12-19-2019

Essays on Imperfect Competition in the Labor Market

Sydnee Christian Caldwell
Massachusetts Institute of Technology

Essays on Imperfect Competition in the Labor Market

Sydnee Caldwell

This dissertation consists of three chapters that examine the impact of imperfect competition in the labor market on wage growth and wage inequality. Chapter 1, coauthored with Nikolaj Harmon, investigates the link between an individual's wages and her outside labor market opportunities. To overcome the fact that many factors that shift an individual's outside opportunities also impact her productivity at her current job, we develop a novel identification strategy that generates within-individual (and within-firm-by-occupation) variation in workers' information about their outside options. This strategy, which we implement using monthly employer-employee data from Denmark, exploits the fact that individuals often learn about job opportunities through their social networks. Using this strategy we find that changes in workers' outside labor market opportunities are reflected in both mobility and wage growth, even for individuals who do not switch firms.

The reduced-form results in Chapter 1 are inconsistent with perfect competition and provide empirical support for a large class of on-the-job search models where workers are able to leverage outside offers into wage increases at the incumbent firm (Cahuc, Postel-Vinay, and Robin 2006; Postel-Vinay and Robin 2002). The results in this chapter therefore suggest that policies that improve workers' ability to receive or accept job offers while employed can raise wages. These policies could include restrictions on firms' use of noncompete or nonsolicitation agreements or harsh punishments for firms' use of no-poach clauses. The results further suggest that industrial policies that promote "good jobs" may, through improving workers' outside options, benefit a broad set of workers—not simply the workers who get those jobs (Acemoglu 2001).

Chapter 2, coauthored with Oren Danieli, investigates the role that cross-sectional differences in individuals' outside options play in generating between-group wage inequality. We use a two-sided matching model to micro-found a measure of workers' outside options, which we call the outside options index. We then use German administrative data to estimate this index and use two sources of variation: 1) the introduction of high-speed trains and 2) a standard shift-share instrument to identify the elasticity between our index and wages. When we combine our measure of options with this elasticity, we find that roughly one-third of the gender wage gap in Germany can be explained by differences in options, mostly the result of differences in effective labor market size.

Chapter 3, coauthored with Emily Oehlsen, investigates whether, in the absence of commuting costs, monopsonistic firms have an incentive to pay women less. The chapter uses data from a series of experiments, conducted in collaboration with a national ride-share company, to estimate the elasticity of men's and women's labor supply to both the market (Frisch elasticities) and the firm. These elasticities are sufficient to calculate the firm's optimal gender wage gap in settings where hours are flexible and settings where they are not. We find that women are twice as elastic as men are to the market: in response to a 10 percent increase in wages, they work 7 percent more compared to 3 percent more. This does not reverse at the firm level; both men and women have firm-specific elasticities of between 2 and 4, and there is no evidence that women are less elastic than men. These elasticities suggest optimal wage-markdowns of between 20 and 33 percent. The results also indicate that, in the absence of commuting costs, firms have no incentive to pay equally productive women less than their male counterparts.

Together, the second and third chapters show that there are differences in the option sets that men and women face in the labor market, and that these differences contribute to the gender wage gap. However, these chapters also suggest that much of the outside options gap is the result of differences in willingness or ability to commute. To that end, policies that make it easier for women—particularly those with children—to commute or to access child care near their desired place of work may be an effective tool in combating the gender gap.

Chapter 1

Outside Options, Bargaining, and Wages: Evidence from Coworker Networks

(with Nikolaj Harmon)

There is growing evidence that imperfect competition and frictions in the labor market have a significant impact on the wage distribution (Barth et al. 2016; Card, Heining, and Kline 2013; Card et al. 2016). In such a labor market, workers' wages depend not only on their productivity, but on the characteristics of the firm for which they work and on the characteristics of the firms for which they could have worked.[1] A growing number of papers have documented that equally skilled workers earn different amounts depending on where they work (see, e.g. Barth et al. 2016; Card, Heining, and Kline 2013; Song et al. 2015). However, to date, there is little empirical evidence on the link between workers' outside options and their wages.[2] If two workers at a firm are equally productive, does the worker with better opportunities at other firms—or better information about these opportunities—earn more? Can workers renegotiate their wage with their current firm if they receive an outside offer?

