Estimating Labor Market Power*

JOSÉ AZAR, STEVEN BERRY, AND IOANA MARINESCU

September 18, 2019

Abstract

How much power do employers have to suppress wages below marginal productivity? It depends on the firm-level 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 in the largest occupations. We use a nested logit model of worker's utility for applying to jobs with varying wages and characteristics, including distance from the potential worker's home. We account for the endogeneity of wages by using several different instrumental variable strategies. We find that failing to instrument results in implausibly low elasticities, whereas plausible instruments result in more elastic estimates. Still, the implied market-level labor supply elasticity is about 0.6, while the firm-level labor supply elasticity is about 5.8. This implies that workers produce about 17% more than their wage level, consistent with employers having significant market power even for the largest occupational 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 (Azar, Marinescu and Steinbaum, 2017; Benmelech, Bergman and Kim, 2018; Hershbein, Macaluso and Yeh, 2019; Rinz, 2018; Lipsius, 2018; Abel, Tenreyro and Thwaites, 2018; Martins, 2018). There is also an ongoing literature attempting to directly estimate firm-level labor supply elasticities (Staiger, Spetz and Phibbs, 2010; Falch, 2010; Ransom and Sims, 2010; Matsudaira, 2013; Webber, 2015; Azar, Marinescu and Steinbaum, 2019). Perfect competition would imply that these elasticities are infinite, but the existing literature often finds extremely low elasticities (Manning, 2011) —for example Dube et al. (2018) reports elasticities in the range of 0.1—which would imply an extremely high degree of firm labor market power.

There is also a literature that differentiation in the desirability of jobs is one source of labor market power. For example, Card et al. (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 job vacancies.

In particular, we model job application choice using a nested logit framework. As in Berry (1994), each choice (here, a job) has both observed and unobserved product characteristics, leading to a endogenous price (here, wage). However, unlike in Berry (1994) or Berry, Levinsohn and Pakes (1995), we have information that matches some characteristics of applicants to the jobs that they apply to. A long literature, including McFadden, Talvitie and Associates (1977) and Berry, Levinsohn and Pakes (2004), notes the advantages of this kind of "micro" level data on choices. In particular, Berry and Haile (2016*b*) and Berry and Haile (2016*a*) note that such micro data allows us to cleanly identify substitution patterns, such as workers' substitution across job classifications. This leaves the issue of endogenous wages. Identifying the effect of wages will require instrumental variables. These variables must shift wages without being correlated with unobserved job characteristics.

We define a job market as a commuting zone by 6-digit SOC and allow workers to apply both within and across markets. 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 17%. 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.

Modeling job applicants as choosing to apply to differentiated jobs has benefits and costs relative to existing approaches to labor supply. Importantly for us, the approach provides a standard and well-understood framework, borrowed from IO, for modeling the kind of job-level differentiation suggested by Card et al. (2018). However, IO product markets do not typically have the kind of search, matching and bargaining characteristics that are sometimes thought typical of labor markets. In the model, as in Card et al. (2018), idiosyncratic match characteristics are restricted to familiar i.i.d. logit errors. Further, the model features posted wages rather than bargained wages. Posted wages may actually be a good description for many low and medium skilled jobs, but not for other jobs. However, treating jobs as having posted wages with no explicit search costs may be more realistic in the context of applications to posted jobs on an internet job site. This leaves the question, though, of how to understand the relationship between job application elasticities and standard labor supply elasticities. In the discussion below, we directly estimate the application elasticities and then use a standard stylized model to convert these into labor supply elasticities.

In the remainder of the paper, we discuss data and estimation, then the empirical results and their implications for elasticities. We close with an illustrated use of the elasticities in determining whether occupation/commuting zone markets ought to be considered as relevant antitrust markets.

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 26% 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 the SOCs used in Azar, Marinescu and Steinbaum (2017). We drop from the sample markets that didn't have at least two vacancies every week during the second quarter of 2012, or for which there was no quarter with two vacancies in different zipcodes (otherwise, we cannot identify the distance coefficient for that market). We end up with 844 commuting zone-occupation pairs.

[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 (indexed by j) 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 ($j = 0$), i.e. either 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.[2]

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 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).

In equation (1), the expression $\nu_{imt}(\lambda_m) + \lambda_m \epsilon_{ijmt}$ contains the error terms. The term ϵ_{ijmt} is a random "match value" for the match between worker i and a particular job j. The random utility of the outside option is $u_{i0mt} = \epsilon_{i0mt}$, which implies that the utilities of the jobs in the focal market are interpreted as relative to the utility of the outside option. 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, 1997). 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 $g = 0$. The nest $g = 0$ contains the outside option $j = 0$.

[2] Some workers apply to multiple vacancies in a market, in which case we treat each application as a separate decision.

Finally, the critical term δ_j is McFadden's "alternative-specific constant" (or "mean utility") which captures the average level of utility of a job vacancy. It is defined as:

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

where 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, because they are a very strong predictor of posted wages (Marinescu and Wolthoff, 2016). 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. The unobservable may also include the perceived stringency of (unobserved to us) job qualifications. In the content of job applications, a job that is too hard to get is job that is not desirable to apply for. It is natural to assume that the unobservable job attributes are correlated with the wage. For example, a firm offering good benefits and stable employment may offer lower wages, whereas jobs with stringent qualifications may have high wages.

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 – including the outside option –, 0 otherwise) on her distance to the vacancy (z_{ijm}). The coefficient on z_{ijm} identifies γ_m/λ_m, and the fixed effects identify δ_j/λ_m. The estimates for λ_m from the top level of the nested logit, which we describe in the next subsection, will then allow us to identify δ_j and γ_m.

We also obtain the predicted probability of choosing each vacancy conditional on applying to a job in the market, $s_{ijmt|g}$, which we will use for the elasticity calculations.

The inclusive parameter capturing user i's utility of applying to the best job in 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.

2.4 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 pooling observations from all markets; an observation is a job application. 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 the probability that they apply to market m (equation 4) *times* the probability that they apply to a particular job j in that market conditional on applying in that market (equation 3):

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

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 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

As in Berry (1994), we then run a regression of the estimated 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. Note that the steps of estimating the nested logit from the application level data has already uncovered the "substitution" parameter λ_m and so the parameters of δ are all that remains for us to learn. We also note that, while only vacancies with a posted wage are used for the wage equation, all vacancies are used in the nested logit model to estimate δ_j.

2.6 Wage and Nesting Parameters Flexibility

We estimate three versions of the model, with different levels of flexibility for the wage and nesting parameters. In the first version, we treat both the wage coefficient α and the nesting parameter λ as constant across markets. In a second version of the model, we allow for the α and λ parameters to be different for low-skill and high-skill occupations, and for rural and urban commuting zones. In particular, we classify an occupation into the "low-skill" group if the mean BLS OES hourly wage in 2012 for its 3-digit SOC code is below the median (18.16 in 2012 dollars), and into the "high-skill" group if it is above the median. We classify a commuting zone as "rural" if its population density is below the median (520 people per square mile[3]), and "urban" if its population density is above the median. We call this version, with heterogeneity across two wage categories by two population density categories, the "intermediate" version of the model.

Our third version allows for the α and λ parameters to vary freely across 6-digit SOC occupations, as well as continuously with population density (using a quadratic in population density). We call this the "full" version of the model.

In all three versions, for the top level logit we include the following controls: a dummy for whether the applicant is in the same commuting zone as the job vacancy, market (CZ$\times$SOC) fixed effects, and year-week fixed effects. For the bottom level multinomial logit we include the log of distance between the applicant's zipcode and the zipcode of the job vacancy.

[3] A density of at least 1,000 people per square mile is one of the criteria for an urbanized area as defined by Census: https://www2.census.gov/geo/pdfs/reference/GARM/Ch12GARM.pdf

For the wage equation that relates the job fixed effects to job characteristics (including the wage), we include market fixed effects, firm size measured as the log number of employees, and firm size squared. In some specifications we also include job title fixed effects.

2.6.1 Instruments for the wage

Since the wage is endogenous, we need to instrument for it. We consider three broad strategies to construct instruments for the wage. One is to use variation in the choice set of job seekers, the so-called "BLP" instruments, as in Berry, Levinsohn and Pakes (1995), Nevo (2000) and Berry and Haile (2014). Another set of instruments is based on the idea that multi-location firms may try to limit within-firm cross-location variation in wages. Reasons for this could have to do with fairness norms or with the ease of setting wage policy at the firm level. For all sets of instruments, we have to consider whether they are likely to be correlated with unobserved (to us) job characteristics.

To construct instruments that measure the variation in possible choice sets across markets, we look at the characteristics of other job vacancies in the market, such as the number of vacancies, and the size of the other firms that post in the market. If the model is correct, these are excluded because they do not enter the formula for the utility of user i for product j (see Berry and Haile (2014)). However, if job characteristics can respond relatively quickly to changes in demand, and to changes in the jobs posted by other firms, then the exogeneity restriction could be violated. This is an important drawback of these instruments in the context of labor market vacancies, as the vacancies change constantly and rapid responses may be possible.

