Essays on Imperfect Competition in the Labor Market

by

Sydnee Christian Caldwell

B.A., University of California, Berkeley (2012)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2019

Signature redacted

Author ....

Department of Economics
May 1, 2019

Signature redacted

Certified by ....

Daron Acemoglu
Elizabeth and James Killian Professor of Economics
Thesis Supervisor

Signature redacted

Certified by ....

Joshua Angrist
Ford Professor of Economics
Thesis Supervisor

Signature redacted

Certified by ....

David Autor
Ford Professor of Economics
Thesis Supervisor

Signature redacted

Accepted by ....

Ricardo Caballero
Chairman, Department Committee on Graduate Theses

Essays on Imperfect Competition in the Labor Market

by

Sydnee Christian Caldwell

Submitted to the Department of Economics
on May 1, 2019, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis consists of three chapters on imperfect competition in the labor market. The first chapter (joint with Nikolaj Harmon) explores the relationship between an individual's wages and the quality of her opportunities at other firms (her outside options). 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 linked employer-employee data from Denmark, exploits the fact that individuals often learn about labor market opportunities through their social networks.

We find that increases in labor demand at former coworkers' current firms increases an incumbent worker's job-to-job mobility and wage growth. Consistent with theory, larger changes are necessary to induce a job-to-job transition than to induce a wage gain. Tests that exploit within-firm or within-firm-and-occupation variation and tests that exploit different subsets of an individual's former coworkers confirm that the results are not driven by unobserved changes in demand for workers' skills. Finally, we use our reduced form moments to identify a structural search model incorporating both posting and bargaining firms. We find that bargaining is more prevalent among high skilled workers.

The second chapter (joint 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" (OOI). The index is similar to those used in the industrial organization literature to measure concentration (e.g. the Herfindal-Hirschman Index, the HHI). We then use German administrative data to estimate this index and use two sources of quasi-random 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 these two ingredients, we find that roughly 1/3 of the gender wage gap in Germany can be explained by differences in options, mostly the result of differences in effective labor market size (commuting costs).

The third chapter (joint with Emily Oehlsen) asks whether, in the absence of commuting costs, firms with market power have an incentive to pay women less than men. We use data from a series of experiments at Uber where we offered random subsets of male and female drivers higher "wages". Drivers varied both in the size of the wage increase and in whether they could drive for Uber's main competitor, Lyft. These two sources of variation allow us to experimentally identify: (1) Frisch elasticities and (2) firm substitution elasticities. We find that women have Frisch elasticities double those of men on both the intensive and extensive margin. However, unlike the prior literature, we find that women are not less likely to shift between firms in response to changes in relative wages. The results suggest that, at least in the gig economy, firms have little incentive to wage discriminate between men and women based on their labor supply choices.

JEL Codes: J00, J31, J42

Thesis Supervisor: Daron Acemoglu
Title: Elizabeth and James Killian Professor of Economics

Thesis Supervisor: Joshua Angrist
Title: Ford Professor of Economics

Thesis Supervisor: David Autor
Title: Ford Professor of Economics

Acknowledgments

I am indebted to my advisors Daron Acemoglu, Josh Angrist, and David Autor for their guidance and encouragement over the past six years.

I could not have hoped for a better main advisor than Daron. Daron sets an impossible standard for what it means to be a researcher, teacher, advisor, and mentor. I came out of every conversation with him more excited about research than before; every paper I have written is measurably better as a result of his advice. I am incredibly grateful for Daron's support and patience over the past several years.

I was very fortunate to have had the opportunity to learn from and work with Josh Angrist. One of the lucky breaks I got at MIT was when Josh asked me if I would like to work with him on a new project at Uber. I learned a lot from Josh about how to design and run an experiment, how to write a paper, and how to ski.

David Autor was the first to contact me after I was admitted to MIT, and has been a constant source of advice ever since. I most appreciate how thoughtful David has been about giving me opportunities to grow as an economist. I am particularly grateful that he gave me my first RA job at MIT and that he asked me to be his TA for 14.662 (and let me prepare my own lecture!).

I am also very thankful to my undergraduate advisor, David Card, who introduced me to economics. Most of what I know about labor economics I learned during long talks with Dave in the CLE. This dissertation would not exist if it were not for his advice and moral support.

My fellow graduate students made the PhD a much more enjoyable experience. There are too many to name, but this list especially includes: Sarah Abraham, Alex Bartik, Josh Dean, Ellora Derenoncourt, Alex Fahey, Alex He, Peter Hull, Cara Nickolaus, Brendan Price, Vira Semenova, Elizabeth Setren, Mahvish Shaukat, Marco Tabellini, Roman Andres Zarate, Vira Semenova, David Silver, and Ariel Zucker. Some of my friends were also coauthors: Oren Danieli, Nikolaj Harmon, Scott Nelson, Emily Oehlsen, and Daniel Waldinger deserve extra thanks.

Outside of MIT, I was lucky to count on the support of my family and friends. Danica Shei was always available to help put things in perspective. Holly Teliska and Ivy, Olive, and Erik Hardenbergh deserve an extra thank you for giving me a quiet place to write early drafts of my job market paper. I am most grateful to my husband, Brendan, for endless love and support over the past ten years. Thanks for following me from Berkeley to Cambridge – and back again.

I thank the National Science Foundation and Brendan Bulik-Sullivan for financial support.

This thesis is dedicated to the memory of Cara Nickolaus, a kind friend and brilliant economist.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

List of Figures

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

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

[1] JOINT WITH NIKOLAJ HARMON

1.1 Introduction

There is growing evidence that imperfect competition and frictions in the labor market have a significant impact on the wage distribution (Card et al., 2013; Barth et al., 2016; Card et al., 2016c). In such a labor market, workers' wages depend not only on their productivity, but on the character-

[1]Caldwell thanks Daniel le Maire for many generous discussions regarding the data used in the paper. Caldwell also thanks Alex Bartik, John Coglianese, Oren Danieli, Joshua Dean, Ellora Derenoncourt, Amy Finkelstein, Victoria Gregory, Jonathan Hazell, Alex He, Peter Hull, Simon Jäger, Lawrence Katz, Patrick Kline, Joseph Mullins, Scott Nelson, Alan Olivi, Harry Pei, Raffaele Saggio, David Silver, Isaac Sorkin, Marco Tabellini, Will Rafey, John van Reenen, Daniel Waldinger, Jaya Wen, Román Andrés Zárate, Nathan Zorzi, and seminar participants at Berkeley, Booth, Columbia, Harvard, HBS, Kellogg, Michigan, MIT, Stanford GSB, and Wharton for thoughtful discussions and comments. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1122374 (Caldwell). Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

istics of the firm they work at and on the characteristics of the firms they could have worked at.[2] However, to date, there is little empirical evidence on the link between workers' outside options and their wages. 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 options and her wages is important both for distinguishing between different models of wage-setting and for understanding how recent developments in the labor market, including the use of no-poach agreements and the rise in labor market concentration, will impact wages. However, examining this link empirically is challenging both because outside options are not observed in standard datasets 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 that workers learn about job opportunities through their social networks (Granovetter, 1973; Ioannides and Datcher Loury, 2004).[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 the same occupation, and even within the same firm-and-occupation group, we are able to exploit differences in information

[2]This is explicit in models where wages are determined by bargaining between an individual and a firm or a union and a firm (Pissarides, 2000; Acemoglu, 2001; Farber, 1986). 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, 2003a).

[3]Similar facts were presented in prior work by Myers and Shultz (1951), Rees (1966), and Rees and Shultz (1970).

between workers in the same narrow skill group.

The data come from a new monthly linked employer-employee database covering the universe of employees at Danish companies. While wages in Denmark were historically set by union bargaining, firms today have considerable latitude to negotiate wages with individual employees (Dahl et al., 2013). Our data cover the period post-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 search model where firms renegotiate wages with workers that receive outside offers. 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 both through public sources and through their, individual-specific, social networks. This allows us to derive a measure of outside options that we can take to the data.

We then test the key predictions of this theory 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, non-parametrically, 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, that are correlated with unobserved determinants of job-to-job mobility and wage growth.

We present non-parametric evidence that confirms both predictions: (1) changes in workers'

[4]Our data cover the period 2008-2016. Most wage decentralization occurred in the 1990s.

information about their outside opportunities lead to mobility and wage growth, and (2) 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 an additional ten new positions at an individual's former coworkers' current firms results in a fifteen percent higher probability the worker makes a job-to-job transition that month.[5] The same change translates to approximately $50 more 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 of 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% increase in base pay.

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 non-parametric 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. We also show that the results are 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 to see all workers within the same firm and occupation see equal wage growth.

We decompose our measure of outside options into portions coming 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, consistent with our information-transmission story. In particular, coworkers who are still working in the same (of five) administrative region matter more, as do coworkers the individual worked with more recently. We also

[5]Because our data are monthly, the base rate is low: roughly one percent of workers make a job-to-job transition each month.

construct similar measures of outside options based on an individual's future coworker network. If our results were driven by unobserved demand shocks, we would expect these measures to have a significant impact on both mobility and wage growth. We would also expect that adding these measures as controls to our baseline regression would reduce our estimates. We do not find support for either of these predictions.

Wage effects are largest for higher skilled workers and for workers with more specialized skills. We divide workers into eight broad occupation groups and re-estimate the effects within each group.[6] We find that the impact on workers with specialized skills (professionals) is double that of workers in the middle skill group (technicians), and nearly five times that of workers in the least skilled group. Because workers with specialized skills also have higher baseline earnings, these impacts on translates into substantially larger effects on the level of earnings. However, there are impacts on mobility and earnings for all but the least skilled workers. Within each skill group, women benefit less than men.

In addition, both job-mover and stayers appear to benefit. We find that individuals who stay at their current firm obtain roughly 20% the earnings gain of job movers. Posting models—including 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.

Our reduced form results only indicate that some firms and workers engage in wage renegotiation; they do not indicate that all firms renegotiate. Some firms may be able to commit not to renegotiate wages with employees who receive outside offers (Postel-Vinay and Robin, 2004; Doniger, 2015). To assess the extent to which firms negotiate with different groups of workers, we use our reduced form estimates to identify a structural search model incorporating on-the-job search, information transmission through networks, and a mass of posting firms. The model is based on Flinn and Mullins (2017); our estimates contribute to a small literature on the empirical

[6]The 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. These groups are based on the broad International Standard Classification of Occupation (ISCO) codes.

relevance of wage-posting and bargaining (Hall and Krueger, 2012; Doniger, 2015).

We estimate this model, separately for different skill groups, using simulated method of moments. Our estimates indicate that wage renegotiation is more common among high skilled workers. Using these parameters, we estimate that a 50% reduction in the arrival rate for employed workers would lead to a significant reduction in wage growth. For high skilled workers a larger portion of this is due to decreased on-the-job bargaining; for lower skilled workers this is mostly due to decreased mobility. Overall, the results indicate that changes in the labor market that hamper workers' ability to obtain or use outside offers may have meaningful impacts on wage growth.

1.1.1 Related Literature

This paper contributes to several distinct literatures. In particular, outside options are a key ingredient in macroeconomic search and bargaining models, which assume that individual workers negotiate–and potentially renegotiate–their wage with their employers (Pissarides, 2000; Postel-Vinay and Robin, 2002; Cahuc et al., 2006). In some of these models, the worker's outside option is the value of non-employment. In models where employed workers can renegotiate wages with their current firm, the outside option is typically the best outside offer the worker has received. Bargaining on the basis of outside offers rationalizes many macroeconomic phenomena including wage dispersion (Hornstein et al., 2011), mismatch (Hagedorn et al., 2017; Lise and Robin, 2017), and wage cyclicality for job switchers (Gertler et al., 2016). A large literature uses structural models of renegotiation on the basis of outside offers to measure the determinants of wage growth (see, e.g., Bagger et al., 2014b; Lise et al., 2016; Jarosch, 2015).[7]

Only a handful of papers have directly examined the role of workers' outside options in wage-setting and none, to our knowledge, have used individual-level variation. Beaudry et al. (2012) use cross-city variation in the growth of different industries to show that there are sectoral linkages in

[7]Within the search literature, this paper is most related to work by Lamadon (2014), who investigated the transmission of both firm- and worker- productivity shocks to wages using a directed, competitive search model estimated on Swedish matched employer-employee data. That paper used the correlation in wage growth between an individual and his current coworkers (who experience the same firm shocks) to separately identify worker- and firm- productivity shocks in the context of a directed search model.

wages, consistent with bargaining models (see also Bidner and Sand, 2016; Fortin and Lemieux, 2015). Our approach is more similar to that used by Hagedorn and Manovskii (2013). They use a proxy for the number of offers an individual has received since starting a job—based on the vacancy-to-unemployment rate—to test the predictions of spot market models.[8] Contemporaneous work by Jäger et al. (2018) finds that wages are do not respond to changes in the value of non-employment, suggesting that other wage-setting protocols, including the one investigated in this paper, may be more relevant. This paper differs from the prior literature in its focus on whether firms negotiate with individual workers on the basis of changes in the worker's opportunities at other firms.[9]

This paper is also related to recent work that has found that idiosyncratic changes in firm rents impact the wages of workers at those firms (Abowd and Lemieux, 1993; Van Reenen, 1996; Card et al., 2014; Kline et al., 2018; Mogstad et al., 2017). The small rent-sharing elasticities reported in this literature (.05-.15) suggest that workers may be able to capture a large portion (85-95%) of changes in the value of their outside option (Card et al., 2016c). This is because, in simple bargaining models, workers' wages are a weighted average of the rents produced in the match and the workers' outside options. However, this one-to-one relationship might break down if changes in outside options are not verifiable, or if firms are able to commit not to renegotiate (Hall and Lazear, 1984; Manning, 2003a). The results in this paper suggest workers capture a much smaller portion of changes in their outside options than most rent-sharing estimates imply.

This paper also contributes to a rapidly growing literature on information transmission through social networks (Simon and Warner, 1992; Hensvik and Skans, 2016; Giorgi et al., 2016; Bailey et al., forthcoming; Beaman, 2016; Gee et al., 2017; Glitz and Vejlin, 2018). A recent series of papers shows that recently displaced workers and workers entering a new labor market use information obtained from their former coworkers (Glitz, 2013; Saygin et al., 2014), classmates (**?**), family

[8]One key difference between this paper and that paper is that we generate a measure of the arrival rate that varies across workers with identical patterns of employment/non-employment.

[9]Caldwell and Danieli (2018) investigate the role of outside options in generating between-group wage inequality. The authors use an assignment model to derive an index of workers' outside options, analogous to concentration indices often used in the industrial organization literature. They then use the cross-sectional distribution of workers to estimate this index.

members (Kramarz and Skans, 2014), and neighbors (Bayer et al., 2008; Schmutte, 2014) to find job opportunities. This paper shows that currently employed workers also use this information. It is most related to Shue (2013); that paper shows that an individual's wages are more related to those of his/her randomly assigned Harvard Business School section-mates than to his/her different-section classmates (2013). The data in that paper contain pay for CEOs and other top executives; we focus on workers who are unable to set their own pay and must, instead, bargain with their employer.[10]

The rest of the paper proceeds as follows: Section 1.2 develops a theoretical model that incorporates both on-the-job search and a mix of firm wage-setting strategies. It then uses this model to derive empirical predictions, to derive a measure of workers' outside options, and to explain the key identifying assumption. Section 1.3 describes the institutional features of the Danish labor market and the administrative population-based registers we use. Section 1.4 explains the empirical strategy and maps variables in our data onto the theoretical objects described in Section 1.2. Section 1.5 presents reduced form results on mobility and earnings and Section 1.6 explores heterogeneity in these results. Section 1.7 uses the reduced form to identify bargaining parameters and to estimate the extent of bargaining, relative to posting, for workers in different skill groups. Section 1.8 concludes.

1.2 Outside Options and Wages

We start by developing a continuous time search model with bargaining and on-the-job search. Our model is based on Flinn and Mullins (2017) and is standard in all but two respects.

First, rather than assuming that all firms and workers bargain over wages, the model allows for two types of firms: those that post wages and those that negotiate and renegotiate wages with workers. This relaxation makes our model more general, since the other leading model of wage-setting under imperfect competition (monopsony) features wage-posting. Further, the speed with

[10]This paper is also related to a recent series of papers that use linked employer-employee data to investigate the importance of an individual's coworkers in determining wage growth (Cornelissen et al., 2017; Herkenhoff et al., 2018; Jarosch et al., 2018). These papers have shown that individuals appear to learn from their more productive colleagues, and that this learning is reflected in their wages.

which various changes in the labor market—changes in concentration or the enforcement of no-poach agreements—will impact wages depends on whether workers have to switch firms in order to benefit from changes in their outside options.

Second, we depart from the model in Flinn and Mullins (2017) and allow workers to learn about job opportunities both through public sources and through their social networks. This is key for our reduced form strategy. For simplicity, we first derive the model assuming all workers—conditional on employment status—learn about jobs at the same rate. In Section 1.2.4 we show how the addition of social networks allows us to generate a measure of outside options that we can take to the data.

1.2.1 Model Setup

Workers vary in ability a and firms vary in productivity θ. A worker of ability a matched with a firm of type θ produces $a\theta$. Time is continuous and both firms and workers are risk neutral and discount the future at rate ρ. Matches are dissolved exogenously at rate δ and workers receive ab while unemployed. The parameter b reflects both the value of unemployment benefits and the value of non-work time.[11] Search is undirected and workers learn about new job opportunities at rate λ^E while employed and λ^U while unemployed.

There are two types of firms: posting (P) and renegotiating (R). Posting firms commit ex ante to a wage schedule and do not renegotiate with workers who receive an outside offer. They post wage premia; a firm that posts w pays the worker wa. Renegotiating firms bargain with workers, both at the beginning of the employment relationship and when one of the parties receives a credible outside offer (Cahuc et al., 2006). Renegotiation is costless and occurs only by mutual consent. We follow the prior literature in assuming that worker-firm bargaining at renegotiating firms follows the infinite-horizon alternating offer game in Rubinstein (1982). When an employed worker receives an outside offer, the incumbent and outside firm engage in competition over that worker. By definition,

[11] It is standard to assume that the value of non-employment is proportional to ability. This is reasonable given that most unemployment benefits are based, at least in part, on a worker's wages. However, a direct implication of this assumption is that all workers have the same reservation firm type when unemployed.

posting firms do not adjust their bids. Then, if the winning firm is a renegotiating firm, the worker and firm bargain over wages. The worker uses the maximum value she could have obtained at the losing firm as her outside option.[12]

We close the model in Appendix 1.11. We assume that, when deciding whether to post a vacancy, firms draw both θ and a vacancy type (P or R). This is somewhat simpler from the setup in Flinn and Mullins (2017), and more suited to the counterfactuals we consider in Section 1.7.

1.2.2 Value Functions

We next derive the value functions for workers who: (1) work at renegotiating firms, (2) work at posting firms, or (3) are unemployed. Workers and renegotiating firms bargain over how to split the total surplus produced in the match, $T_R(a\theta)$. The value the worker receives, $V_R(a, \theta, \overline{w})$, depends both on the productivity of the match and the last outside option she used for bargaining, $\bar{w}$. A worker at a posting firm earning w obtains value $V_P(w)$. V_U denotes the value function for an unemployed worker.

It is useful to first state a result from Flinn and Mullins (2017):

Lemma 1. *A worker who receives the total surplus created by the match θ at a renegotiating firm (type R) has the same value as a worker earning θ at a posting firm (type P). That is,*

$$T_R(a\theta) = V_R(a\theta, a\theta) = V_P(a\theta)$$

Proof. See Appendix 1.11. □

The intuition behind this is simple: once a worker at a renegotiating firm receives the full surplus of her match, her wage can no longer adjust at that firm. She will receive the full surplus of the match whenever her last bargaining position was θ, the match productivity.[13] The worker's

[12] If the losing bid came from a renegotiating firm, this is the total value that would have been produced in the match. If the losing bid came from a posting firm, this is the total value the worker would have received at that firm. This may not be equal to the total value of the match.

[13] This is a direct implication of the bargaining protocol. It also implies the worker's wages satisfy $\omega(\theta, \theta) = \theta$.

mobility decisions will be the same as those of a worker at a posting firm earning θ; like that worker, her wage will not adjust at the current firm. This result is useful because it means there is a sufficient statistic that governs workers' mobility patterns and wage growth: the maximum wage they could earn at a firm. At posting firms this is simply the offered wage; at renegotiating firms, this is θ.

Renegotiating Firm Because we have assumed transferable utility, the total surplus of a match between a worker and a renegotiating firm is the sum of the value to the worker and to the firm. At a renegotiating firm of productivity θ this total surplus is:

$$\begin{aligned} \rho T_R(\theta) &= \theta + \underbrace{\delta\left(V_U - T_R(\theta)\right)}_{\text{unemployment}} + \\ &\underbrace{\lambda^E p_R \int \beta\left[T_R(x) - T_R(\theta)\right]^+ dF_\theta(x)}_{\text{poached by renegotiating firm}} + \underbrace{\lambda^E(1-p_R)\int \left[V_P(x) - T_R(\theta)\right]^+ d\Phi(x)}_{\text{poached by posting firm}} \end{aligned} \tag{1.1}$$

The first term is the output of the match, θ. With probability δ, the worker loses her job. The third and fourth terms measure the additional surplus the worker gets if she meets a firm that outbids her current firm.[14] We use F_θ and Φ to refer to the distributions of offers from renegotiating and posting firms, respectively. The worker receives an offer from a renegotiating firm with probability $\lambda^E p_R$ and an offer from a posting firm with probability $\lambda^E(1-p_R)$. She moves to that firm if (1) the outside firm is a more productive renegotiating firm or if (2) the outside firm is a posting firm that offers her more than she produces with her current firm. If she moves to a new renegotiating firm, she uses the value produced in her current match as her outside option and obtains a fraction β of the rents produced in the new match $(T_R(x) - T_R(\theta))$. We refer to the parameter β as workers' bargaining power. It measures the proportion of rents a worker is able to obtain in bargaining. The free entry condition ensures that if the worker is poached, the value to the incumbent firm is zero.

The firm and worker bargain over how to split this total surplus $T_R(\theta)$. The worker's value

[14]We use the notation $[a]^+ = \max\{a, 0\}$.

function depends not only on the total value produced but on her last bargaining position. If her last offer came from a renegotiating firm, this is

$$V_R(\theta, x) \quad = \quad \underbrace{T_R(x)}_{\text{last outside option}} \quad +\beta \underbrace{(T_R(\theta) - T_R(x))}_{\text{surplus}}$$

Similarly, if her last offer came from a posting firm, this is:[15]

$$V_R(\theta, x) \quad = \quad \underbrace{V_P(\bar{w})}_{\text{last outside option}} \quad +\beta \underbrace{(T_R(\theta) - V_P(x))}_{\text{surplus}}$$

In order to achieve this split, workers and firms agree on wages $\omega(\theta, x)$. The worker's value function depends only on her current wages $(\omega(\theta, x))$, not the total value produced by the match. With probability δ she is unemployed next period. With probability λ^E she receives an outside offer. With probability $\lambda^E(1 - p_R)$ that offer is from a posting firm. If the outside offer comes from a firm that is sufficiently good, she is poached. For more moderate values, she renegotiates her wage at her current firm. Her value function is given by $V_R(\theta, x) =$

$$\begin{aligned}
&\{\omega(\theta, x) + \delta V_U \\
&\lambda^E p_R \left(\underbrace{\int_\theta^\infty \left[T_R(\theta) + \beta\left(T_R(x) - T_R(\theta)\right)\right] dF_\theta(x)}_{\text{poached}} + \underbrace{\int_x^\theta \left[\left(T_R(x) + \beta\left(T_R(\theta) - T_R(x)\right)\right] dF_\theta(x)}_{\text{renegotiate}} \right) + \\
&\lambda^E(1 - p_R) \left(\underbrace{\int_\theta^\infty V_P(x) d\Phi(x)}_{\text{poached}} + \underbrace{\int_x^\theta \left[V_P(x) + \beta\left(T_R(\theta) - V_P(x)\right)\right] d\Phi(x)}_{\text{renegotiate}} \right) \}/ \\
&\left(\rho + \delta + \lambda^E p_R(1 - F_\theta(x))\right) + \lambda(1 - p_R)(1 - \Phi(x))\big)
\end{aligned}$$

[15]This expression illustrates the symmetric relationship between the importance of rents and options mentioned in the introduction. When $T_R(\theta) = \theta$ and $V_R(\theta, w) = w$, $w = (1 - \beta)\bar{w} + \beta\theta$. If rents pass through at a rate β, options pass through at a rate $1 - \beta$.

Posting Firm The value function of a worker at a posting firm depends only on the wage she receives at that firm:

$$\begin{aligned}\rho V_P(w) \quad = \quad & w + \underbrace{\delta\left(V_U - V_P(w)\right)}_{\text{unemployment}} + \\ & \underbrace{\lambda^E p_R \int \beta\left[T_R(x) - V_P(w)\right]^+ dF_\theta(x)}_{\text{poached by renegotiating firm}} + \underbrace{\lambda^E(1-p_R)\int\left[V_P(x) - V_P(w)\right]^+ d\Phi(x)}_{\text{poached by posting firm}}\end{aligned} \tag{1.2}$$

The first term is her wage, w. With probability δ the match is dissolved and the worker becomes unemployed. She receives an offer from a renegotiating (posting) firm with probability λp_R^E (posting firm: $\lambda^E(1-p_R)$). She is poached if (1) the outside firm is a more productive renegotiating firm or (2) the outside firm is a posting firm that offers her more her current firm. The β in the third term reflects the fact that if she moves to a renegotiating firm, she will obtain a fraction β of the rents produced by that match and will use her current value function $V_P(w)$ as her fallback option.

Unemployment The value function for an unemployed worker, V_U, also has a simple recursive formula:

$$\rho V_U = \underbrace{b}_{\text{benefits}} + \underbrace{\lambda^U p_R \int \beta\left[T_R(x) - V_U\right]^+ dF_\theta(x)}_{\text{offer from renegotiating firm}} + \underbrace{\lambda^U(1-p_R)\int\left[V_N(x) - V_U\right]^+ d\Phi(x)}_{\text{offer from posting firm}}$$

If a worker is unemployed, she receives benefits b this period. She receives an offer from a renegotiating firm with probability $\lambda^U p_R$. Her wage upon accepting employment at that firm is set such that she receives V_U plus a fraction, β, of the match surplus $(T_R(x) - V_U)$. She receives an offer from a posting firm with probability $\lambda^U(1-p_R)$.

1.2.3 Reduced Form Predictions

The continuous time model corresponds to the limit of an analogous discrete time model. To derive the reduced form predictions, we consider what happens in a period of length $t = 1$. In a slight abuse of notation, we use h_θ to refer to the combined distribution of match productivities θ and wage offers w. We also assume that λ is small enough so that the probability of receiving multiple offers in this period of time is negligible. This gives us the following results:

Claim 2. The probability a worker makes a job-to-job transition this period is

$$\Pr(\text{move}) \approx \underbrace{\lambda^E e^{-\lambda^E}}_{\text{arrival rate}} \underbrace{\left[\int_\theta^\infty h_\theta(\theta')d\theta'\right]}_{\text{prob better firm}} = \tilde{\lambda}\tilde{H}(\theta)$$

where $\tilde{H}(\theta) = (1 - H_\theta(\theta))$ and θ is her current employer's type.

Proof. See Appendix 1.11 for details. □

The intuition behind Claim 2 is simple: the probability an individual makes a job-to-job transition this period is simply the probability she receives an offer multiplied by the probability that offer came from a firm that was willing to pay her more than her current firm would match. By construction, the probability a worker at a posting firm sees a wage change is identical to the probability she makes a job-to-job transition. However, workers at renegotiating firms may see wage changes, even if they do not move firms.

Claim 3. The probability a worker at a renegotiating firm of type θ sees a wage change is:

$$\Pr(\text{wage change}) \approx \underbrace{\lambda^E e^{-\lambda^E}}_{\text{arrival rate}} \underbrace{\left[\int_q^\infty h_\theta(\theta')d\theta'\right]}_{\text{prob better offer}} = \tilde{\lambda}\tilde{H}(w) > \tilde{\lambda}\tilde{H}(\theta)$$

where q is her last bargaining position.

Proof. See Appendix 1.11 for details. □

The probability the worker sees a wage change is the probability she receives an offer, multiplied by the probability that the offer came from a firm that is better than the last offer she used in renegotiation. This is always weakly higher than the probability of making a job-to-job transition because the outside firm doesn't have to be willing to outbid her current firm.

The two predictions are summarized in Figure 1-1. Offers are ranked according to the maximum value a worker could receive. When a worker at a renegotiating firm receives an outside offer (Panel A), one of three things will occur:

1. **Worker is Poached:** If the outside offer is sufficiently good, the outside firm will 'win' during competition with the incumbent firm. This happens if:

$$\begin{aligned} T(\theta') &> T(\theta) \\ V_P(w) &> T(\theta) \end{aligned}$$

2. **Wage Renegotiation:** If the outside firm loses to the incumbent firm, but would have offered the worker more than her last outside option, the worker will renegotiate her wage with her current firm. This will happen if:

$$\begin{aligned} T(\theta') &\in [T(w'), T(\theta)] \\ V_P(w) &\in [T(w'), T(\theta)] \end{aligned}$$

where w' is the worker's last bargaining position.