The link between an individual's outside labor market opportunities and her wages is important both for distin-

guishing between different models of wage-setting and for understanding how recent developments in the labor market, including the use of no-poach and nonsolicitation agreements (which reduce workers' ability to receive offers while employed), will impact wages. However, examining this link empirically is challenging both because outside options are not observed in standard data sets and because most factors that shift workers' outside options also shift their productivity in their current job. This is a problem because changes in productivity at the incumbent firm should impact wages, even if the labor market is perfectly competitive.

This paper overcomes these challenges by combining a novel identification strategy that exploits changes in workers' information about their outside opportunities with rich administrative data that contain high-frequency (monthly) wage data and detailed measures of workers' skills. The empirical strategy is motivated by a large literature, pioneered by Granovetter, that documents how workers learn about job opportunities through their social networks (Granovetter 1973; Ioannides and Datcher Loury 2004; Topa 2011).[3] We create measures of a worker's information about outside opportunities by weighting firm-specific changes in labor demand by each worker's unique coworker network. These networks consist of the set of individuals a worker has worked with in the recent past but is no longer working with. They allow us to identify which new positions an individual is likely to hear about. Because networks vary across workers within an occupation, and even within a firm-and-occupation group, we are able to exploit differences in information between workers within a narrow skill group.

We implement this approach using a new monthly linked employer-employee database covering the universe of employees at Danish firms. While wages in Denmark were historically set by union bargaining, firms today have considerable latitude to negotiate wages with individual employees (Dahl, Le Marie, and Munch 2013). Our data cover the period after decentralization.[4] The data contain information on individuals' monthly earnings and hours worked, and on their six-digit industry and occupation.

We start by deriving our measure of outside options from a standard on-the-job search model where firms renegotiate wages with workers that receive outside offers (Postel-Vinay and Robin 2002; Cahuc, Postel-Vinay, and Robin 2006). The model allows us to illustrate the two key predictions of this class of models. First, workers who receive outside offers from more productive firms leave. Second, workers who receive outside offers from less-productive firms that dominate their current position renegotiate. We modify the model to allow workers to learn about job opportunities through both public sources and their individual-specific social networks to derive a measure of outside options that we can take to the data.

We then test the key predictions of this model by regressing indicators for mobility and measures of wage growth on our individual- and time-specific measures of outside options. Our baseline measure weights the number of new positions at each firm by an individual's exposure to that firm through their coworker network. The identifying assumption is that, conditional on the included covariates, unobserved determinants of individual mobility or wage growth are uncorrelated with time-varying labor demand at an individual's former coworkers' current firms. In order to focus on variation in outside options over time for a given worker, we include worker fixed effects in all of our specifications. We also control, nonparametrically, for month- and (four-digit) industry-specific demand shocks. The primary threat to validity, which we address through a series of distinct tests, is that the coworker networks proxy for specific types of skills, and that there are unobserved month-specific changes in demand for these skills, which are correlated with unobserved determinants of job-to-job mobility and wage growth.

We present nonparametric evidence that confirms both predictions of the theoretical model: changes in workers' information about their outside opportunities lead to mobility and wage growth, and larger changes are necessary to induce a job-to-job transition than to induce a wage change. Virtually all of the increased mobility is the result of moves to firms where the worker has a former coworker. This is consistent with the idea that workers learned about the opportunity through their former colleagues. We find that an additional 10 new positions at an individual's former coworkers' current firms results in a 15 percent higher probability that the worker makes a job-to-job transition that month.[5] The same change translates to an approximate $50 increase in earnings over the course of the year. However, most individuals do not renegotiate: the impact on whether an individual sees an earnings gain is less than a percentage point. If all the gains were associated with gains for workers who were driven to renegotiate (see a positive earnings change), the average full-time worker would see an 11 percent increase in base pay.

Both job-stayers and job-movers see higher wages in response to changes in their outside options; job-stayers obtain roughly 20 percent the earnings gain of job-movers. Posting models—such as monopsony models—would predict a ratio of zero: wages do not adjust unless the individual switches firms. Spot market models where wages freely fluctuate in response to changes in demand for a worker's skill would predict a ratio of one. We are able to reject both of these extremes.