We consider two variants of instruments based on the job characteristics: first a simple version instrumenting only with the number of vacancies in the market, and then a version instrumenting with the 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.[4] In the latter case, since we have more than one instrument, we estimate the equation using 2-step GMM.

Second, given the potential endogeneity of BLP-style instruments in our context, we consider alternative identification strategies. One alternative set of instruments is a variant of the "Hausman instru-

[4]There is a tradition of constructing the BLP instruments as sums of other products' characteristics, as in Berry, Levinsohn and Pakes (1995), although this ignores arguments for various non-linear functions that may better approximate optimal instruments, as discussed in Berry, Levinsohn and Pakes (1999), Reynaert and Verboven (2014) and Gandhi and Houde (2017).

ments" often used in the industrial organization literature (Hausman, Leonard and Zona, 1994). 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. The Hausman idea of instruments based on prices in other markets relies on costs being correlated across locations, but demand shocks that are not. In the context of wages, the similar condition for the estimation of labor supply is that firms' within-market wage decisions reflect a common (across markets) firm-level marginal productivity (labor demand), but not common shocks to unobserved job characteristics.

Alternatively, if firms set wages with some effort to reduce cross-location wage inequality, variation in competitive labor market conditions across the markets where a firm operates can lead to correlated wages, within-firm across markets. This will lead to other market wages that predict own-market wages but are naturally excluded from own-market labor supply decisions. DellaVigna and Gentzkow (2019) discuss and document firm-level retail pricing decisions that hold prices constant across markets within firm. Similar forces could be at work with wages. Of course, the empirical relevance of other-market within-firm wages can be tested in the data.

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). Excluding the same CZ or same SOC makes the instrument somewhat more plausible.

Our first version of Hausman instruments can easily be criticised on the grounds that firm-level wage decisions across locations may involve simultaneous changes to unobserved benefits and working conditions. If so, a firm's wages in other locations may reflect the unobserved quality of the job in this location. We therefore use a set of instruments that go one step further, 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. To implement this strategy, we first obtain predicted wages for each market by running a regression of wages on CZ-SOC fixed effects and year-week fixed effects, and obtaining the corresponding predicted values for the wages. We then take averages of the predicted wages across vacancies posted by the same firm in other markets. As with the Hausman instruments, we calculate two versions: one using the average predicted wage for vacancies by the same firm in all other CZ

$\times$ SOC markets, and one – plausibly more exogenous – excluding markets that are in the same CZ or the same SOC from the average predicted wage calculation. Intuitively, this third set of instruments amounts to saying that if a firm faces high-wage competition in other markets, then it also pays higher wages in the current market.

Our preferred specification is this last one, the instrument based on the average *predicted* wage for vacancies by the same firm in all other CZ $\times$ SOC markets excluding markets that are in the same CZ or the same SOC from the average predicted wage calculation. We also run a version of this specification with job title fixed effects, which improves the wage prediction, but also creates estimation issues due to the large number of instruments.

One important caveat about the instruments based on wages in other markets is that some types of firms, which hire some classes of occupations, may engage in across-location wage-setting, but in other cases firms may set wages only based on own-market conditions. Thus, our Hausman-like instruments may work better in some occupational markets in others, a possibility that suggests further research focusing on subsets of firms and occupations where instruments can be more closely tailored to specific conditions.

2.7 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

3.1 Basic model

In the basic model (Table 1 and 2), wage coefficients and nesting parameters are constrained to be the same for all markets.

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, in the rest of Table 1, we instrument the wage with a series of different BLP style instruments, instrumenting the wage with characteristics of competing jobs. In columns 3 and 4 (without and with job title fixed effects), we instrument the wage with the number of vacancies in the market posted by competitor firms. This yields much larger wage coefficients than the OLS estimates and suggests that OLS suffers from significant downward bias. In columns 5 and 6, 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 wage coefficient is smaller but still much bigger than the OLS estimate. It is possible that the very high wage coefficients found when instrumenting 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.

In Table 2, we instrument the wage with a series of different more Hausman-like instruments. Instrumenting the wage by the average wage of the firm in other markets leads to a positive estimate of the wage effect (column 1), while further controlling for job titles increases the magnitude of the estimate (column 2), 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 1-2 vs. 3-4). The wage estimates with this instrument are also positive and slightly higher in magnitude than in columns 1-2. Instead of the actual wages posted by the firm itself in other markets, we then use as instrument the predicted wage for the firm in other markets, either including the same CZ and same SOC (Table 2, columns 5-6) or excluding these (columns 7-8). This strategy yields higher wage coefficients than in columns 1-4 but typically smaller than for the BLP style instruments in Table 1.

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 Tables 1 and 2; the median is reported because means are driven by outliers (distributions of estimates are reported for the "full" model in Figures 4 and 6). 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 in Table 1 lead to implausibly low elasticities. The IV estimates lead to a median market-level elasticity between 0.12 and 1.1, and a firm-level elasticity between 1.3 and 7.8. Generally, including job titles yields higher elasitcity estimates and BLP instruments yield higher elasticities than Hausman-style instruments.

3.2 Heterogeneity

We now allow the wage coefficients and the nesting parameters to vary across labor markets. Otherwise, the estimation strategy stays the same.

In Tables 3 and 4, we report estimates of the intermediate model, which allows for different wage coefficients and nesting parameters for low-skill and high-skill occupations and for urban and rural commuting zones. As before, OLS yields elasticities that are implausibly low. For the intermediate model, the BLP style instruments have limited relevance, especially when also including job title fixed effects (see F-stat in Table 3). The Hausman style instruments (Table 4) are more relevant, but the F-stat is still too low for specifications with job title fixed effects. In terms of the magnitudes of the elasticity estimates, the Hausman style instruments yield similar magnitudes to those in Table 2.

How do elasticities differ between low and high skill occupations and between urban and rural commuting zones? Figures 1 and 2 report elasticities from our preferred specification, column 6 in Table 4. It is clear that rural commuting zones have lower market and firm-level elasticites than urban commuting zones, both for low and high skilled occupations. These lower elasticites may reflect lower job opportunities in less densely populated areas and is consistent with reduced-form elasticity estimates in Azar, Marinescu and Steinbaum (2019). On the other hand, differences in elasticities between low and high skilled occupations are very small for firm-level elasticities, while low skill occupations have higher market-level elasticities than high skill occupations.

We then move to an even more flexible model, the "full" model, which allows for different wage coefficients and nesting parameters by SOC-6 occupation and by (a quadratic in) population density. With this model like with the intermediate model, rural commuting zones have lower elasticities than urban commuting zones. The median firm-level elasticities of this full model are very similar to the median firm-level elasticities of the intermediate model (Figure 2). The median market-level elasticities are slightly higher in the full model, especially for high skill occupations Figure 1).

In contrast to what we found for the intermediate level, low skill occupations always have *lower* firm-level and market-level elasticities than high skill occupations in the full model. This difference between the intermediate and full model is mostly driven by the high skill occupations having a higher elasticity in the full model as opposed to the intermediate model, while the low skill occupations have similar elasticities in both models. This suggests that there may be substantive heterogeneity between high skill occupations that the intermediate model may not be flexible enough to pick up.

Based on Figures 1 and 2, we conclude that rural commuting zones have lower elasticities than urban commuting zones. Low and high skill workers have similar elasticities, though the exact ranking depends on the model, with a more flexible model yielding lower elasticities for low skill occupations. We cannot blindly assume that low skill workers can more easily move to other jobs because specific skills do not matter. Therefore, there is no strong reason to believe that a SOC-6 by commuting zone market definition is decidedly too narrow for low skill as opposed to high skill workers.

We now give more detail on the heterogeneity of elasticities by occupation and by population density in the full model.

In Figures 3 and 4, we show the range of elasticity estimates for each SOC-6 for the full model. The Figures exclude two outlier occupations ("Industrial engineers" and "other" small occupations) whose elasticities are show in appendix Figures 7 and 8. Most occupations have positive median elasticities, but there are a few occupations with negative median elasticities. While a negative elasticity is implausible, such negative estimates are likely to occur when we slice the data finely; for example, Farber (2015) finds negative labor supply elasticities for some individual taxi drivers. Rankings of occupations by market-level elasticities are somewhat different from rankings by firm-level elasticities. For market-level elasticities (Figure 3), sales representatives have the highest elasticity with a median slightly above 0.5. Truck drivers, whether heavy or light, have among the lowest market-level elasticities, suggesting that there are few suitable job substitutes for truckers in other labor markets (other occupation and/or other commuting zone). For firm-level elasticties, administrative assistant types of occupations occupy the top three spots for the most elastic, suggesting that these type of jobs tend to be more substitutes for each other within their market than other types of jobs. At the bottom of the firm elasticity scale are again light truck drivers, who also have the lowest market-level elasticity.

In Figure 5, we plot the median market-level elasticity as a function of the population density. Market-level elasticities increase with population density, especially when log population density is above 6.5, i.e. population density is above 665 per square mile; remember that the density threshold that is part of the definition of an urbanized commuting zone is 1000. So the median market level elasticity increases particularly fast with population density once we move into more urbanized commuting zones.