3. **No Change:** If neither of these conditions is met, the worker stays at her current firm and continues to earn her current wage. Renegotiation only occurs by mutual consent; the worker will not initiate wage renegotiation if it would lead to a wage cut.

Panel B shows that, for workers at posting firms, outside offers can only lead to job-to-job transitions.

This figure shows the key empirical predictions. We should see a positive relationship between our measure of outside options and both job-to-job mobility and wage growth. We should also see effects on earnings through a greater portion of the outside options distribution. This is because outside offers only need to dominate whatever a worker last used for negotiation, not the maximum wage that firm would be willing to pay.

1.2.4 Information Transmission Through Networks

Standard search models assume that all workers in a given labor market face the same job arrival rate. Suppose instead that workers learn about job opportunities through both public sources—which are common to all workers—and through their own social networks. Search through both sources is undirected.

We can then decompose the probability a worker receives an outside offer into two components:

$$\lambda^P + \alpha \underbrace{\int s(x)v(x)dx}_{\text{social ties}} \tag{1.3}$$

where the first term, λ^P, is the arrival rate of offers through public sources. The second term measures the arrival of offers through networks. We assume that the probability a worker hears about one of the $v(x)$ offers at firms of type x scales with the number of people they know at that firm, $s(x)$.[16] The parameter α is the joint probability of learning about an opening through social ties and receiving an offer.

The probability a worker makes a job-to-job transition or sees a wage change depends on the probability the worker receives a 'good enough' offer. The probability a worker receives an offer better than θ has a similar expression:

$$\underbrace{\lambda^P \int_\theta^\infty p(x)dx}_{\text{public sources}} + \alpha \underbrace{\int_\theta^\infty s(x)v(x)dx}_{\text{social ties}} \tag{1.4}$$

In section 1.5 we test whether individuals are more likely to move or earn more in periods when

[16]We define $v(x)$ such that $\int v(x)dx = 1$. We assume that workers do not strategically pick jobs in order to gain access to better social networks, and that firms do not hire workers to take advantage of their network. In the empirical work we both control for time-invariant individual heterogeneity—which could include a propensity to strategically move to new firms—and for the number of connections in an individual's network.

they were more likely to receive an outside offer through one of their connections. Our reduced form measure of an individual's outside options, Ω_{it}, is based on the expression for arrival rates in equation 1.3. We construct

$$\Omega_{it} = \underbrace{\sum_j \text{Share Coworkers}_{ijt}}_{\text{coworker network}} \times \underbrace{s_{jt}}_{\text{firm demand}} \times \underbrace{\omega_{jt}}_{\text{firm quality}}$$

We weight firm-specific measures of labor demand (s_{jt}) by the share of an individual's former coworkers' at that firm. Our baseline specification sets $\omega_{jt} = 1 \; \forall j, t$. In this case, Ω_{it} is simply a proxy for the arrival rate of offers through the individual's social network. In some specifications, we attempt to measure the probability a worker received a 'good' outside offer by weighting changes in firm demand by the different measures of firm quality. We provide more details on how we construct this measure in Section 1.4.3.

1.3 Setting and Data

1.3.1 Institutional Setting

Several features of the Danish labor market make it a good setting for studying the relationship between outside options and wages. First, job-to-job and occupational mobility rates are high by European standards, and are comparable to those in the United States (Botero et al., 2004; Groes et al., 2014).[17] This flexibility is the result of the Danish "flexicurity" system, which combines low firing costs with a generous social safety net (Andersen and Svarer, 2007).

Second, while wages in Denmark were historically set by sector-level bargaining agreements, wages today are mostly set at the firm level or are negotiated between individual workers and firms (Dahl et al., 2013).[18] Further, private sector collective agreements do not typically cover man-

[17]Appendix Figure 1-17 shows that relative to other OECD countries (those marked in blue have readily available linked employee-employer registers), hiring and separation rates are high.

[18]During the 1990s there was a shift towards the use of collective agreements that specified only general working conditions: working hours, rules regarding hours flexibility, and minimum wages. There are four wage-setting regimes

agers, executives, or university graduates. Instead, these workers negotiate employment conditions individually. Unions still play a significant role in setting sector-level minimum wages, which generally apply to inexperienced or new workers. Denmark has no national minimum wage. Danish unions are also very important in organizing unemployment insurance; most unemployment insurance funds are associated with one of the unions.

One key difference between the American and Danish labor markets is that multiple job-holding is common in Denmark, as it is in most Nordic countries. In 2015 eight percent of Danish workers worked for more than one employer (Pouliakas, 2017). The incidence over an eight year period is significantly higher. Throughout our analysis we focus on the population of single job-holders.

1.3.2 Primary Data Sources

We combine three types of administrative data: (1) a monthly employer-employee register, (2) person-level demographic registers, and (3) firm-level Customs and Trade registers. We provide more information on the data in Appendix 1.12.

Most of our data come from a new monthly employer-employee database, known as the BFL (or e-Income) database. Danish firms must report individual hours and earnings to the Danish Customs and Tax Authority on a monthly basis. The BFL register contains the data from these reports. In addition to information on hours and earnings, the data also contain six-digit industry and occupation codes. We have data from January 2008, the start of the register, through March 2016. To construct coworker networks for the first three years of our sample, we supplement these data with a separate monthly employer-employee register (MIA), which contains monthly information on place of employment from 1999-2008.

in Denmark. In the standard rate system, agreements determine wages for most workers in a sector. In the minimum wage and minimum pay systems, these agreements only specify wage floors; only very inexperienced workers earn the minimum rate. Under the fourth system, there are no centrally bargained minimum wage rates. Rather, wages are negotiated at the plant or firm level. There is scope for individual-firm bargaining in all but the fourth system. See Dahl et al. (2013) and http://www.ac.dk/media/768440/den-danske-model.pdf ("The foundation and dynamics of the Danish labour market") for more information.

The data are well-suited for our analysis because they contain high-frequency (monthly) earnings data and because, unlike most employer-employee datasets, they contain firm-reported measures of hours worked. The hours data allow us to examine whether changes in monthly earnings are driven by changes in hours worked, or changes in hourly earnings.

We use unique person identifiers to link the employer-employee data to demographic registers that contain information on age, sex, country of origin, education, and household characteristics. We collapse the education codes in our registers to nine broad level codes and eleven broad field codes following the International Standard Classification of Education (ISCED) codes. We use these codes to distinguish between workers with the same level of education but different skills.

We use firms' unique identifiers to link our employer-employee data to the Danish Foreign Trade Statistics Register. For each firm and month between January 2004 and December 2015 we have the value (in Danish Kroner) of imports and exports disaggregated by product and by origin (imports) or destination (exports). The original data are reported at the eight-digit Combined Nomenclature level; we aggregate flows to the six-digit Harmonized System. We use these data in section 1.5.6 when we consider measures of outside options based on world demand for each firm's products.

1.3.3 Descriptive Statistics and Sample Restrictions

Workers Column 1 of Table 1.1 provides descriptive information on the set of workers who appear at least once in the BFL data between January 2008 and March 2016. The average worker in our sample (weighted by months in sample) is nearly forty years old and has annualized earnings of around $40,000, before taxes. Nearly half of individuals are married or in a registered partnership; more are cohabiting. About a third of the workers have a college degree.

Before constructing coworker networks we restrict the sample to Danish citizens who work in firms with between 2 and 1000 employees. We exclude non-Danish citizens both because our demographic data (especially our education measures) are most complete for Danish citizens and because our data include all employees of Danish firms, including those not residing in Denmark.

We exclude connections formed in firms with more than one thousand employees because it is unlikely that an individual knows all of her coworkers in a very large firm. This type of restriction is common in the networks literature.[19]

We impose two additional restrictions to generate our regression sample. First, we focus on prime-age workers between the ages of 25 and 60. Very young (under 25) workers will not have had enough time in the labor market to develop a network; older (over 60) workers are likely close to retirement. Second, we focus attention on workers who are, over the sample period, single job holders. This is a relatively significant restriction given the prevalence of multiple job-holding in Denmark. However, this restriction allows us to remove a significant portion of part-time workers whose earnings fluctuations likely reflect both changes in hourly wages and changes in hours worked. Further, the theoretical framework is about single job holders. Column 3 of Table 1.1 shows the impact of these restrictions. These workers are more likely to be married or in a couple. They also have higher average annualized earnings because they are more likely to be working full-time.

Firms Workers in our sample are spread across 352,010 distinct firms (tax identifiers). Column 1 of Table 1.2 shows that the average firm has eleven employees, though there is substantial variance: the standard deviation is over two hundred. Most firms have a single establishment and over half are located in one of two regions: the Capital Region and Central Denmark. These regions contain Copenhagen and Aarhus, Denmark's first and second largest cities, respectively. Column 2 shows that most firms fall within the network sample: they have between two and one thousand employees, on average, throughout our sample period. Most (>99%) of the firms that are excluded from our network sample are single-employee firms. There is much less variation in firm size within the network sample. The average firm has eight employees and the standard deviation is thirty-two.

[19]For instance, Hensvik and Skans (2016) only consider firms with less than 500 employees and Eliason et al. (2017) only consider connections formed in firms of fewer than 100 employees. Saygin et al. (2014) include connections formed in all firms with fewer than 3000 workers but consider a smaller set of workers: those involved in mass layoffs. Glitz and Vejlin (2018) include all of an individual's coworkers from the prior ten years but focus only on workers who were hired in a given year. A different set of papers has defined an individual's network by her set of same-citizenship peers (see, e.g. Dustmann et al., 2015).

Column 3 shows that most firms are neither importers nor exporters; only fourteen percent of firms appear in the trade register. However, because the average firm in this register is double the size of the average firm in the full sample, these firms cover a substantial portion of employment. Most exporting firms export a single product. Because our firm size threshold is generous, most firms in the trade register fall within the network sample; less than fifteen percent are excluded (Column 3 versus Column 4). Most firms are located in the capital region and Central Denmark; the fraction of employment located in these two regions is even greater. Figure 1-18 presents a map of the five administrative regions.

1.3.4 Earnings Outcomes

We examine the impact of changes in workers' outside options on changes in five measures of earnings. The monthly register data contain two measures of monthly earnings: a broad measure, which includes income derived from benefits (e.g. contributions to retirement accounts or fringe benefits) and a narrow measure which captures post-mandatory-contribution take-home pay. We look at log changes in both measures. However, we prefer the broad measure because it does not respond to changes in ATP (mandatory pension) contributions that arise due to changes in legislation or changes in hours worked.[20] Further, using the broad measure allows us to account for the fact that some workers may to with their employer over retirement contributions or fringe benefits. Our third measure is log hourly earnings. While both earnings measures are available for all workers, hours are imputed for roughly a quarter of the sample. We focus on the subset of observations with firm-reported hours.

Finally, we use the panel component of our data to construct the fourth and fifth measures: "bonus pay" and "base pay". We identify bonuses by looking for one-month increases in earnings that are followed by a decrease of approximately the same magnitude. We define base pay as difference between total monthly earnings and any bonuses. We provide more details on how we construct the earnings measures in Appendix 1.12.

[20]Individuals pay different rates based on which of four bins their monthly hours falls in: 0-38, 39-77, 78-116, and 117+ hours.

1.4 Empirical Strategy

Our empirical strategy exploits the fact that individuals often learn about new job opportunities through their former coworkers. The logic is simple: individuals should be more likely to hear about job opportunities when their former coworkers' firms are expanding more relative to other periods.

1.4.1 Graphical Illustration of Empirical Strategy

The strategy can be explained in three pictures. First, Figure 1-2 shows how this approach identifies variation in outside options within an individual over time. Panel A shows the individual's network: each collection of dots represents a firm and each dot represents a worker. Each blue dot represents one of the individual's former coworkers. Our identification strategy relies on the assumption that the worker is less likely to hear about job openings in Panel B—where she does not know anyone at the expanding firm (displayed in red)—than in Panel C.

In some of our analysis we exploit variation in information about outside labor market opportunities between workers in the same firm (or firm and occupation). Workers who join a firm at the same time will have different networks if they moved to that firm from different firms. Figure 1-3 shows that within-firm variation also arises due to differences in firm tenure. The figure shows two workers—marked in blue and purple—at firm A in period 1. In the next period, the purple worker moves to firm C, and is replaced by the red worker from firm B. In the final period, firm C expands. Our within-firm analysis relies on the fact that the red worker at firm A is less likely to hear about this expansion than the blue worker. This is because the blue worker has a shared history with someone at firm C. Panel B shows that it may not be the case that the worker with longer tenure at the firm has access to 'better' information.

Figure 1-4 illustrates how we divide up the monthly panel to create (1) coworker networks and (2) time-specific shocks. In each period we use the prior thirty-six months to generate an individual's network. We then look at how firm-specific shocks between t=0 and t=1, weighted by this network, translate into mobility or wage growth in the same period. Specifically, we examine whether an individual is at a new firm or earns more in period 1 than they did in period 0.

1.4.2 Coworker Networks

Individual i's coworker network in month t consists of all workers she worked with in the prior three years who are now at different firms.[21] There are two key restrictions. First, we only include connections in firms with between 2 and 1000 people. In large firms, it is unlikely that a worker knows, or shares information with, all of their former coworkers. Second, we exclude connections that were formed more than three years ago both because older connections are likely to be less informative and because, without a fixed window, network size or quality would vary mechanically over the sample window. We examine the robustness to these restrictions in Section 1.5.7.

We also remove connections who are working at firms individual i worked at in the past three years so that the network does not vary mechanically with mobility. Specifically, if a worker moves from firm A to firm B, we do not include her former coworkers at firm A in the network, unless they move to another firm. This is important: if we did not do this workers who switched firms would, mechanically, see a large increase in network size. Our measure of outside options for that worker would also be heavily weighted towards the firm they just left.

Column 1 of Table 1.3 shows that, on average, workers have 156 connections, which connect them to 60 distinct firms. However, the distribution is very skewed and the median worker has only 60 former coworkers. Columns 2 and 3 compare the networks of male and female workers and show that women do not appear to have significantly weaker networks by any metric: the number of connections, the number of industries, or the average value added per worker at a connection's firm. However, because Danish firms—like firms in most countries—are somewhat segregated by gender, women's networks primarily consist of other women (Card et al., 2016a; Hellerstein et al.,

[21] Specifically, for each month t, we construct the bipartite adjacency matrix A^t where $A^t_{ij} = 1$ whenever i and j worked together (at the same firm, at the same time) in the previous thirty-six months and $A^t_{ij} = 0$ otherwise. We can then rewrite our measure of outside options using network notation:

$$\sum_{c \in N} \frac{A_{ict}}{\sum_{n \in N} A_{int}} \times s^{+}_{jt} \times \omega_{\psi(c,t),t}$$

Note that we only consider first-degree connections.

2008).

1.4.3 Measuring Outside Options

Our measure of outside options is motivated by the theoretical model described in Section 1.2. We create individual- and time-specific measures of outside options by weighting time-varying measures of firm-specific labor demand by each individual's coworker network. For each individual i and month t we construct:

$$\Omega_{it} = \sum_j \text{Share Coworkers}_{ijt} \times \underbrace{s_{jt}}_{\text{firm demand}} \times \underbrace{\omega_{jt}}_{\text{firm quality}}$$

where s_{jt} is a measure of firm labor demand and the ω_{jt} are firm-quality weights.[22]

Our baseline measure uses the number of new positions at an individual's former coworkers' firms as the measure of firm demand, and weights all firms equally: $s_{jt} = (E_{j,t} - E_{j,t-1})^+$ and $\omega_{jt} = 1 \ \forall j, t$. We focus on the number of new positions, rather than the overall number of hires, which reflects changes in both labor demand and churn. To prevent a mechanical correlation between the change in employment at these firms and an individual's own job-to-job mobility decisions, we use a "leave-out" version where we do not include new positions created for individual i.

While our baseline specification does not use firm quality weights, the model in Section 1.2 suggests that the probability an individual moves or renegotiates scales with both the total number of offers and the probability that an offer comes from a sufficiently good firm. In practice, it is difficult to determine which firms are likely to be more attractive because workers may have preferences over non-wage characteristics (Sorkin, 2018) and because firm wage premia may change

[22] Because the weighting functions vary by individual (leading our measure of outside options to vary across workers within a firm), this is somewhat different from a standard "Bartik"-style instrument. There is an ongoing debate on the identifying assumptions behind these instruments (see, e.g. Borusyak et al., 2018; Goldsmith-Pinkham et al., 2018; Adão et al., 2018; Jaeger et al., 2018). Our identifying assumption, described below, is similar to the "exogeneity of shocks" assumption in Borusyak et al. (2018). In particular, we do not require that the shares—the coworker networks—be randomly assigned.

in response to firm- or market-shocks. In specifications presented in the Appendix we weight positions by the mean wage at the firm, scaled by average wages: $\omega_{jt} = \frac{w_j}{\bar{w}}$. In Appendix Section 1.13.2 we consider the impact of new positions at more and less productive firms, where productivity is measured using value added per worker.

Within an individual, variation in Ω_{it} is driven by changes in firm demand, not changes in network composition. Table 1.16 shows that the number and characteristics of individuals in a worker's coworker network are highly autocorrelated, with autocorrelations above .9, even after a year. The number of hires and new positions at a firm are significantly less autocorrelated.

In Section 1.5.7 we show that our results are robust to different definitions of (1) an individual's coworker network, (2) firm demand s_{jt} and (3) firm weights ω_{jt}. In section 1.5.6 we consider measures of s_{jt} based on world demand for each firm's products.

1.4.4 Estimating Equations and Identifying Assumptions

The main estimating equation is:

$$y_{it} = \gamma\Omega_{it} + X_{it} + c_{it} + \alpha_i + \alpha_{kt} + \epsilon_{it} \tag{1.5}$$

where y_{it} is either an indicator for mobility or one of the five measures of wage growth described in Section 1.3.4. Our measure of an individual's outside options is Ω_{it} and the key coefficient is γ. We control for c_{it}, the number of coworkers in an individual's network, and for individual (α_i) and industry-by-time (α_{kt}) fixed effects. The individual fixed effects allow us to account for non-random sorting of individuals into firms and networks. Our estimates exploit month-to-month changes in labor demand at the firms in an individual's network. The industry-by-month controls absorb variation in demand for specific skills. In our theoretical framework, these changes in demand correspond to variation in the arrival rate of offers through public sources, λ^P. We two-way cluster our standard errors at the person and firm level to account for individual differences in mobility preference and for correlation between the wage growth of employees within a firm.[23]

The identifying assumption is that, conditional on the included covariates, changes in labor demand at an individual's former coworkers' current firms are uncorrelated with unobserved determinants of mobility or wage growth: $E[\epsilon_{it}\Omega_{i,t}|c_{it}, \alpha_i, \alpha_{kt}] = 0$. The primary concern is that there are unobserved changes in the demand for a worker's skill that are correlated with Ω_{it} but not captured by our industry-by-time controls.[24] These changes in demand would lead to both mobility and wage growth in either a competitive model or in a bargaining model. In the next section we perform several distinct empirical exercises that support the identifying assumption.

1.5 Impacts on Mobility and Earnings

This section presents the main reduced form results. We find that in periods when a worker's former coworkers' firms are expanding that worker is more likely to make a job-to-job transition or to see an earnings gain, even if she does not move. Further, while the mobility results are driven by observations with values of Ω_{it} in the top decile, workers with more moderate values of Ω_{it} also see earnings gains. These are exactly the empirical predictions presented in Figure 1-1.

The results suggest that firms pay workers less than their marginal product and that workers learn about job opportunities through their former coworkers. Some workers take these new opportunities; others use this information to renegotiate their wages at their current firm. The findings match the predictions of the on-the-job search and bargaining model in Section 1.2, but are inconsistent with both (1) the frictionless neoclassical model and (2) models of wage-setting where firms post wages and commit not to renegotiate (e.g. monopsony models).

[23]A future version of this paper will use Conley standard errors to account for correlation in wage growth across the network. This is difficult to implement given the computing resources available, as there are $N \times T$ distinct networks, each of which is an $1 \times N$ matrix (Conley, 1999; Conley and Topa, 2003). Here N is over 1 million, and T is nearly 100.

[24]In the theoretical model described in Section 1.2, these correspond to unobserved changes in the arrival of offers through public sources, which are correlated with Ω_{it}.

1.5.1 Graphical Evidence Supporting Theoretical Predictions

We start by presenting non-parametric evidence on the relationship between Ω_{it} and job-to-job mobility decisions and earnings growth. We find empirical support for both theoretical predictions depicted in Figure 1-1: (1) changes in outside options positively impact both mobility and wage growth and (2) larger changes in outside options are necessary to induce a job-to-job transition than to induce a wage change.

Mobility The top-left panel of Figure 1-5 plots the raw probability an individual makes a job-to-job transition in a given period by the quality of their outside options. The probability an individual makes any transition is low because of the high frequency (monthly) nature of our data.[25] However, there is a clear positive relationship between an individual's outside options (as measured by Ω_{it}) and their probability of making a job-to-job transition. This is primarily driven by high-value options: those in the top decile. This is consistent with the model in Section 1.2, where an outside offer only induces a job-to-job transition if the outside firm is willing to outbid the incumbent firm.

The top right panel shows that this increased job-to-job mobility is the result of individuals moving to firms where one of their former coworkers works. Each job-to-job transition can be divided into one of three categories based on whether the move is to (1) a coworker-connected firm, (2) an unconnected but in-sample firm, or (3) an out-of-sample firm. Coworker-connected firms are firms where one of the individual's former coworkers currently works. Out-of-sample firms are firms with more than one thousand employees. The null impact on unconnected firms is consistent with information transmission–without a connection one of these firms, the individual is no more likely to hear about the openings than anyone else. The null impact on out-of-sample firms suggests that there is little selection in or out of our sample.

These patterns do not emerge simply because highly mobile workers have stronger networks and larger values of Ω_{it}. The bottom panel of Figure 1-5 shows that the main results hold after par-

[25]This probability is around 2% in our sample. This is similar to the 1.5-2% number reported in Bagger et al. (2014b). About half of these transitions are job-to-job transitions, which do not include an intervening spell of unemployment or non-employment.

tialling out individual fixed effects and industry-by-time fixed effects, and controlling for network size. The figure on the left shows a clear positive relationship between the residual probability an individual makes a job-to-job transition and the residualized outside options. The bottom right panel shows that the results are driven by moves to connected firms. There is no such relationship between the residualized outside options and the residual probability an individual leaves their firm and is not immediately employed at another firm.

Earnings Figure 1-6 shows that there is also a positive relationship between an individual's change in log earnings and Ω_{it}. As before we regress both the outcome—changes in log monthly earnings—and Ω_{it} on individual and industry-by-time fixed effects and on a linear control for the number of coworkers in an individual's network. We plot the mean residuals of our earnings outcomes by percentile of the residual options distribution. The top panel focuses on log wages; the bottom panel focuses on log earnings. The data show a clear positive relationship between earnings changes and outside options in both cases. As with mobility, only large shocks are important; there is no impact on observations below the sixtieth percentile.

The theoretical model in Section 1.2 suggests that larger changes in outside options are necessary to induce a job-to-job transition than to induce a wage change. This is because, in order to induce a worker to switch jobs, the outside firm has to beat the *maximum* the incumbent firm would be willing to pay. In order to induce a wage change, the outside firm only needs to beat the worker's last outside offer. A comparison of Figures 1-5 and 1-6 shows that, while only the top decile of (residualized) Ω_{it} impacts job-to-job mobility, values of Ω_{it} in the top three deciles lead to changes in earnings. This is exactly the pattern predicted by the model in Section 1.2 (see Figure 1-1).

The primary concern with the interpretation of our estimates is that they might reflect unobserved changes in demand for a worker's skills rather than the pure effect of information about outside options, holding skill demand constant. Appendix Figure 1-19 shows that the main graphs are unchanged by the addition of occupation-by-time fixed effects.

1.5.2 Basic Results

We next turn to quantifying the impacts on mobility and wage growth. To scale our results, we examine how a ten unit change in Ω_{it} (approximately one standard deviation) affects the percent chance an individual makes a job-to-job transition or affects the average worker's annual earnings.

Mobility

Panel A of Figure 1-7 displays estimates of γ from equation 1.5. The outcome variable is an indicator for whether the individual made a job-to-job transition. We scaled γ so it indicates the impact of a 10-unit increase in Ω_{it} on the percent chance an individual will make a job-to-job transition this month. The estimate at the far left is our preferred specification, which includes both individual and industry-by-time fixed effects, and controls for the number of connections an individual has. These covariates allows us to control for time-invariant differences in the quality of an individual's network and for the fact that some individuals are more mobile than others.

The baseline estimate at the far left indicates that a ten-unit change in Ω_{it} (roughly one standard deviation) leads to a 15% higher probability an individual will move to a new firm this month. This is consistent with the idea that individuals stay in contact with and discuss labor market opportunities with their former coworkers. When new vacancies arise at these former coworkers' firms, the worker is more likely to hear about the vacancy than the average worker. In some cases, the worker receives an offer and decides to move.

One way to assess whether the identifying assumption—that changes in Ω_{it} are uncorrelated with unobserved changes in demand for the worker's skill—is satisfied is to examine how our estimate of γ changes when we add more detailed non-parametric controls for changes in demand of different skill groups (Altonji et al., 2005). If Ω_{it} reflected changes in demand for a worker's skill, we would expect γ to fall as we added more controls.

The remaining estimates presented in this figure show that our estimate is, in fact, stable when including different non-parametric controls for changes in demand for certain skill groups. The second specification adds time-varying demographic controls—indicators for whether the individual is married or has children—to our baseline specification. The third and fourth specifications use occupation-by-time and industry-by-occupation-by-time fixed effects instead of industry-by-time fixed effects. The stability of our estimates suggests the results are driven by information transmission through social networks, not unobserved demand shocks.

Table 1.4 presents the raw (unscaled) estimates of γ for each mobility outcome. The fact that the impact on job-to-job mobility is larger than the impact on whether a worker makes any transition (including to non-employment), suggests that some workers use information from their coworkers to avoid short periods of un- or non-employment. This result is consistent with prior work that has shown that newly unemployed or displaced workers use information from their former coworkers to find new employment (Saygin et al., 2014; Glitz and Vejlin, 2018).

The second and third rows show that the increase in job-to-job mobility is entirely driven by increases in moves to coworker-connected firms. There is only a small displacement effect: some individuals who would have moved anyway are more likely to move to a firm where they know a coworker, than to an unconnected firm. However, there is not a significant impact on whether the worker moves to an out-of-sample (i.e. large) firm.

Earnings

Panel B of Figure 1-7 displays estimates of γ from equation 1.5 where the outcome variable is the change in log monthly earnings.We have scaled γ so it indicates the impact of on a ten unit change in Ω_{it} on the average worker's annual earnings. The baseline estimate at the far left indicates that a ten unit change in Ω_{it} is associated with a \$50 increase in annual earnings.

As with job-to-job mobility, the estimates are stable across a number of specifications, including different non-parametric controls for changes in demand for different industry or occupation groups. This suggests that, to the extent that there may be omitted changes in demand for a worker's skill, they are not correlated with our measure Ω_{it}. The final column presents the baseline specification estimated on the sample of job-stayers, those who are at the same firm as in the prior month. Not surprisingly, the estimate is nearly identical to the baseline estimate at the far left. Our data are monthly; the vast majority of our sample ($\approx 99\%$) are job-stayers.

Table 1.5 displays estimates of γ from equation 1.5 for different earnings outcomes (rows) and for each specification presented in Figure 1.5 (columns). The fact that we find a larger impact on narrow earnings—which do not include benefits—than our baseline measure suggests that most of this increase is due to changes in take home pay. The third row shows that, among workers with non-imputed hours, there are changes in hours worked. We see significant impacts on both base pay and bonuses. Appendix Table 1.17 shows that we see similar impacts for job-stayers.

While a \$50 average increase may seem low, this estimate masks the fact that most workers do not see any earnings gains in a given month. If all of the gains were driven by individuals who were driven to renegotiate with their firm (and not gains by those would have renegotiated anyway), the impact on compliers' earnings would be

$$\frac{\beta_{\Delta \log y}}{\beta_{1\{\Delta \log y > 0\}}}$$

Focusing on full time workers, this implies an 11% increase in base pay. Because workers who know they will have a chance to renegotiate their wages with their employer are probably especially likely to seek out information about outside opportunities, this is likely an overestimate.[26]

1.5.3 Exploiting Within-Firm Differences in Information

Even within an occupation or industry-by-occupation group, there may be substantial variation in workers' skills. For instance, software engineers may differ in their knowledge of Python, Julia, or C and these skills may be valued by different firms. In some cases skill variation within industries or occupations may be the result of training received at certain sets of firms (e.g. learning how to format code a certain way). As a result, individuals with a shared work history may have skills that are similar in ways we cannot observe.

We can address the concern that our industry and occupation controls are not sufficient to

[26] This is based on coefficients from the baseline regression. The specification with firm-occupation-time fixed effects yields a ratio of 13%.

absorb time-variation in the demand for workers' skills by adopting a within-firm identification strategy. As discussed in Section 1.4, workers within the same firm or firm-and-occupation group may have different networks due to differences in tenure at that firm and at other firms. Figure 1-8 presents estimates of γ that exploit this variation; Panels A and B present results for mobility and earnings, respectively. The first estimate in each panel presents the baseline specification. The second specification adds firm fixed effects; the third replaces the industry-by-time fixed effects with firm-by-occupation-by-time fixed effects. While the standard errors increase, we cannot reject that the earnings and mobility results are the same as in our baseline specifications. The raw coefficients for both mobility and earnings are presented in Tables 1.4 and 1.5, respectively.