Several distinct pieces of evidence suggest that our results are not driven by unobserved changes in demand for workers' skills. First, we show that the estimates are stable when adding more detailed nonparametric controls for changes in demand for different occupation or skill groups. These controls are based on different combinations of our industry, occupation, and education fixed effects. Second, we show that the results are also robust to adopting a within-firm

identification strategy that exploits variation in coworker networks that emerges from differences in tenure at the current firm and at past firms. The evidence is most consistent with worker-initiated renegotiation, not firm-initiated raises. If the earnings changes were the result of firms learning about the market price of their workers' skills, we would expect all workers within the same firm and occupation to see equal wage growth.

In order to further show that our results are driven by changes in workers' information, and not changes in workers' productivity, we decompose our measure of outside options into portions that come from different subsets of an individual's former coworkers. We find that the changes in earnings are driven by changes in labor demand at the firms of closely connected former coworkers: those who work in the same administrative region and those whom the individual worked with in the more recent past. Placebo tests exploiting an individual's future coworkers—coworkers the individual has not met yet—tell a similar story. If the results were driven by unobserved demand shocks, we would expect measures constructed using these coworkers to have a similar impact on mobility and wage growth. We would also expect that adding these measures as controls to our baseline regression would shrink our estimates.[6]

In the last part of the paper we divide workers into eight broad occupation groups and reestimate the effects within each group.[7] We find that the impact on workers in the highest skill group (professionals) is double that of workers in the middle skill group (technicians), and nearly five times that of workers in the least skilled group. Within each skill group, women benefit less than men. Because workers in higher skill groups also have higher baseline earnings—and men have higher baseline earnings than women—this heterogeneity translates into substantial differences in earnings. The heterogeneity does not appear to be driven by differences in the quality of our measure for different groups of workers: the impacts on mobility do not differ across groups.

To identify whether the heterogeneity across skill groups is due to differences in wage-setting strategies or differences in workers' bargaining power, we then use our reduced-form estimates to identify a structural search model incorporating on-the-job search and a mass of posting firms (Flinn and Mullins 2017). Intuitively, the impact on whether there is a wage change is informative about the fraction of firms that post wages; the impact on wages identifies workers' bargaining power. We estimate this model using a simulated method of moments. We find that the reduced-form heterogeneity is largely driven by differences in wage-setting strategies, not bargaining power: wage renegotiation (posting) is more common among high- (low-) skilled workers. Using our estimated parameters, we show that a reduction in the arrival rate for employed workers would lead to a significant reduction in wage growth. For high-skilled workers, much of this is due to decreased on-the-job bargaining; for lower-skilled workers this is mostly due to decreased mobility.

Chapter 2

Outside Options in the Labor Market

(with Oren Danieli)

In almost every model of the labor market, wages depend on a worker's outside options: the amount of compensation she could receive from different employers. In a perfectly competitive labor market, an equally attractive outside option always exists, and competition between identical employers sets compensation at the marginal product. However, in reality, a worker's next best option could require a different combination of her skills, could involve different working hours, or could be located in a different city. The number of outside options could be systematically lower for some workers because of the health of their local labor market, because they are unwilling or unable to commute, or because their skills are valuable only for a few employers or industries. Such differences could have significant implications for their incomes.

A key challenge for empirical research on this topic is that a worker's outside option set is not typically observed. Even within the same firm and occupation, workers may face different options because of their specific set of skills, their preferences, or their constraints. As a result, little is known about which workers have better outside options and what role options play in generating wage inequality.

The first contribution of this paper is to develop an empirical procedure to uncover a key latent parameter in most wage-setting models: the value of an individual's option set. We show how this latent parameter can be derived from the cross-sectional concentration of similar workers across jobs. If similar workers are concentrated in a certain region, industry, occupation, or other job characteristics, then the worker's options are more limited. We quantify this concentration in a single "outside options index" (OOI). We show that in a matching model of heterogeneous workers and jobs, this OOI is a sufficient statistic for the effect of outside options on compensation, when holding productivity constant. We then estimate the OOI for every worker using administrative matched employer-employee data from a 1 percent representative sample of workers in Germany. Examining the distribution of the OOI, we find which workers' characteristics are associated with better outside options. Next, we quantify the impact on wages by estimating the elasticity between the OOI and wages using two quasi-random sources of variation in the OOI, which holds workers' productivity constant: the introduction of high-speed commuter-rails and a shift-share ("Bartik") instrument.

Our second contribution is to show that differences in outside options explain a substantial portion—30 percent—of the gender wage gap in Germany. This gender difference is driven entirely by differences in willingness to commute or move.

