m} s_{igmt}(1 - s_{igmt}). \tag{19}$$

Firm-level Elasticities There are F firms. Firm f posts a subset $\mathcal{F}_f$ of the job vacancies in market m in week t (which might also be just one vacancy, or no vacancies). Note that firm f could also post jobs in

other CZ $\times$ SOC markets, but we do not take this into account in our analysis because our model treats any application outside the market as part of the outside option. The probability that user i applies to a job posted by firm f in market m in week t is

$$s_{ifmt} = \sum_{j \in \mathcal{F}_f} s_{ijmt}. \tag{20}$$

The slope of the probability that user i applies to a job vacancy posted by firm f with respect to the log wage of job k posted by firm f is a sum of own-wage elasticities and cross-wage elasticities:

$$\frac{\partial s_{ifmt}}{\partial \log w_k} = \sum_{j \in \mathcal{F}_f} \frac{\partial s_{ijmt}}{\partial \log w_k}. \tag{21}$$

The slope of the expected share of firm f with respect to the log wage of job k posted by firm f is a sum of own-wage elasticities and cross-wage elasticities:

$$\frac{\partial s_{fmt}}{\partial \log w_k} = \sum_{j \in \mathcal{F}_f} \frac{\partial s_{jmt}}{\partial \log w_k}. \tag{22}$$

The total slope of the expected share of firm f with respect to a simultaneous change in the log wage for all of its own vacancies is

$$\sum_{k \in \mathcal{F}_f} \frac{\partial s_{fmt}}{\partial \log w_k} = \sum_{k \in \mathcal{F}_f} \sum_{j \in \mathcal{F}_f} \frac{\partial s_{jmt}}{\partial \log w_k}. \tag{23}$$

The total elasticity of the expected share of firm f with respect to a simultaneous increase in the wage for all of its own vacancies is obtained by dividing the last equation by the expected share of firm f in market m in week t:

$$\frac{1}{s_{fmt}} \sum_{k \in \mathcal{F}_f} \frac{\partial s_{fmt}}{\partial \log w_k} = \frac{1}{s_{fmt}} \sum_{k \in \mathcal{F}_f} \sum_{j \in \mathcal{F}_f} \frac{\partial s_{jmt}}{\partial \log w_k} \tag{24}$$

$\frac{\partial s_{jmt}}{\partial \log w_k}$ can be found in equation 12 multiplied by s_{jmt} for the own elasticity, and equation 13 multiplied by s_{jmt} for the cross elasticity.

3 Results

In Table 1, we first report the impact of wages on the job-specific utility δ_j using OLS. Surprisingly, the impact of wages on job utility is negative (column 1). This counterintutive result is explained by Marinescu and Wolthoff (2016) as reflecting important job heterogeneity within a SOC-6 occupation. For example, in the accountant and auditors occupation, "senior accountant" job titles pay more but receive fewer applications than "junior accountant" job titles. Therefore, in column 2, we control for job title fixed effects, which restores the expected positive relationship between wages and the utility of a job.

Wages are likely to be endogeneous because they are correlated with unobserved job characteristics. Therefore, we instrument the wage with a series of different Hausman-style instruments in the rest of Table 1. Instrumenting the wage by the average wage of the firm in other markets leads to a positive estimate of the wage effect (column 3), while further controlling for job titles increases the magnitude of the estimate of the wage, just like in OLS. We then instrument the wage with the average wage of the firm in other markets, but excluding the same SOC6 and the same commuting zone. This instrument is slightly less relevant than the wages of the same firm in all other markets, but still has a very high F-stat (compare columns 3-4 vs. 5-6). The wage estimates with this instrument are also positive and slightly higher in magnitude than in columns 3-4.

Using the structure of the nested logit, we can derive the market-level application elasticity, the firm-level elasticity and the vacancy-level elasticity. The medians of these elasticities are all reported at the bottom of Table 1; the median is reported because means are driven by outliers. In all cases, the firm-level elasticity is only slightly lower than the vacancy level elasticity: this is because most firms do not post multiple vacancies in the same week in the same SOC6 by commuting zone labor market. The OLS estimates lead to implausibly low elasticities. The IV estimates lead to a median market-level elasticity of about 1, and a firm-level elasticity roughly between 2 and 3. Of these, specification 6 is our preferred specification because it controls for job title heterogeneity and the instrument is less likely to be endogenous. In specification 6, the median market-level elasticity is 1.16 and the median firm-level elasticity[3] is 2.83.