We present one additional specification for earnings, focusing on job stayers. The fact that we obtain similar estimates in this sample bolsters the case that earnings growth is the result of worker-initiated renegotiation. An alternative interpretation of our findings is that managers learn about the 'market value' of their employees and raise wages accordingly. This story inconsistent with our finding that, within a firm and occupation, job stayers with more ties to coworkers at expanding firms see larger wage gains. This specification also addresses the concern that Ω_{it} is correlated with unobserved time-specific shocks to an individual's ability a their current firm. Table 1.17 presents the full set of estimates for job-stayers.

1.5.4 Exploiting Different Groups of Coworkers

We can provide further evidence that our results are driven by changes in workers' information about their outside opportunities by decomposing Ω_{it} into the portions coming from different subsets of coworkers. Some subsets of coworkers are more likely to be sources of information than others.

Same- and Different-Region Coworkers

Our first test is based on the geographic location of an individual's former colleagues. We would expect an individual's same-region coworkers to be a more valuable source of information for two reasons. First, individuals are more likely to be in contact with their former colleagues who work in the same geographic area. Second, assuming there are costs to moving, individuals are more likely to obtain actionable information from their same-region coworkers: information about jobs they would likely take. In both cases we would expect changes in labor demand at an individual's nearby coworkers' firms to matter more. By contrast, if our estimates reflected changes in demand for a worker's skill, both sets of coworkers would be roughly equally valuable.

We run regressions of the form

$$y_{it} = \gamma^{\mathrm{IN}}\Omega_{it}^{\mathrm{IN}} + \gamma^{\mathrm{OUT}}\Omega_{it}^{\mathrm{OUT}} + c_{it}^{\mathrm{IN}} + c_{it}^{\mathrm{OUT}} + \alpha_{kt} + \alpha_i + \epsilon_{it} \tag{1.6}$$

where $\Omega_{it}^{\mathrm{IN}}$ and $\Omega_{it}^{\mathrm{OUT}}$ are based on an individual's same-region or different-region former coworkers. These networks are based on the five Danish administrative regions shown in the map in Figure 1-18. Workers are assigned to regions based on the location of the firm they worked for in the prior period. Most workers live in the same region in which they work. Note that, by design, individuals without former coworkers in both their own region and in other regions are excluded. This primarily excludes individuals with very few connections.

We report the results in Figure 1-9. Panel A reports estimates of $\{\gamma^{\mathrm{IN}}, \gamma^{\mathrm{OUT}}\}$ from regressions where the dependent variable is an indicator for whether the worker made a job-to-job transition; Panel B reports analogous results from regressions where the dependent variable is the change in log monthly earnings. Each regression controls for the number of connections the worker has in the same region (c_{it}^{IN}) and in other regions (c_{it}^{OUT}), and for individual and industry-by-time fixed effects.

Both panels clearly show that changes in demand at a worker's same-region coworkers' firms have a significant and positive impact on whether the worker moves to a new firm or experiences earnings growth. New positions at the individual's different-region coworkers' firms have a much smaller effect. For each outcome, we can soundly reject equality of the two coefficients. This is exactly what we expect if individuals are more likely to lose contact with their former coworkers who move to, or start working in, different regions. If, by contrast, Ω_{it} simply reflected changes in the demand for a worker's skill, we would find that $\gamma^{\text{IN}} = \gamma^{\text{OUT}}$. The effect of an individual's different-region coworkers is more precisely estimated than that on an individual's same-region coworkers. This reflects the fact that Ω_{it}^{OUT} is constructed using former coworkers from four regions, relative to a single region. Table 1.6 presents coefficients for the pooled sample of men and women.

The fact that a worker's outside-region coworkers impact mobility, but not earnings is consistent with the idea that workers may not be able to use outside offers at geographically distant firms as leverage with their current employers. Because our data are monthly, earnings impacts are driven entirely by job-stayers. Employers may doubt a worker's willingness to relocate and may, as a result, not see an outside offer from a distant firm as a credible threat. This may also explain why women see lower earnings gains than men: they are less likely to move (top panel).

The main concern with this test is that, if mobility across regions is low, the coworkers who work in different regions may have a different set of skills from those who stay in the same region. This seems somewhat unlikely in the Danish context: conditional on a job-to-job transition, roughly half of workers start working in a different region (Kristoffersen, 2016). This partially reflects the fact that Denmark is a small country: Denmark's two largest cities, Aarhus and Copenhagen, are just a three hour drive apart. Further, Table 1.14 shows that a sizable fraction of an individual's former coworkers now work in different regions. A related concern is that labor markets might be very local: even within an industry, firms in different regions may produce different products or use different combinations of workers' skills. We think that this is somewhat less of a concern than in it would be in other contexts, because Denmark is not very large.

Past versus Future Coworkers

Our second test exploits the fact that some coworkers are more or less likely to provide the worker with information, because of when they worked together. The logic of this test is simple: because workers may lose contact with their former colleagues over time, coworkers a worker worked with in the more distant past are likely to be less valuable sources of information. Further, while an individual's future coworkers likely have similar skills in ways we can and cannot observe, they are less likely to be a source of information in the current period (because they have not yet worked together).[27] These coworkers therefore give us another way to control for changes in demand for a worker's skills. We construct distinct networks comprising individuals the worker worked with (1) 4-5 years ago, (2) 2-3 years ago, and (3) 1 year ago and workers the worker will work with in (4) 1 year, and (5) 2-3 years. In some specifications we divide the third network into coworkers an individual worked with in the past six months and coworkers an individual worked with between 6 and 12 months ago. We describe how we create measures of Ω_{it} for each of these groups in Appendix 1.12.6.

If our results were driven by information transmission, we would expect to see three patterns. First, changes in labor demand at firms at their more recent former coworkers would matter more than changes at firms of coworkers they worked with in the more distant past. Second, changes at an individual's *future* coworkers' firms would not significantly impact wage growth. Third, the coefficients on past measures of Ω_{it} from the "short" regression—that includes only shocks to an individuals former coworkers' firms—should be equal to those in the "long" regression that adds controls for shocks to individual's future coworkers' firms. By contrast, if the results were driven by unobserved demand shocks, we would expect the coefficients on Ω_{it} to fall significantly when we added controls for changes in demand at an individual's future coworkers' firms.

Table 1.7 presents estimates of γ^n from equation 1.7 (columns 2 and 6) and equation 1.8 (remaining columns) for two outcomes: job to job mobility and changes in log monthly earnings. Column 1 of Table 1.7 confirms that the average number of coworkers an individual has each year varies in proportion to the number of years included in the network.

$$y_{it} = \gamma^n \Omega^n_{it} + c^n_{it} + \delta X_{it} + \alpha_i + \alpha_{kt} + \epsilon_{it} \quad (1.7)$$

$$y_{it} = \sum_n (\gamma^n \Omega^n_{it} + c^n_{it}) + \delta X_{it} + \alpha_i + \alpha_{kt} + \epsilon_{it} \quad (1.8)$$

[27] These future colleagues may still be a source of information if they are connected to the worker in other ways–e.g. through family or education networks.

We find empirical support for all three predictions. First, the coefficients in column 2 (also presented graphically in Figure 1-10) show that the effects on mobility decline monotonically as we move from examining the effects of her prior year coworkers to those of the coworkers she last worked with 4-5 years ago. When we include each of these former coworker networks in a single regression (column 3), the same pattern emerges, though some of the standard errors increase, reflecting the fact that there is overlap in the firms covered by each network. Columns 6 and 7 present analogous results for changes in log earnings monthly. Second, changes in demand at her future coworkers' firms are much less important and have no impact on wage growth.[28] Third, the coefficients in the short regressions (columns 3 and 7) are not significantly different from those that include the future network controls (column 4 and 8).

The main concern with this falsification exercise is that an individual's skills may also change over time. While a combination of unobserved demand shocks and rapid changes in skill could explain some of our results, it would not explain the fact that an individual's prior year coworkers influence her wage growth, while her future year coworkers do not.

1.5.5 Dynamics

We next examine how the effects play out over time. We plot estimates of γ from models of the form

$$y_{i,t+j} \quad = \quad \gamma\Omega_{it} + \gamma_j \Omega_{i,t+j} + c_{ij} + \alpha_i + \alpha_{k,t+j} + \epsilon_{it} \tag{1.9}$$

The coefficients describe how variation in Ω_{it} impacts mobility decisions and wage growth in subsequent months, after controlling for variation in outside opportunities in those periods.

Mobility Figure 1-11 shows that moving from an average of one vacancy per former coworker to two vacancies per former coworker increases the probability that an individual will make a job-to-

[28]While there is a statistically significant impact of changes in demand at an individual's next-year coworkers' firms on mobility (not earnings), the impact is less than a third of the size of that for her prior year coworkers, and is smaller than the impact of her 4-5 year removed coworkers.

job transition this period. There is a somewhat negative impact on mobility the next month, but no impact in subsequent months. There is no impact on whether individuals exit to non-employment or unemployment or on whether they move to an unconnected or out-of-sample firm. The results are driven entirely by moves to connected firms. This indicates that, after controlling for the value of Ω_{it} in a given period, there is no additional impact of past values of Ω_{it} (past outside options).[29]

Earnings Figure 1-12 shows the dynamic effects of Ω_{it} on an individual's base pay (bottom panel) and bonuses (top panel). The top panel shows that there is an immediate impact on idiosyncratic 'bonus' pay. There is a negative impact on bonus pay in future periods, after controlling for future values of Ω_{it}. The bottom panel shows that individuals' base pay takes two months to adjust. This is what we would expect to see if it takes a while to negotiate with one's boss for a raise.

Figure 1-13 presents estimates of γ from a regression of the change in log base pay in period $t+j$ (relative to period t) on Ω_{it} and on that period's value of $\Omega_{i,t+j}$:

$$y_{i,t+j} - y_{i,t} = \gamma\Omega_{it} + \gamma_j\Omega_{i,t+j} + c_{ij} + \alpha_i + \alpha_{k,t+j} + \epsilon_{it} \tag{1.10}$$

This figure shows that changes in base pay do not revert in the short run; four months later, the worker is still earning more.

1.5.6 Exploiting Trade-Induced Changes in Labor Demand

We can also address the concern that our results are driven by unobserved changes in the demand for workers' skills by examining changes in labor demand that are driven by changes in global demand for each firm's exports. Changes in world demand for different products may lead firms to expand employment or raise wages for incumbent workers (Hummels et al., 2014; Garin and Silverio, 2018).

[29]The fact that the moves occur very quickly is likely due to the fact that our measure of outside options uses realized hires or positions created. These reflect vacancies that were posted one or two months prior. This also reflects the fact that firm labor demand is highly serially correlated; firms that expand in one month often expand in the next. Each regression controls for that period's value of Ω_{it}, which enters significantly. We do not find such fast adjustment with the trade-based measures used in Section 1.5.6.

Instrument Because realized changes in a firm's exports may be confounded by changes in firm productivity or changes in local conditions, we follow the prior literature and use world export demand to construct a measure based on predicted firm-level exports. We construct $\hat{\Omega}_{it}^{\text{trade}}$ in three steps. First, we use data from the first six years of our administrative trade register (2004-2009) to calculate the share of Danish exports of each six-digit product p accounted for by each firm j, π_p^j. Fixing the product shares using pre-period data ensures that our measure of demand for a firm's exports does not respond to changes in firm productivity. Second, we weight monthly measures of total world exports of each product (less exports from Denmark) from COMTRADE by these firm-product weights. Our COMTRADE data begin in 2010 and run through the end of our register.

The first two steps are similar to those used in prior work.[30] In a third step we weight firm-specific measures of log predicted exports by each individual's coworker network. Because most individuals do not work in firms covered by the trade register, we use weights based on the fraction of former coworkers who are in exporting firms. All of our regressions control for the total number of workers in an individual's network that are in exporting firms (c_{it}^{trade}). More information on this instrument is provided in Appendix 1.12.7.

Reduced Form Table 1.8 presents estimates of γ from the regression:

$$y_{it} = \gamma^{\text{trade}}\hat{\Omega}_{it}^{\text{trade}} + \beta X_{it} + c_{it}^{\text{trade}} + \alpha_{kt} + \alpha_i + \epsilon_{it} \qquad (1.11)$$

for different outcomes y_{it}. All of the regressions control for individual fixed effects to control for time-invariant differences in network quality and trade-register coverage. They also include controls for the number of coworkers in an individual's network and the share covered by the trade register. Standard errors are two-way clustered by individual and firm. Note that the sample in this table differs from that in our usual tables, because individuals who do not have any former

[30] In particular, Hummels et al. (2014) also used fixed product shares to weight world demand for exports. However, they also used measures of transportation costs and import demand. They find a negative relationship between changes in log (annual) earnings and changes in log predicted imports (an offshoring index) at the firm level. Our work differs from theirs in that we focus on the pass-through to workers at different firms.

coworkers in firms covered by the trade register are excluded.

Columns 1 and 2 show that $\hat{\Omega}_{it}^{\text{trade}}$ is positively related both to (1) a measure of $\Omega_{it}^{\text{trade}}$ computed using actual (not predicted) firm exports and (2) the measure of outside options used in the prior sections. The second result is consistent with earlier research: firms that experience increased demand for their exports expand the size of their labor force.

Columns 4-6 present the reduced form. These columns show that there is also a relationship between $\hat{\Omega}_{it}^{\text{trade}}$ and both mobility and wage growth: when an individual's former coworkers' firms see more demand for their exports, that individual is more likely to move or see an increase in earnings. Panel B shows that this relationship is robust to the inclusion of firm-occupation-time controls. These controls allow us to account for the fact that changes in product demand at an individual's former coworkers' firms is likely correlated with changes in demand for the exports at the individual's own firm.

1.5.7 Robustness Checks

We have conducted a number of robustness checks to ensure that the results in the previous sections are not driven by choices we made in constructing our measure of outside options or in constructing our regression sample. In one set of checks we verified that we obtain similar results when using different definitions of an individual's coworker network. When constructing these networks we had to specify both a window of time over which to define the network, and a firm size cutoff. Appendix Tables 1.21 and 1.22 present estimates of γ from equation 1.5 using three alternate definitions. The first two columns of each table use the baseline firm-size threshold but use a two-year or five-year window. The third column uses the baseline window of time but excludes connections formed in firms with more than 500 workers, rather than 1000 workers.

We also test whether our results are robust to including former coworkers who move to large firms. We removed these coworkers because new positions at very large firms are likely to be known to all workers. We find that relaxing this does not meaningfully impact the results. One reason for this is that many workers do not have any connections to these firms. Figure 1-20 shows that adding versions of Ω_{it} calculated among the set of former coworkers who now work at large firms ($\Omega_{it}^{\text{large}}$) does not change the main estimate. The coefficient on $\Omega_{it}^{\text{large}}$ is orders of magnitude smaller, and is not stable across specifications. This is exactly what we would expect to see if information about new positions at these firms is common to all workers.

In a second set of checks, we verified that our qualitative findings also emerge when considering alternative measures of firm shocks (s_j) or alternative weighting functions (ω_j). Appendix Tables 1.19 and 1.20 compare our baseline estimates of equation 1.5 to those computed using four alternative measures. The first pair (columns 1 and 2 of each table) is based on the number of new positions at each firm; column 1 presents our baseline measure and column 2 presents the same measure, weighted by the mean wage at each firm. The second pair (columns 3 and 4) uses the leave-out number of hires, rather than the change in employment. While it is difficult to compare magnitudes across columns, the qualitative patterns are the same.

We have run a number of additional checks in addition to those reported in this paper. First, in results not reported here we show that there are similar effects on mobility when we include multiple job-holders, or when we relax our definition of multiple job-holding. However, for individuals disposed to multiple job-holding, changes in information about outside options also lead an individual to obtain an additional job (at a coworker connected firm). We have also verified that the qualitative earnings results are similar for the full worker sample when we sum earnings across all firms an individual works for in a given month. It is hard to interpret these results as our theory is about single job holders. Second, in a separate analysis we also interacted our measure of outside options with year fixed effects. We did not find any systematic differences across years of our sample. Third, we have verified that we obtain similar results using data from an annual employer-employee register (Danish IDA register). Those results are noisy, however, reflecting the fact that we have a very short panel of annual data. We describe some of these checks at more length in Appendix 1.12.

1.6 Heterogeneity and Mechanisms

We find that the impacts on earnings are primarily driven by changes in hourly earnings, not changes in hours worked and that both job movers and job stayers benefit. We also find that earnings impacts are concentrated among workers in the top half of the skill distribution. The results suggest that firms may not renegotiate wages with low skilled workers who receive outside offers. Within skill groups, women gain less than men.

1.6.1 Hours versus Hourly Earnings

We first examine whether the results are driven by changes in hours worked or changes in hourly earnings. For this analysis we focus on the subset of observations non-imputed hours.[31] We use the accounting identity

$$d \log y \;=\; d \log w + d \log h$$

The overall impact on log earnings depends on how both hours and hourly earnings change.

Figure 1-14 plots the ratio of the coefficient γ estimated from regressions of log hours (numerator) and log earnings (denominator). As usual, each regression controls for industry-by-time fixed effects, individual fixed effects, and the number of connections included in Ω_{it}.

The estimate at the far left shows that most—more than three quarters—of the impact on monthly earnings is the result of changes in earnings per hour, not changes in hours worked. However, there is heterogeneity across groups. While changes in hours worked explain only 14% of the impact for college-educated workers, they explain nearly 28% of the impact for non-college workers.

[31]We require both this month's hours and the previous month's hours to be reported by the firm. This excludes observations for workers who move between firms with different reporting statuses.

1.6.2 Movers and Stayers

We next examine the relative returns for job-stayers and movers. While this is a descriptive exercise, it is a useful one. Models where individuals cannot use outside offers to renegotiate wages (e.g. posting models) would predict a ratio of 0. Models where wages perfectly reflect the price of a worker's skill (and there are no 'match' effects) would predict a ratio of 1.

Figure 1-15 presents estimates of $\frac{\gamma^S}{\gamma^M}$ from

$$\Delta \log w_{it} = \gamma^S \Omega_{it} \times Stay_{it} + \gamma^M \Omega_{it} \times \text{Move}_{it} + \beta X_{it} + \alpha_i + \alpha_{jt} + \alpha_t \times \text{Move} + \epsilon_{it} \quad (1.12)$$

The baseline specification controls includes all of our baseline controls, as well as time-varying differences in the value of staying or moving. For this exercise, we focus on the subset of workers with non-overlapping job spells, in order to make sure that the earnings changes for movers reflect a full month's pay. The baseline estimate, presented at the far left, shows that, on average, stayers capture 20% the gain of movers. We can firmly reject zero. The remaining columns add additional controls. The second column adds time-varying demographic controls; the remaining columns add combinations of industry and occupation fixed effects. The point estimate remains stable across a variety of specifications.[32]

1.6.3 Heterogeneity: Skill Groups and Gender

It is not surprising that some workers, such as academic economists, use outside offers as leverage to obtain a raise. However, the implications of many labor market policies depend on the nature of competition in the market for relatively homogeneous workers. In order to examine whether this link between options and wages is important for workers throughout the skill distribution, we examine heterogeneity across different occupations.

We divide workers into 8 categories, corresponding to broad ISCO (International Standard

[32]Even if movers and stayers with the same value of Ω_{it} were equally likely to have heard about an offer through their former coworkers, those that chose to move should have—according to the model in Section 1.2—received better offers. The relative return for movers and stayers is likely a lower bound on the fraction of rents workers are able to capture.

Classification of Occupation) codes: (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) elementary occupations.[33] We then estimate our baseline regression (equation 1.5) separately for workers in each group.

Figure 1-16 presents estimates of γ for each occupation group and gender. We produced these estimates by interacting Ω_{it} with indicators for whether the worker is male or female. Panel A presents results for mobility; Panel B presents results for earnings. We find that, while there is variation in magnitudes, the impacts on mobility are significant and positive for all subgroups. This suggests that workers throughout the wage distribution use information they obtain from their former coworkers to find new labor market opportunities. Women within each skill group are slightly less responsive than men but the differences are not significant.

We find that the earnings effects are largest for high skilled workers: those in the 'professional' category. There is no impact on assembly workers, manual-skilled workers, or craftsmen. Panel B of Table 1.9 presents estimates of these parameters, scaled to represent the impact of a ten unit change in Ω_{it} on a worker's annual earnings. Because more skilled workers typically earn more, scaling the parameters by mean earnings magnifies the differences between groups.

Because our earnings estimates are driven by job-stayers, our results suggest that wage renegotiation and bargaining is an important channel of wage growth only for workers in the top half of the skill distribution. By contrast, workers in the lower half of the skill distribution are more likely to be in jobs where wage renegotiation is less important for wage growth. In section 1.7, we show that this reflects differences in the probability that these individuals are at firms that renegotiate wages (not simply lower bargaining power).

Within skill groups that see earnings gains, women's earnings respond less than men's do. The results are in line with recent research showing that women obtain a smaller portion of changes in firm rents (Card et al., 2016b; Kline et al., 2018). There are several possible mechanisms. For instance, we would see this pattern if women are less likely to initiate wage renegotiation in

[33]There are 10 broad ISCO categories. We do not have data on workers in the armed forces and very few workers are classified under the "skilled agricultural, forestry, and fishery workers" category.

response to an outside offer (Bowles et al., 2007; Babcock and Laschever, 2009). We would also see this pattern if women are equally likely to initiate wage renegotiation, but are less successful in bargaining.

1.7 Structural Parameters

Finally, we use the reduced form estimates to identify the structural model described in Section 1.2. The model allows us to estimate two key parameters of interest: (1) worker's bargaining power β and (2) the fraction of offers from 'posting' firms $(1 - p_R)$. This allows us to determine whether the heterogeneity we observed in Section 1.6.3 was the result of lower skilled workers having lower bargaining power or being less likely to work in firms that are willing to renegotiate. This is important both for distinguishing between classic models of wage setting under imperfect competition (i.e. monopsony and search models), and for determining how changes in the labor market will influence workers. Greater values of β mean greater pass through from options to wages. However, higher fractions of wage-posting firms mean that workers are only able to see wage gains if they switch jobs. We estimate the model separately for each of the occupation groups described in Section 1.5. In Section 1.7.4 we use these estimates to examine how a decrease in the arrival rate for employed workers would impact both the overall level of wages and the level of wage growth.

1.7.1 Setup

We make a number of parametric assumptions before taking the model described in Section 1.2 to the data. First, we follow prior work and fix the monthly discount rate at $\rho = \dfrac{1}{1 + .0050} \approx .995$ (Bagger et al., 2014b). Second, we allow posting and renegotiating offers to come from two distinct distributions. We assume that both are log normally distributed with means and variances (μ_P, σ_P) and (μ_R, σ_R), respectively.

Workers face different job arrival rates when they are employed and unemployed. The mean arrival rate is λ^U for an unemployed worker and λ^E for an employed worker. There is also variation in arrival rates across workers with the same employment status because some workers have better access to information than others. This varies both within an individual worker over time and across individuals within a time period. If a fraction s of an individual's network is expanding, she faces the arrival rates:

$$\begin{aligned} \lambda^E(s) &= \lambda^E + \tilde{s}\alpha^E \\ \lambda^U(s) &= \lambda^U + \tilde{s}\alpha^U \end{aligned}$$

where the α are scaling parameters and $\tilde{s}$ is the deviation of her information quality from that of the average worker. We assume that $\tilde{s}$ is drawn from a normal distribution with mean zero. We cannot separately identify the variance of $\tilde{s}$ and α.

Table 1.10 lists the 12 parameters we estimate. There are four parameters governing the offer distributions and four parameters governing the job arrival rates. The remaining parameters are: the exogenous job destruction rate, the value of non-employment, the fraction of offers coming from "posting" firms, and the bargaining power parameter.

1.7.2 Estimation Strategy

We estimate the model using simulated method of moments. The strategy finds values of the structural parameters that minimize the distance between a set of observed moments and the same moments calculated from a simulated version of the model. We use ξ to denote the true value of the parameters in our model. Our estimate of ξ minimizes the weighted distance between the simulated moments (given ξ) and the observed moments S_N:

$$\hat{\xi} = \arg\min_x (S_N - S(x))' W (S_N - S(x)) \tag{1.13}$$

The method is intuitive, but computationally intensive. For each guess of the parameters we simulate a panel of worker histories with 20,000 workers and 100 periods. We then calculate moments implied by this panel. Some of these moments are simple means; others are coefficients from linear regressions using variables in our simulated panel. We then calculate the weighted distance between these simulated moments and those observed in our data. More details are provided in Appendix 1.12.8.

Moments We identify the parameters in Table 1.10 using three sets of moments. The first two sets are standard in the literature and are based on transition rates and moments of the log wage (and log wage change) distribution. The third set comes from our reduced form estimates. Table 1.11 lists the full set of moments.

1. **Transition rates** Monthly job-to-job, employment to non-employment, and non-employment to employment transition rates.
2. **Log Wages**
   (a) Mean and residual variance of log wages[34]
   (b) Mean and variance of log wage changes for job stayers
   (c) Mean log wage change associated with a job-to-job transition[35]
   (d) Quantiles of log wage change distribution: 25th, 50th, 75th, and 95th percentiles
3. **Regression coefficients** We also match the coefficient on Ω_{it} from regressions with the following dependent variables:
   (a) $1\{\text{Job-to-Job Transition}\}$, estimated on the sample of employed workers
   (b) $1\{\text{U2E Transition}\}$, estimated on the sample of un- and non-employed workers
   (c) $\Delta \log y_{it}$, estimated on the sample of job-stayers
   (d) $1\{\Delta \log y_{it} > 0\}$, estimated on the sample of job-stayers

In our model there is no non-employment; individuals are either employed, or searching for work. We also do not separate non-employment and unemployment spells in our data. Any transition where an individual was not present at a firm in one month but was in the following month is counted as a U2E transition.

[34] We do not match overall variance of log wages in order to avoid estimating the ability distribution. Instead we match the variance after removing individual fixed effects. This is the same as the approach taken by Jarosch (2015).

[35] A prior version of this paper did not attempt to match this moment. We found that matching this improved the overall fit of the model.

We calculate these moments using the subset of workers we observe working full-time jobs (or, in the case of unemployed workers, those whose last job was a full-time job). For earnings, we focus on the base pay measure, which does not contain annual bonuses, severance pay, or other one-time payments that we can identify. Both of these choices are motivated by the fact that our model does not allow for changes in labor supply (or month-to-month fluctuations in hours worked) and the fact that we have not allowed for measurement error in earnings.

Identification We can consider the sources of variation that are used to identify each parameter in Table 1.10. First, the three transition moments—U2E, J2J and E2U—provide the variation necessary to identify δ, λ^E and λ^U. The employment-to-unemployment transition rate provides information about δ; the job-to-job transition rate provides information about the mean job arrival rate for employed workers, λ^E; and the unemployment-to-employment transition rate provides informative about the mean arrival rate for unemployed workers, λ^U. The reduced form coefficients from the two mobility regressions identify the two α parameters.

The regression coefficients for wages are informative about the key parameters of interest: β and p_R. Intuitively, larger values of p_R decrease both the probability job-stayers see wage gains and the average size of these gains. For a fixed value of p_R, larger values of β increase the value of each outside offer, and increase the size of the wage gains. Finally, the quantiles of the log wage change distribution provide more information about the offer distributions and about our bargaining parameters β and p_R.

1.7.3 Posting and Bargaining

Table 1.10 presents parameters we estimated for the full sample.[36] There are two main results.

First, a substantial portion of offers—more than fifty percent—come from firms that would not be willing to renegotiate wages. Table 1.10 shows that low skilled workers are more likely to be in jobs where they cannot renegotiate their wage.[37] This is not simply a feature of the Danish labor

[36] Standard errors are a work in progress.

[37] Note that we do not plot $1 - p_R$, but the equilibrium fraction of workers at posting firms. This is somewhat more

market. While low skilled workers are more likely to be in firms or sectors that do not negotiate wages individually (Dahl et al., 2013), Appendix Figure 1-22 shows that we find similar results when we break down responses from the United States-based survey analyzed in Hall and Krueger (2012).[38] This suggests that models that feature wage posting (e.g. monopsony models) may be more appropriate for lower skilled workers; models that allow for individual-firm bargaining and renegotiation may be more appropriate for higher skilled workers.

Second, job search through networks appears to be more important for non-employed than employed workers: $\alpha_0/\alpha_1 > 1$ for both groups of workers. This is consistent with prior work that has shown that new labor market entrants are more likely to rely on their family networks for employment in economic downturns (Kramarz and Skans, 2014).

1.7.4 Impact of a Fall in the Job Arrival Rate

Many current policy debates center on regulations that impact workers' ability to receive and take offers from other firms. For instance, changes in the enforcement of non-poach or non-compete clauses directly impact workers' ability to move between firms; changes in antitrust enforcement influence workers ability to receive outside offers, by changing the number of outside firms. Our model allows us to investigate the mechanisms through which these developments would impact different groups of workers.