We start by outlining a static model of the labor market that illustrates how, with two-sided heterogeneity, differences in outside options lead to differences in compensation, even for equally productive workers. Our model is based on the classic Shapley and Shubik (1971) assignment game—a two-sided matching model with transfers. Compensation in this setting is set to prevent workers from moving to their outside options; because of heterogeneity, this will be below their full productivity in the first-best option.

We derive a sufficient statistic from this model, the OOI, which summarizes the impact of options on compensation. It measures the quantity of relevant jobs for a given worker. If a worker gets access to more similar jobs, their compensation would increase by exactly the increase in OOI times a constant elasticity, even though their productivity remains constant. The OOI depends on two factors: the supply of jobs and worker flexibility (i.e., a worker's ability or willingness to take jobs in more places, more occupations, or more industries). Workers with more relevant jobs, as captured by the OOI, will have on average a better outside option, and will also be able to sort into better matches, conditional on their productivity.

We show that the OOI is equal to a standard concentration index: workers with more options are those who, in equilibrium, are found in a greater variety of jobs. Under standard assumptions on the distribution of match quality (Choo and Siow 2006; Dupuy and Galichon 2014), the OOI is equal to the entropy index. This index, with a negative sign, is used in the industrial organization literature as a measure of market concentration (Tirole 1988), similar to the Herfindhal-Hirschman Index (HHI), which has also been used to measure concentration in labor markets (Azar, Marinescu, and Steinbaum 2017; Benmelech, Bergman, and Kim 2018). In contrast to most concentration indices, our index is not measured on a specific dimension, such as occupation or industry. Instead, workers' and firms' characteristics are allowed to vary continuously; this allows for the fact that some workers may have employment opportunities in different occupations or industries or in different geographic areas. Options are estimated in equilibrium, based on matches we observe in cross-sectional data. To isolate the effect of more options from the effect of productivity, the OOI is calculated without using any information on wages or wage offers.

We then develop a method to estimate the OOI for each worker in the labor market, which is computationally feasible even in large data sets. The OOI is a function of the joint probability of every worker to be in every job. Our method estimates this probability using the cross-sectional distribution of similar workers. We show that this problem can be translated into a logistic regression framework. We then use the fast implementation of logistic regressions to estimate the probabilities for every worker-job combination. From those probabilities we can directly calculate the OOI for each worker.

We use the OOI to analyze the impact of outside options on inequality, starting with identification of which workers have better outside options. Specifically, we estimate the OOI for every worker in a representative sample of German workers in 2014 using administrative linked employer-employee data. In order to validate our measure, we show that the OOI predicts which workers are less affected by a mass layoff: workers with better outside options recover more quickly from a displacement. Because we do not use wages to calculate the OOI, there is not a mechanical link between the OOI and wages.

We use two sources of variation in options to estimate the elasticity between the OOI and wages: 1) the introduction of high-speed commuter rail stations (Heuermann and Schmieder 2018), and 2) a standard industry shift-share ("Bartik") instrument (Beaudry, Green, and Sand 2012). These instruments allow us to identify the elasticity between our OOI and wages. The first instrument uses the introduction of new commuter rail in small German towns. These stations, located along existing routes, effectively increased the labor market size for workers in small cities who happened to live on routes between major German cities. Prior work showed that the exact choice of town was largely driven by political considerations (Heuermann and Schmieder 2018). The second instrument uses differences in exposure to industry growth trends between local labor markets. We compare workers who work in the same industry but have outside options in different industries because they reside in different parts of the country. Both instruments yield a similar semielasticity of roughly 0.17–0.32 between the OOI and wages. Further, while the new train stations have the largest impact on the options of highly skilled workers, the elasticity does not vary across education groups.[8]

Combining this elasticity with the estimated distribution of the OOI, we find that differences in outside options lead to wage inequality. Differences in options lower compensation for women (immigrants) by 6 (8) percentage points. This explains roughly 30 percent (88 percent) of the overall gap in Germany. We also find large effects that graduates of higher-secondary education have greater outside options, which increases their compensation by 7 percentage points (one-fourth of the total return).[9]

Finally, the last part of the paper examines why certain groups of workers have better options than others. We find that heterogeneity in the ability or willingness to commute can account for the full gender gap in outside options. We also find that without their higher willingness to work at more distant jobs, highly educated workers would actually have fewer options. This is likely because their skills tend to be more industry specific.