Finally, in Figure 6, we show the relationship between population density and the distaste for distance calculated using the nesting parameters. Workers in more densely populated commuting zones are less likely to apply to jobs far away from their zip code of residence than workers in less densely populated areas. This could at first blush suggest that commuting zones are too narrow a definition of the labor market for rural areas. However, the very definition of a commuting zone takes into account the fact that people in rural areas are more willing to commute over long distances. In fact, the market level elasticities are *smaller* in less densely populated areas (Figure 5), suggesting that if anything commuting zones in rural areas may be too broad to define a labor market.

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 separation 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 separation 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 is roughly 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 our preferred specification 7 from Table 4, the median markdown is 0.17, implying that workers' marginal productivity is about 17% higher than their wages. Overall, our results show that, even in the largest occupations and therefore arguably most competitive labor markets, wage markdowns are substantial, with a median of 17%.

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 monopsonist 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 Figure 3, we can perform the hypothetical monopsonist test for each occupation. 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 all occupations, indicating that the hypothetical monopsonist would find it profitable to decrease wages by 5%.

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. In our preffered specification, the firm-level application elasticity is 2.88, which is only slightly lower than the vacancy-level elasticity. The market-level application elasticity is 0.30. 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 0.17, meaning that workers' productivity is 17% greater than their wages, indicating substantial employer market power.

References

Abel, Will, Silvana Tenreyro, and Gregory Thwaites. 2018. "Monopsony in the UK." Centre for Macroeconomics (CFM) 1827.

Azar, José, Ioana Marinescu, and Marshall I Steinbaum. 2017. "Labor market concentration." National Bureau of Economic Research.

Azar, José, Ioana Marinescu, and Marshall Steinbaum. 2019. "Measuring Labor Market Power Two Ways." *AEA Papers and Proceedings*, 109: 317–321.

Benmelech, Efraim, Nittai Bergman, and Hyunseob Kim. 2018. "Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?"

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile prices in market equilibrium." *Econometrica: Journal of the Econometric Society*, 841–890.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1999. "Voluntary Export Restraints on Automobiles: Evaluating a Strategic Trade Policy." 89(3): 189–211.

Berry, Steven, James Levinsohn, and Ariel Pakes. 2004. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Vehicle Market." *Journal of Political Economy*, 112(1): 68–105.

Berry, Steven T. 1994. "Estimating discrete-choice models of product differentiation." *The RAND Journal of Economics*, 242–262.

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

Berry, Steven T., and Philip A. Haile. 2016*a*. "Identification in Differentiated Products Markets." *Annual Review of Economics*, 8: 27–52.

Berry, Steven T., and Philip A. Haile. 2016*b*. "Nonparametric Identification of Multinomial Choice Demand Models with Heterogeneous Consumers." Yale University.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics*, 36: S13–S70.

Cardell, N. Scott. 1997. "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity." *Econometric Theory*, 13(2): 185–213.

DellaVigna, Stefano, and Matthew Gentzkow. 2019. "Uniform Pricing in US Retail Chains." NBER Working Paper 23996.

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

Falch, Torberg. 2010. "The Elasticity of Labor Supply at the Establishment Level." *Journal of Labor Economics*, 28(2): 237–266.

Farber, Henry S. 2015. "Why you Can't Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers." *The Quarterly Journal of Economics*, 130(4): 1975–2026.

Gandhi, Amit, and Jean-Francois Houde. 2017. "Measuring Substitution Patterns in Differentiated Products Industries." University of Wisconsin-Madison.

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.

Hausman, Jerry, Gregory Leonard, and J Douglas Zona. 1994. "Competitive analysis with differenciated products." *Annales d'Economie et de Statistique*, 159–180.

Hershbein, Brad, Claudia Macaluso, and Chen Yeh. 2019. "Concentration in U.S. local labor markets: evidence from vacancy and employment data." 40.

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.

Lipsius, Ben. 2018. "Labor Market Concentration does not Explain the Falling Labor Share." Job Market Papers pli1202.

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.

Martins, Pedro S. 2018. "Making their own weather? Estimating employer labour-market power and its wage effects." Queen Mary, University of London, School of Business and Management, Centre for Globalisation Research 95.

Matsudaira, Jordan D. 2013. "Monopsony in the Low-Wage Labor Market? Evidence from Minimum Nurse Staffing Regulations." *The Review of Economics and Statistics*, 96(1): 92–102.

McFadden, Daniel, A. Talvitie, and Associates. 1977. *Demand Model Estimation and Validation.* Berkeley CA:Institute of Transportation Studies.

Nevo, Aviv. 2000. "A Practitioner's Guide to Estimation of Random Coefficients Logit Models of Demand." *Journal of Economics & Management Strategy*, 9(4): 513–548.

Ransom, Michael R, and David P. Sims. 2010. "Estimating the Firm's Labor Supply Curve in a "New Monopsony" Framework: Schoolteachers in Missouri." *Journal of Labor Economics*, 28(2): 331–355.

Reynaert, Mathias, and Frank Verboven. 2014. "Improving the performance of random coefficients demand models: The role of optimal instruments." *Journal of Econometrics*, 179(1): 83–98.

Rinz, Kevin. 2018. "Labor Market Concentration, Earnings Inequality, and Earnings Mobility." Center for Economic Studies, US Census Bureau Working Paper 2018-10.

Staiger, Douglas O, Joanne Spetz, and Ciaran S Phibbs. 2010. "Is there monopsony in the labor market? Evidence from a natural experiment." *Journal of Labor Economics*, 28(2): 211–236.

Webber, Douglas A. 2015. "Firm market power and the earnings distribution." *Labour Economics*, 35: 123–134.

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 in the basic model: OLS and BLP-style instruments.
Estimation is by OLS in the first two columns, and by 2SLS in columns (3) to (6). The sample includes jobs in 844 CZ×SOC labor markets over the period between April and June of 2012. The job-specific utilities are estimated using a nested logit model. The elasticity numbers represent medians across all observations at the job × year-week level. Standard errors are clustered by CZ × SOC.
Data source: CareerBuilder.com

	Dependent variable: δ_j, job-specific utility					
	OLS		IV: Number of Vacancies		IV: BLP Instruments	
	(1)	(2)	(3)	(4)	(5)	(6)
Log Wage	-0.0163***	0.0194***	0.887***	1.316***	0.647***	0.352***
	(0.00387)	(0.00600)	(0.186)	(0.449)	(0.133)	(0.115)
Log Employees	0.00146	0.000451	-0.0734***	-0.0998***	-0.0535***	-0.0253***
	(0.00159)	(0.00198)	(0.0161)	(0.0338)	(0.0115)	(0.00936)
(Log Employees)^2	-0.000710***	-0.000728***	0.00459***	0.00596***	0.00318***	0.000990
	(0.000128)	(0.000162)	(0.00114)	(0.00227)	(0.000823)	(0.000633)
CZ × SOC FE	✓	✓	✓	✓	✓	✓
Job Title FE		✓		✓		✓
Observations	16,481	12,139	16,481	12,139	16,481	12,139
R-squared	0.044	0.052	-4.569	-6.747	-2.439	-0.396
Median Market-Level Elasticity	-0.0135	0.0161	0.734	1.088	0.535	0.291
Median Firm-Level Elasticity	-0.144	0.172	7.839	11.63	5.712	3.112
Median Vacancy-Level Elasticity	-0.147	0.176	8.017	11.89	5.843	3.183
Kleibergen-Paap F-stat			33.69	13.35	17.70	11.92

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 in the basic model: Hausman-style instruments

Estimation is by 2SLS. The sample includes jobs in 844 CZ×SOC labor markets over the period between April and June of 2012. The job-specific utilities are estimated using a nested logit model. The elasticity numbers represent medians across all observations at the job × year-week level. Standard errors are clustered by CZ×SOC.