[3]To calculate the median, elasticities are weighted by the predicted vacancy share.

In Table 2, we turn to an alternative set of instruments. First, we use BLP style instruments, instrumenting the wage with characteristics of competing jobs. In columns 1 and 2 (without and with job title fixed effects), we use as instrument for the wage the number of vacancies in the market posted by competitor firms. This yields elasticities that are about double the ones found with the Hausman-style instruments. In columns 3 and 4, we add to the instrument list: instead of just the number of vacancies posted by competitors in the market, we also include the sum of the log number of employees of competitor firms, and the sum of square of the log number of employees of competitor firms. In these specifications, the elasticities are similar to the one estimated with Hausman-style instruments, especially when controlling for job titles (column 4). It is possible that the very high elasticities found when instrumented with the number of vacancies alone are biased by the endogeneity of the number of vacancies: indeed, unlike the number of product varieties in a market, the number of vacancies can be adjusted at high frequency and is thus more likely to be endogeneous.

Finally, in the last two columns (5-6) of Table 2, we use as instruments the market conditions faced by the posting firm in other markets. Here, we use the state unemployment for other markets where the firm posts. The instrument is weak with an F-test well below 10 and a coefficient on the wage that is not statistically significant, but the estimates of the elasticity are similar to the Hausman-style and BLP-style instruments.

Having shown a variety of specifications, we think that specification 6 in Table 1 is perhaps the most credible because the instrument (wages of the same firm in other markets, excluding same SOC and same commuting zone) is more likely to be exogenous yet has a strong first stage.

We now turn to describing the heterogeneity in estimated elasticities across occupations (Table 3) based on our preferred specification 6 in Table 1. The SOC-6 occupation by commuting zone estimates are collapsed at the occupation level weighing by the predicted vacancy share. The occupations are sorted in the order of increasing firm-level elasticity. There is substantial variation in the median firm-level elasticity, which ranges from 1.25 for industrial engineers to 30.80 for "secretaries and administrative assistants (except legal, medical and executive)". The elasticity for secretaries and administrative assistants is an outlier though, and the next highest elasticity is 5.61 for bookkepping, accounting and auditing clerks. Interestingly, the markets with the highest elasticities tend to be less skilled, suggesting that the wage is a more important component of the utility of a low skill job. The median market-level

elasticity is similar across occupations, at around 1.2; the similarity is driven by the fact that the market-level elasticity is mostly determined by the wage coefficient α (see equation 16 and following).

Overall, our preferred specifications yield a firm-level wage elasticity of applications of about 2, and a market-level elasticity of about 1. Furthermore, low skill occupations tend to have somewhat higher firm-level elasticities, consistent with wages being more important to low skill workers.

4 Implications for the labor supply elasticity and antitrust policy

4.1 From the application elasticity to the labor supply elasticity

Our econometric model gives us an estimate of the elasticity of job application with respect to wages. How does that translate into an elasticity of firm employment with respect to wages, i.e. a labor supply elasticity? We will use two approximations to calculate the labor supply elasticity based on our estimated application elasticity: (i) the labor supply elasticity is roughly equal twice the recruitment elasticity (increase in hires in response to an increase in wages), and (ii) the application elasticity is roughly equal to the recruitment elasticity.