Chapter 3

Monopsony and the Gender Wage Gap: Experimental Evidence from the Gig Economy

(with Emily Oehlsen)

The third chapter also investigates the relationship between imperfect competition in the labor market and between-group wage inequality. When the labor market is not perfectly competitive, firms are not price takers: in order to recruit or retain more workers, they must offer higher wages (see surveys in Ashenfelter, Farber, and Ransom [2010]; Bhaskar, Manning, and To [2002]; Boal and Ransom [1997]; Manning [2003]). Firms have an incentive to pay higher wages to workers that are harder to recruit or retain (workers that are more elastic to the firm), even if they are no more productive than other workers.

The idea that imperfect competition in the labor market could lead to a gender wage gap dates back to Joan Robinson's 1933 book, in which she coined the term monopsony. Women may earn less than men if they are, on average, less willing to leave their employer in response to changes in firm and market conditions (Card et al. 2016).[10] This could be true if women face smaller effective labor markets due to discrimination or commuting costs, or if they are less aggressive or strategic about taking advantage of new opportunities (Babcock and Laschever 2009; Card, Cardoso, and Kline 2016).

In this chapter we use data from a series of randomized experiments conducted at Uber to produce new evidence on the elasticity of men and women's labor supply, to both the firm and the market. We use these elasticities to test whether gender differences in firm-specific elasticities might contribute to a gender wage gap. The key advantage of the Uber setting is that we are able to generate exogenous variation in wages in a setting where it is clear what the alternative firm is: Lyft.

The first part of the chapter outlines a theoretical model that allows workers to adjust both how much they work (participation and hours) and for whom (firm substitution). The model illustrates that when hours are flexible, the optimal wage markdown depends on both the traditional firm substitution/recruitment elasticity and how responsive workers' total hours are to changes in wages. The first elasticity measures the extent to which workers join or leave individual firms in response to changes in relative wages. The second measures the extent to which workers increase their overall labor supply (at the expense of leisure) in response to wage changes. We show that these two sets of elasticities are sufficient to calculate the optimal wage gap in a setting where hours are flexible and in a setting where they are not.

We structured our experiments so that we were able to estimate each of these elasticities. In these experiments we offered random subsets of male and female drivers the opportunity to have 25–39 percent higher wages for a week. While some drivers had access to a competing ride-share company (Lyft), others did not. Some of this variation comes from the fact that Lyft temporarily left the Houston market. We conducted one experiment while Lyft was out of this market and another experiment after Lyft had returned. Some of this variation comes from the fact that some Uber drivers are, because of the age of their vehicle, ineligible to drive for Lyft. Drivers in our experiments received offers by e-mail and text message, as well as through the Uber application itself. Because drivers were required to opt-in in order to receive the wage increase, we are able to account for inattention as a possible confounder (Mas and Pallais 2018).

We first use data from the experiment conducted when Lyft was out of the market to estimate Frisch elasticities for men and women. While these elasticities serve as a baseline for our analysis of firm substitution, they are also of independent interest as they are a key component of most business cycle models. And despite the large volume of research on male and female labor supply, there is little quasi-experimental or experimental evidence that intensive or extensive margin Frisch elasticities differ by gender (Killingsworth and Heckman 1986; McClelland and Mok 2012).[11]

We find that women have Frisch (market-level) elasticities double those of men. In response to a 10 percent increase in wages, female drivers work 7 percent more hours ($\varepsilon = 0.7$) while male drivers work only 3 percent more hours ($\varepsilon = 0.3$). The results are not driven by baseline differences in usual hours worked or by differences in age. Our estimate of the Frisch elasticity for men is similar to the estimates presented in prior studies of taxi drivers (Farber 2005, 2015), but is somewhat smaller than estimates in similar experiments (Fehr and Goette 2007). We argue that this may be due, in part, to the fact that it is typically difficult to observe part-time workers shifting hours across firms or platforms. The extensive margin elasticities are modest, even among our sample of marginally attached drivers. In response to a 10 percent increase in wages, women are at most 2 percentage points more likely to drive (an elasticity of at most 0.18), relative to a single percentage point for men (at most 0.09).[12]

To assess firm substitution, we then compare these market-level Frisch elasticities to estimates from two similar experiments where only a subset of the drivers cannot drive for Lyft. Because the estimates from these experiments are not precise, we also use data from a large Uber-run promotion we call the "individual driver bonus" to corroborate our findings. Using both sets of data we find that drivers with the opportunity to work for competing platforms are significantly more elastic. This likely reflects the fact that some of their hours do not come from leisure but from Lyft. The gap between the elasticities for drivers who could and could not access Lyft is particularly pronounced for younger drivers, who are likely more adept with the technologies Uber and Lyft use.