We re-estimate the model, assuming a 50% reduction in λ_1 and α_1. We then compare the old and new steady states, ignoring transition dynamics. Note that this is a partial equilibrium exercise. The assumption that posted wages do not change in response to a change in the on-the-job arrival rate is somewhat unrealistic. However, it is a useful benchmark if firms that post wages may be unable or unwilling to change these wages in the short run. Recent research suggests that there

informative. It is, in general, lower than $1 - p_R$ because firms that are willing to renegotiate wages do a better job of retaining workers.

[38] We use question 34D from that survey, which asked job seekers: "When you were offered your (current/previous job), did your employer take-it-or leave-it offer or was there some bargaining that took place over the pay?". We calculate the fraction of workers in each occupation who reported that there was some bargaining over pay. More details are provided in Section 1.13.3.

may in fact be substantial downward nominal rigidity in posted wages for new hires (Hazell and Taska, 2018).

Table 1.12 presents the main results. The first column shows that, in response to a 50% reduction in the arrival rate of offers for employed workers, mobility falls by less than 50%. This is because workers correctly anticipate that they will be less likely to receive offers while employed, they impose a higher bar on jobs that they accept out of unemployment. The second column shows that equilibrium wages are now lower. While the absolute magnitude is small (<1%) for both groups, the absolute magnitude of the change in the arrival rate was also small (~1%). Wage growth is significantly lower for both groups.

1.8 Conclusion

This paper uses a novel empirical strategy to show that changes in an individual's information about their outside labor market opportunities lead to job mobility and wage growth. The results are consistent with search and bargaining models where firms renegotiate wages with workers who receive outside offers. The results are inconsistent with both a competitive neoclassic model, and with models where all firms commit not to renegotiate workers' wages (pure posting models). They also suggest that bargaining is important for a wide range of workers, not just those at the very top of the skill distribution.

The reduced form results have several immediate policy implications. First, our finding that workers are able to leverage changes in their information about labor market opportunities into increased pay suggests that pay transparency policies—which give workers information about what they could receive at other firms—may be an effective way to promote wage growth. Second, the results are consistent with recent arguments that increases in labor market concentration or changes in regulations that restrict worker mobility may have detrimental impacts on wages (Council of Economic Advisors, 2016; Krueger, 2017; Ashenfelter and Krueger, 2018). Finally, our results suggest that there are strategic interactions in firm-wage setting. As a result, policies that encourage productive firms to open a plant in a labor market—through local tax breaks or other incentives—may be an effective way to boost wages of all workers (Acemoglu, 2001; Green, 2015). In addition, policies that nominally target the wages at some firms (e.g. large firms) may benefit workers throughout the economy.

One limitation of this paper is that we have limited information about an individual's information set. Future research could use new sources of data from social or employment networking sites to better identify the former coworkers an individual interacts with. A particularly interesting direction would be to work with such a platform to directly vary workers' information about labor market opportunities. A different direction would be to gather direct evidence on the frequency and nature of wage renegotiation, perhaps by collecting survey data analogous to that in Hall and Krueger (2012) for a large sample of employed workers.

1.9 Tables and Figures

Figure 1-1: Main Theoretical Predictions

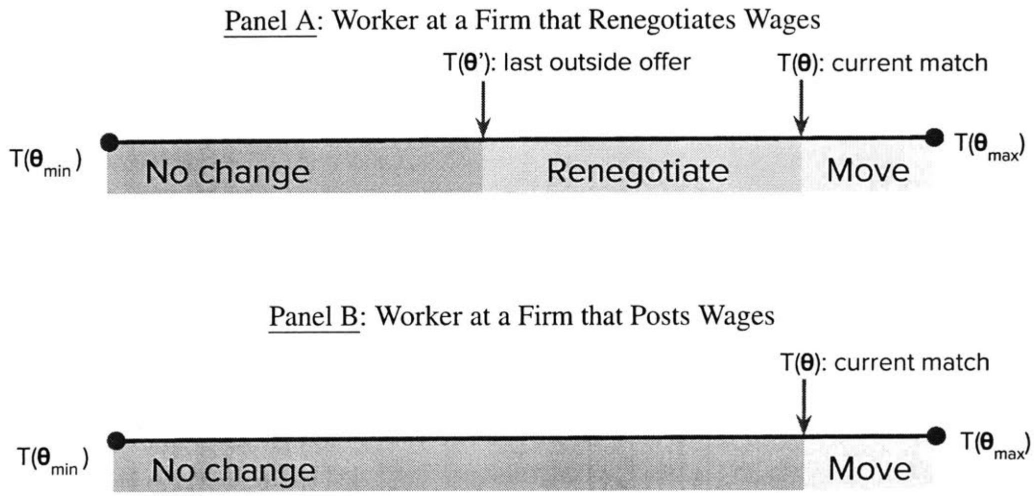

Note: This figure illustrates the main theoretical predictions in Section 1.2. Offers are ranked according to the maximum wage a worker could receive. For offers from renegotiating firms, this is the total value produced by the match; for offers from posting firms, this is the posted wage. Panel A shows what will happen to a worker at a renegotiating firm who receives an outside offer. If the outside offer is higher than the total value of her current match, $T(\theta)$, she will move to the new firm. If the offer is lower than the total value, but is better than whatever she last used to negotiate ($T(w')$, she will renegotiate with her firm for a raise. If the offer is lower than this, she will not initiate renegotiation. Panel B shows that, for workers at posting firms, outside offers can only lead to job-to-job mobility.

Figure 1-2: Variation in Outside Options over Time

Panel A: Network Structure

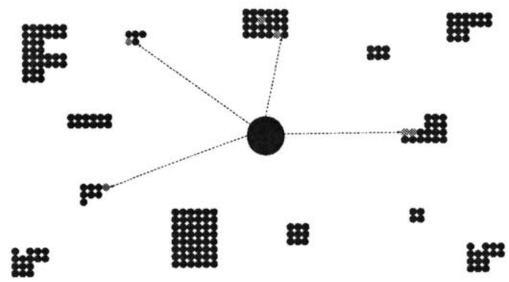

Panel B: Only Unconnected Firms are Hiring

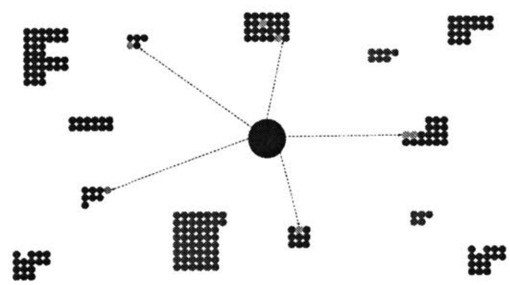

Panel C: Connected and Unconnected Firms are Hiring

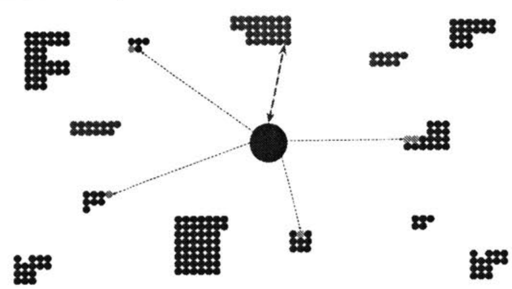

Note: This figure illustrates the identification strategy. Panel A shows the network structure. The big blue dot in the middle represents worker i. Each collection of dots represents a firm; each dot within a collection is a worker. The blue dots are workers that worker i has worked with in the past. Panels B and C depict a scenario where some of the firms (marked in red) expand. In Panel B, worker i does not have any former coworkers at the expanding firms; in Period C she does. Our identification strategy assumes that worker i is more likely to hear about job openings in the situation presented in Panel C than the situation in Panel B.

Figure 1-3: Variation in Networks Within a Firm

Panel A: Incumbent Worker Has Better Information

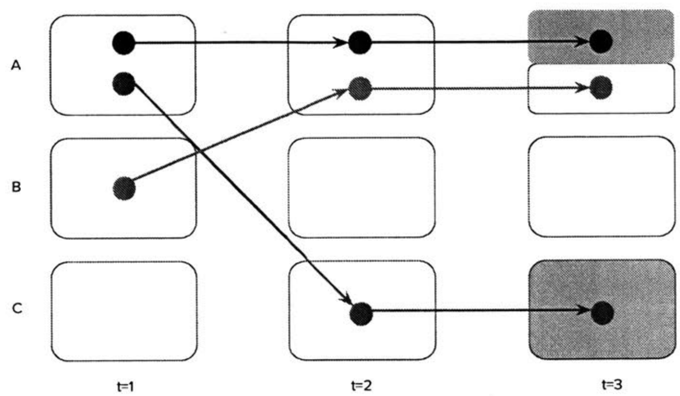

Panel B: New Worker Has Better Information

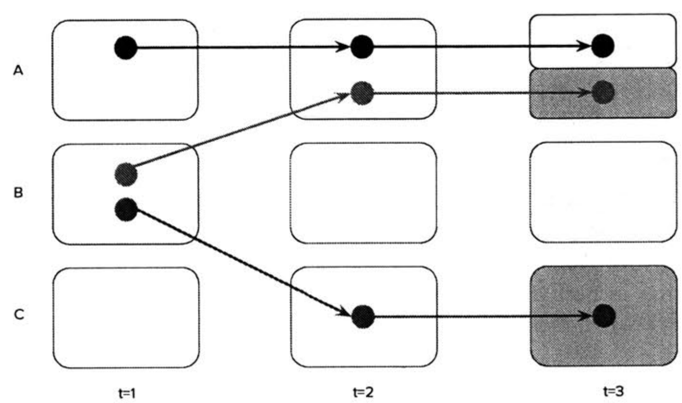

Note: Coworker networks can vary between workers in the same firm both due to their history at other firms and due to differences in tenure at their current firm. This figure shows how networks vary between workers in the same firm due to differences in tenure. Panel A shows an example where the incumbent (blue) worker has better information than a new worker (red). In the first period, the blue worker works with the purple worker at firm A; the red worker is alone at firm B. In period two, the red worker moves to firm A and the purple worker moves to firm C. In the third period the blue worker's coworker network will include the purple worker (firm C) and the red worker's will not. Panel B shows a similar example where the worker with less tenure at firm A (red) is more closely connected to firm C.

Figure 1-4: Graphical Depiction of Timing

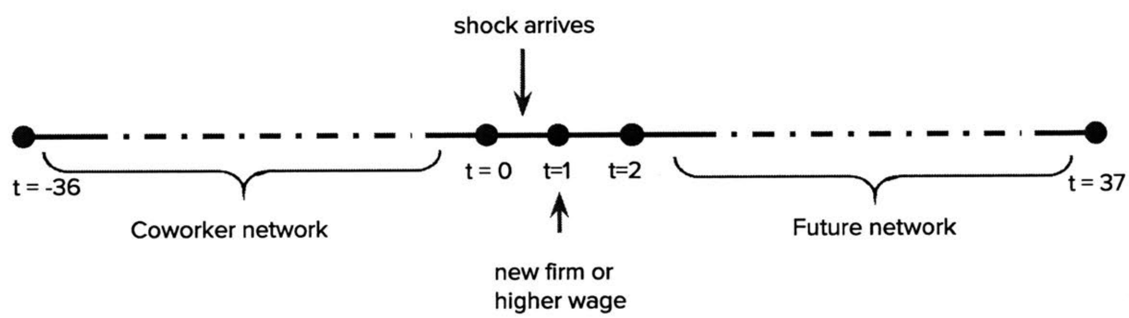

Note: This figure shows the timing of the shocks and coworker networks. For each month, we use data from the previous 36 months to construct the coworker network (excluding a worker's prior firms). We use changes in employment from last period (period 0) to this period (period 1) to construct the firm-specific shock. We look at mobility decisions and earnings changes from period 0 to period 1. We use data from the next 36 months (starting in period 2) to construct the future coworker network. We exclude current coworkers from this network when there is overlap.

Figure 1-5: Impact of Outside Options on Probability of Moving

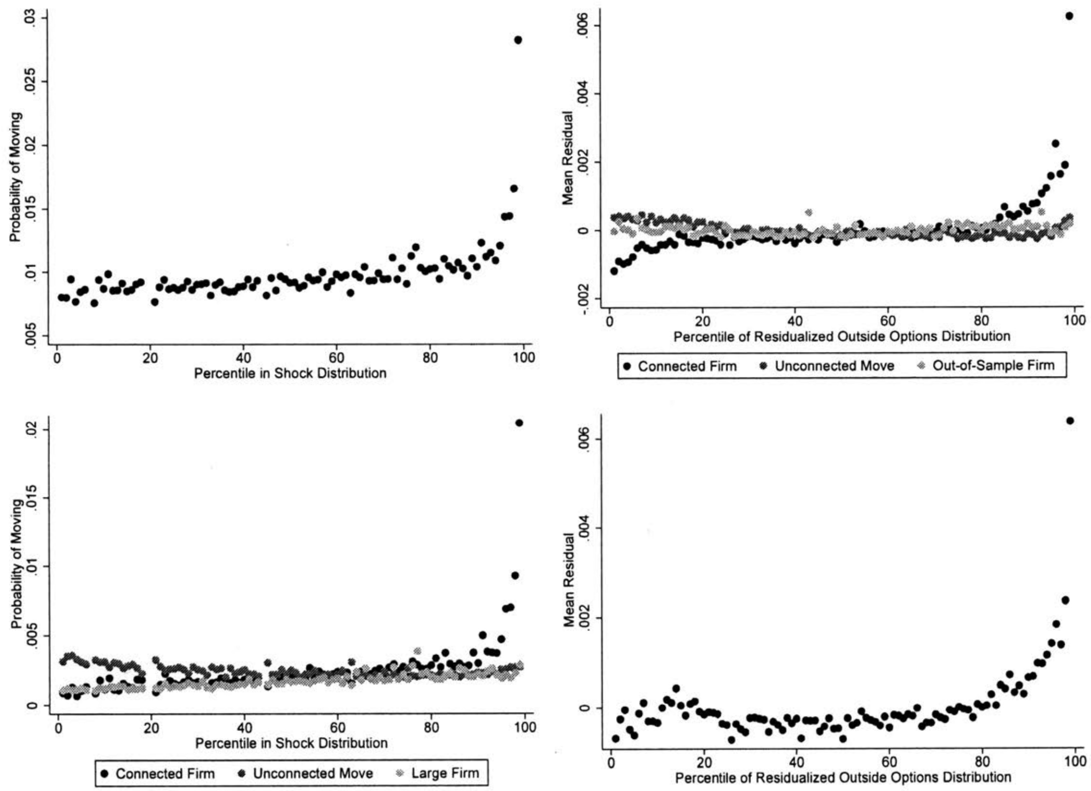

Note: This figure shows how the probability of making a transition depends on Ω_{it}. The percentiles and probabilities in the top panel are computed from the raw data. The percentiles and probabilities in the bottom panel are computed after partialling out individual and four-digit industry-by-time fixed effects. An individual makes a job-to-job transition if they are working at a different firm this month than they were working at last month. A connected move is a job-to-job transition to an in-sample firm where one of the individual's former coworkers works. An unconnected move is a job-to-job transition to an in-sample firm where an individual does not know any employees. An out-of-sample move is a job-to-job transition to a firm whose average employment exceeds 1000 over the sample period.

Figure 1-6: Impact of Outside Options on Changes in Log Earnings and Hourly Earnings

Panel A: Hourly Earnings

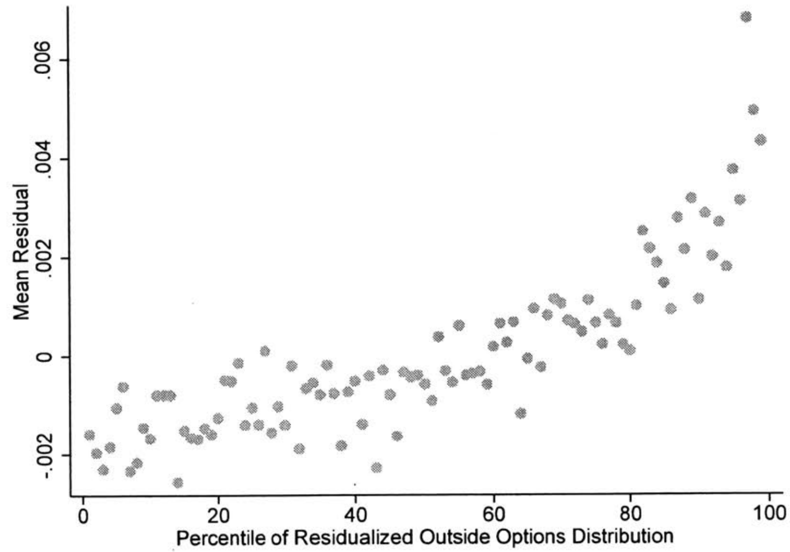

Panel B: Monthly Earnings

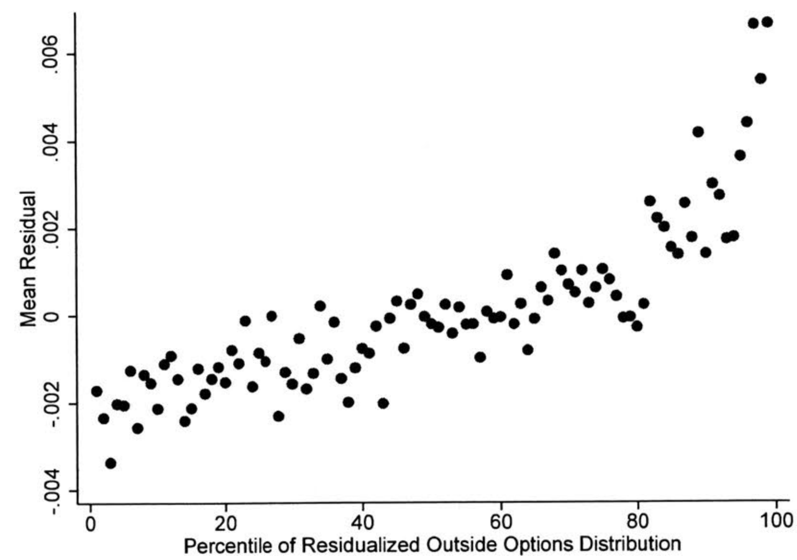

Note: This figure shows how the the change in log earnings in period t depends on the average hiring rates at an individual's former coworkers' firms between $t-1$ and t. The dependent variable in Panel A is log wages and the dependent variable in Panel B is log monthly earnings. The percentiles and earnings changes are computed after partialling out individual and four-digit industry-by-time fixed effects.

Figure 1-7: Reduced Form Results

Panel A: Job-to-Job Mobility

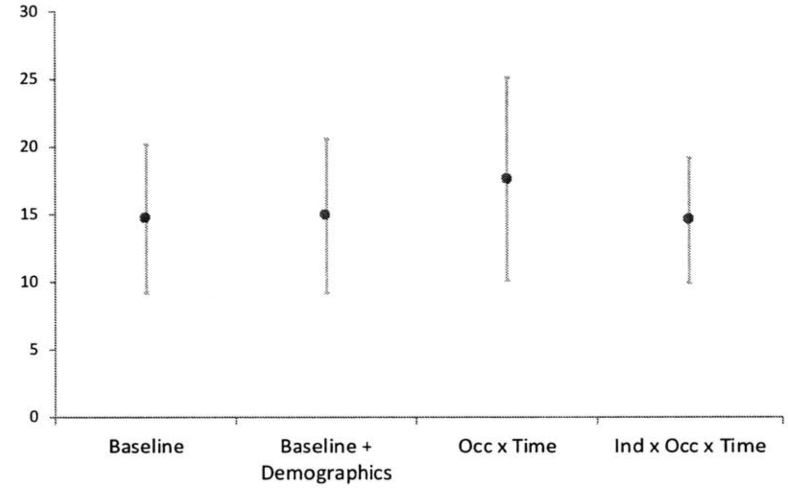

Panel B: Change in Log Monthly Earnings

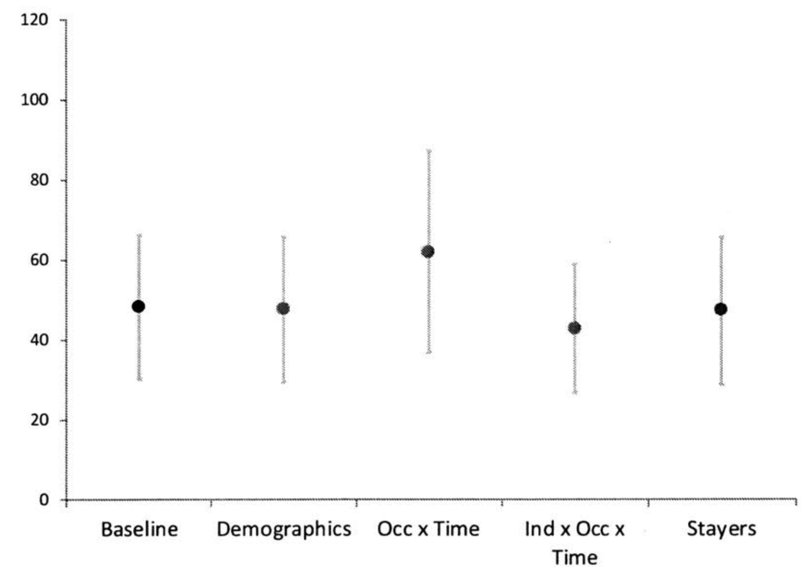

Note: This figure plots scaled estimates of γ from equation 1.5. The outcome variable in Panel A is an indicator for whether the individual made a job-to-job transition. The figure plots the percent impact of a ten-unit change in Ω_{it} on the probability an individual made a job-to-job transition. The outcome variable in Panel B is the change in log monthly earnings. We scale these coefficients to represent the average impact (in 2016 USD) on the average worker's annual earnings. Each regression controls for worker fixed effects, four digit industry-by-time fixed effects, and the number of connections in an individual's network. Additional controls are as listed on the x-axis. Whiskers indicate 95% confidence intervals. Standard errors are two-way clustered by individual and firm.

Figure 1-8: Within-Firm Results

Panel A: Job-to-Job Mobility

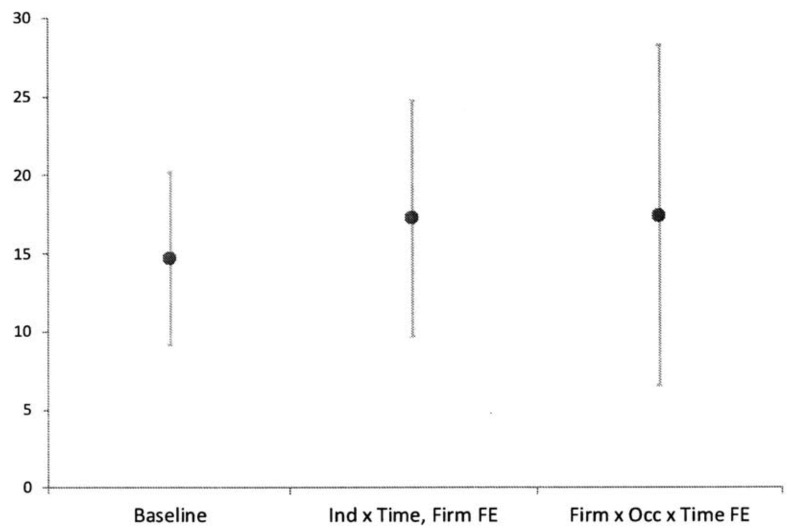

Panel B: Change in Log Monthly Earnings

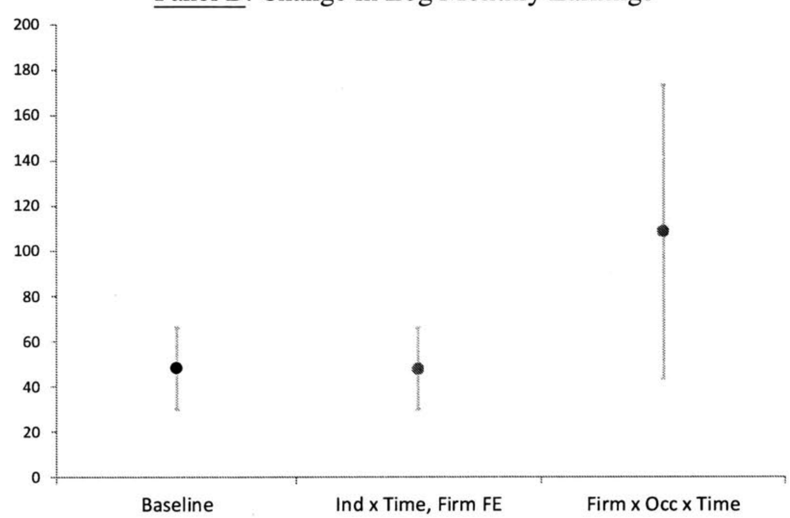

Note: This figure plots the impact of a 10 unit increase in Ω_{it} on the probability of making a job-to-job transition (panel A) or on the change in log monthly earnings (panel B). Each regression controls for worker fixed effects and the number of connections in an individual's network. Additional controls are as indicted. Whiskers indicate 95% confidence intervals. Standard errors are two-way clustered by individual and firm. Raw coefficients are reported in Tables 1.4 and 1.5.

Figure 1-9: Same and Different Region Coworkers

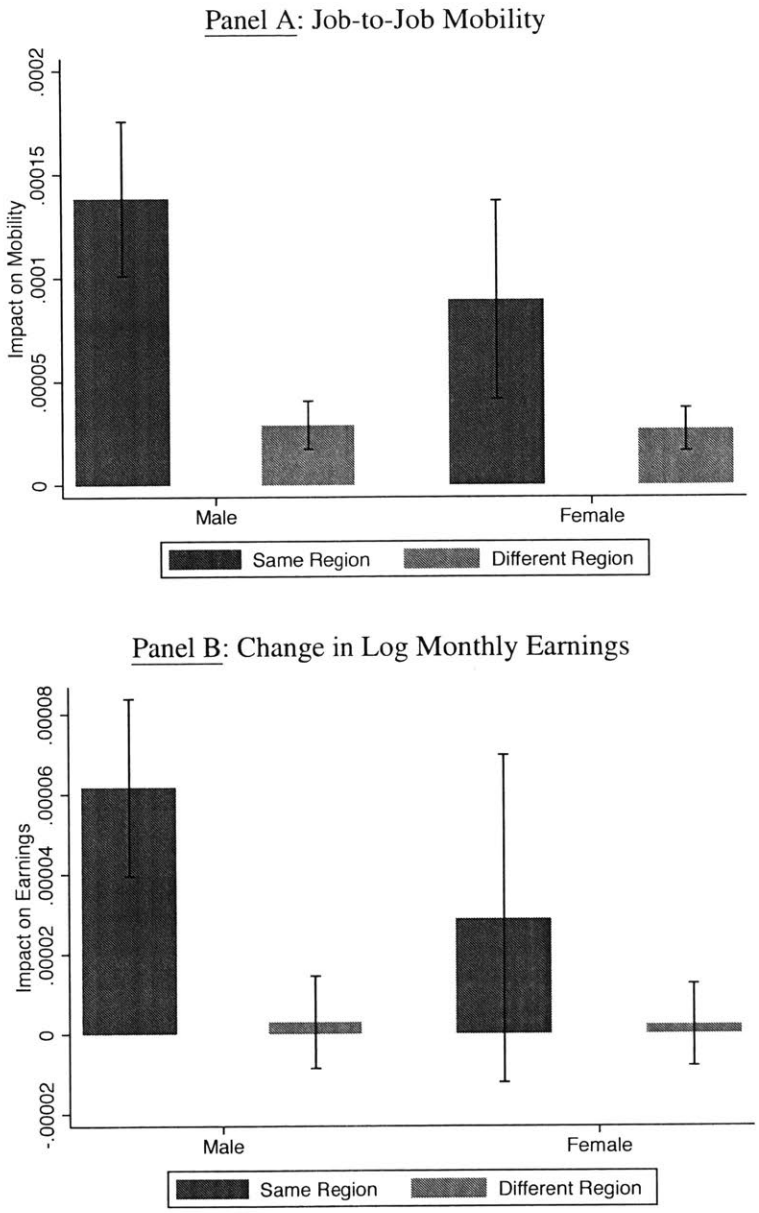

Note: This figure compares the mobility and earnings response to $\Omega_{it}^{\mathrm{IN}}$ and $\Omega_{it}^{\mathrm{OUT}}$, which measure the average number of new positions created among an individual's same-region and different-region coworkers. Each regression controls for worker fixed effects, four-digit industry-by-time fixed effects, and includes linear controls for the number of coworkers in $\Omega_{it}^{\mathrm{IN}}$ and $\Omega_{it}^{\mathrm{OUT}}$. Individuals are not included in these regressions if they do not have any former coworkers working in the same region or in any of the other four regions. Whiskers indicate 95% confidence intervals. Standard errors are two-way clustered by individual and firm.