Data source: CareerBuilder.com

	Dependent variable: δ_j, job-specific utility							
	IV: Average Wage of Same Firm in Other Markets		IV: Average Wage of Same Firm in Other Markets (Excluding Same CZ and Same SOC)		IV: Average Predicted Wage of Same Firm in Other Markets		IV: Average Predicted Wage of Same Firm in Other Markets (Excluding Same CZ and Same SOC)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Wage	0.117***	0.210***	0.148***	0.231***	0.254***	0.528***	0.314***	0.619***
	(0.0178)	(0.0355)	(0.0277)	(0.0540)	(0.0343)	(0.0803)	(0.0470)	(0.122)
Log Employees	-0.0125***	-0.0184***	-0.0205***	-0.0250***	-0.0258***	-0.0466***	-0.0391***	-0.0639***
	(0.00248)	(0.00377)	(0.00373)	(0.00604)	(0.00386)	(0.00751)	(0.00565)	(0.0122)
(Log Employees)^2	0.000306	0.000569**	0.000893***	0.00103***	0.00121***	0.00240***	0.00210***	0.00344***
	(0.000187)	(0.000270)	(0.000262)	(0.000398)	(0.000276)	(0.000508)	(0.000386)	(0.000779)
CZ × SOC FE	✓	✓	✓	✓	✓	✓	✓	✓
Job Title FE		✓		✓		✓		✓
Observations	13,865	10,368	11,781	8,851	13,860	10,366	11,777	8,849
R-squared	-0.054	-0.072	-0.108	-0.111	-0.345	-0.823	-0.534	-1.155
Median Market-Level Elasticity	0.0969	0.173	0.122	0.191	0.210	0.437	0.260	0.512
Median Firm-Level Elasticity	1.035	1.852	1.308	2.042	2.241	4.667	2.773	5.471
Median Vacancy-Level Elasticity	1.058	1.895	1.337	2.089	2.292	4.774	2.836	5.596
Kleibergen-Paap F-stat	508.3	146.3	186.7	52.92	229.3	82.36	157.3	43.78

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3. Impact of wages on utility derived from a job vacancy in the intermediate model: OLS and BLP-style instruments
Estimation is by OLS in the first two columns, and by 2SLS in columns (3) to (6). The sample includes jobs in 844 CZ×SOC labor markets over the period between April and June of 2012. The job-specific utilities are estimated using a nested logit model. The elasticity numbers represent medians across all observations at the job × year-week level. Standard errors are clustered by CZ × SOC. We classify CZs into rural and urban based on whether the population density is below or above our sample median of 520 persons per square mile. We classify occupations into low-skill and high-skill based on whether their 3-digit SOC category's average hourly wage according to the BLS Occupational Employment Statistics data is above or below our sample median of 18.83 USD per hour.
Data source: CareerBuilder.com

	Dependent variable: δ_j, job-specific utility					
	OLS		IV: Number of Vacancies		IV: BLP Instruments	
	(1)	(2)	(3)	(4)	(5)	(6)
Log Wage × Rural CZ × Low-Skill SOC	-0.00942	0.0268*	0.851***	1.618	0.388***	-0.0733
	(0.0101)	(0.0144)	(0.214)	(2.225)	(0.0879)	(0.230)
Log Wage × Urban CZ × Low-Skill SOC	-0.00461	0.0261	2.365**	11.04	1.153**	-1.556**
	(0.0120)	(0.0168)	(1.194)	(30.06)	(0.515)	(0.723)
Log Wage × Rural CZ × High-Skill SOC	-0.0202***	0.0111	0.504***	1.275	0.365***	0.281*
	(0.00596)	(0.00784)	(0.128)	(1.623)	(0.0835)	(0.161)
Log Wage × Urban CZ × High-Skill SOC	-0.0171***	0.0142**	1.655***	3.806	1.400***	1.705*
	(0.00533)	(0.00669)	(0.511)	(2.781)	(0.425)	(0.882)
Log Employees	0.000556	0.000713	-0.0894***	-0.238	-0.0573***	-0.0133
	(0.00171)	(0.00223)	(0.0240)	(0.440)	(0.0138)	(0.0253)
(Log Employees)^2	-0.000671***	-0.000795***	0.00563***	0.0150	0.00345***	0.000230
	(0.000139)	(0.000183)	(0.00163)	(0.0288)	(0.000973)	(0.00173)
CZ × SOC FE	✓	✓	✓	✓	✓	✓
Job Title FE		✓		✓		✓
Observations	16,455	12,119	16,455	12,119	16,455	12,119
R-squared	0.046	0.055	-11.778	-141.075	-5.693	-6.783
Median Market-Level Elasticity	-0.0136	0.0121	0.791	1.504	0.361	0.236
Median Market-Level Elasticity (Rural)	-0.0155	0.0102	0.464	1.142	0.313	0.215
Median Market-Level Elasticity (Urban)	-0.0131	0.0125	1.452	3.342	1.080	1.306
Median Market-Level Elasticity (Low-Skill)	-0.00658	0.0211	0.769	1.462	0.351	-0.0663
Median Market-Level Elasticity (High-Skill)	-0.0154	0.0103	1.089	2.505	0.922	1.122
Median Firm-Level Elasticity	-0.141	0.124	6.951	14.37	4.115	2.944
Median Firm-Level Elasticity (Rural)	-0.175	0.124	5.588	13.03	3.183	2.434
Median Firm-Level Elasticity (Urban)	-0.138	0.124	14.31	33.05	11.34	13.81
Median Firm-Level Elasticity (Low-Skill)	-0.0594	0.185	6.933	13.18	3.161	-0.597
Median Firm-Level Elasticity (High-Skill)	-0.149	0.120	10.02	23.05	8.484	10.33
Median Vacancy-Level Elasticity	-0.146	0.125	6.971	14.41	4.125	3.042
Median Vacancy-Level Elasticity (Rural)	-0.194	0.125	5.661	13.23	3.501	2.693
Median Vacancy-Level Elasticity (Urban)	-0.144	0.124	14.46	33.27	11.79	14.35
Median Vacancy-Level Elasticity (Low-Skill)	-0.0675	0.192	6.964	13.24	3.175	-0.600
Median Vacancy-Level Elasticity (High-Skill)	-0.150	0.123	11.72	26.95	9.917	12.07
Kleibergen-Paap F-stat			2.151	0.0337	1.489	0.860

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 4. Impact of wages on utility derived from a job vacancy in the intermediate model: Hausman-style instruments.
Estimation is by 2SLS. The sample includes jobs in 844 CZ×SOC labor markets over the period between April and June of 2012. The job-specific utilities are estimated using a nested logit model. The elasticity numbers represent medians across all observations at the job × year-week level. Standard errors are clustered by CZ×SOC. We classify CZs into rural and urban based on whether the population density is below or above our sample median of 520 persons per square mile. We classify occupations into low-skill and high-skill based on whether their 3-digit SOC category's average hourly wage according to the BLS Occupational Employment Statistics data is above or below our sample median of 18.83 USD per hour.
Data source: CareerBuilder.com

	Dependent variable: δ_j, job-specific utility							
	IV: Average Wage of Same Firm in Other Markets		IV: Average Wage of Same Firm in Other Markets (Excluding Same CZ and Same SOC)		IV: Average Predicted Wage of Same Firm in Other Markets		IV: Average Predicted Wage of Same Firm in Other Markets (Excluding Same CZ and Same SOC)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log Wage × Rural CZ × Low-Skill SOC	0.0998***	0.129***	0.159***	0.196***	0.237***	0.349***	0.355***	0.514***
	(0.0274)	(0.0451)	(0.0431)	(0.0699)	(0.0483)	(0.0922)	(0.0651)	(0.115)
Log Wage × Urban CZ × Low-Skill SOC	0.205***	0.327***	0.278***	0.482**	0.634***	1.077***	0.507***	0.887***
	(0.0600)	(0.114)	(0.0945)	(0.213)	(0.156)	(0.388)	(0.132)	(0.287)
Log Wage × Rural CZ × High-Skill SOC	0.0887***	0.166***	0.145***	0.218***	0.152***	0.392***	0.256***	0.547***
	(0.0176)	(0.0373)	(0.0258)	(0.0579)	(0.0280)	(0.0852)	(0.0402)	(0.136)
Log Wage × Urban CZ × High-Skill SOC	0.149***	0.294***	0.179***	0.291***	0.297***	0.680***	0.406***	0.889***
	(0.0373)	(0.0707)	(0.0479)	(0.0853)	(0.0713)	(0.139)	(0.110)	(0.311)
Log Employees	-0.0133***	-0.0167***	-0.0241***	-0.0272***	-0.0271***	-0.0439***	-0.0439***	-0.0663***
	(0.00261)	(0.00411)	(0.00412)	(0.00685)	(0.00402)	(0.00780)	(0.00615)	(0.0143)
(Log Employees)^2	0.000324	0.000418	0.00109***	0.00112**	0.00124***	0.00216***	0.00237***	0.00356***
	(0.000199)	(0.000302)	(0.000290)	(0.000451)	(0.000289)	(0.000526)	(0.000426)	(0.000929)
CZ × SOC FE	✓	✓	✓	✓	✓	✓	✓	✓
Job Title FE		✓		✓		✓		✓
Observations	13,845	10,354	11,760	8,839	13,840	10,352	11,756	8,837
R-squared	-0.071	-0.114	-0.163	-0.221	-0.545	-1.311	-0.730	-1.680
Median Market-Level Elasticity	0.0916	0.158	0.137	0.201	0.208	0.374	0.301	0.520
Median Market-Level Elasticity (Rural)	0.0771	0.127	0.125	0.175	0.140	0.315	0.234	0.449
Median Market-Level Elasticity (Urban)	0.130	0.248	0.157	0.255	0.261	0.597	0.351	0.722
Median Market-Level Elasticity (Low-Skill)	0.0902	0.117	0.144	0.177	0.214	0.316	0.319	0.465
Median Market-Level Elasticity (High-Skill)	0.0979	0.194	0.133	0.205	0.196	0.448	0.267	0.585
Median Firm-Level Elasticity	0.994	1.855	1.520	2.380	1.929	4.401	2.883	6.012
Median Firm-Level Elasticity (Rural)	0.816	1.437	1.302	1.889	1.679	3.400	2.752	4.742
Median Firm-Level Elasticity (Urban)	1.285	2.383	1.553	2.521	2.580	5.894	3.464	7.201
Median Firm-Level Elasticity (Low-Skill)	0.813	1.051	1.294	1.593	1.932	2.845	2.879	4.185
Median Firm-Level Elasticity (High-Skill)	1.002	1.875	1.534	2.402	1.801	4.435	2.887	6.173
Median Vacancy-Level Elasticity	1.001	1.870	1.554	2.433	1.941	4.427	2.902	6.132
Median Vacancy-Level Elasticity (Rural)	0.851	1.589	1.395	2.090	1.708	3.761	2.828	5.246
Median Vacancy-Level Elasticity (Urban)	1.299	2.477	1.566	2.541	2.598	5.944	3.526	7.487
Median Vacancy-Level Elasticity (Low-Skill)	0.817	1.055	1.301	1.600	1.940	2.858	2.902	4.208
Median Vacancy-Level Elasticity (High-Skill)	1.053	2.084	1.560	2.444	2.105	4.815	2.902	6.297
Kleibergen-Paap F-stat	33.43	8.322	25.28	2.899	8.049	2.766	12.71	4.173