We use the elasticities for drivers who could and could not drive for Lyft to compute implied firm substitution elasticities for men and women. We find mean elasticities between two and four. These estimates are in line with other recent estimates of firm-specific elasticities, and suggest optimal markdowns of between 20 and 33 percent.[13] Our low elasticities reflect the fact that, even in this setting, switching between firms is not trivial.

In contrast to prior nonexperimental work, we do not see any significant differences in firm-specific elasticities between men and women (Hirsch, Schank, and Schnabel 2010; Ransom and Oaxaca 2010; Webber 2016). Our results suggest that, even if gig economy firms wield monopsony power, they do not have any incentive to pay women less. This is true even in a world where hours are not flexible.

We view these estimates as a lower bound on the extent to which monopsonistic firms outside of the gig economy might be incentivized to pay women less than men. Our results suggest that women are no less strategic about taking advantage of the opportunity to earn higher wages. However, in other contexts, women may face higher commuting costs or constraints, which could result in lower firm-specific elasticities, and thus lower wages. The results in Chapter 2 suggest that the impact of these constraints can be sizable.

Notes

1. This is explicit in models where wages are determined by bargaining between an individual and a firm or a union and a firm (Acemoglu 2001; Farber 1987; Pissarides 2000). It is implicit in models with posting; in these models, the wage a firm chooses to post depends on the wages chosen by other firms (Burdett and Mortensen 1998; Manning 2003).
2. Prior work has largely exploited industry- and region-level variation (see, e.g., Beaudry, Green, and Sand 2012; Bidner and Sand 2016; Fortin and Lemieux 2015; Hagedorn and Manovskii 2013). Contemporaneous work by Jäger et al. (2018) shows that there is no link between the value of nonemployment and wages.
3. Similar facts were presented in prior work by Myers and Shultz (1951), Rees (1966), and Rees and Shultz (1970).
4. Our data cover the period 2008–2016. Most wage decentralization occurred in the 1990s.
5. Because our data are monthly, the base rate is low: roughly 1 percent of workers make a job-to-job transition each month.
6. In a separate set of robustness checks we show that we obtain similar qualitative results when exploiting measures that are based on changes in world demand for the products exported by an individual's former coworkers' firms (Garin and Silverio 2018; Hummels et al. 2014).
7. The eight groups are 1) managers, 2) professionals, 3) technicians and associate professionals, 4) clerical support workers, 5) service and sales workers, 6) craft and related trade workers, 7) plant and machine operators, and 8) assembly workers.
8. The fact that highly educated workers' options are most affected is not surprising; tickets on these trains are fairly expensive.
9. The level that grants a certificate allowing college admission.
10. Similarly, search models predict that workers with lower arrival rates of job offers earn less in equilibrium (Black 1995).
11. While a few studies have exploited temporary wage variation in settings where workers can freely choose their hours, the populations in these studies are predominantly male (Farber 2005, 2015; Fehr and Goette 2007; Oettinger 1999; Stafford 2015). Though most (more than 85 percent) of Uber drivers are male, we structured our experiment to include roughly equal numbers of male and female drivers and to include both part- and full-time drivers (Hall and Krueger 2015).
12. These elasticities are significantly smaller than those typically used to calibrate dynamic models; these models typically assume an elasticity greater than 1 (Chetty et al. 2013; King and Rebelo 1999).
13. Dube et al. (forthcoming) use a bunching estimator to derive labor supply elasticities from administrative wage data and the CPS. They report estimates of two and three (Table 3, Panel B) for moderate values of optimization frictions.

References

Acemoglu, Daron. 2001. "Good Jobs Versus Bad Jobs." *Journal of Labor Economics* 19(1): 1–21.

Ashenfelter, Orley C., Henry Farber, and Michael R. Ransom. 2010. "Labor Market Monopsony." *Journal of Labor Economics* 28(2): 203–210.

Azar, José, Ioana Marinescu, and Marshall I. Steinbaum. 2017. *Labor Market Concentration*. Technical report. Cambridge, MA: National Bureau of Economic Research.