The first assumption, that the labor supply elasticity is roughly equal twice the recruitment elasticity is discussed in Manning (2011). Let $R(w)$ be the number of new hires (recruitment) as a function of the wage, and let $s(w)$ be the number of separations as a function of the wage. If a firm's employment level $N(w)$ does not change (steady state), we must have $N(w) = R(w)/s(w)$, and this implies that the labor supply elasticity ϵ is $\epsilon = \epsilon_{Rw} - \epsilon_{sw}$, where ϵ_{Rw} is the recruitment elasticity and ϵ_{sw} is the separation elasticity. If $\epsilon_{Rw} \approx -\epsilon_{sw}$, then $\epsilon \approx 2\epsilon_{Rw}$. This rough equality between the recruitment and the separation elasticity is supported by empirical evidence: Dube et al. (2018) use their own estimates and other estimates from the literature to show that the recruitment elasticity (Horton, Rand and Zeckhauser, 2011) and the retention elasticity (Ho et al., 2016; Hsieh and Kocielnik, 2016; Yin, Suri and Gray, 2018) are very close to each other.

The second assumption, that the wage elasticity of applications is the same as the wage elasticity of hiring, cannot be checked directly in our data but can be checked by comparing estimates of the application elasticity with estimates of the hiring elasticity. Dube et al. (2018) and Manning (2011) report estimates of the recruitment and the retention elasticites, and they are typically between 0 and 2, which

is consistent with the application elasticity estimated by Marinescu and Wolthoff (2016), as well as the results from this paper. In summary, the wage elasticity of applications is of the same order of magnitude as recruitment and retention elasticities.

Overall, it seems reasonable to assume that the application elasticity roughly is equal to the hiring elasticity, which is equal to half of the labor supply elasticity. Under these assumptions, we can calculate the labor supply elasticity by multiplying our estimated application elasticities by two.

Under this assumption, we can calculate the markdown using our estimates of the wage elasticity of applications. Based on specification 6 from Table 1, the median markdown is 0.18, implying that workers' marginal productivity is about 18% higher than their wages. Table 3 breaks down the markups by occupation. The highest markup is for industrial engineers, whose productivity is 40% above their wages, while the lowest markdown is for secretaries and administrative assistants, whose productivity is 2% above their wages. Generally, wage markdowns seem to be higher for higher skilled occupations. Overall, our results show that, even in the largest and therfore arguably most competitive labor markets, wage markdowns are substantial, with a median of 18%.

4.2 From labor supply elasticities to market definition

Since 1982, the horizontal merger guidelines have included the hypothetical monopolist test to determine whether a product market could be profitably monopolized. The idea of the hypothetical monopolist test is to use as the relevant antitrust market the smallest market for which a hypothetical monopolist that controlled that market would find it profitable to implement a "small significant non-transitory increase in price" (SSNIP).

In the 1982 horizontal merger guidelines, there were no specific instructions about how this SSNIP test could be applied, but an influential paper soon defined a methodology (Harris and Simons, 1991): critical loss analysis. Analogously, the hypothetical monopsonist test would suggest as the relevant antitrust market the smallest labor market for which a hypothetical monopsonist that controlled that labor market would find profitable to implement "small significant non-transitory reduction in wages" (SSNRW).

Consider a simple model of monopsony, with a constant value of marginal product of labor given by a, a wage w which depends on the employment level of the monopsonist L. The profits of the monop-

sonist are

$$\pi(L) = (a - w)L.$$

If the monopsonist changes wages by Δw, and this generates a change in labor supply ΔL, the change in profits is

$$\Delta\pi = \Delta L \times (a - w - \Delta w) - \Delta w \times L.$$

Thus, the SSNRW is profitable for the monopsonist if and only if

$$\Delta L \times (a - w - \Delta w) > \Delta w \times L.$$

Dividing on both sides by wL, we obtain

$$\frac{\Delta L}{L} \times \left(\underbrace{\frac{a - w}{w}}_{\text{Markdown } \mu} - \frac{\Delta w}{w} \right) > \frac{\Delta w}{w}.$$

Rearranging terms (and taking into account that the change in wage is negative, which changes the direction of the inequality):

$$\frac{\Delta L / L}{\Delta w / w} < \frac{1}{\mu - \Delta w / w}.$$

Since the left-hand side is approximately the elasticity of labor supply, which we denote η, we have that the critical elasticity (see Harris and Simons (1991) for the corresponding concept in the product market) for the wage reduction to increase profits is:

$$\eta \approx \frac{1}{\mu - \Delta w / w}.$$

The antitrust practice typically considers a 5% increase in price (for at least a year) as the SSNIP. Therefore, we will consider a 5% "small significant non-transitory reduction in wages" (SSNRW). The market is too broad if the actual labor supply elasticity is less than the critical elasticity.