Figure 1-10: Impacts of Coworkers From Different Time Horizons: Separate Regressions

Panel A: Job-to-Job Transition

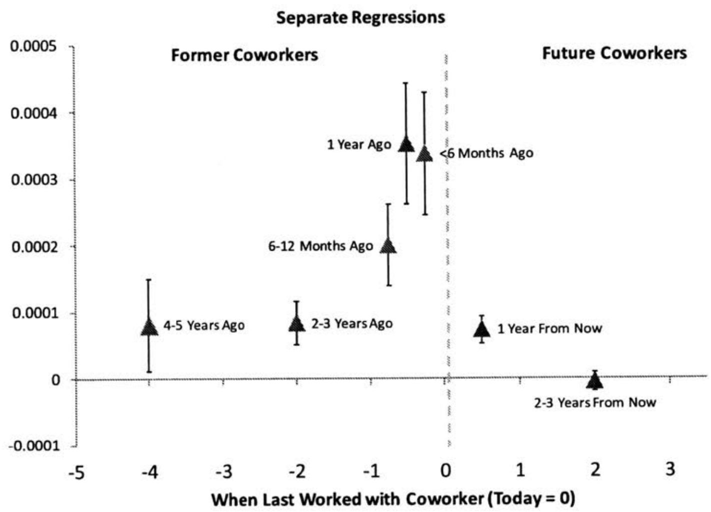

Panel B: Change in Log Monthly Earnings

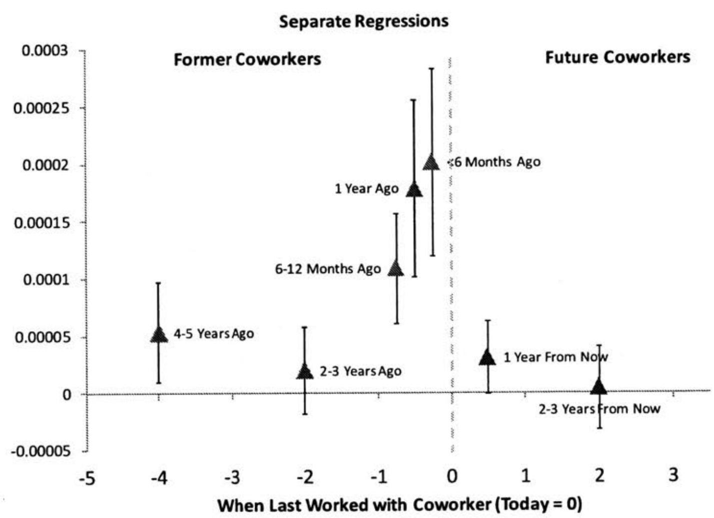

Note: This figure shows how the impact of Ω_{it} varies based on the length of time since the worker worked with his/her former coworkers or the length of time before the worker starts working with his/her future coworkers. Each figure reports estimates from separate regressions of the outcome variable on each network, as described by in equation 1.7. Each regression controls for worker fixed effects, four digit industry-by-time fixed effects, and the number of connections in the included network. Whiskers indicate 95% confidence intervals. Standard errors are two-way clustered by individual and firm. Table 1.7 presents analogous results from regressions which include prior and future networks in the same regression.

Figure 1-11: Dynamics: Probability of Moving

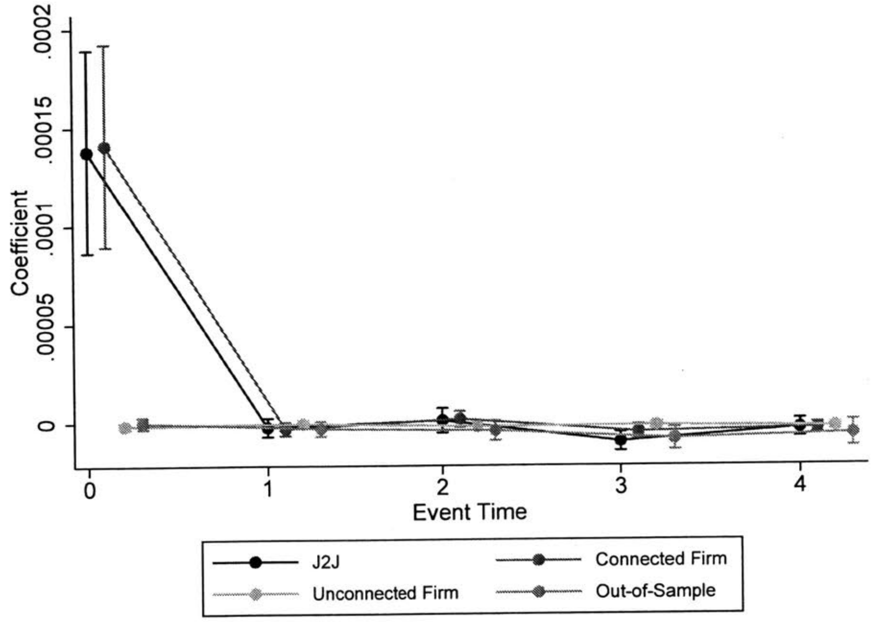

Note: This figure shows how the probability of making a job-to-job transition or moving to a coworker-connected firm depends on their value of Ω_{it} at $t = 0$. The sample includes all individuals that are in the network sample at time $t = 0$. Each dot represents a separate regression. Each regression controls for worker fixed effects, four digit industry-by-time fixed effects, and the number of connections in the individual's network. We also control for that period's value of Ω_{it}. Whiskers indicate 95% confidence intervals. Standard errors are two-way clustered by individual and firm.

Figure 1-12: Dynamics: Earnings

Panel A: Base Pay

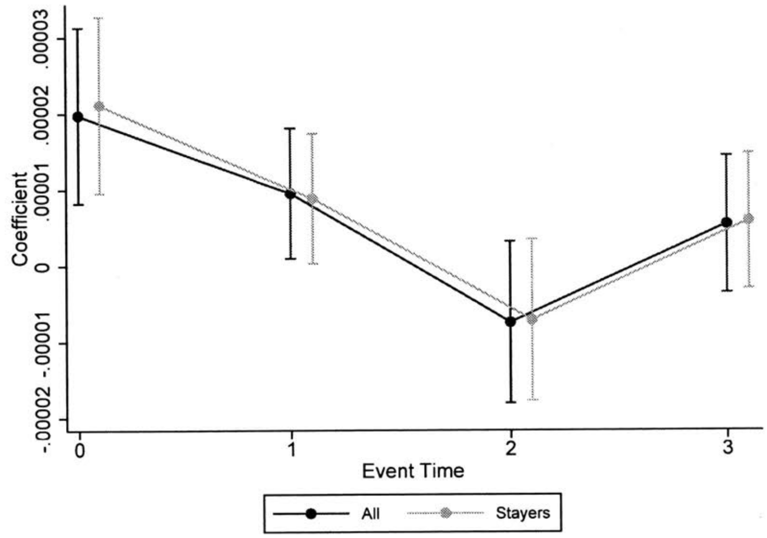

Panel B: Bonus Pay

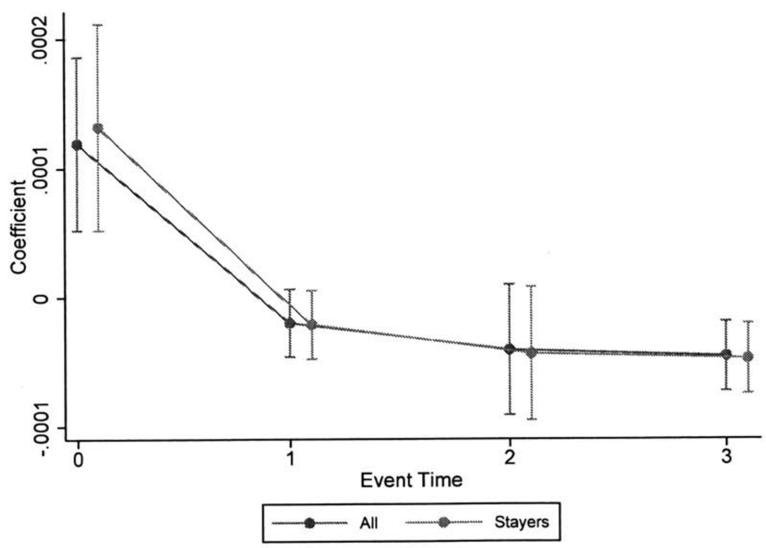

Note: This figure shows how earnings changes depend on their value of Ω_{it} at $t = 0$. The first panel presents results for base pay, which excludes bonuses. The second focuses on bonus pay. Each dot represents a separate regression. Each regression controls for worker fixed effects, four digit industry-by-time fixed effects, and the number of connections in an individual's network. We also control for that period's value of Ω_{it}. Stayers are workers that did not change firms at $t = 0$. Earnings are in kroner. Whiskers indicate 95% confidence intervals. Standard errors are two-way clustered by individual and firm.

Figure 1-13: Long-Run Impacts on Base Pay

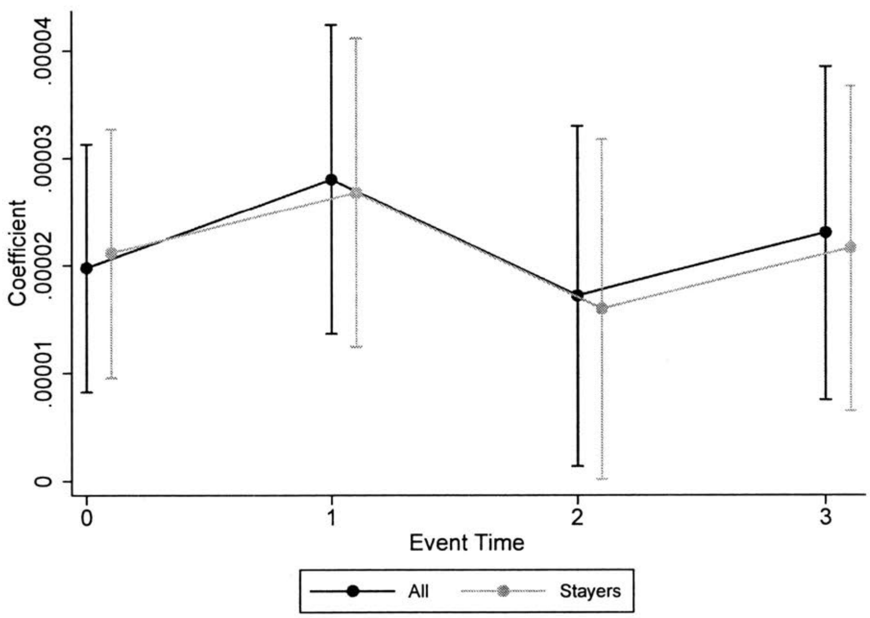

Note: This figure shows how the relationship between earnings at period $t + k$ and earnings at period $t - 1$ depends on the value of Ω_{it} at $t = 0$. Our measure of earnings is "base pay": earnings without bonuses. Each regression controls for worker fixed effects, four digit industry-by-time fixed effects, and the number of connections in an individual's network. Stayers are workers that did not change firms at $t = 0$. Earnings are in kroner. Whiskers indicate 95% confidence intervals. Standard errors are two-way clustered by individual and firm.

Figure 1-14: Mechanisms: Hours versus Hourly Earnings

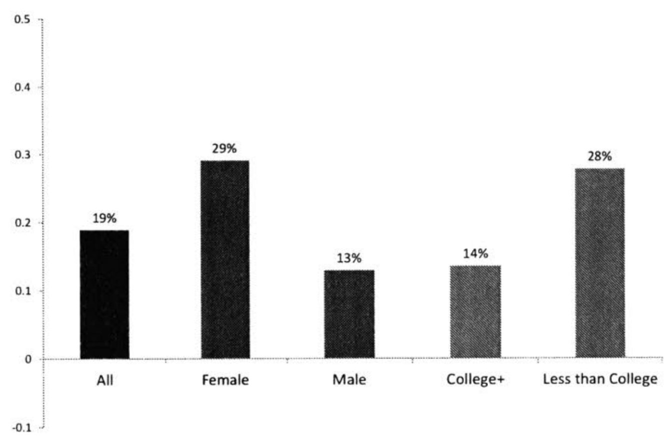

Note: This figure shows what portion of our earnings results are driven by changes in hours worked. We use the accounting identity: $d\log y = d\log h + d\log w$. For each demographic group we present the ratio of the coefficients from equation 1.5. The numerator comes from a regression where the outcome is the change in log hours. The denominator comes from a regression where the outcome is the change in log earnings. The sample differs from that in Table 1.5 because both the earnings and hours regressions only include observations for workers who had non-imputed hours both this month and in the prior month. The results are discussed in Section 1.6.1.

Figure 1-15: Reduced Form Evidence on Bargaining: Returns for Movers and Stayers

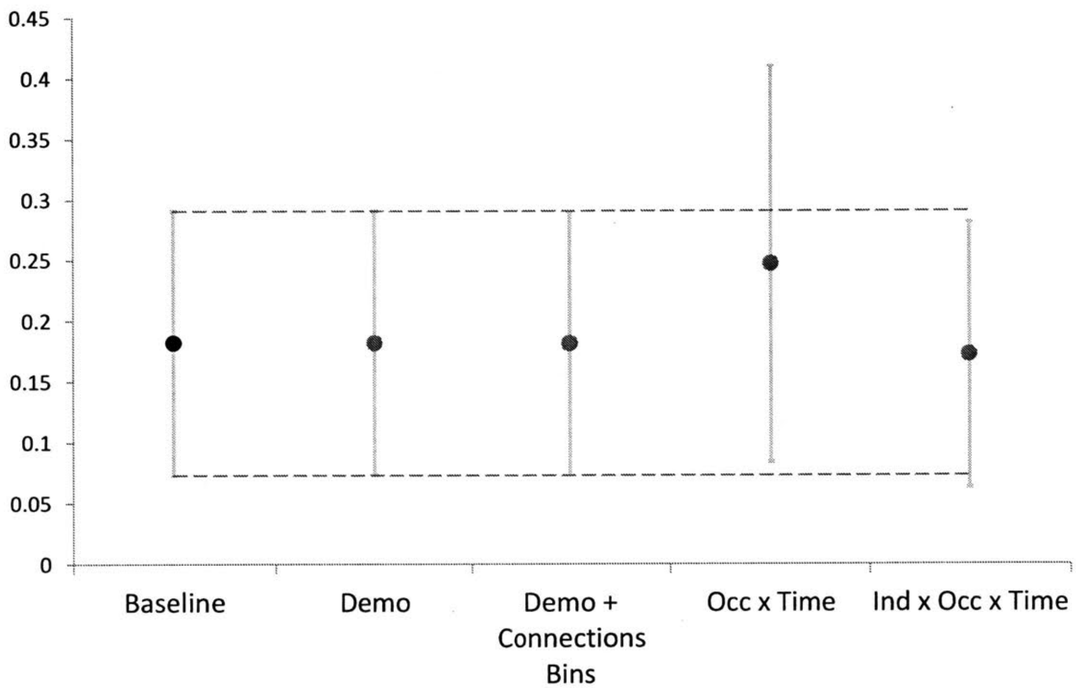

Note: This figure shows how the ratio of returns to information for job stayers and movers changes across different specifications. Each estimate comes from regressions of equation 1.12. Our sample includes the subset of workers with non-overlapping job spells. The demographic controls are indicators for whether the worker has kids or is married. In the third estimate, we replace the linear control for the number of connections with indicators for deciles of the connections distribution. The fourth estimate includes four-digit occupation-by-time fixed effects instead of the industry-by-time fixed effects. The fifth estimate include four-digit industry by two-digit occupation by time fixed effects.

Figure 1-16: Heterogeneity by Occupation

Panel A: Any Transition

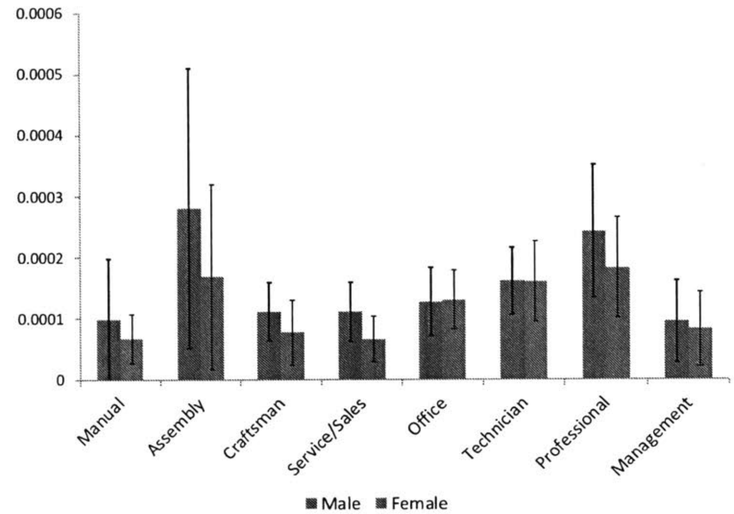

Panel B: Change in Log Monthly Earnings

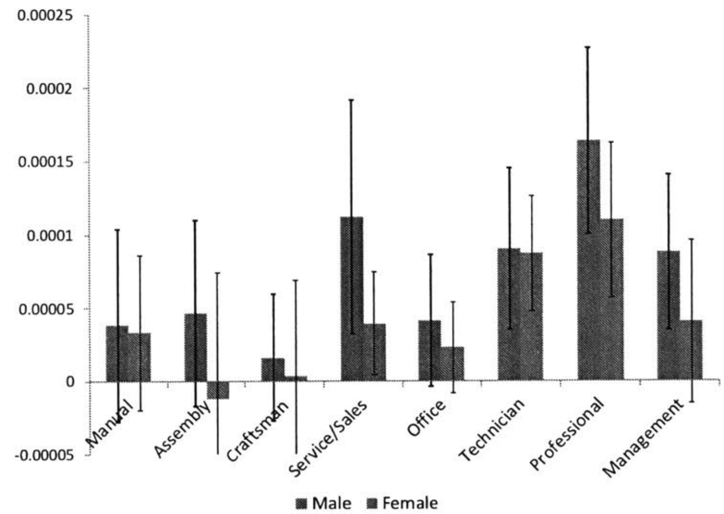

Note: This figure shows how mobility and wage responses differ across occupation groups. We group workers according to broad ISCO (International Standard Classification of Occupations) codes. We then estimate equation 1.5 within each occupation. The dependent variable is as indicated in each panel. Each regression controls for individual fixed effects, four-digit industry-by-time fixed effects and a linear control for the number of connections in an individual's network. Standard errors are two-way clustered by individual and firm. We do not have sufficient data on workers in the military (ISCO 10) or in agricultural occupations (ISCO 6). Table 1.9 presents coefficients for the pooled (both male and female) sample.

Table 1.1: Descriptive Statistics: Workers

		Regression Sample		
	All	All	Male	Female
	(1)	(2)	(3)	(4)
Worker-Month Observations	270115520	72130704	43351245	28779459
Workers	3809303	1096761	657086	439675
Demographics				
Danish	91%	1	1	1
	(0.29)	--	--	--
Age	39.89	43.21	43.26	43.13
	(13.55)	(9.50)	(9.56)	(9.40)
Female	0.49	0.40	0.00	1.00
	(0.50)	(0.49)		
Married	0.49	0.57	0.55	0.60
	(0.50)	(0.50)	(0.50)	(0.49)
In a Couple	0.64	0.74	0.73	0.76
	(0.48)	(0.44)	(0.44)	(0.43)
Has Children	0.51	0.54	0.51	0.58
	(0.50)	(0.50)	(0.50)	(0.49)
College +	0.39	0.31	0.27	0.37
	(0.49)	(0.46)	(0.45)	(0.48)
Employment				
Annual Earnings (2016 USD)	$42,650	$53,591	$58,648	$46,015
	(70,592)	(77,747)	(93,935)	(42,261)
Number of Firms	3.41	3.05	3.05	3.05
	(2.89)	(1.28)	(1.32)	(1.23)
Number of Industries	2.69	2.21	2.22	2.20
	(1.96)	(1.45)	(1.52)	(1.33)
Number of Months	95.00	83.12	83.49	82.57
	(44.63)	(23.18)	(23.24)	(23.07)

Note: The first entry in each row is the mean. The standard deviation is in parentheses. The regression sample includes Danish single-job-holders who are working in firms with between 2 and 1000 employees. Annual earnings are computed using the "broad" income measure and includes fringe benefits and mandatory retirement contributions. The number of industries is calculated using 4-digit NACE codes. The number of months in column 1 is the number of firm-month observations; our regression sample contains only single job-holders, who are at a maximum of one firm each month. More details on the variables are provided in the Appendix 1.12.

Table 1.2: Descriptive Statistics: Firms

	All		Trade Register	
	All	Network	All	Network
	(1)	(2)	(3)	(4)
Number of Firms	352010	272346	49574	43423
Employment	11.07	8.41	41.68	21.11
	(224)	(32)	(544)	(63)
Establishments	1.19	1.12	1.78	1.35
	(4.64)	(1.24)	(11.00)	(2.48)
Firm Accounts				
In Accounting Data	69%	76%	89%	93%
	(0.46)	(0.43)	(0.32)	(0.26)
Revenue (1000 2016 USD)	2.74	2.57	9.58	7.90
	(57)	(49)	(121)	(99)
Value Added/Worker	73.99	74.53	102.24	105.48
	(2049)	(2200)	(4285)	(4458)
Trade Data				
In Trade Register	14%	16%	100%	100%
	(0.35)	(0.37)		
Importer			90%	90%
			(0.30)	(0.29)
Exporter			53%	56%
			(0.50)	(0.50)
Number of Products			3.32	3.29
			(16.81)	(14.78)
Annual Export Value (1000 2016 USD)			279.14	218.39
			(3709.46)	(1907.75)
Location				
Capital Region	33%	32%	36%	35%
	(0.47)	(0.47)	(0.48)	(0.48)
Central Denmark	22%	23%	24%	24%
	(0.42)	(0.42)	(0.42)	(0.43)
North Denmark	11%	11%	10%	10%
	(0.31)	(0.31)	(0.30)	(0.30)
Zealand Region	14%	14%	10%	10%
	(0.34)	(0.34)	(0.31)	(0.31)
Southern Denmark	21%	21%	21%	21%
	(0.41)	(0.41)	(0.40)	(0.41)

Note: This table presents descriptive statistics on the firms in our sample. The first column includes all firms included in our data. The second column restricts to the set of firms in our network sample: those with more than 1 and and fewer than 1000 employees. We used the average number of employees over the sample period to define the network sample. Standard deviations are in parentheses. We calculate the number of establishments at each firm by linking our observations to the annual IDA panel. The firm accounting variables come from the accounting register (FIRE). The trade variables come from the UHDM register. For each firm we calculate the mean number of (six-digit) products each firm exports, across all months the firm is in the trade register. In order to comply with Statistics Denmark privacy regulations, we calculated the median number of products after taking a ten-firm moving average of the data.

Table 1.3: Characteristics of Coworker Networks

	All (1)	Male (2)	Female (3)	College (4)	Less Than College (5)
Number of Connections	156	146	172	175	148
	(278)	(252)	(312)	(313)	(261)
	[58.9]	[56.9]	[62.4]	[72.2]	[53.5]
Characteristics of Connections					
Fraction Female	39%	29%	55%	47%	36%
	(0.26)	(0.21)	(0.24)	(0.23)	(0.26)
Fraction College+	34%	29%	40%	52%	26%
	(0.25)	(0.24)	(0.25)	(0.24)	(0.20)
Mean Age	43.2	43.6	42.7	43.4	43.1
	(7.4)	(7.1)	(7.9)	(7.3)	(7.5)
Fraction in Trade Register	35%	36%	34%	36%	35%
	(0.28)	(0.28)	(0.28)	(0.32)	(0.26)
Connected Firm Characteristics					
Mean Value Added Per Worker	560.5	567.9	549.4	608.2	539.6
	(2201.1)	(1841.9)	(2653.7)	(3592.7)	(1147.0)
Mean Hourly Earnings (kroner)	215.3	217.1	212.5	228.7	209.4
	(38.4)	(37.5)	(39.6)	(45.1)	(33.4)
Fraction Female	39%	33%	48%	45%	37%
	(0.16)	(0.14)	(0.15)	(0.14)	(0.16)
Dispersion					
Number of Industries	31.9	31.5	32.4	32.4	31.6
	(36.4)	(35.9)	(37.2)	(34.7)	(37.1)
Number of Firms	60.5	58.9	62.8	63.5	59.2
	(92.8)	(89.8)	(97.0)	(92.9)	(92.7)
Observations	1096764	657089	439675	334269	762219

Note: This table describes the characteristics of the coworker networks. Each individual's (time-varying) coworker network consists of individuals he/she has worked with in the past three years. We provide more details on how we constructed these networks in Section 1.4. The first entry in each row is the mean. Standard deviations are reported in parentheses; medians are reported in brackets. To comply with Statistics Denmark's privacy regulations, we computed the medians after taking a 10-person moving average.

Table 1.4: Impact of Outside Options on Mobility

					Within-Firm Analysis	
	Baseline	With Demographic Controls	Occupation-by-Time	Industry-Occupation-by-Time	Baseline + Firm FE	Firm-Occupation-Time FE
	(1)	(2)	(3)	(4)	(5)	(6)
Any Transition	1.255 ***	1.399 ***	1.521 ***	1.255 ***	1.468 ***	1.430 ***
	(0.255)	(0.273)	(0.350)	(0.217)	(0.351)	(0.485)
Job-to-Job	1.378 ***	1.400 ***	1.654 ***	1.368 ***	1.618 ***	1.637 ***
	(0.263)	(0.273)	(0.360)	(0.223)	(0.361)	(0.520)
Connected Firm	1.408 ***	1.431 ***	1.656 ***	1.406 ***	1.642 ***	1.685 ***
	(0.263)	(0.273)	(0.352)	(0.222)	(0.354)	(0.521)
Connected Industry	-0.014	-0.013	0.004	-0.018 *	-0.010	-0.020
	(0.010)	(0.010)	(0.012)	(0.009)	(0.012)	(0.015)
Unconnected Firm	-0.015 *	-0.015 *	-0.005	-0.018 **	-0.009	-0.017
	(0.009)	(0.009)	(0.008)	(0.009)	(0.008)	(0.012)
Out-of-Sample Firm	-0.001	-0.003	0.000	-0.001	-0.005	-0.011
	(0.015)	(0.015)	(0.018)	(0.014)	(0.018)	(0.023)
Observations	57922601	55697201	57923303	57595436	57919273	49148828
Individual FE	X	X	X	X	X	X
Additional Controls	Industry-Period FE	Industry-Period FE	Occupation-Period FE	Industry-Occupation-Period FE	Industry-Period FE, Firm FE	Firm-Occupation-Period FE
Level of Occupation Codes	N/A	N/A	2-digit	2-digit	N/A	4-digit
Level of Industry Codes	4-digit	4-digit	N/A	4-digit	2-digit	N/A

Note: This table presents estimates of γ from equation 1.5. Outcomes vary by row; specifications vary by column. All regressions control for individual fixed effects and for the number of connections in an individual's network. Standard errors are two-way clustered at the individual and firm level. Coefficients are scaled by 10000, for readability. A transition is any observation where an individual is not where they were in the prior month: either at a different firm, or at no firm. A job-to-job transition occurs when the individual is, by the first of the month, at a new firm. Connected (unconnected) firms are those in the network sample where the individual has (does not have) a former coworker. Out of sample firms are firms with more than 1000 employees. Levels of significance: *10%, ** 5%, and *** 1%.

Table 1.5: Impact of Outside Options on Earnings

					Within-Firm Analysis	
	Baseline	With Demographic Controls	Occupation-by-Time	Industry-Occupation- by-Time	Baseline + Firm FE	Firm-Occ - Period FE
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Log Earnings	0.703 ***	0.694 ***	0.904 ***	0.624 ***	0.699 ***	1.581 ***
	(0.134)	(0.135)	(0.188)	(0.120)	(0.135)	(0.481)
	56134045	54689099	56134776	55806641	56130355	47695803
Δ Log Earnings (Narrow)	0.998 ***	0.975 ***	1.228 ***	0.915 ***	0.995 ***	2.209 ***
	(0.185)	(0.187)	(0.240)	(0.168)	(0.186)	(0.653)
	56153349	54707318	56154079	55825981	56149670	47711986
Δ Log Hours	0.197 ***	0.113 **	0.219 **	0.209 ***	0.197 **	0.044
	(0.076)	(0.055)	(0.099)	(0.073)	(0.077)	(0.166)
	50199730	49027776	50200508	49870846	50195900	43109059
Δ Log "Base Pay"	0.231 ***	0.224 ***	0.299 ***	0.224 ***	0.225 ***	0.307 **
	(0.056)	(0.054)	(0.066)	(0.054)	(0.057)	(0.122)
	55861195	54417929	55861927	55533703	55857498	47453513
Bonus/Base Pay	0.991 ***	1.104 ***	1.248 ***	0.778 **	0.994 ***	2.057 ***
	(0.307)	(0.356)	(0.357)	(0.303)	(0.303)	(0.731)
	57063082	54888239	57063791	56735772	57058999	48403944
Individual FE	X	X	X	X	X	X
Demographic Controls		X				
Additional Controls	Industry-Period FE	Industry-Period FE	Occupation-Period FE	Industry-Occupation-Period FE	Industry-Period FE, Firm FE	Firm-4-digit-Occupation-Period FE

Note: This table presents estimates of γ from equation 1.5. Outcomes vary by row; specifications vary by column. All regressions control for individual fixed effects and for the number of connections in an individual's network. Additional controls are indicated in the relevant column. Coefficients are scaled by 10000, for readability. Earnings are in kroner. Standard errors are two-way clustered at the individual and firm level. We provide more information about how we decompose the raw earnings measures into base pay and bonuses in Appendix 1.12.4. Table 1.17 presents results for job-stayers. Levels of significance: *10%, ** 5%, and *** 1%.