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

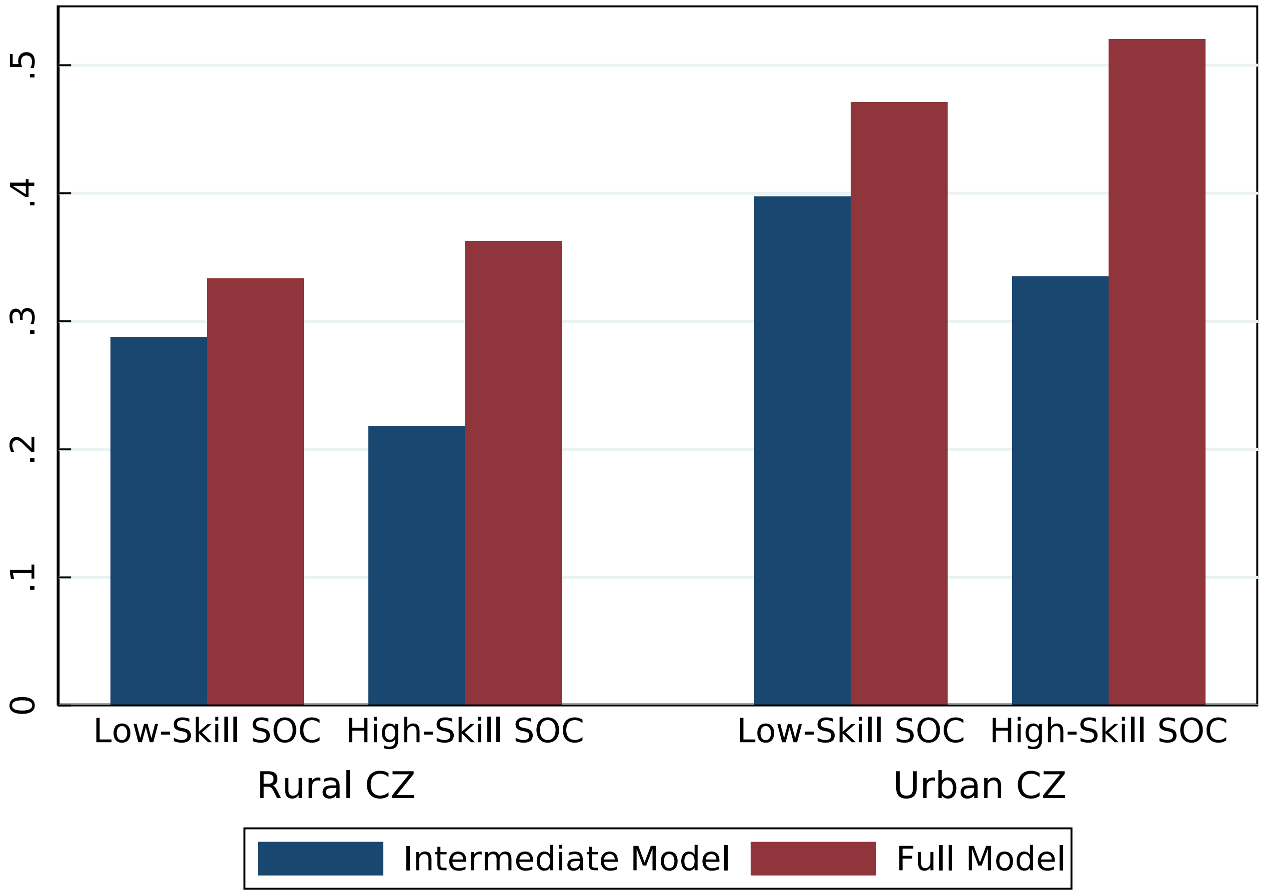

Figure 1. Distribution of Market-Level Elasticities by Rural-Urban and Skill Classification. This figure shows the median market-level elasticity of applications by low/high skill level, and by urban-rural classification for our "Intermediate Model" and for the "Full Model". The intermediate model allows for heterogeneity in the wage and nesting parameters by low/high skill level, and by urban-rural classification. The full model is more flexible and allows the wage and nesting coefficients to vary for each 6-digit occupation, as well as a quadratic polynomial in log population density. The results are for the specification that includes both CZ × SOC fixed effects and job-title fixed effects, and instruments the wage using the average predicted wage of the same firm in other markets (excluding the same CZ and same SOC). The interactions of the wage are instrumented using interactions of the average predicted wage for the same firm in other markets, interacted with the corresponding variables.

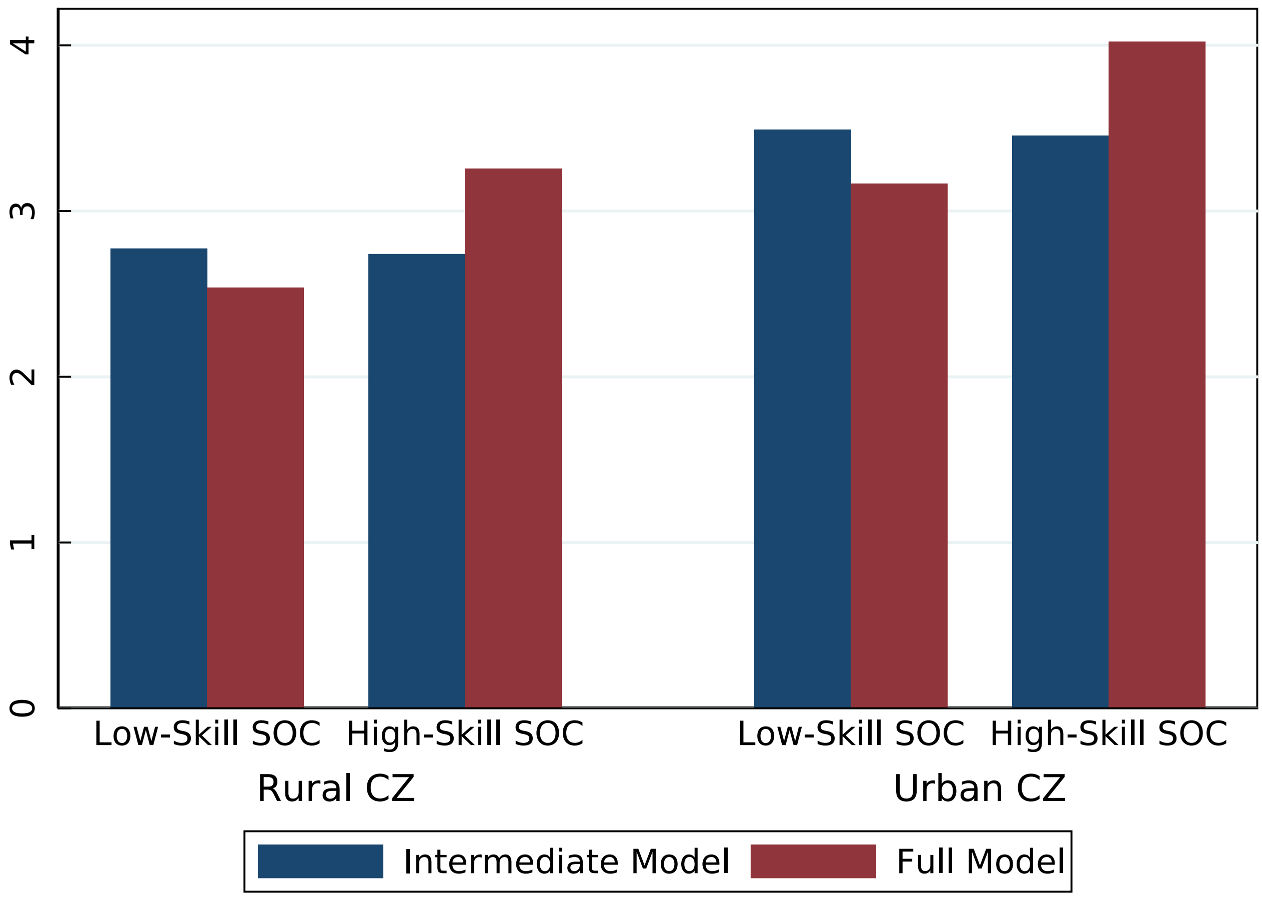

Figure 2. Distribution of Firm-Level Elasticities by Rural-Urban and Skill Classification. This figure shows the median firm-level elasticity of applications by low/high skill level, and by urban-rural classification for our "Intermediate Model" and for the "Full Model". The intermediate model allows for heterogeneity in the wage and nesting parameters by low/high skill level, and by urban-rural classification. The full model is more flexible and allows the wage and nesting coefficients to vary for each 6-digit occupation, as well as a quadratic polynomial in log population density. The results are for the specification that includes both CZ $\times$ SOC fixed effects and job-title fixed effects, and instruments the wage using the average predicted wage of the same firm in other markets (excluding the same CZ and same SOC). The interactions of the wage are instrumented using interactions of the average predicted wage for the same firm in other markets, interacted with the corresponding variables.