If we have data on the firm-level and the market-level elasticity, then the inverse firm level elasticity can be used to estimate μ (by the Lerner rule), and the market-level elasticity can be used to estimate η. It is clear that the hypothetical monopsonist would face the market-level elasticity. But is it incoherent

then to estimate μ from the firm-level elasticity of a non-monopsonist? It would be if we were talking about the monopsony profit-maximizing, because then the markdown should be $1/\eta$. However, we are merely asking if the profit would increase from the status quo of the non-monopsonist.

Based on our estimates in Table 1, column 6, we perform the hypothetical monopsonist test for each occupation: we use $1/(2 * median firm - level wage elasticity of applications)$ to calculate the markdown μ. We report the results in Table 3. The median market-level labor supply elasticity can be calculated as $2 * median market level elasticity$. The market-level elasticity is smaller than the critical elasticity for almost all occupations, indicating that the hypothetical monopsonist would find it profitable to decrease wages by 5%. In the case of industrial engineers, the market-level labor supply elasticity is slightly larger than the critical elasticity, indicating that this market may not be a relevant antitrust market. Overall, these results indicate that, for almost all occupations, a monopsonist could profitably reduce wages by 5% and therefore these 6-digit SOC form relevant antitrust markets.

5 Conclusion

Using data from a large online job board, we estimate the wage elasticity of applications at the vacancy-level, firm-level, and at the market-level where a labor market is defined by a SOC6 occupation and a commuting zone. We estimate workers' choice among differentiated jobs using a nested logit model. A series of instrumental variable estimators yield fairly consistent results. The firm-level application elasticity is about 2, which is only slightly lower than the vacancy-level elasticity. The market-level application elasticity is about one. Multiplying these elasticities by two, one obtains an estimate of the labor supply elasticity. Our elasticity estimates also allow us to determine that a hypothetical monopsonist would find it profitable to decrease wages by 5% is most markets, implying that most occupations at the SOC-6 level are relevant antitrust markets. Our findings imply that the median wage markdown is 0.18, meaning that workers' productivity is 18% greater than their wages, indicating substantial employer market power even in the largest labor markets.

References

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri. 2018. "Monopsony in Online Labor Markets." National Bureau of Economic Research Working Paper 24416.

Harris, Barry C., and Joseph J. Simons. 1991. "Focusing Market Definition: How Much Substitution is Necessary The DOJ Merger Guidelines Part I: Market Definition 12 Research in Law and Economics 207 (1989)." *Journal of Reprints for Antitrust Law and Economics*, 21: 151–172.

Ho, Chien-Ju, Aleksandrs Slivkins, Siddharth Suri, and Jennifer Wortman Vaughan. 2016. "Incentivizing High Quality Crowdwork." *SIGecom Exch.*, 14(2): 26–34.

Horton, John J., David G. Rand, and Richard J. Zeckhauser. 2011. "The online laboratory: conducting experiments in a real labor market." *Experimental Economics*, 14(3): 399–425.

Hsieh, Gary, and Rafał Kocielnik. 2016. "You Get Who You Pay for: The Impact of Incentives on Participation Bias." 821–833. San Francisco, California, USA:ACM Press.

Manning, Alan. 2011. "Imperfect competition in the labor market." *Handbook of labor economics*, 4: 973–1041.

Marinescu, Ioana, and Roland Rathelot. 2018. "Mismatch Unemployment and the Geography of Job Search." *American Economic Journal: Macroeconomics*, 10(3): 42–70.

Marinescu, Ioana, and Ronald Wolthoff. 2016. "Opening the Black Box of the Matching Function: the Power of Words." National Bureau of Economic Research Working Paper 22508. DOI: 10.3386/w22508.

Yin, Ming, Siddharth Suri, and Mary L. Gray. 2018. "Running Out of Time: The Impact and Value of Flexibility in On-Demand Crowdwork." *CHI '18*, 430:1–430:11. New York, NY, USA:ACM. event-place: Montreal QC, Canada.