Table 1.6: Impacts by Region of Former Coworker

	Mobility		Earnings	
	Job-to-Job	Connected Move	Change in Log Earnings	Change in Log Wages
	(1)	(2)	(3)	(4)
		A. Full Sample		
Same-Region Coworkers	1.140 ***	1.134 ***	0.470 ***	0.310 ***
	(0.216)	(0.214)	(0.128)	(0.076)
Different-Region Coworkers	0.275 ***	0.127 ***	0.025	0.022
	(0.053)	(0.046)	(0.046)	(0.043)
Observations	48806147	48806147	47367747	47145131
		B. Male Workers		
Same-Region Coworkers	1.383 ***	1.373	0.617 ***	0.327 ***
	(0.191)	(0.189)	(0.113)	(0.108)
Different-Region Coworkers	0.289 ***	0.133	0.029	0.062
	(0.059)	(0.049)	(0.059)	(0.058)
Observations	30394074	30394074	29396003	29254327
		C. Female Workers		
Same-Region Coworkers	0.894 ***	0.894 ***	0.287	0.280 ***
	(0.244)	(0.242)	(0.209)	(0.078)
Different-Region Coworkers	0.266 ***	0.127 ***	0.022	-0.031
	(0.053)	(0.049)	(0.052)	(0.048)
Observations	18410492	18410492	17970161	17889214

Note: This table presents estimates of γ^{IN} and γ^{OUT} from equation 1.6. Each regression controls for worker fixed effects, four digit industry-by-time fixed effects, and the number of connections an individual has in the same region and in other regions. We assign individuals to regions based on the location of their firm in the prior period. Coefficients are scaled by 10000, for readability. Earnings are in kroner. Individuals who do not have former coworkers in both their own region and outside regions are excluded, by design. Standard errors are two-way clustered by individual and firm. Levels of significance: *10%, ** 5%, and *** 1%. Some of the coefficients are plotted in Figure 1-9.

Table 1.7: Impacts by when An Individual Last Worked with the Coworker

		Job to Job Transition				Change in Log Earnings			
	Connections	Separate	Single Regression			Separate	Single Regression		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Former Coworkers									
1 Year Ago	22.2	3.522***	5.285***	5.281***		1.784***	2.357***	2.364***	
		(0.462)	(0.717)	(0.714)		(0.393)	(0.628)	(0.630)	
1-6 Months Ago	11.2	3.370***			4.794***	2.011***			2.847***
		(0.470)			(0.751)	(0.417)			(0.706)
7-12 Months Ago	11.0	2.003***			2.149***	1.085***			0.976***
		(0.312)			(0.440)	(0.245)			(0.369)
2-3 Years Ago	44.88	0.833***	0.789**	0.761**	0.829	0.552***	0.213	0.217	0.289
		(0.166)	(0.380)	(0.376)	(0.513)	(0.194)	(0.386)	(0.385)	(0.515)
4-5 Years Ago	61.76	0.804**	0.871	0.878	0.929	0.532**	0.909***	0.910***	0.710*
		(0.354)	(0.607)	(0.598)	(0.801)	(0.223)	(0.338)	(0.338)	(0.418)
Future Coworkers									
1 Year From Now	23.85	0.725***		0.598***	0.606***	0.308*		-0.062	-0.262
		(0.105)		(0.180)	(0.233)	(0.162)		(0.256)	(0.323)
2-3 Years From Now	56.06	-0.052		-0.880***	-1.286***	0.040		0.006	-0.052
		(0.073)		(0.194)	(0.261)	(0.185)		(0.299)	(0.355)
Observations		Varies	14670466	14670466	12373445	Varies	14325491	14325491	12109820

Note: This table presents estimates of γ^n from equation 1.7. Each row contains coefficients from a separate regression. Each regression controls for worker fixed effects, four digit industry-by-time fixed effects, and the number of connections in each included network. Standard errors are two-way clustered by individual and firm. We exclude the first two years and final three years of our regression sample so that network quality does not vary across years of our sample. Levels of significance: *10%, ** 5%, and *** 1%.

Table 1.8: Trade-Based Measures

	First Stage		Reduced Form			
			Transitions		Earnings	
	Log Exports	New Positions	Job to Job	Connected Firm	Δ Log Hourly Earnings	Δ Log Earnings
	(1)	(2)	(3)	(4)	(5)	(6)
	A. Baseline: Industry-Time Controls					
Predicted Log Exports	0.306 ***	0.012 ***	0.161 *	0.183 **	0.291 **	0.386 ***
	(0.032)	(0.003)	(0.090)	(0.080)	(0.133)	(0.147)
	23730053	23730053	23730053	23730053	23109099	23186630
	B. Within Firm: Firm-Occupation-Time Controls					
Predicted Log Exports	0.598 ***	0.014 ***	0.026	0.124 *	1.002 ***	1.021 **
	(0.131)	(0.004)	(0.070)	(0.065)	(0.386)	(0.417)
	20586717	20586717	20586717	20586717	20078092	20134138

Note: This table presents estimates of γ from equation 1.11. Outcome variables vary by column. The outcome variable in the first column is a measure of Ω_{it} based on realized measures of firms' exports. The outcome variable in the second column is our baseline measure Ω_{it}. The third and fourth columns present mobility results. The fifth and sixth columns present earnings results. Coefficients are scaled by 10000, for readability. Earnings are in kroner. Standard errors are two-way clustered by individual and firm. We provide details on how we construct $\Omega_{it}^{\text{trade}}$ in Appendix 1.12. Levels of significance: *10%, ** 5%, and *** 1%. Note that the sample differs from other tables because a large fraction of workers do not have any coworkers in exporting firms.

Table 1.9: Heterogeneity by Occupation

	Manual	Assembly	Craftsman	Service/Sales	Office	Technician	Professional	Managers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	A. Coefficients							
Job-to-Job Mobility	0.897 **	2.575 **	1.087 ***	0.816 ***	1.302 ***	1.612 ***	2.137 ***	0.915 ***
	(0.396)	(1.002)	(0.229)	(0.180)	(0.235)	(0.278)	(0.470)	(0.311)
Δ Log Earnings	0.366	0.344	0.151	0.643 ***	0.280 *	0.888 ***	1.368 ***	0.743 ***
	(0.266)	(0.267)	(0.208)	(0.228)	(0.146)	(0.201)	(0.247)	(0.224)
	B. Scaled Impact							
Scaled by Annual Earnings	$16	$17	$8	$28 ***	$14 *	$54 ***	$91 ***	$67 ***
	($12)	($13)	($9)	($10)	($7)	($12)	($16)	($20)
	5284780	5585636	9552430	5753531	7218323	10473806	11960350	4038416

Note: This table shows mobility and wage responses differ across occupations. We group workers according to broad ISCO (International Standard Classification of Occupations) codes. We then estimate equation 1.5 separately within each occupation. Each regression controls for individual fixed effects, four-digit industry-by-time fixed effects and a linear control for the number of connections in an individual's network. Standard errors are two-way clustered by individual and firm. Coefficients are scaled by 10000, for readability. Earnings outcomes in Panel A are in kroner. We do not have sufficient data on workers in the military (ISCO 10) or in agricultural occupations (ISCO 6). Panel B estimates the impact of a 10 unit change in Ω_{it} on a worker's annual earnings (in dollars). We calculate annual earnings separately for each group. Levels of significance: *10%, ** 5%, and *** 1%. Figure 1-16 plots coefficients similar to those in Panel A, for each sex.

Table 1.10: Parameters

Parameters	Description	Manual	Professionals
β	Bargaining Power	.8733	0.8008
$1 - p_R$	Fraction of Offers from Renegotiating Firms	.3094	0.5125
δ	Exogenous Job Destruction Rate	0.0246	0.0269
λ_U	Outside Offer Arrival Rate for Unemployed Workers	0.0374	0.0593
α_U	Connected-Offer Arrival Rate for Unemployed Workers	0.0009	0.0068
λ_E	Outside Offer Arrival Rate for Employed Workers	0.0170	0.0291
α_E	Connected-Offer Arrival Rate for Employed Workers	0.0007	0.00002
$b, \sigma_P, \sigma_R, \mu_P, \mu_R$	Other parameters	—	—

Note: The table above displays the parameters that we estimate in Section 1.7. We allow these parameters to vary by skill group. We do not estimate the discount rate ρ, but instead fix it at $1/(1 + .005) \approx .995$.

Table 1.11: Moments

Number	Description
3	Mean transition rates: J2J, U2E, E2U
2	Mean and variance of (residual) log wage changes for job-stayers
1	Mean log wages
1	Residual variance of log wages
2	Mean and variance of (residual) log wage changes
1	Mean log wage gain associated with a job to job transition
2	Regression coefficients for mobility: J2J and U2E
2	Regression coefficients for earnings of job-stayers: $\Delta \log y$ and $1\{\Delta \log y > 0\}$

Note: This table lists the moments used to estimate the model in Section 1.7. We include a set of moments for each distinct labor market we consider.

Table 1.12: Impact of a Decreased Arrival Rate

Group	$\%\ \Delta$Mobility	$\%\Delta$ Wages	$\%\ \Delta$ Wage Growth
Professionals	-41%	-.8%	-32%
Manual Skills	-46%	-.3%	-43%

Note: This table shows how a reduction in the job arrival rate influence wages and wage growth for two main groups of workers: professionals and manual skilled workers.

1.10 Appendix Tables and Figures

Figure 1-17: Labor Market Flexibility in OECD Countries

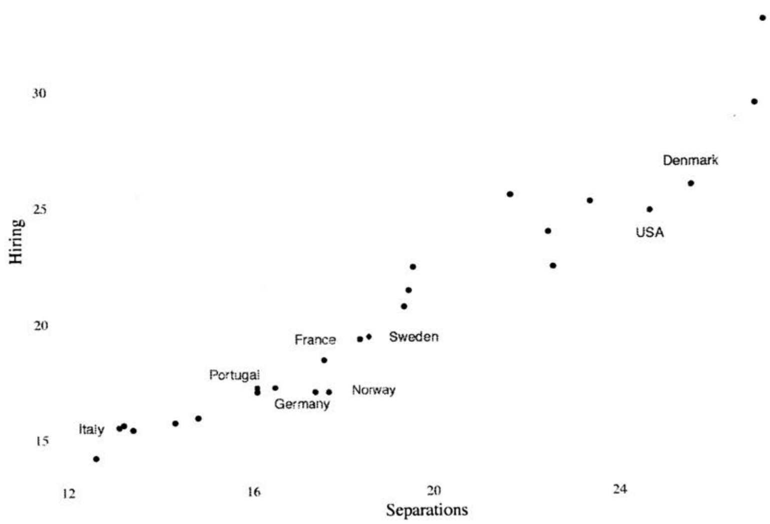

Note: This figure plots mean hiring and separation rates for OECD countries using data from OECD (2004). The original data are adjusted for industrial composition. The years used vary by country. For more details, see OECD (2004).

Figure 1-18: Map of Danish Administrative Regions

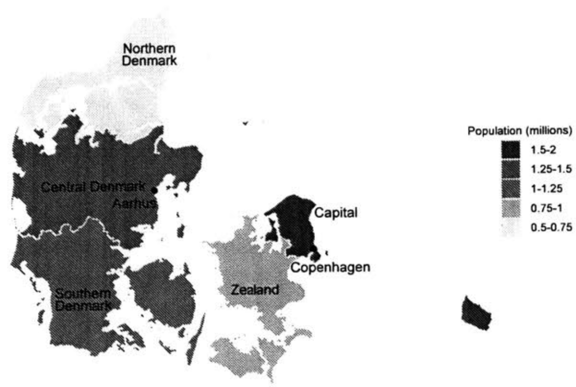

Note: This figure shows the five administrative regions in Denmark. We construct same-region and different-region coworker networks on the basis of these regions. The population data are taken from Statistic Denmark's Statbank.

Figure 1-19: Robustness: Controlling for Occupation-Time FE

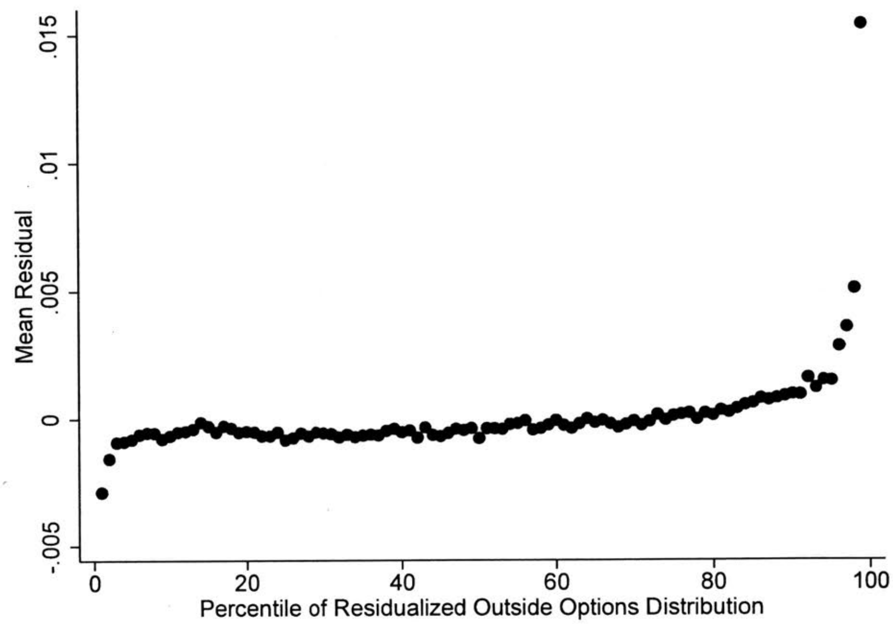

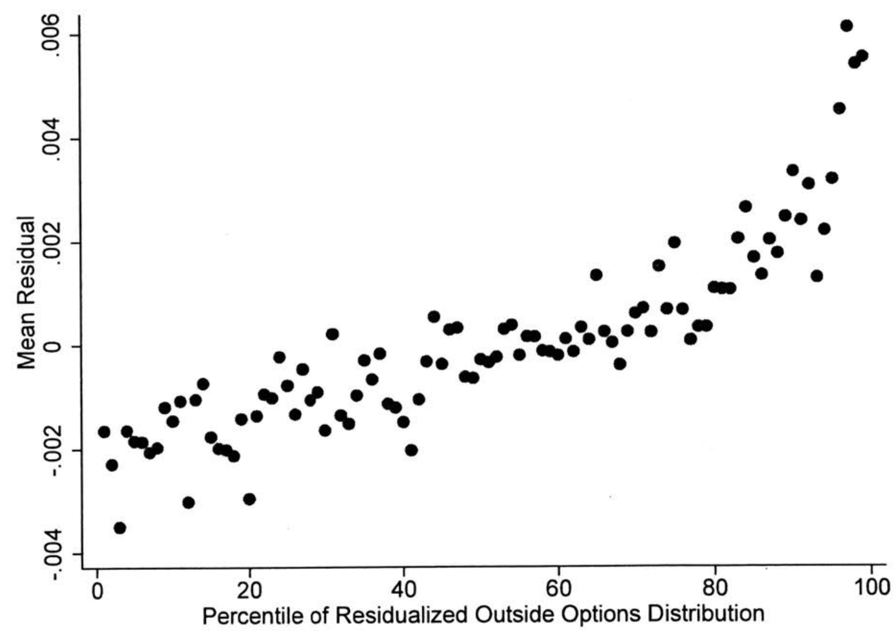

Note: This figure shows how the probability of making a transition or the average change in monthly earnings depends on Ω_{it}. We first residualize both the dependent variables and Ω_{it} on individual fixed effect and industry-by-time and occupation-by-time fixed effects.

Figure 1-20: Robustness: Value of Connections at Larger Firms

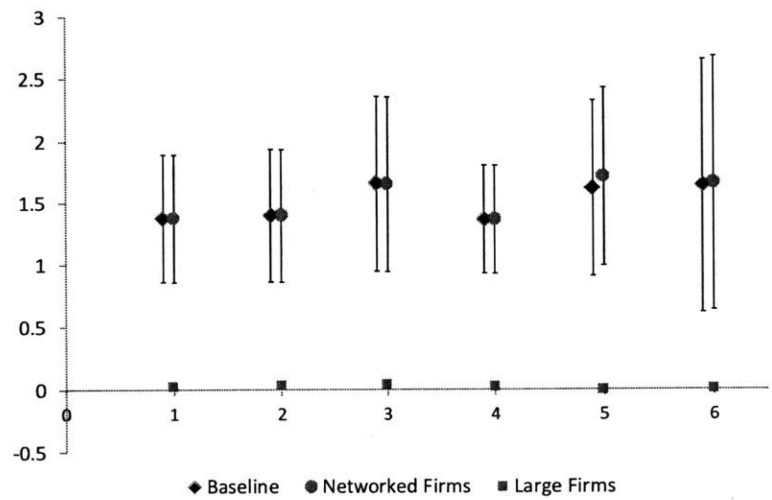

Note: This figure shows that adding measures of Ω_{it} based on connections at large firms (more than 1000 employees) does not change our estimates of γ. The outcome variable is an indicator for whether the worker made a job-to-job transition and coefficients are scaled as in Figure 1-7. The six specifications correspond to those in Table 1.4. The black dot presents our baseline estimates. The blue dots show how the estimate of γ changes when we include measures based on connections at large firms (more than 1000) employees as a separate regressor ($\Omega_{it}^{\text{large}}$). The red squares show the coefficient on $\Omega_{it}^{\text{large}}$ in this regression. Standard errors two-way are clustered by worker and firm.

Figure 1-21: Quality of Trade Prediction

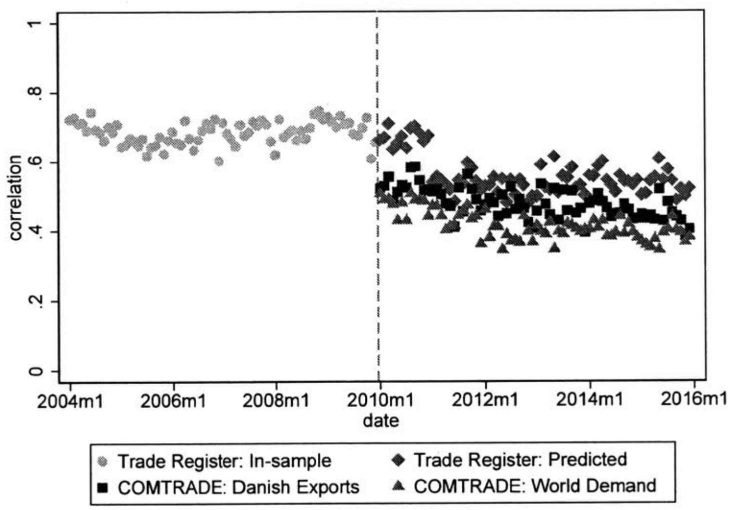

Note: This figure assesses the quality of our firm-level predicted trade measures. We use data from 2004-2009 to fix each firm's share of total Danish exports of each product. The light blue dots show the correlation between the predicted exports–based on total Danish exports of each product as reported in the administrative register– and the firm's actual exports. The dark blue dots show the correlation between the actual exports and those predicted using Danish exports in COMTRADE. The red triangles show the correlation between actual firm exports those predicted using world (minus Denmark) exports in COMTRADE. This is the measure used in $\Omega_{it}^{\text{trade}}$. More details are provided in Section 1.12.7.

Figure 1-22: Posting and Bargaining in the United States: Hall and Krueger (2012)

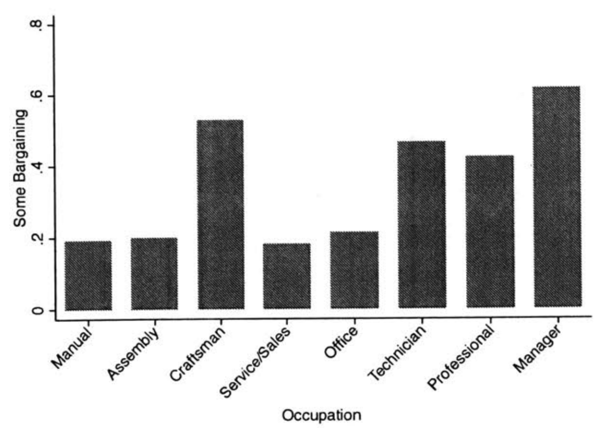

Note: This figure uses survey data from Hall and Krueger (2012) to plot the mean fraction of workers in each occupation group who engage in bargaining at the start of a job spell. More details are provided in Section 1.13.3.

Figure 1-23: Impacts by Quality of Outside Option

Panel A: Mobility

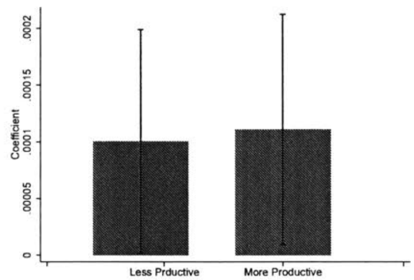

Panel B: Change in Log Monthly Earnings

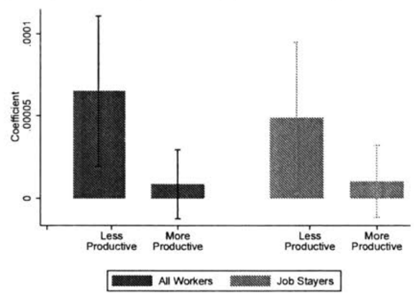

Panel C: Indicator for Wage Growth

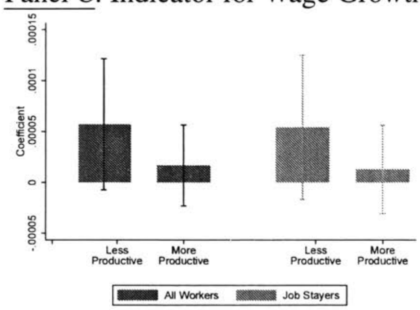

Note: This figure shows how mobility and wage responses differ based on the source of the outside option. We group into twenty vigintiles by their mean value added per worker (in real 2016 kroner) over the sample period. We have these data for roughly three quarters of the firms in our sample. We then estimate equation 1.14 and plot estimates of γ^{ABOVE} and γ^{BELOW}. Panel A presents estimates where the outcome is an indicator for whether the individual made a job to job transition. Panel B shows results for changes in log monthly earnings; Panel C shows results for whether there was an earnings change. Each regression controls for the number of connections in each network and for individual and four-digit industry-by-time fixed effects. Standard errors are two-way clustered by individual and firm.

Table 1.13: Constructing Network Sample

	Workers (1)	Worker-Months (2)
Spells Covering 1st of Month	3809303	248252752
Danish Workers	3295211	225944464
Work at a Firm of <=1000 Workers At Least Once	2785351	190373936
Between 25 and 60 Years Old	2111834	142660144
Single Job-Holder	1842082	126885936
At a Firm of <=1000 Workers	1096764	60491824

Note: This table shows how each of our sample restrictions changes the number of observations and individuals in our data. Each row represents an additional restriction, relative to the row above. These restrictions are described in Section 1.3.3 and in Appendix 1.12.

Table 1.14: Network Dispersion Across Regions

	Capital Region (1)	Central Denmark (2)	North Denmark (3)	Zealand Region (4)	Southern Denmark (5)
Capital Region	73%	8%	3%	9%	8%
Central Denmark	15%	65%	6%	3%	11%
North Denmark	11%	15%	65%	2%	7%
Zealand Region	32%	6%	2%	53%	7%
Southern Denmark	15%	12%	3%	3%	66%

Note: This table shows the dispersion of an individual's coworker network across different regions of network. The rows indicate the individual's region of work. The columns indicate the region each of their coworkers lives in. Each row sums to 100%.

Table 1.15: Transition Rates

	All (1)	Men (2)	Women (3)	College (4)	No College (5)
Make a Job to Job Transition	1.0%	1.0%	0.9%	1.0%	0.9%
		Types of Job to Job Transitions			
Connected Firm	57.4%	59.9%	52.9%	50.6%	60.3%
Connected Industry	12.8%	13.7%	11.2%	12.2%	13.1%
Unconnected Firm	12.1%	12.3%	11.6%	11.8%	12.2%
Out-of-Sample firm	17.7%	14.1%	24.3%	25.5%	14.4%

Note: This table shows the raw probability an individual would make each type of transition each period. These transitions are defined in Appendix 1.12.

Table 1.16: Autocorrelation of Network Characteristics

	Lag											
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Number of Connections	0.993	0.987	0.980	0.973	0.967	0.960	0.952	0.942	0.933	0.923	0.913	0.904
Female	0.987	0.977	0.970	0.962	0.955	0.948	0.941	0.935	0.928	0.921	0.915	0.909
College+	0.987	0.978	0.970	0.963	0.956	0.950	0.943	0.936	0.930	0.923	0.917	0.911
Age	0.981	0.967	0.956	0.945	0.935	0.926	0.916	0.906	0.896	0.886	0.877	0.868
Characteristics of Coworkers' Firms												
Mean Value Added (Time-varying)	0.121	0.105	0.058	0.078	0.024	0.033	0.034	0.029	0.025	0.513	0.025	0.039
Mean Hourly Earnings (Time-varying)	0.908	0.869	0.847	0.822	0.802	0.786	0.764	0.748	0.743	0.726	0.711	0.708
Mean Fraction Female	0.970	0.950	0.936	0.923	0.912	0.903	0.893	0.885	0.877	0.869	0.862	0.856

Note: This table shows that the characteristics of an individual's network remain stable over time. The first row looks at the number of coworkers in the network. The remaining rows look at the correlation between the average (across all coworkers) characteristics of an individual's network in a given month.

Table 1.17: Impact of Outside Options on Earnings: Job Stayers

	All				Full-Time			
	Baseline		Within Firm-Occupation		Baseline		Within Firm-Occupation	
	(1)		(2)		(3)		(4)	
Δ Log Earnings	0.688	***	1.520	***	0.719	***	1.673	***
	(0.137)		(0.490)		(0.171)		(0.461)	
	54266264		46185273		15280379		12549605	
Δ Log Earnings (Narrow)	0.732	***	1.794	***	0.745	***	1.691	***
	(0.144)		(0.551)		(0.174)		(0.465)	
	54284385		46200647		15280506		12549728	
Δ Log Hours	0.101		-0.118					
	(0.516)		(0.130)					
	48718166		41905015					
Δ Log "Base Pay"	0.239	***	0.253	**	0.085	***	0.103	
	(0.056)		(0.123)		(0.033)		(0.069)	
	54001519		45950270		15281563		12549852	
Bonus/Base Pay	1.104	***	2.353	***	1.229	***	2.111	***
	(0.302)		(0.810)		(0.466)		(0.771)	
	54571865		46406029		15290342		12557319	
Individual FE	X		X		X		X	
Additional Controls	Industry-Period FE		Firm-Occupation-Period FE		Industry-Period FE		Firm-Occupation-Period FE	

Note: This table presents estimates of γ from equation 1.5. Outcomes vary by row; specifications vary by column. All regressions control for individual fixed effects and for the number of connections in an individual's network. Additional controls are indicated in the relevant column. Coefficients are scaled by 10000, for readability. Earnings are in kroner. Standard errors are two-way clustered at the individual and firm level. We explain how we decompose raw earnings measures into base pay and bonuses in Appendix 1.12.4. The sample includes only job stayers: those who are at the same firm as in the prior month. Table 1.5 presents results for all workers. Levels of significance: *10%, ** 5%, and *** 1%.

Table 1.18: Impacts on U2E Transitions

	>=1 Month (1)		>=2 Months (2)		>=3 Months (3)	
1 Year Ago	28.748	***	32.039	***	38.461	***
	(2.315)		(2.464)		(3.355)	
2-3 Years Ago	10.662	***	11.184	***	15.999	***
	(1.532)		(2.107)		(4.005)	
4-5 Years Ago	12.033	***	12.314	***	9.333	**
	(1.586)		(2.085)		(3.756)	
Observations	611809		392682		174068	

Note: This table presents estimates of γ from equation 1.7 for workers who are not currently at a firm. Each column presents a separate regression. In addition to the listed covariates, each regression controls for individual and four-digit industry-by-time fixed effects, and for the number of connections in each included network. The first column contains all individuals who were non-employed in the prior month; the remaining columns condition on remaining non-employed for 2 or 3 months. Standard errors are two-way clustered by individual and firm. Levels of significance: *10%, ** 5%, and *** 1%.

Table 1.19: Impact on Mobility: Alternative Measures of Outside Options

	Positions		Hires	
	Baseline	Weighted	Unweighted	Weighted
	(1)	(2)	(3)	(4)
Any Transition	1.255 ***	0.539 ***	1.731 ***	0.640 ***
	(0.255)	(0.175)	(0.312)	(0.180)
Job-to-Job	1.378 ***	0.570 ***	1.800 ***	0.658 ***
	(0.263)	(0.175)	(0.318)	(0.180)
Connected Firm	1.408 ***	0.579 ***	1.792 ***	0.658 ***
	(0.263)	(0.175)	(0.315)	(0.179)
Unconnected Firm	-0.014	-0.005	-0.009	-0.004
	(0.010)	(0.003)	(0.009)	(0.004)
Out-of-Sample Firm	-0.001	-0.001	0.013	0.002
	(0.015)	(0.006)	(0.015)	(0.005)
Observations	57922601	57922601	57922601	57922601

Note: This table presents estimates of γ from equation 1.5 for different measures of Ω_{it}. The rows correspond to different mobility outcomes. Each regression controls for individual and four-digit industry-by-time fixed effects. Coefficients are scaled by 10000, for readability. Standard errors are two-way clustered by individual and firm. Levels of significance: *10%, ** 5%, and *** 1%.