Babcock, Linda, and Sara Laschever. 2009. *Women Don't Ask: Negotiation and the Gender Divide*. Princeton, NJ: Princeton University Press.

Barth, Erling, Alex Bryson, James C. Davis, and Richard Freeman. 2016. "It's Where You Work: Increases in Earnings Dispersion across Establishments and Individuals in the US." *Journal of Labor Economics* 34: S67–S97.

Beaudry, Paul, David A. Green, and Benjamin Sand. 2012. "Does Industrial Composition Matter for Wages? A Test of Search and Bargaining Theory." *Econometrica* 80(3): 1063–1104.

Benmelech, Efraim, Nittai Bergman, and Hyunseob Kim. 2018. *Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?* Technical report. Cambridge, MA: National Bureau of Economic Research.

Bhaskar, Venkataraman, Alan Manning, and Ted To. 2002. "Oligopsony and Monopsonistic Competition in Labor Markets." *Journal of Economic Perspectives* 16(2): 155–174.

Bidner, Chris, and Ben Sand. 2016. "Job Prospects and Pay Gaps: Theory and Evidence on the Gender Gap from US Cities." Discussion Paper No. 16-14. Burnaby, British Columbia: Simon Fraser University.

Black, Dan A. 1995. "Discrimination in an Equilibrium Search Model." *Journal of Labor Economics* 13(2): 309–334.

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

Burdett, Kenneth, and Dale T. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39(2): 257–273.

Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin. 2006. "Wage Bargaining with On-The-Job Search: Theory and Evidence." *Econometrica* 74(2): 323–364.
Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women." *Quarterly Journal of Economics* 131(2): 633–686.
Card, David, Ana Rute Cardoso, Jörg Heining, and Patrick Kline. 2016. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics* 53(1): S13–S70.
Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *Quarterly Journal of Economics* 128(3): 967–1015.
Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber. 2013. "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-analysis of Extensive Margin Elasticities." *NBER Macroeconomics Annual* 27(1): 1–56.
Choo, Eugene, and Aloysius Siow. 2006. "Who Marries Whom and Why." *Journal of Political Economy* 114(1): 175–201.
Dahl, Christian M., Daniel Le Maire, and Jakob R. Munch. 2013. "Wage Dispersion and Decentralization of Wage Bargaining." *Journal of Labor Economics* 3 (3): 501–533.
Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri. Forthcoming. "Monopsony in Online Labor Markets." *American Economic Review: Insights*.
Dupuy, Arnaud and Alfred Galichon. 2014. "Personality Traits and the Marriage Market," *Journal of Political Economy* 122(6): 1271–1319.
Farber, Henry S. 1987. "The Analysis of Union Behavior." In *Handbook of Labor Economics*, Vol 2. Orley Ashenfelter and R. Layard, eds. Amsterdam: Elsevier, pp. 1039–1089.
———. 2005. "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers." *Journal of Political Economy* 113(1): 46–82.
———. 2015. "Why You Can't Find a Taxi in the Rain and Other Labor Supply Lessons from Cab Drivers." *Quarterly Journal of Economics* 130(4): 1975–2026.
Fehr, Ernst, and Lorenz Goette. 2007. "Do Workers Work More If Wages Are High? Evidence from a Randomized Field Experiment." *American Economic Review* 97(1): 298–317.
Flinn, Christopher, and Joseph Mullins. 2018. "Firms' Choice of Wage-Setting Protocols in the Presence of Minimum Wages." Working paper. Minneapolis: University of Minnesota.
Fortin, Nicole M., and Thomas Lemieux. 2015. "Changes in Wage Inequality in Canada: An Inter-provincial Perspective." *Canadian Journal of Economics/Revue Canadienne D'économique* 48(2): 682–713.
Garin, Andrew, and Filipe Silverio. 2018. "How Does Firm Performance Affect Wages? Evidence from Idiosyncratic Export Shocks." Harvard University Working Paper. Cambridge, MA: Harvard University.
Granovetter, Mark. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78(6): 1360–1380.
———. 1995. *Getting a Job: A Study of Contacts and Careers*. Chicago: University of Chicago Press.
Hagedorn, Marcus, and Iourii Manovskii. 2013. "Job Selection and Wages over the Business Cycle." *American Economic Review* 103(2): 771–803.
Hall, Jonathan V., and Alan B Krueger. 2015. "An Analysis of the Labor Market for Uber's Driver-Partners in the United States." Princeton University Industrial Relations Section Working Paper No. 587. Princeton, NJ: Princeton University.
Heuermann, Daniel F., and Johannes F. Schmieder. 2018. "The Effect of Infrastructure on Worker Mobility: Evidence from High-Speed Rail Expansion in Germany." *Journal of Economic Geography* 19(2): 335–372.
Hirsch, Boris, Thorsten Schank, and Claus Schnabel. 2010. "Differences in Labor Supply to Monopsonistic Firms and the Gender Pay Gap: An Empirical Analysis Using Linked Employer-Employee Data from Germany." *Journal of Labor Economics* 28(2): 291–330.
Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang. 2014. "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data." *American Economic Review* 104(6): 1597–1629.
Ioannides, Yannis M., and Linda Datcher Loury. 2004. "Job Information Networks, Neighborhood Effects, and Inequality." *Journal of Economic Literature* 42(4): 1056–1093.
Jäger, Simon, Benjamin Schoefer, Sammy Young, and Josef Zweimulle. 2018. "Wages, Bargaining, and the Nonemployment Option: Evidence from Unemployment Insurance." Paper presented at conference, "Wages, Bargaining, and the Nonemployment Option: Evidence from Unemployment Insurance," held at University of California, Berkeley, April 5.
Killingsworth, Mark R., and James J. Heckman. 1986. "Female Labor Supply: A Survey." In *Handbook of Labor Economics*, Vol. 1. Orley Ashenfelter and R. Layard, eds. Amsterdam: Elsevier, pp. 103–204.
King, Robert G., and Sergio T. Rebelo. 1999. "Resuscitating Real Business Cycles." In *Handbook of Macroeconomics*, Vol. 1. J.B. Taylor and M. Woodford, eds. Amsterdam: Elsevier, pp. 927–1007.
Manning, Alan. 2003. *Monopsony in Motion*. Princeton, NJ: Princeton University Press.
Mas, Alexandre, and Amanda Pallais. 2018. "Labor Supply and the Value of Non-Work Time: Experimental Estimates from the Field." *American Economic Review: Insights*. 1(1): 111–126.
McClelland, Robert, and Shannon Mok. 2012. "A Review of Recent Research on Labor Supply Elasticities." Working Paper No. 2012-12. Washington, DC: Congressional Budget Office.
Myers, Charles A., and George P. Shultz. 1951. *The Dynamics of a Labor Market*. New York: Prentice-Hall.
Oettinger, Gerald S. 1999. "An Empirical Analysis of the Daily Labor Supply of Stadium Vendors." *Journal of Political Economy* 10(2): 360–392.
Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory*. Cambridge, MA: Massachusetts Institute of Technology Press.
Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica* 70(6): 2295–2350.
Ransom, Michael R., and Ronald L. Oaxaca. 2010. "New Market Power Models and Sex Differences in Pay." *Journal of Labor Economics* 28(2): 267–289.
Rees, Albert. 1966. "Information Networks in Labor Markets." *American Economic Review* 56(1/2): 559–566.
Rees, Albert, and George Pratt Shultz. 1970. *Workers and Wages in an Urban Labor Market*. Chicago: University of Chicago Press.
Robinson, Joan. 1933. *The Economics of Imperfect Competition*. London: Palgrave Macmillan.

Shapley, Loyd S., and Martin Shubik. 1971. “The Assignment Game I: The Core.” *International Journal of Game Theory* 1(1): 111–130.
Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. 2015. *Firming Up Inequality*. Technical report. Cambridge, MA: National Bureau of Economic Research.
Stafford, Tess M. 2015. “What Do Fishermen Tell Us That Taxi Drivers Do Not? An Empirical Investigation of Labor Supply.” *Journal of Labor Economics* 33(3): 683–710.
Tirole, Jean. 1988. *The Theory of Industrial Organization*. Cambridge, MA: Massachusetts Institute of Technology Press.
Topa, Giorgio. 2011. “Labour Markets and Referrals.” In *Handbook of Social Economics*, Vol. 1. Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, eds. Amsterdam: Elsevier, pp. 1193–1221.
Webber, Douglas. 2016. “Firm-Level Monopsony and the Gender Pay Gap.” *Industrial Relations* 55(2): 323–345.