Table 1. Impact of wages on utility derived from a job vacancy.
Controls are the log of firm employment, and the square of the log of firm employment. To calculate the median elasticity, elasticities are weighted by the predicted vacancy share.
Data source: CareerBuilder.com

	Dependent variable: δ_j, job-specific utility					
	OLS		IV: Average Wage of Same Firm in Other Markets		IV: Average Wage of Same Firm in Other Markets (Excluding Same CZ and Same SOC)	
	(1)	(2)	(3)	(4)	(5)	(6)
Log Wage	-0.0706***	0.0780**	0.839***	1.364***	1.064***	1.453***
	(0.0256)	(0.0330)	(0.171)	(0.330)	(0.293)	(0.484)
Observations	9,369	6,973	7,621	5,749	6,487	4,914
Controls	Yes	Yes	Yes	Yes	Yes	Yes
CZ x SOC Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Job Title Fixed Effects		Yes		Yes		Yes
Median Market-Level Elasticity	-0.0563	0.0622	0.669	1.088	0.849	1.159
Median Firm-Level Elasticity	-0.137	0.152	1.632	2.654	2.071	2.827
Median Wage Markdown	-3.650	3.289	0.306	0.188	0.241	0.177
Median Vacancy-Level Elasticity	-0.140	0.155	1.661	2.701	2.108	2.878
Kleibergen-Paap F-stat			243.2	147.7	117.5	78.84

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 2. Impact of wages on utility derived from a job vacancy.
Controls are the log of firm employment, and the square of the log of firm employment. To calculate the median elasticity, elasticities are weighted by the predicted vacancy share. The BLP instruments are characteristics of competitors: number of vacancies, sum of the log number of employees of competitor firms, sum of square of the log number of employees of competitor firms.
Data source: CareerBuilder.com.

	Dependent variable: δ_j, job-specific utility					
	IV: Number of Vacancies		IV: BLP Instruments		IV: Average Unemployment Rate of Same Firm in Other Markets	
	(1)	(2)	(3)	(4)	(5)	(6)
Log Wage	3.453***	4.493**	2.218***	0.837	0.787	1.414
	(1.103)	(2.126)	(0.818)	(0.516)	(1.236)	(1.106)
Observations	9,369	6,973	9,369	6,973	7,844	5,901
Controls	Yes	Yes	Yes	Yes	Yes	Yes
CZ x SOC Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Job Title Fixed Effects		Yes		Yes		Yes
Median Market-Level Elasticity	2.754	3.584	1.769	0.668	0.628	1.128
Median Firm-Level Elasticity	6.718	8.743	4.316	1.629	1.531	2.750
Median Wage Markdown	0.074	0.057	0.116	0.307	0.327	0.182
Median Vacancy-Level Elasticity	6.839	8.901	4.394	1.659	1.559	2.800
Kleibergen-Paap F-stat	28.43	10.07	14.38	9.670	1.495	3.939

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3. Elasticities and wage markdowns by occupation.
The estimates are based on specification 6 in Table 1.
Data source: CareerBuilder.com.

Occupation	Median Firm-Level Elasticity	Median Wage Markdown	Median Market-Level Elasticity	Median Critical Elasticity
Industrial Engineers	1.25	0.40	1.22	2.22
Computer User Support Specialists	2.21	0.23	1.20	3.62
First-Line Supervisors of Production and Operating Workers	2.29	0.22	1.20	3.73
Medical Secretaries	2.40	0.21	1.11	3.87
Accountants and Auditors	2.42	0.21	1.19	3.89
Financial Analysts	2.72	0.18	1.21	4.27
Registered Nurses	2.72	0.18	1.28	4.28
Executive Secretaries and Executive Administrative Assistants	2.77	0.18	1.14	4.34
Customer Service Representatives	3.10	0.16	1.11	4.74
Bookkeeping, Accounting, and Auditing Clerks	5.61	0.09	1.19	7.18
Secretaries and Administrative Assistants, Except Legal, Medical, and Executive	30.80	0.02	1.22	15.10