Table 1.20: Impact on Earnings: Alternative Measures of Outside Options

	Positions		Hires	
	Baseline	Weighted	Unweighted	Weighted
	(1)	(2)	(3)	(4)
Δ Log Earnings	0.703 ***	0.243 ***	0.147 **	0.069 **
	(0.134)	(0.052)	(0.072)	(0.031)
Δ Log Earnings (Narrow)	0.998 ***	0.333 ***	0.488 ***	0.164 ***
	(0.185)	(0.067)	(0.139)	(0.047)
Δ Log Base Pay	0.231 ***	0.070 ***	0.100 ***	0.035 **
	(0.056)	(0.018)	(0.038)	(0.015)
Bonus/Base Pay	0.991 ***	0.361 ***	0.334	0.134
	(0.307)	(0.095)	(0.329)	(0.094)
Observations	56134045	56134045	56134045	56134045

Note: This table presents estimates of γ from equation 1.5 for different measures of Ω_{it}.. The rows correspond to the earnings outcomes described in Section 1.12.4. Each regression controls for individual and four-digit industry-by-time fixed effects. Coefficients are scaled by 10000, for readability. Earnings are in kroner. Standard errors are two-way clustered by individual and firm. Levels of significance: *10%, ** 5%, and *** 1%.

Table 1.21: Impact on Mobility: Robustness to Alternate Network Definitions

	Baseline		Alternative Network Definitions: Past Two Years		Past Five Years		Past Three Years	
	(1)		(2)		(3)		(4)	
Any Transition	1.255	***	1.735	***	1.220	***	3.195	***
	(0.255)		(0.269)		(0.269)		(0.554)	
Job-to-Job	1.378	***	1.925	***	1.401	***	3.542	***
	(0.263)		(0.274)		(0.274)		(0.545)	
Connected Firm	1.408	***	1.993	***	1.405	***	3.691	***
	(0.263)		(0.270)		(0.270)		(0.543)	
Connected Industry	-0.014		-0.008		-0.001		-0.043	**
	(0.010)		(0.011)		(0.011)		(0.021)	
Unconnected Firm	-0.015	*	-0.051	***	-0.010		-0.040	**
	(0.009)		(0.010)		(0.010)		(0.019)	
Out-of-Sample Firm	-0.001		-0.009		0.006		-0.066	***
	(0.015)		(0.018)		(0.018)		(0.023)	
Firm-Size Cutoff	1000		1000		1000		500	
Observations	57922601		58819990		44047438		57769832	

Note: This table presents estimates of γ from equation 1.5 using different network definitions. The rows correspond to different mobility outcomes. Each regression controls for individual and four-digit industry-by-time fixed effects. Standard errors are two-way clustered by individual and firm. Coefficients are scaled by 10000, for readability. Levels of significance: *10%, ** 5%, and *** 1%.

Table 1.22: Impact on Earnings: Robustness to Alternate Network Definitions

	Baseline	Alternative Network Definitions: Past Two Years	Past Five Years	Past Three Years
	(1)	(2)	(3)	(4)
Δ Log Earnings	0.703 ***	1.362 ***	0.777 ***	1.736 ***
	(0.134)	(0.132)	(0.132)	(0.240)
Δ Log Earnings (Narrow)	0.998 ***	1.847 ***	1.050 ***	2.444 ***
	(0.185)	(0.186)	(0.186)	(0.399)
Δ Log Base Pay	0.231 ***	0.300 ***	0.373 ***	0.464 ***
	(0.056)	(0.063)	(0.063)	(0.106)
Bonus/Base Pay	0.991 ***	2.209 ***	0.929 ***	3.481 ***
	(0.307)	(0.354)	(0.354)	(1.065)
Firm Size Cutoff	1000	1000	1000	500
Observations	56134045	57002153	42723997	55986364

Note: This table presents estimates of γ from equation 1.5 using different network definitions. The rows correspond to the earnings outcomes described in Section 1.12.4. Each regression controls for individual and four-digit industry-by-time fixed effects. Standard errors are two-way clustered by individual and firm. Coefficients are scaled by 10000, for readability. Earnings are in kroner. Levels of significance: *10%, ** 5%, and *** 1%.

Table 1.23: Mobility: Exploiting Annual Data

	Change in Employment		Hires		Pct Change in Employment		Positions, Scaled by VA	
	(1)		(2)		(3)		(4)	
Job-to-Job	0.031	***	0.044	***	0.014	**	13.229	***
	0.012		0.017		0.007		4.462	
Connected Firm	0.044	***	0.071	***	0.026	***	16.829	***
	0.012		0.018		0.006		3.547	
Connected Industry	-0.008	**	-0.013	**	-0.003		-0.675	
	0.003		0.005		0.002		1.310	
Unconnected Firm	-0.004		-0.014	**	-0.009	***	-2.924	**
	0.004		0.006		0.002		1.308	
Observations	4205495		4205495		4203842		4205495	

Note: This table replicates our main mobility results using annual data from the IDA (the integrated database for labor market research). Coefficients are scaled by 10000, for readability. Earnings are in kroner. More details are provided in Appendix 1.13.

Table 1.24: Earnings: Exploiting Annual Data

	Change in Employment		Hires		% Change in Employment		Positions, Scaled by VA	
	(1)		(2)		(3)		(4)	
Δ Log Earnings	0.028	***	0.028	***	0.007	**	2.651	
	(0.007)		(0.010)		(0.003)		(2.040)	
Observations	2948997		2948997		2947863		2948997	
Δ Log Daily Earnings	0.013	**	0.003		0.004		8.315	***
	(0.006)		(0.008)		(0.003)		(1.932)	
Observations	2183376		2183376		2182526		2183376	

Note: This table replicates our main earnings results using annual data from the IDA (the integrated database for labor market research). Coefficients are scaled by 10000, for readability. More details are provided in Appendix 1.13.

Table 1.25: Results By Relative Productivity of Outside Firm

	Full Sample						Job Stayers			
	Job to Job Transition (1)		Change in Log Earnings (2)		$1\{\Delta$ Log Earnings>0$\}$ (3)		Change in Log Earnings (4)		$1\{\Delta$ Log Earnings>0$\}$ (5)	
					A. Baseline					
More Productive	1.131	**	0.089		0.159		0.103		0.127	
	(0.526)		(0.107)		(0.204)		(0.111)		(0.222)	
Less Productive	1.029	**	0.655	***	0.561	*	0.490	**	0.539	
	(0.504)		(0.234)		(0.329)		(0.235)		(0.363)	
Observations	22361250		21639845		22361250		20863189		21309619	
					B. Controlling for Same-Vigintile Impact					
More Productive	1.110	**	0.085		0.166		0.102		0.128	
	(0.519)		(0.106)		(0.203)		(0.111)		(0.222)	
Less Productive	1.004	**	0.650	***	0.571	*	0.487	**	0.541	
	(0.504)		(0.234)		(0.329)		(0.235)		(0.363)	
Observations	22361250		21639845		22361250		20863189		21309619	

Note: This table shows mobility and wage responses differ based on the productivity of the outside firm. We group firms into vigintiles based on their mean value added per worker (in real terms) over the sample period. We are able to do this for roughly 75% of the firms in our dataset. We then construct measures of Ω_{it} using only firms from higher and lower productivity firms (with strict equality). Panel A presents estimates of γ^{ABOVE} and γ^{BELOW} from equation 1.14. Panel B adds controls for Ω_{it} based on coworkers in the same vigintile. Because not all individuals have coworkers in the same vigintile, we replace missing values with 0's, and include a dummy for whether an individual has any coworkers in the same vigintile. Each regression controls for individual fixed effects, four-digit industry-by-time fixed effects, vigintile fixed effects, and a linear control for the number of connections in an individual's network. Standard errors are two-way clustered by individual and firm. Coefficients are scaled by 10000, for readability. Levels of significance: *10%, ** 5%, and *** 1%.

Table 1.26: Model Fit

Description	Target (full sample)	Model (full sample)
J2J Transition Rate	0.0071	.0126
U2E Transition Rate	0.0374	0.0389
E2U Transition Rate	0.0189	0.0198
Mean log monthly earnings (kroner)	10.45	11.27
Residual variance of log monthly earnings	.1623	.6354
Mean change in log earnings	0.0015	.0085
Variance of log earnings changes	0.1129	0.1124
Mean log earnings gain during J2J transition	0.141	0.162
γ^{J2J}: Regression coefficient for J2J mobility	0.000118	0.000159
γ^{U2E}: Regression coefficient for U2E mobility	.000169	.000879
Regression coefficient for earnings of job-stayers: $\Delta \log y$	.00000406	.0000106
Regression coefficient for earnings of job-stayers: $1\{\Delta \log y > 0\}$	0.000027	0.000134

Note: This table examines the fit of our estimates for model described Section 1.7. Note that in order to comply with Statistics Denmark's privacy regulations, we do not present quantiles. Standard errors are a work in progress.

1.11 Theoretical Appendix

In this section we provide the proofs referenced in the text. Many of the results follow those in Flinn and Mullins (2017).

1.11.1 Main Text Proofs

Lemma. *1 A worker who receives the total surplus created by the match θ at a renegotiating firm (type R) has the same value as a worker earning θ at a posting firm (type P). That is,*

$$T_R(\theta) = V_P(\theta)$$

Proof. This proof follows that in Flinn and Mullins (2017). Let Φ be the endogenous distribution of offers from posting firms. A worker who is at a posting firm earning w has the following value function:

$$\begin{aligned}(\rho+\delta)V_N(w) \quad = \quad & w + \underbrace{\lambda_E p_R \int \beta\left[T_R(x) - V_N(w)\right]^+ dF_\theta(x)}_{\text{better offer: renegotiating firm}} \\ & + \underbrace{\lambda_E(1-p_R)\int \left[V_N(x) - V_N(w)\right]^+ d\Phi(x)}_{\text{better offer: posting firm}} \\ & + \quad \underbrace{\delta V_U}_{\text{unemployment}}\end{aligned}$$

The total surplus created by a match θ at a bargaining firm is similar.

$$
\begin{aligned}
(\rho+\delta)T_R(\theta) &= \underbrace{\theta + \lambda_E p_R \int \beta\left[T_R(x) - T_R(\theta)\right]^+ dF_\theta(x)}_{\text{better offer renegotiating firm}} \\
&\quad \underbrace{+\lambda_E(1-p_R)\int \left[V_N(x) - T_R(\theta)\right]^+ d\Phi(x)}_{\text{better offer posting firm}} \\
&\quad + \underbrace{\delta V_U}_{\text{unemployment}}
\end{aligned}
$$

This period the match produces θ. With probability $\lambda_E p_R \int \beta\left[T_R(x) - T_R(\theta)\right]^+ dF_\theta(x)$ the worker gets an offer from a bargaining firm with higher overall match surplus. The probability depends on the arrival rate, the fraction of posting firms, and the density of offers that are better. When this occurs, the worker leaves the firm. She is able to use $T_R(\theta)$ as her outside option, and gains a fraction β of the match surplus. Because there is free entry, the firm's value is 0. With probability $\lambda_E(1-p_R)\int\left[V_N(x) - T_R(\theta)\right]^+ d\Phi(x)$ the worker receives a more attractive offer from a posting firm. Her new value at that firm is simply $V_N(\theta)$. Again, the firm's value is zero, by free entry. Both equations will hold if $V_N(\theta) = T_R(\theta)$. □

1.11.2 Firm Value Functions

If a vacancy of match quality θ offers wage w, the expected discounted profit is the probability the vacancy is filled, multiplied by the discounted stream of profits. The probability the vacancy is filled is:

$$
\mathscr{H}(w) = \underbrace{\frac{\lambda_U M_U}{\lambda_U M_U + \lambda_E M_E}}_{\text{unemployed}} \times \underbrace{1}_{\text{accept offer}} + \underbrace{\frac{\lambda_E M_U}{\lambda_U M_U + \lambda_E M_E}}_{\text{employed}} \times \underbrace{G(w)}_{\text{accept offer}}
$$

If the firm meets an unemployed worker, the vacancy is filled with probability 1. The firm meets such a worker with probability $\frac{\lambda_U M_U}{\lambda_U M_U + \lambda_E M_E}$. The numerator of this expression is the search intensity of unemployed workers. The denominator is the aggregate search intensity. The worker meets an employed worker with probability $1 - \frac{\lambda_U M_U}{\lambda_U M_U + \lambda_E M_E}$.. The offer is accepted only if the worker is currently at a firm that would be willing to pay her at most w. We use $G(w)$ to denote this function.

The expected discounted profits associated with the offer are $\theta - w$, divided by the expected length of the match. This is $\rho + \delta + \lambda_E \left\{ f_R \tilde{F}_\theta(\theta) + (1 - f_r)\tilde{\Phi}(\theta) \right\}$. The match is exogenously dissolved at rate δ. The final term measures the probability of an endogenous separation.

We use $J(\theta, w)$ to denote the value to a firm of opening this type of vacancy and paying w:

$$J(w, \theta) = \mathscr{H}(w) \times \frac{\theta - w}{\rho + \delta + \lambda_E \left\{ f_R \tilde{F}_\theta(\theta) + (1 - f_R)\tilde{\Phi}(\theta) \right\}}$$

The firm chooses w to maximize this expression. A key contribution of Flinn and Mullins (2017) is to show that, in this setting, the firm's optimal wage offer is given by a deterministic function $\omega(\theta)$ (denoted $\varphi(\theta)$ in that paper) that is monotonically increasing in θ, lower semi-continuous, and almost everywhere differentiable. We do not repeat the proofs here, and direct the reader to that paper. that paper also derives the differential equation that describes this wage offer function.

It is important to note that when we take the model to the data, we simply allow the offers from posting and bargaining firms to come from different distributions. We do not directly estimate the link between the productivity of posting firms and the wages that they post.

1.11.3 Closing the Model

Flinn and Mullins (2017) close the model by assuming that posting and bargaining firms face different costs of posting a vacancy, and that the costs of posting a vacancy are increasing in the measure of each type. However, for some values of costs, there is no solution.

We take a different approach. We assume that when a firm wants to open a vacancy, it takes a random draw from the (exogenously given) productivity distribution. With probability $1 - p_R$, the firm can post wages for the vacancy. With probability p_R the firm has to bargain with workers. This probability is exogenously given. When we estimate the model, we allow this to vary across skill groups.

This is somewhat simpler than the set-up in Flinn and Mullins (2017) because it allows for a single free-entry condition. We assume the marginal cost of posting a vacancy is c. Then the free entry condition is given by:

$$\begin{aligned} c &= q(\kappa)p_R \int_{\theta^*}^{\theta^{\max}} \int_{\theta^*}^{\theta^{\max}} (1-\beta)\left[T_R(\theta) - T(x)\right]^+ dF_\theta d\hat{G}(x) \\ &\quad + q(\kappa)(1-p_R) \int_{\theta^*}^{\theta^{\max}} J(\theta, \omega(\theta)) dF_\theta d\hat{G}(x) \end{aligned}$$

Firms enter until the cost of posting a vacancy is equal to the expected benefit. This depends on $q(\kappa)$, the rate at which workers meet vacancies, and on the expected rents associated with the match. The rate at which workers meet vacancies depends on the total number of vacancies, and the form of the matching function. When wages are set by bargaining, the firm gets to keep $(1-\beta)$ of the difference between the rents produced in the match, and the rents produced by the hypothetical match the worker used for bargaining. When wages are posted, the firm's value function is $J(\theta, \omega(\theta))$ as defined above.

1.11.4 Equilibrium

The equilibrium of the model is characterized by a set of flow equations and three steady state distributions:

1. the distribution of workers across renegotiating firms $G(x)$
2. the distribution of 'last best' offers for workers at each type of renegotiating firm $H(q|x)$
3. the distribution of workers across posting firms

Each of these can be derived using the relevant balance condition. We follow Flinn and Mullins (2017) to derive the equilibrium distributions. We omit worker ability for simplicity.

Unemployment

In equilibrium, the flow rates in and out of unemployment must balance. Each period a fraction δ of workers are displaced from their jobs. Using M_U to denote the fraction of workers who are unemployed, we can write

$$\begin{aligned}
\underbrace{\delta(1-M_U)}_{\text{Entering}} &= \underbrace{M_U\left(\lambda_U\left(p_R\left[1-F_\theta(\theta^*)\right]+(1-p_R)\left[1-\Phi(\theta^*)\right]\right)\right)}_{\text{Exiting}} \\
M_U &= \frac{\delta}{\delta+\lambda_U\left(p_R\left[1-F_\theta(\theta^*)\right]+(1-p_R)\left[1-\Phi(\theta^*)\right]\right)}
\end{aligned}$$

Each period, a mass δ of workers who were unemployed last period, become unemployed. The probability a worker who was unemployed last period becomes employed is the probability she meets a vacancy, multiplied by the probability that vacancy exceeds the value of unemployment.

Contact Rates

In equilibrium, the contact rates are determined via a standard matching function. Unemployed workers come across a vacancy with probability λ_U and employed workers come across a vacancy with probability λ_E.

We define

$$\kappa = \frac{v}{\lambda_U M_U + \lambda_E M_E}$$

to be the market tightness measure. This is the ratio of the number of vacancies to the (search-intensity weighted) number of searchers. If we assume the matching function is Cobb-Douglass, $\lambda_U = \kappa^\gamma$ and the rate at which workers arrive at vacancies is $q(\kappa) = \kappa^{\gamma-1}$.

Distribution of Workers Across Firm Types

We can derive the distribution of workers across firm type by looking at the steady state relationship. $G(x)$ measures the fraction of workers at firms of match quality x or below.

$$
\begin{aligned}
dG(x) &= \underbrace{M_E G(x)}_{(1)} \left\{ \underbrace{\delta + \lambda p_R \tilde{F}_\theta(x) + \lambda(1-p_R)\tilde{\Phi}(x)}_{(2)} \right\} - \\
&\quad \underbrace{M_U \lambda}_{(3)} \left\{ \underbrace{p_R \left(F_\theta(x) - F_\theta(b)\right) + (1-p_R)\left(\Phi(x) - \Phi(b)\right)}_{(4)} \right\}
\end{aligned}
$$

The first line measures the flow out of these firms. The measure of workers currently at firms of quality x or lower is given by expression (1): the mass of employed workers multiplied by the fraction at these types of firms. The second term gives us the probability a worker is no longer at one of these firms. This can occur if they are displaced and move into unemployment, or if they receive an offer from a better firm.

The second line measures the flow into these firms. No workers at higher quality firms will ever flow into these firms (directly). The inflow is the product of (3): the measure of unemployed workers who get a job offer and (4): the probability that, conditional on receiving an offer, it is more attractive than unemployment.

In equilibrium inflows equal outflows and $dG(x) = 0$. We can rearrange the above expression to get

$$
G(x) = \frac{M_U \lambda \left(p_R \left(F_\theta(x) - F_\theta(b)\right) + (1-p_R)\left(\Phi(x) - \Phi(b)\right)\right)}{M_E \left(\delta + \lambda p_R \tilde{F}_\theta(x) + \lambda(1-p_R)\tilde{\Phi}(x)\right)}
$$

We use $G(x, R)$ and $G(x, P)$ to denote the distributions of workers at each type of firm and g_r, g_p to denote the corresponding densities.

Distribution of Best Offers for Workers at Renegotiating Firms of Each Match Quality

We can use a similar logic to derive the distribution of best offers received for workers at each type of firm. This is important because it determines the distribution of wages that we observe: at bargaining firms, workers' wages directly depend not only on the type of firm they are at, but on the best offer they have received. The flow equation for workers at type x firms whose last best offer was from a firm of type at most q is:

$$\begin{aligned} dH(q|x) \times g_r(x) &= -\left(\delta + \lambda_E p_R \tilde{F}_\theta(q) + \lambda_E(1-p_R)\tilde{\Phi}(q)\right) H(q|x)g(x,R)M_E \\ &\quad + \lambda_E p_R f_\theta(x) G(q) M_E + \lambda_U p_R f_\theta(x) M_U \end{aligned}$$

The flow rate into this state depends on the probability an employed worker receives an offer from a firm of q or lower or the probability an unemployed worker receives an offer from a firm of type x. The flow rate out of this state depends on (1) the probability the worker is no longer at a firm of type x and (2) the probability a worker receives an offer better than q.

1.12 Data Appendix

1.12.1 Data Sources

Individual Characteristics

Demographic Information First, we obtain basic demographic information from the BEF and FAIN registers: sex, year of birth, and country of origin. The two registers draw from different administrative databases, but together provide nearly complete coverage.

Family Structure Next, we use the annual FAM and FAIN registers to determine whether someone has children in a given year. The annual BEF population register provides a unique identifier for each individual's spouse or partner if the individual in question has a valid person identifier in Denmark. While Statistics Denmark does distinguish between types of couples, they provide a partner ID if two people are living together and fall into one of the following four types:

1. Married couple
2. Registered partnership
3. Non-married or registered couple that live together and have at least one child in common
4. Cohabiting couple: two persons of different sex who live together with no other adults and who have an age gap of less than 15 years

Following Statistics Denmark, we consider anyone with a valid partner ID to be in a couple. We consider anyone in a type 1 or 2 relationship (married or in a registered partnership) to be married. We include these variables in our baseline regression in some specifications.

Education Our education variables come from the UDDA ("Uddannelser") register. This register combines information in the student register– which contains information on registered education in Denmark– and the qualification register. The qualification register in turn combines a number of sources including self-reported and imputed information on immigrants' education and information from professional membership registers (e.g. engineering associations). We focus on an individual's highest completed education, and use Statistics Denmark's own crosswalks to convert the detailed education codes (for information from the student register) to ISCED codes.

Occupation We use the BFL and RAS registers to code individuals' occupations. Most observations in the BFL register contain a six-digit occupation code ("DISCO"). We use the first four digits (prior studies including Groes et al. 2014 note that these are roughly equivalent to three-digit SOC codes) or the first two digits. Each table notes which aggregation level we employ. There are a small number of observations that do not have a valid occupation code. For these observations, we supplement the data with information from the RAS ("Registerbaserede arbejdsstyrkestatistik") register.

In Danish register data (both BFL and RAS) there is a break in the occupation coding in 2010.[39] Before this the register data report the codes based on the 1988 coding (which changes over time); after this the register data use the 2008 codes. This does not impact our main analysis because we typically rely on occupation by time fixed effects, allowing the coding to vary by period. To construct the occupation based networks we assign workers to a single occupation based on their most frequently reported occupation over the 2010-2016 time period. To construct firm by occupation by time fixed effects or industry by occupation by time fixed effects, we use 2-digit occupation codes.

Firm Characteristics

We obtain a number of firm characteristics from the BFL, IDA, RAS, FIRM, and FIRE registers.

Industry We obtain information on industry from the BFL. The BFL includes six-digit industry codes for each firm in the data. The six digit industry variables are too detailed for our purposes. For instance, they distinguish between stores that sell women's clothing and men's clothing, and

[39] This break is not indicated in any of the official online descriptions of the registers.

stores that sell both men and women's clothing. In all of our analysis we use the first four digits (which correspond to the NACE code) or first two digits.

Region We use data from IDA (the integrated database for labor market research) to assign each firm to one of Denmark's five administrative regions: (1) the capital region, (2) Southern Denmark, (3) Northern Denmark, (4) Central Denmark, and (5) the Zealand region. We are able to assign most firms to regions using data from the firm-level panel. For the small number of firms with multiple establishments in multiple regions, we use the region where the firm has the greatest number of employees. There are a small number of firms that do not have a district listed in the firm-level panel. For these observations, we use information from the worker panel to assign the firm to the region where the greatest number of employees live. Appendix Figure 1-18 shows a map of the five regions.

Value Added We use data from a firm accounting register (FIRE/FIRM) to calculate value added per worker. It is straightforward to calculate value-added following the procedure in Bagger et al. (2014a). This is the same procedure Statistics Denmark uses to produce national accounts.

We group firms into vigintiles based on mean value added per worker (in real terms) over the sample period.

1.12.2 Monthly Series

As discussed in section **??**, our monthly earnings and hours data come from the administrative monthly employment for employees ("BFL") register. Danish firms are required by law to report wages paid at least once a month to the Danish Customs and Tax Administration daily. Firms also typically report the number of hours worked. The hours data should be of reasonably high quality because the obligatory payments to the Danish supplementary pension fund (ATP) depend on hours worked.[40] Statistics Denmark compiles these data into the BFL register based on the information

[40]The cost is nominally shared by the employee and employer; payment is taken as a payroll deduction. There are four different bins corresponding to full-time, 2/3 time, 1/3 time, and less than 1/3 time. In terms of monthly hours the bins are: 0-38, 39-77, 78-116, and 117+ hours. In terms of weekly hours the bins are: 0-8, 9-17, 18-27, and 27+.

provided to the Customs and Tax Authority. In cases where the firm does not report any hours, or Statistics Denmark considers the data 'invalid or improbable', hours are imputed. The data imputation flag indicates that approximately 15% of the hours data are imputed.

This register contains all employees in Danish registered companies, regardless of whether the employee lives in Denmark or abroad. The register reports data at the person-month-firm level. A key strength of our data is that for each observation there is start date and an end date (number within the month). Individuals who transition between firms within a month will have earnings observations at both firms that month, with the end date at one firm preceding the start date at the second.

In order to create a monthly series, we restrict attention to observations that span the first of each month. This means that moves that occur mid-month will only be captured by the following month's data. We then re-scale earnings, when necessary, so that they are equal to a full month's work. The table below illustrates this. Here, the individual worked at firm A in the first month and for half of the second month. Then, she switched to firm B, where she continued to work in month 3. Our final dataset lists this individual as working at the first firm in the first and second months and the second firm in the third month. Earnings during the second month are multiplied by two so they are equal to a full month's work. If we did not perform this adjustment, we would over-state the wage increase the individual got when he/she switched firms. Based on this dataset, we will first observe that a move has occurred in period 3.

Person	Period	Start	End	Firm	Raw Earnings	Adjusted	Monthly Panel
1	1	1	31	A	2000	2000	Yes
1	2	1	15	A	1000	2000	Yes
1	2	16	31	B	1500	3000	No
1	3	1	31	B	3000	3000	Yes

One unique feature of the BFL is that individuals who leave the firm for a short period of time (less than 45 days) but return are reported as still working at the firm, but receiving no earnings. This may occur if the worker is receiving training, or is on a short medical leave. This means that short disappearances and reappearances from the firm will not be counted in the firm-level hiring shock.

There are a very few person-firm-month observations (less than half of a percent) where an individual has two records for a single employer. In ninety-six percent of these problematic observations, there are two records for the individual-firm that month. Most of these observations appear to be for salaried workers. One of the observations corresponds to their salary earnings (160 hours per month when hours are reported); the second observation appears to be the result of additional hours worked, likely due to over-time payments. We create a separate earnings variable: total earnings and total hours.

Comparison with other Administrative Registers Relative to the standard Danish employer-employee dataset (IDA), our data are unique in providing monthly (not annual) data and in containing all employment spells for all workers. The IDA only includes data for workers employed as of the last week of November in the reference year and does not contain start or end dates for each worker-firm-year observation. To deal with these shortcomings, some researchers have used other registers to construct a weekly 'spells' dataset with high frequency information on workers' employment status (but not earnings) (see, e.g. Bagger et al., 2014b). The key advantage of the BFL data is the addition of earnings.

Our data also have several advantages, relative to other linked employer-employee datasets. The inclusion of hours data (not hours bins) is unique, relative to administrative registers in Germany and the United States. The fact that we observe all firms is an advantage relative to industry datasets in the United States. Other countries with administrative linked employer-employee data–including Portugal and Italy–have very rigid labor markets. It is not clear results in those contexts would be relevant for the United States.

1.12.3 Sample Selection Criteria

There are five million people in Denmark. The labor force participation rate is around 65%. Between a fifth and a quarter of workers are part-time. There are nearly three million workers in our dataset. Starting from the raw BFL dataset we make five primary sample restrictions. First, we drop individuals who are not Danish citizens. We do this because the BFL includes observations on all employees of Danish firms, including those who do not work in Denmark. In addition, our demographic information is most complete for Danish citizens.

Second, we drop individuals who never work at firms with fewer than 1000 people. This removes about 15% of the sample. These individuals are not used to construct the networks and do not appear in our regression sample.

Next we remove observations for workers who are younger than 25 or older than 60. This step also removes workers for whom we do not have valid birth year data.

Finally, we restrict our sample to individuals who, over the course of our sample, never work in more than one job at a time. We identify a multiple job holder as someone who works in two distinct firms (tax identifiers) on the first of the month and has positive earnings in both firms. In some cases we see individuals with a one-month long spell at a firm that overlaps another, longer spell. We include these individuals in our main sample. If we did not do this, our final sample would include roughly 800,000 workers. Because we allow for these one-time slips, our final sample includes slightly over one million workers.