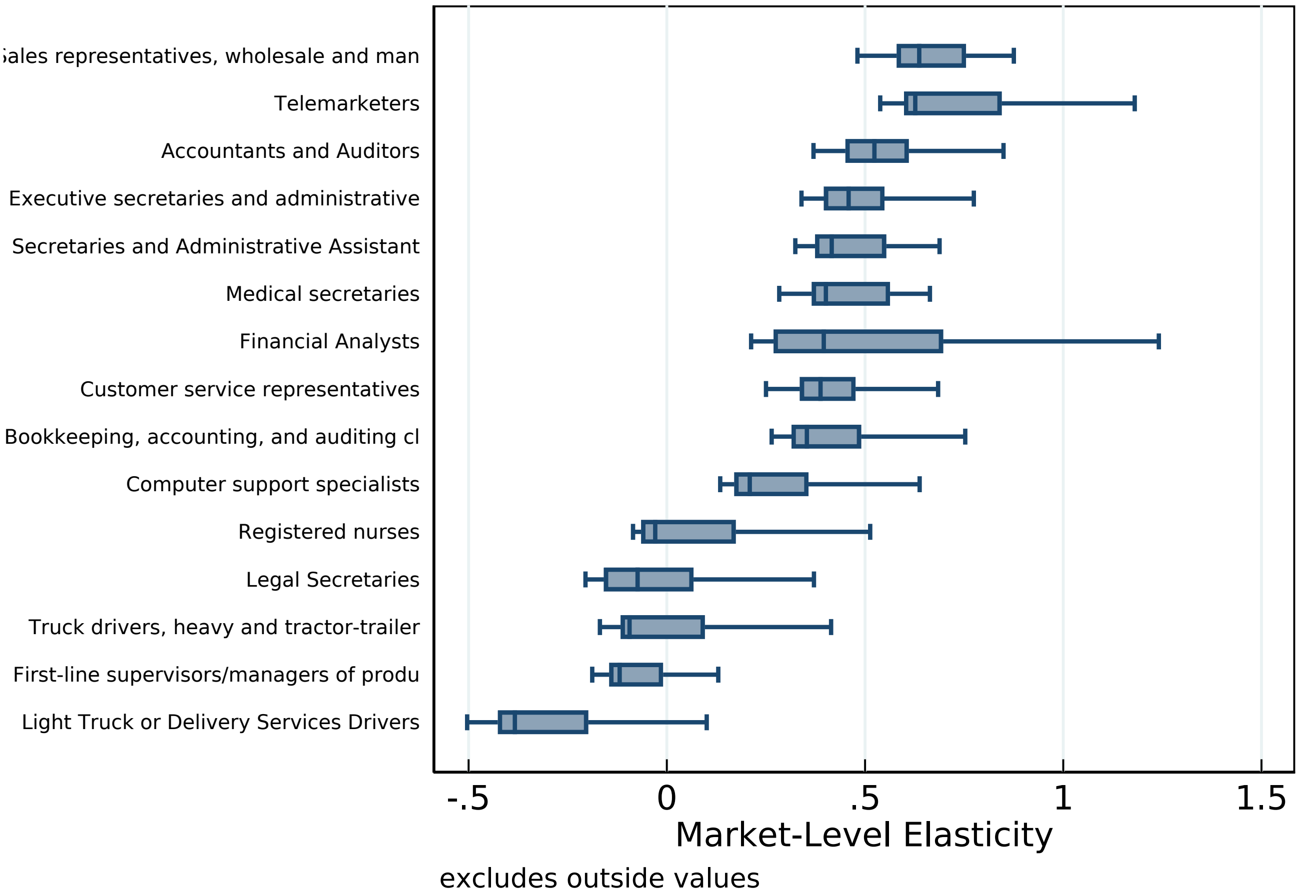

Figure 3. Distribution of Market-Level Elasticities by 6-digit SOC. This figure shows the median market-level elasticity of applications by 6-digit SOC for the "Full Model". The full model is allows the wage and nesting coefficients to vary for each 6-digit occupation, as well as a quadratic polynomial in log population density. The results are for the specification that includes both CZ × SOC fixed effects and job-title fixed effects, and instruments the wage using the average predicted wage of the same firm in other markets (excluding the same CZ and same SOC). The interactions of the wage are instrumented using interactions of the average predicted wage for the same firm in other markets, interacted with the corresponding variables. The figure excludes "Industrial Engineers" and the "Other SOCs" category (grouping SOCs with less than 100,000 observations at the application level) because the range of elasticities was too wide for visualization. Graphs with these categories are included in the appendix.

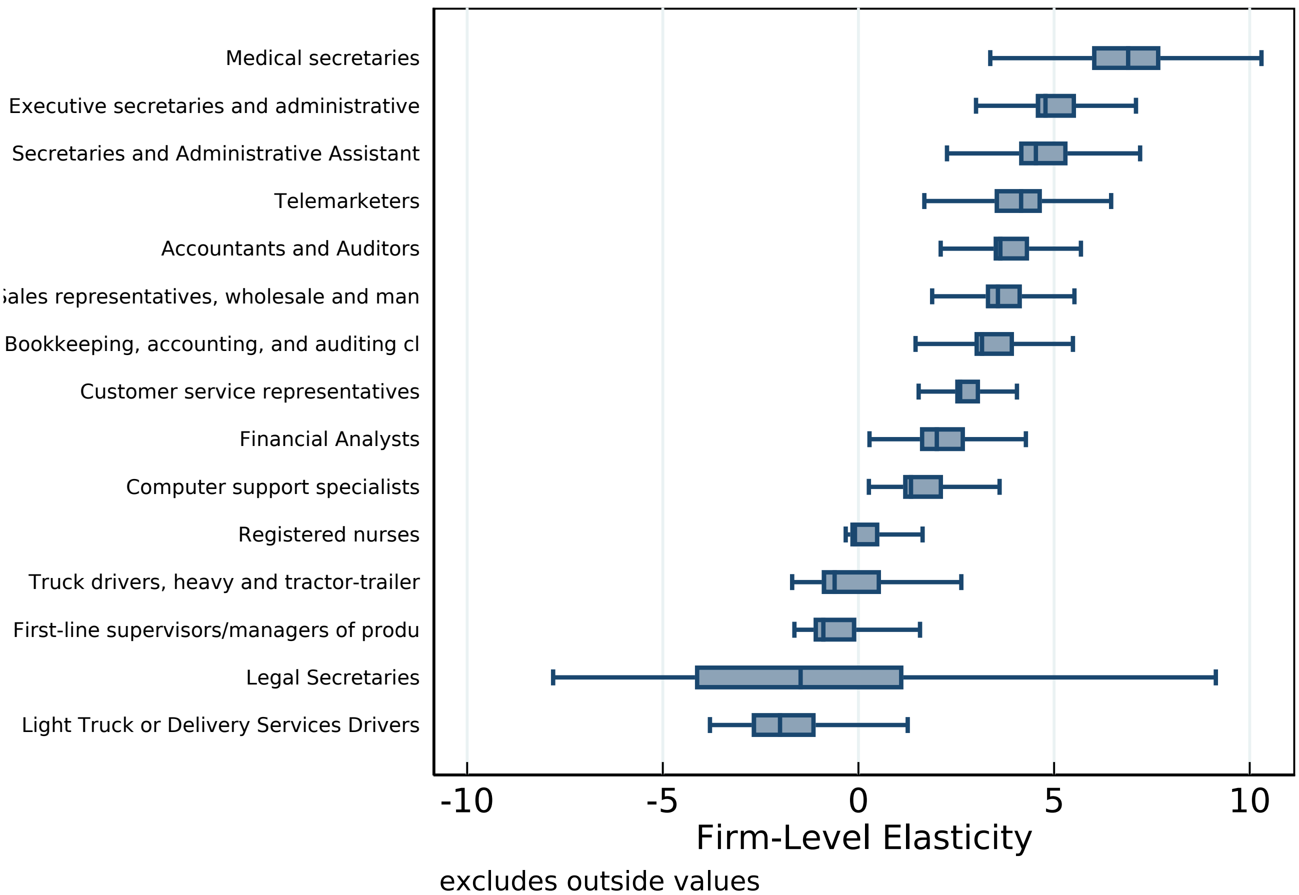

Figure 4. Distribution of Firm-Level Elasticities by 6-digit SOC. This figure shows the median firm-level elasticity of applications by 6-digit SOC for the "Full Model". The full model is allows the wage and nesting coefficients to vary for each 6-digit occupation, as well as a quadratic polynomial in log population density. The results are for the specification that includes both CZ × SOC fixed effects and job-title fixed effects, and instruments the wage using the average predicted wage of the same firm in other markets (excluding the same CZ and same SOC). The interactions of the wage are instrumented using interactions of the average predicted wage for the same firm in other markets, interacted with the corresponding variables. The figure excludes "Industrial Engineers" and the "Other SOCs" category (grouping SOCs with less than 100,000 observations at the application level) because the range of elasticities was too wide for visualization. Graphs with these categories are included in the appendix.

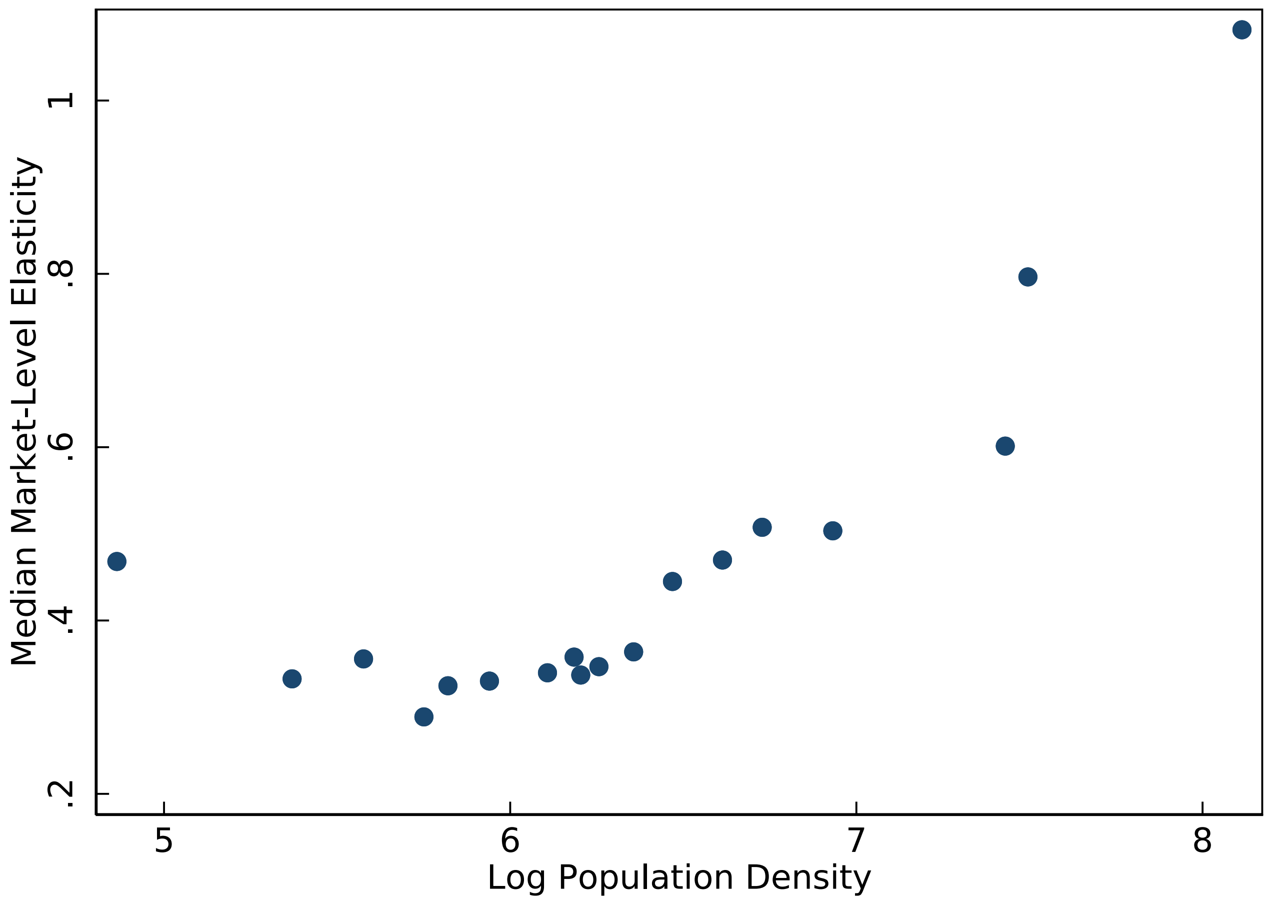

Figure 5. Median Market-Level Elasticities by Population Density. This figure shows the median market-level elasticity of applications by population density for the "Full Model". The full model is allows the wage and nesting coefficients to vary for each 6-digit occupation, as well as a quadratic polynomial in log population density. The results are for the specification that includes both CZ $\times$ SOC fixed effects and job-title fixed effects, and instruments the wage using the average predicted wage of the same firm in other markets (excluding the same CZ and same SOC). The interactions of the wage are instrumented using interactions of the average predicted wage for the same firm in other markets, interacted with the corresponding variables.

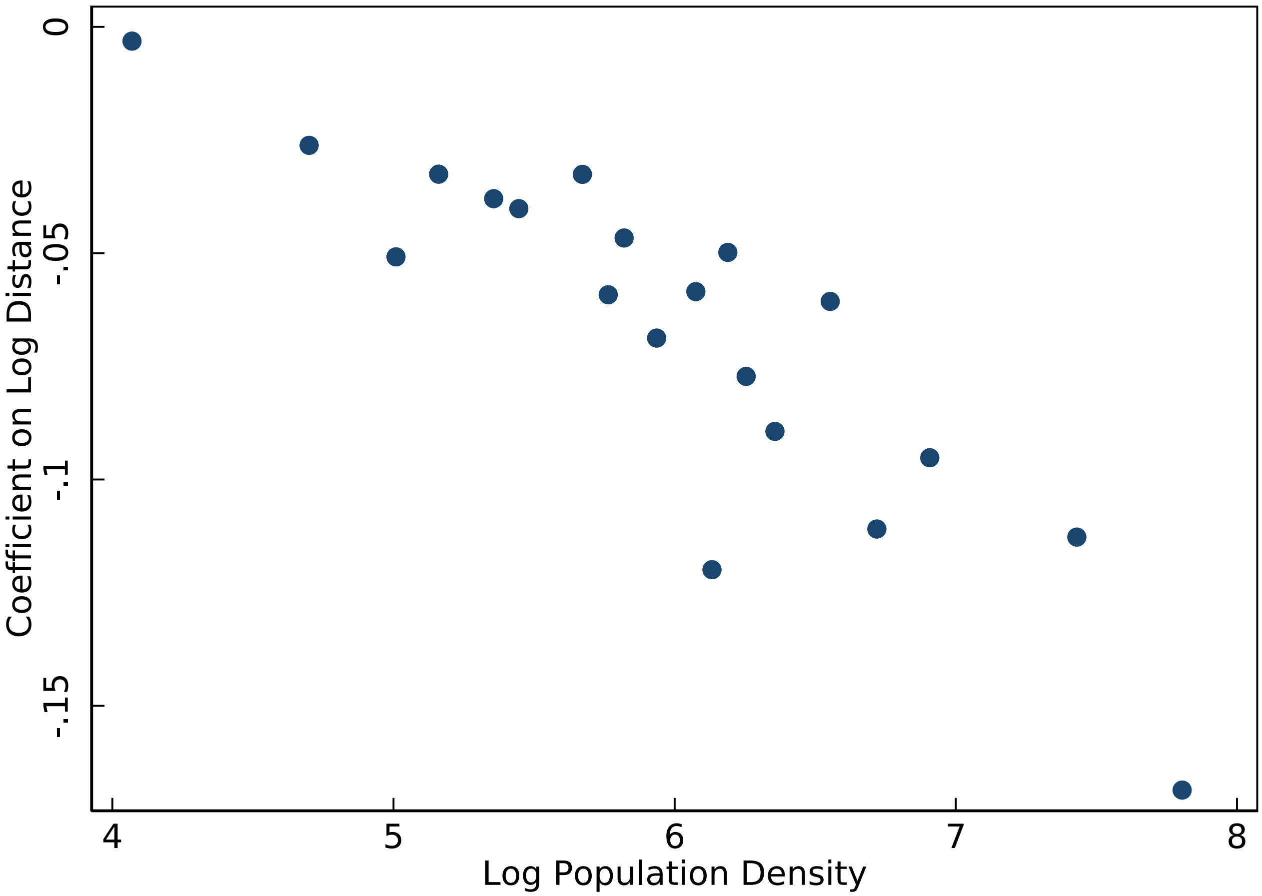

Figure 6. Distance coefficient by population density. This figure shows a binned scatter of the median distance coefficient (which varies by CZ × SOC market) by log population density. The distance coefficients are calculated using the nesting parameters from the full model, as the coefficients obtained from estimating the multinomial logit models in equation (3) for each market, times the corresponding nesting parameter based on the estimation of equation (4) using the full (i.e. most flexible) model.

A Appendix Figures and Tables

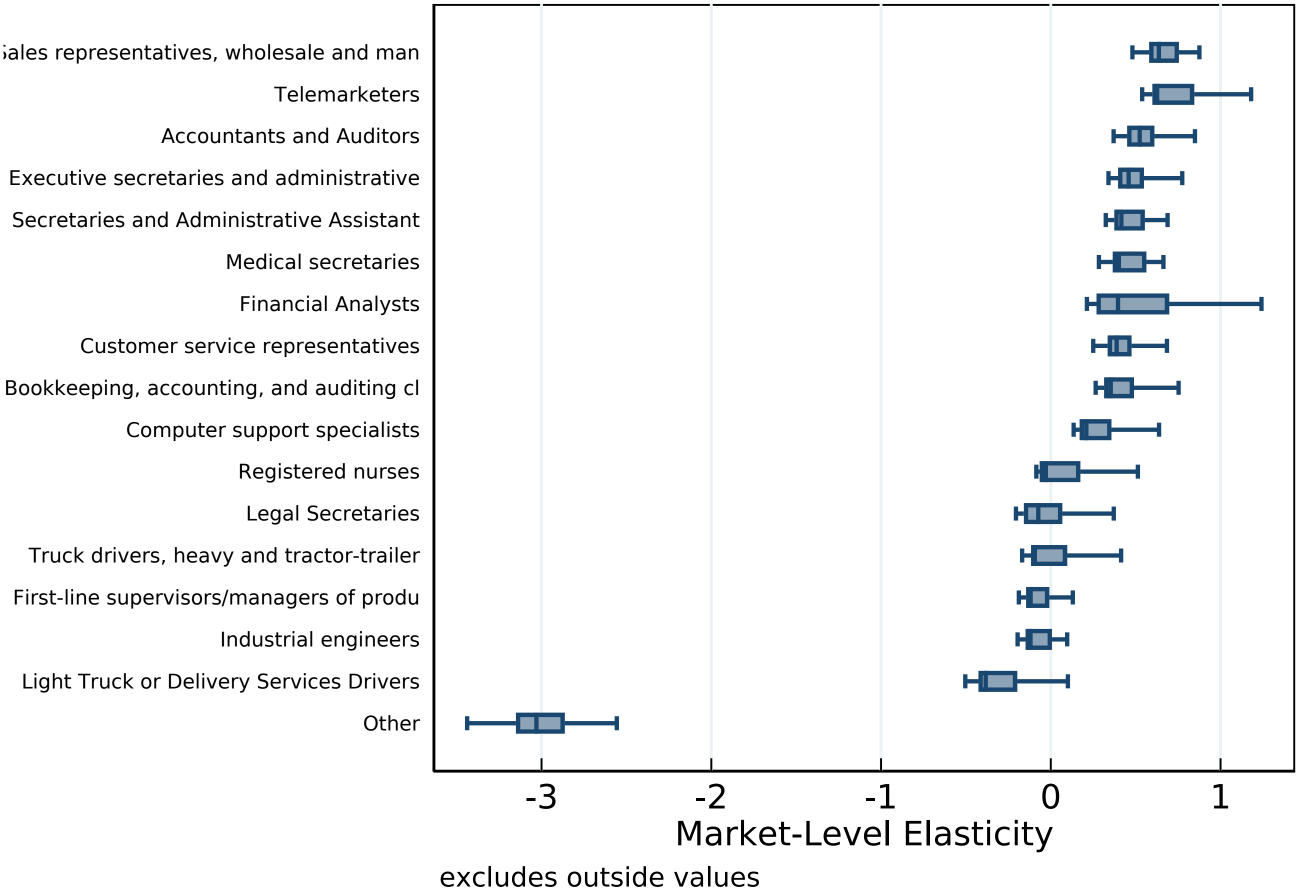

Figure 7. Distribution of Market-Level Elasticities by 6-digit SOC. This figure shows the median market-level elasticity of applications by 6-digit SOC for the "Full Model". The full model is allows the wage and nesting coefficients to vary for each 6-digit occupation, as well as a quadratic polynomial in log population density. The results are for the specification that includes both CZ $\times$ SOC fixed effects and job-title fixed effects, and instruments the wage using the average predicted wage of the same firm in other markets (excluding the same CZ and same SOC). The interactions of the wage are instrumented using interactions of the average predicted wage for the same firm in other markets, interacted with the corresponding variables.

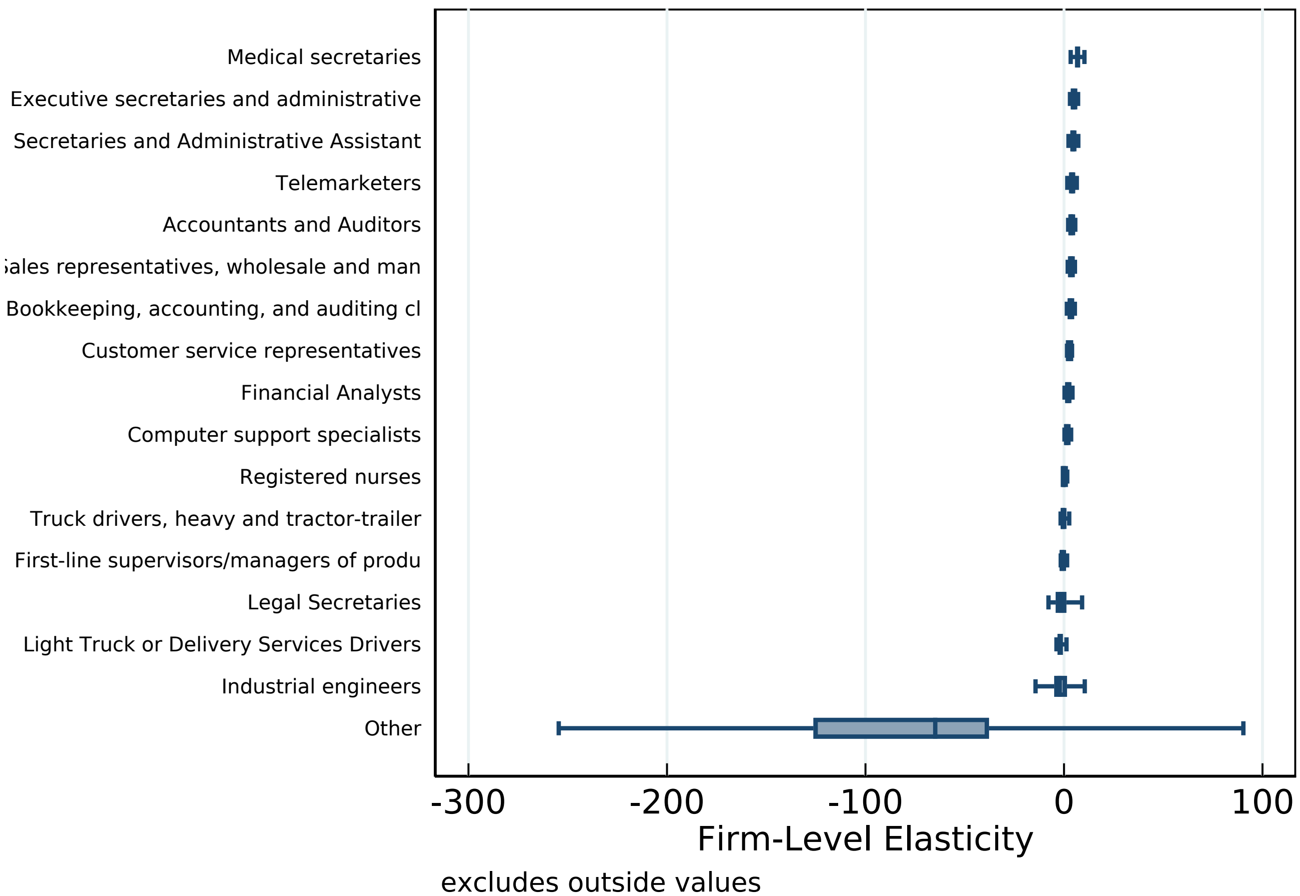

Figure 8. Distribution of Firm-Level Elasticities by 6-digit SOC. This figure shows the median firm-level elasticity of applications by 6-digit SOC for the "Full Model". The full model is allows the wage and nesting coefficients to vary for each 6-digit occupation, as well as a quadratic polynomial in log population density. The results are for the specification that includes both CZ $\times$ SOC fixed effects and job-title fixed effects, and instruments the wage using the average predicted wage of the same firm in other markets (excluding the same CZ and same SOC). The interactions of the wage are instrumented using interactions of the average predicted wage for the same firm in other markets, interacted with the corresponding variables.