We think that our measure of multiple job holding is somewhat conservative. However, we see job-to-job transition rates in this sample that are similar to those reported using other Danish registers. In practice, this restriction removes a large number of part-time workers. We have verified that our analysis is robust to including only the roughly 800,000 workers who meet the more stringent definition of multiple job-holding and to including a broader set of workers.

Finally, our main regression sample focuses on workers who are currently in firms with below 1000 workers. This is somewhat different from the sample of workers who *ever* work in one of these firms over the eight year period. We track whether workers move in or out of the sample but, once they leave, do not continue to follow them. This is because these workers cease to accumulate connections. Our results suggest that the length of time since an individual worked with a connection matters for information transmission. Table 1.13 shows how each restriction impacts the size of the analysis sample.

1.12.4 Earnings and Hours Data

The earnings data are reported electronically directly by the firms to the Danish Customs Authority. We scale all earnings and hours to the monthly level using the start and end dates in the register. There are substantial period-to-period changes in hourly earnings in the raw data. This arises for several reasons.

1. **Severance Pay:** Workers in Denmark are often eligible for severance pay when they leave a job. Including this in our measure of earnings would lead us to think that many workers see nominal wage *decreases* upon switching firms.

2. **Fringe Payments:** Second, the broad earnings measure includes payments for fringe benefits, including pay for housing or telephones, pay for vacation, and some contributions to retirement plans. While some of these may be smoothed over a worker's tenure, others may be reported lumpily by the firm. It is hard for us to, without additional data, remove these from a firm's base pay.

3. **Overtime:** Third, hourly earnings may change in response to changes in hourly worked that result in overtime payments.

4. **Annual Bonuses:** Finally, because we examine monthly data, annual bonus payments will lead to large monthly changes in income. If a worker receives a 15% annual bonus in December, for instance, her hourly earnings would more than double that month. She would then see her pay cut in half the following month. A worker who received a 20% annual bonus would see her earnings increase by 200% and then fall by a similar amount.

Earnings Measure

Most of our analysis focuses on a measure of monthly earnings that is processed in two ways:

1. **Severance Pay:** We remove severance pay. White collar workers in Denmark are sometimes eligible for severance pay upon dismissal, depending on their tenure at the firm. They are typically entitled to 1 month of pay after 10 years of tenure or 3 months of pay after 20 years of tenure. Firms may also elect to pay severance pay beyond that required by law. We do not want to include this in our baseline earnings measures as it will lead us to spuriously conclude that an individual received a raise in their last month at the firm.

   In cases where an individual appears to have been given severance pay upon termination, we re-code their final month's earnings with their earnings in the prior month. Specifically, we identify workers whose earnings more than doubled in their final month of employment at the firm (not due to changes in hours). For these observations we re-code final month's earnings with earnings from the prior month, adjusting for differences in hours as necessary. This is illustrated in the figure below. The earnings numbers are deliberately stylized:

Period	1	2	3	4	5	6
Firm	A	A	A	$\emptyset$	$\emptyset$	B
Raw	5000	5000	20000			4000
Δ Raw		0	15000			
Clean	5000	5000	5000			4000
Δ Cleaned		0	0			

2. **Double Pay:** We correct a small number of observations where an individual is not paid in one month, but receives twice their normal pay the following month. In this case we spread the earnings evenly over the two months.

In practice we still see substantial volatility in monthly earnings, partially arising from what appear to be annual bonuses. Because we have no reason to expect this volatility to be correlated with our measure of outside options, we prefer to use this measure as is, rather than attempt to smooth it in some way. Measurement error on the left hand side will simply inflate our standard errors.

Wages We construct wages by dividing earnings by hours worked. We only use the subset of the data with firm-reported hours. It is important to note that this measure may still vary with hours worked, if individuals cross ATP-contribution thresholds or if they receive overtime.

Real Earnings Most of our analysis is based on changes in nominal wages. In some descriptive tables we report real earnings measures, converted into US dollars. To convert earnings into real 2016 numbers, we use a consumer price index provided by Statistics Denmark. We then use the 2016 exchange rate between the US dollar and the Danish krone to convert numbers into dollars.

Base Pay and Bonuses

In some of our analysis we further process the data in order to investigate the impacts on base pay and bonuses. Our goal is to separate base pay and bonuses as in the following picture:

Raw	100	120	100	100
Earnings	100	100	100	100
Bonus	0	20	0	0

We define bonuses are one month increases in pay that revert the following month. We identify these bonuses by looking for earnings growth that:

1. Lasts one month and:

   (a) Is not driven by changes in wages. We require there to be more than a 7 DKK (~$1) increase in hourly earnings[41]

   (b) Is not permanent. We require that this month's earnings (wages) are more than 70 DKK (7DKK) larger than next month's. This allows us to ignore raises.

[41] In practice the exact value of this cutoff does not matter.

2. In some cases, we see what appears to be both a bonus and a raise. we see bonuses and raises at the same time. In this case we examine how wages change the following month. This is illustrated below:

Raw	100	150	130	130
Earnings	100	130	130	130
Bonus	0	20	0	0

We have done some sensitivity to thresholds of the above. While these decisions do impact the overall distribution of changes, they appear to be uncorrelated with our measure of outside options, and thus do not impact our point estimates.

In practice we think our ability to distinguish between bonuses and base pay is best for salaried workers. For these workers the percentage of observations with a bonus is 6% and the percentage of observations with a raise is 7%. These both seem reasonable: if all individuals received an annual bonus or received raise (even a cost of living adjustment) each year, we would expect to see rates of $\frac{1}{12} \approx 8\%$. The mean bonus is around half of an individual's usual monthly earnings, though there is substantial variation.

1.12.5 Mobility

There are three possible transitions a currently employed worker could make in a given period:

1. Stay: the worker is at the same firm this month as he/she was at last month
2. Exit: the worker was employed last period, but is not employed this period. This could mean unemployment or non-employment.
3. Job to Job Transition: the worker is employed in both periods, but at different firms. We decompose these moves into two types:
   (a) Move to a connected firm ("connected move"): the worker moves to a firm where he/she has a former coworker.

(b) Move to an unconnected firm ("unconnected move"): the worker moves to a firm where he/she does not have a former coworker but that firm still falls within our sample of "not too large" firms

(c) Move to a large firm ("out-of-sample move").

Note that moves that lead to an individual being not employed over the first of the month will be coded as an exit. This could occur if the individual takes time off between ending one job and starting the next, even if the transition was entirely voluntary.

Because our concept of the firm is a tax-identifier, individuals whose firms are involved in a merger or acquisition may see a change in firm codes. We recode the small number of observations that appear to be associated with this type of move as a non-transition. However, we do not include them in our analysis of stayers' earnings. In practice, this did not affect our analysis in any way. Table 1.15 presents descriptive statistics on transitions.

1.12.6 Networks

We use the baseline monthly series to construct an individual's coworker network. We first augment our data with register data from MIA (cleaned in the same way) so that we can construct coworker networks for individuals for the first three years of the BFL series (2008-2011). We then drop observations that correspond to spells in a firm with less than 2 people or more than 1000 people. These are standard restrictions in the empirical networks literature. We then construct the bipartite adjacency matrix A. This is a symmetric matrix where $A_{ij} = 1$ whenever i and j worked together in the past 3 years (excluding the period in question). Because we do not want to include shocks that come from an individual's former firms, we then remove all coworkers who are currently working at these firms (or the individual's current firm).

Note that because we have data from MIA dating back to 2004, we could have, theoretically, extended our analysis of job to job mobility by a few additional years. One issue, however, is that in 2005 and 2007 there were significant changes in some of the establishment identifiers, coincident with a reorganization of Danish municipalities. This is not a problem for how we define the networks because, within a month, we are still able to identify who is working together. However, it does make it more difficult to identify job-to-job transitions.

Future Networks Each individual's future coworker network consists of all individuals she works with in the next year or two-three years who are not currently working at firms she moves to. As we did when constructing an individual's past coworker network, we construct an individual's future coworker network, leaving out connections at firms she moves to in the next three years. As a result, her future coworker network consists of workers who will come join her at her current (or her future) firm, but have, themselves, not moved yet. If we did not do this, the size of the future coworker network would mechanically vary with job to job mobility decisions. We also exclude individuals who are in her former coworker network. The bottom row of Table 1.7 shows that the number of future coworkers in each year is roughly equivalent to the number of prior coworkers in each year. We only construct these networks for observations from January 2008 to December 2013 (60 months) so that the quality of the measure does not vary across periods in our sample.

Computation To generate each coworker network, we must generate and store an N x N matrix. We use MATLAB's sparse matrix packages to efficiently handle the data.

1.12.7 Trade Details

Denmark is a small open economy that is thoroughly integrated into the world market. Its main source of imports and destination of exports is Germany, which it shares a border with. Sweden, Britain, Norway and the United States are also important partners.

We merge our data with monthly bilateral trade flow from COMTRADE in order to compute the predicted value of exports for each firm and month, based on world product demand. We have flows from January 2010 to March 2016 at the six-digit Harmonized System level. For each product and month, we calculate the total value of imports of that product by all countries (less Denmark) from countries other than Denmark. We merge these to the administrative data at the firm-product-month level.

Instrument Construction

Step 1: We use data from 2004-2007 to calculate the average share of exports of product p that is accounted for by firm j. We do this by dividing the value of exports of product p by firm j over this time period by the total value of exports of product p by all firms over this time period:

$$\pi_p^j = \frac{\sum_t \sum_c \exp_{j,p,c,t}}{\sum_j \sum_t \sum_c \exp_{j,p,c,t}}$$

Step 2: We merge our administrative trade data with $\exp_{p,t}^{-1}$ from the product database after aggregating to the HS-6 level. We calculate the total predicted value of exports by weighting the leave-out measure from COMTRADE by these firm-specific product shares:

$$\hat{\exp}_{j,t} = \sum_{p,t} \pi_p^j \sum_c \exp_{\mathrm{p,c,t}}^{-1}$$

Step 3: We weight log predicted exports by an individual's former coworker network. The denominator is based on connections who are in firms covered by the trade register.

Instrument Quality

One concern is that firms may change their product or export destination mix over time. To examine this, we look at the correlation between realized and predicted exports over time. Figure 1-21 shows that the correlation between the observed and predicted measures is high, though declining over the sample period. The light blue dots show the correlation between a firm's (total) exports and that predicted using the product shares π_p^f during the period used to construct the shares. This is high by construction. The blue diamonds show the correlation between a firm's total exports and that predicted by weighting total Danish exports (as reported in the trade register) by π_p^f after the period used to construct the shares. This correlation trends down somewhat over time, but remains around .5 throughout our period, suggesting that firms do not change their product mix too much over time.

The dark blue squares show the correlation between actual firm exports and that predicted by weighting Danish exports in COMTRADE by the firm-level product shares. The correlation is similarly high. We would expect this to differ from the prediction using register data because of differences in reporting thresholds in COMTRADE and in the administrative data. The red triangles show the correlation between actual firm exports and predicted exports based on world demand. The correlation is fairly stable over our sample period.

1.12.8 Estimating the Model

There are five steps to evaluating the objective function specified in equation 1.13, given a set of parameters, a set of empirical moments, and a weighting function.

1. Given a guess of parameters, ξ, solve for $V_P(w)$ and $T_R(\theta)$ by value function iteration

2. Given $V_P(w)$ and $T_R(\theta)$, solve for the wage function $\phi(\theta, w)$, which says what wages a worker at a renegotiating firm of type θ will earn, if her last offer was from a firm whose max offer was w.

3. Simulate the model for 10,000 workers and 100 periods

   (a) Set worker's workers initial conditions using the equilibrium distribution of workers

   (b) For t=2:T

      i. With probability δ, an employed worker is exogenously separated from her employer

      ii. With probability $\lambda^E + \alpha s$ or $\lambda^U + \alpha s$ an employed or unemployed worker receives an offer from an outside firm. With probability p_R this is from a renegotiating firm. The worker decides to move, renegotiate, or stay (with no wage change)

4. Calculate the moments $S(\xi)$ using the simulated data

5. Find the weighted distance between the simulated and empirical moments

A common challenge in the literature is specifying the matrix W, which weights the distances between each of the simulated and observed moments. Given standard regularity assumptions, as $N \to \infty$ the estimator $\hat{\xi}$ is consistent and asymptotically normal for any positive-definite W (Gourieroux et al., 1993). However, in finite samples the weighting matrix often matters. We first attempted to use a diagonal matrix where the i^{th} element is the inverse variance of the i^{th} component of S_N.[42] In practice we, like prior authors, found that this led us to under-weight, and not match, key moments of interest (Jarosch, 2015). including transition rates and the variance of low wage changes. In order to match key moments, including the transition rates and variance of log wage changes, we over-weighted the relevant moments.

[42] For means we use bootstrapped standard errors.

1.13 Supplementary Results

1.13.1 Job Search by the Unemployed and Non-Employed

In this section we present results for unemployed and non-employed workers. Prior work in this literature has shown that workers displaced in a mass layoff use information from their job search networks in order to find new employment (Glitz, 2013; Saygin et al., 2014). Because we have not accounted for selection into unemployment or non-employment, we view this as a purely descriptive exercise. A more comprehensive analysis examining the impacts on workers involved in exogenous separations is beyond the scope of this paper.

We run our standard reduced form regressions on the sample consisting of workers whose last job was at an in-sample firm but who are, as of the first of the month, are not employed at any firm

$$\mathrm{U2E}_{it} = \sum_n \gamma^n \Omega_{it}^n + c_{it} + \alpha_i + \alpha_{jt} + \epsilon_{ijt}$$

The results are presented in Table 1.18. The first column presents the baseline results. Columns 2 and 3 condition on the individual remaining un/non-employed for at least 2 or 3 months. Overall, the results are similar to those presented in Section 1.5: an individual's more recent coworkers matter more for their recover from unemployment/non-employment. One thing to note is that the gradient is less steep: This is consistent with the idea that only an individual's very recent coworkers are likely to proactively give them information. However, when an individual is without work, they reach out to a broader set of former colleagues. Note that we deliberately exclude an individual's future coworkers from this specification: if an individual does not find reemployment, they do not have future coworkers.

1.13.2 Quality of Outside Options

While our baseline measure of Ω_{it} treats all firms equally, the theoretical model in Section 1.2 suggests that both the number and quality of outside options matter. If all firms renegotiated wages, only offers from higher-productivity firms would matter for mobility; only offers from lower-productivity firms would matter for on-the-job wage growth. For workers at posting firms, offers should only impact wages through mobility.

We use firm accounting data to group firms into vigintiles based on mean value added per worker over the sample period. We then construct measures of $\Omega_{it}^{\text{ABOVE}}$ and $\Omega_{it}^{\text{BELOW}}$ using only connections at firms in higher and lower vigintiles (with strict inequality); we construct $\Omega_{it}^{\text{SAME}}$ using connections at firms in the same vigintile. Note that we are only able to do this using the connections that are covered by the accounting register.[43] We then run our baseline regression, replacing Ω_{it} with these measures and adding fixed effects for the vigintile of an individual's current firm (v_{it}).

$$y_{it} = \gamma^{ABOVE}\Omega_{it}^{\text{ABOVE}} + \gamma^{\text{BELOW}}\Omega_{it}^{\text{BELOW}} + c_{it}^{\text{ABOVE}} + c_{it}^{\text{BELOW}} + v_{it} + X_{it} + \alpha_{kt} + \alpha_i + \epsilon_{it} \quad (1.14)$$

Table 1.25 reports estimates of γ^{ABOVE} and γ^{BELOW} from this regression (Panel A) and from regressions that control for the number of new positions created at firms in the same vigintile (Panel B). The coefficients in Panel A are presented graphically in Figure 1-23.

Figure 1-23 shows that higher and lower productivity firms have similar impacts on job-to-job mobility. By contrast, only positions at less productive firms matter for wage growth. The earnings results are exactly in line with the predictions of the model in Section 1.2: workers are able to use outside offers from less productive firms as leverage to renegotiate wages at their current firm.

The fact that workers move to positions at both more and less productive firms is somewhat at odds with a simple model where all firms renegotiate wages.[44] However, it is exactly what we would see if some workers are at firms that have committed not to renegotiate wages (Flinn

[43]The results for this exercise are somewhat less precise, reflecting the fact that we have a smaller sample and use a smaller number of coworkers in each network. Column 2 of Table 1.2 shows that we are able to rank roughly 75% of the firms in our sample. The regressions only include workers who have coworkers both at higher- and lower-productivity firms.

[44]While many search models predict that workers will be willing to accept pay cuts in order to move to more productive firms, on average, workers in our sample see larger wage gains if they move to more productive firms.

and Mullins, 2017).[45] In this case, mobility is not always efficient: workers may move to less productive firms if the incumbent firm is unwilling to renegotiate.

1.13.3 Posting and Bargaining by Occupation: Hall and Krueger (2012) Data

Hall and Krueger (2012) provided evidence on the incidence of wage posting and bargaining among workers in the United States. We used survey data from that paper to construct posting and bargaining rates by occupation.[46] Figure 1-22 plots the unweighted fraction of workers in each broad occupation group that answered "some bargaining over pay" when asked:

> When you were offered your (current/previous job), did your employer make a "take-it-or leave-it" offer or was there some bargaining that took place over the pay?
>
> - Take it or leave it offer
> - Some bargaining over pay

This is question 34D in the survey. Out of 2513 interviewees, there are 1373 workers with valid responses to this question: not all workers were asked the question, and some that were asked refused to respond.

The figure shows that, in the United States, bargaining at the beginning of the job spell is less common among workers in less-skilled occupations. These mirror the results in Section 1.6.3.

[45]There are other possible explanations for this finding. For instance, lower productivity firms could offer more non-wage benefits, as in Sorkin (2018).

[46]The data are available at https://www.acaweb.org/articles%3Fid%3D10.1257/mac.4.4.56. We used a Department of Labor crosswalk to convert the SOC codes in the data to the broad ISCO codes in our Danish data.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

Outside Options in the Labor Market

[1] JOINT WITH OREN DANIELI

2.1 Introduction

In almost every model of the labor market, wages depend on a workers' outside options: the amount of compensation they 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 their skills, could involve different working hours or could be located in a different city. The availability of outside options could be systematically worse for some workers due to 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

[1]This paper benefitted significantly from the advice and feedback of Chiara Farranota, Roland Fryer and Larry Katz. Isaiah Andrews, David Autor, Gary Chamberlain, John Coglianese, Ashley Craig, Ellora Derenoncourt, Ed Glaeser, Nathan Hendren, Simon Jäger, Myrto Kalouptsidi, Robin Lee, Ro'ee Levy, Alan Olivi, Amanda Pallais, Will Rafey, John Roth, Neil Thakral and Dan Svirsky provided helpful comments. This paper uses data from the German Institute for Employment Research (IAB) under the project "Market Power in the Labor Market" (project number 1241/1242). All results based on IAB micro-data have been cleared for disclosure to protect confidentiality. We thank Johannes Schmieder and Daniel Heuermann for generously sharing their data on German train expansions with us.

typically observed. Even within the same firm and occupation, workers may face different options due to 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% representative sample of workers in Germany. Examining the distribution of the OOI, we find what 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, that 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 substantial portions of several widely-discussed wage gaps between different segments of workers. Outside options explain 30% of the gender wage gap in Germany. This gender difference is driven entirely by differences in willingness to commute or move. We also find that differences in outside options account for 88% of the wage gap between German citizens and non-citizens, and about 25% of the high-education wage premium. The availability of more options also increases the wage premium for urban residents. In contrast, differences in outside options reduce inequality between occupations, since high-paying occupations tend to be more specialized and workers in them therefore have fewer options.

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. A direct implication is that workers' compensation is not only determined by what they produce, but also by their ability to produce in more places.

We derive a sufficient statistic from this model, the "outside options index" (OOI), that 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, more industries, etc.). Workers with more relevant jobs, as captured by the OOI, will on average have both a better outside option, and will 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 (?), similar to the Herfindhal-Hirschman Index (HHI), which has also been used to measure concentration in labor markets (Azar et al., 2017; Benmelech et al., 2018). In contrast to most concentration indices, our index is not measured on a specific dimension such as occupation, or industry. Instead, workers with more options are those that are less concentrated across jobs, on all dimensions included in our data set. Options here are estimated in equilibrium, based on matches we actually observe in cross-sectional data. Jobs that the worker will never take in practice because they are less attractive will not enter the OOI nor affect compensation even if the employer is willing to hire. 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 develop a method that estimates the OOI for each worker in the labor market, which is computationally feasible even in large datasets. 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 then 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. Looking across observed workers' characteristics, we find that the OOI is higher for men, German citizens, city residents, more educated and more experienced workers. We also find that higher skill workers such as medical doctors or pilots tend to be more specialized in their current industry, which narrows down their outside options. The OOI also predicts which workers will be 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 quasi-random variation in options, that do not affect productivity, in order to estimate the elasticity between the OOI and wages: the introduction of high-speed commuter rail stations (Heuermann and Schmieder, 2018), and a standard industry shift-share ("Bartik") instrument (Beaudry et al., 2012). These sources of variation allow us to verify that, even if our model is not perfectly specified, there is a link between our outside options index and wages in the data. Our first source of variation focuses on the introduction of new train stations that were constructed along existing routes. These stations effectively increased the labor market size for workers in small German cities that happened to live along the shortest route between two major cities. The second source of variation in outside options 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. We instrument for the growth in outside options in other industries with the national industry trends to exclude the impact of local productivity shocks. Both quasi-random sources of variation yield a similar semi-elasticity of roughly .17-.32 between the OOI and wages.

Combining this elasticity with the estimated distribution of the OOI, we find that differences in outside options tend to increase wage inequality. Differences in options lower compensation for women by six percentage points, explaining roughly thirty percent of the overall gap in Germany. They also account for an eight percentage points difference in compensation between immigrants and natives, which is 88% of the overall gap. We also find large effects on the return to higher-secondary education.[2] Graduates from higher-secondary education have access to more options, which increases their compensation by seven percentage points. This is about a quarter of the total return to higher-secondary education.

Finally, we examine the reasons why workers face different options. We start by examining the parameters that determine workers' options, to understand which ones are most significant. We then use the underlying model to create counterfactual changes to the OOI under different scenarios. These exercises show that the heterogeneity in the ability to commute or move is a key factor in explaining variation in outside options. This factor can account for the full gender gap in outside options. We also find that without their higher willingness to work at more distant jobs, high-educated workers would actually have fewer options. Our analysis suggests that this is likely because their skills tend to be more industry specific.

Related Literature Our paper contributes to at least three distinct literatures. First, we contribute to a large literature on imperfect competition in the labor market by estimating the impact of outside options for every worker in the labor market. While outside options are a key parameter in many labor models, prior work has not focused on estimating the distribution of this parameter across different workers. Most empirical work on imperfect competition has used natural experiments in

[2]The level that grants a certificate allowing college admission.

specific segments of the labor market to show that firms face upward sloping supply curves (see, e.g. Naidu, 2010; Naidu et al., 2016; Ransom and Sims, 2010; Staiger et al., 2010). Beaudry et al. (2012) and Caldwell and Harmon (2018) take a different approach, and provide direct evidence that outside options directly impact workers' earnings, but do not investigate which workers have better options, nor the consequences for between-group wage inequality.[3] Our paper adds to this literature by providing estimates of the distribution of options and by providing descriptive evidence on why some workers have more options than others. By combining this distribution with a causal estimate of the impact of options on wages, we are also able to present the first estimates of the impact of outside options on each individual's wages.

Our theoretical framework textitasizes market imperfections arising from worker and employer heterogeneity. This is similar to the approach taken by Card et al. (2018), and is a standard approach in the industrial organization literature for analyzing market imperfections (for instance, Berry et al., 1995). Recent work by Dube et al. (forthcoming) shows that even small amounts of heterogeneity can generate substantial market imperfections. One difference between our approach and that in the search literature is that we focus on a static equilibrium. While work by Postel-Vinay and Robin (2002) shows that differences in options (as the result of on-the-job search) can impact wage growth during an employment spell, these dynamic considerations are beyond the scope of this paper.

Second, our paper contributes to a small literature on the impact of imperfect labor market competition on between-group wage inequality. Theoretical papers in this literature have argued that some groups such as women or minorities have systematically worse options, enabling their employers to pay them lower wages. These worse options may generate either higher search costs (Black, 1995) or less elastic supply to a particular firm (Robinson, 1933), and can lead to racial or gender wage gaps. Empirical papers in this literature have shown evidence that group differences

[3]Beaudry et al. (2012) show that shocks to one industry "spill over" onto the wages of other industries. Caldwell and Harmon (2018) show that workers with better information about their outside options see greater wage growth. Jäger et al. (2018) focus on a specific outside option– unemployment insurance– and find that changes in UI generosity has little to no effect on workers' wages. Their result fits our finding that what matters for wage-setting is the value of a worker's best alternative to a match. For most workers, this is likely the value of working in another job, not the value of unemployment.

do exist in both labor supply to a firm (Manning, 2003b; Hirsch et al., 2010a; Ransom and Oaxaca, 2010) and in rents (Card et al., 2016a). A key advantage of our setting is that we are able to combine our estimates of group differences in outside options with a causally estimated elasticity between options and wages. This allows us to translate the estimated group differences in options into group differences in wages, and quantify the portion of between-group inequality that can be attributed to imperfections in the labor market.[4]

Finally, our paper contributes to a recent empirical literature on labor market size and concentration, by characterizing workers options using multiple worker and job characteristics at once. Manning and Petrongolo (2017) and Nimczik (2017) develop methods to uncover the size of a workers' labor market based on willingness to commute and on observed firm-firm transitions. Azar et al. (2017) and Benmelech et al. (2018) examine trends in labor market concentrations by calculating Herfindahl-Hirschman indices (HHI's) by occupation/industry, within a geographic area. Hsieh et al. (2013) estimate concentration trends by occupations and demographics such as gender and race using a model similar to ours.

The remainder of the paper proceeds as follows: Section 2.2 outlines the theoretical matching model and derives the Outside Options Index (OOI). Section 2.3 describes the relevant features of the German labor market and the key features of the administrative linked employer-employee data that we use. Section 2.4 explains the empirical procedure of estimating the OOI. Section 2.5 describes the empirical estimates of worker outside options and presents descriptive statistics on their distribution. Section 2.6 estimates the elasticity between the outside options index and wages using two quasi-random sources of variation in options. Section 2.7 analyzes the overall effect on wage inequality. Section 2.8 concludes.

[4] Our setting expands the setting of Bidner and Sand (2016) who quantify the portion of the gender gap that can be attributed to differences in outside options driven solely by differences in access to industries. Our method includes several additional factors, such as differences in commuting costs, that we find to be generating the majority in differences in outside options between genders. We also analyze additional wage gaps beyond gender such as the education, city and citizenship premium.

2.2 A Model of Outside Options and Wages

This section derives a model of a heterogeneous competitive labor market. We use this model to derive the outside options index (OOI), and show it is a sufficient statistic for the impact of outside options on wages. To provide additional intuition for the OOI and its effect in the model, we describe a simple parametric example. We summarize this section by discussing what is and what is not captured in the OOI using the model's assumptions.

2.2.1 Setup

There is a continuum set of workers $\mathcal{I}$ with measure I and a continuum set of one-job firms $\mathcal{J}$ with a measure J which we pin down to 1. If a worker $i \in \mathcal{I}$ works at job $j \in \mathcal{J}$, they produce a value of y_{ij} to the employer and a job-specific amenity valued a_{ij} to the worker. The value of y_{ij} is net of all costs, including capital and amenities. The value for a_{ij} includes all non-pecuniary impacts on worker i's utility including effort, interest, number of vacation days and more. The sum of these two values is the total value of a match, τ_{ij}. This is defined for every potential worker-job pair, even those that are not observed in equilibrium.The value of τ_{ij} is taken as exogenous; all decisions by workers and employers that could affect this value such as investment in capital or human capital and location choices are pre-determined.[5]

Employers and workers decide how to split the total surplus τ_{ij} into worker compensation (ω_{ij}) and employer profits (π_{ij}).

$$\tau_{ij} = \pi_{ij} + \omega_{ij} = y_{ij} + a_{ij}$$

This division is accomplished via a set of transfers (wages) w_{ij}, which allow the worker and employer to divide the total value produced in any way between them:

$$\begin{aligned} \pi_{ij} &= y_{ij} - w_{ij} \\ \omega_{ij} &= a_{ij} + w_{ij} \end{aligned}$$

[5]This is similar to Kreps and Scheinkman (1983) who show how even with competition on prices, pre-determined quantities would deviate from a Bertrand competition.

2.2.2 Equilibrium

We next derive the allocation of workers into jobs and equilibrium wages. We use an equilibrium notion based on cooperative game theory, which is identical to the assignment game, first analyzed by Shapley and Shubik (1971). We assume a static framework with perfect information. There are additional equilibrium concepts that lead to the same result.[6] We use a cooperative framework since it is somewhat more general as it does not make any assumption about how agents reach this equilibrium (e.g. who makes offers).

An allocation is defined as a set $M = \{(i,j)\,|i \in \mathcal{I}, j \in \mathcal{J}\}$ in which no i or j appears twice, so every worker can work only in one job, and every job can hire exactly one worker. For a given allocation M we can define an invertible function on the domain of matched workers $m(i)$ such that $(i, m(i)) \in M$. Note that we do not require all workers and jobs to be in M; some workers can be unemployed and some jobs could be vacant. If a worker is unmatched, she produces u_i, which could be thought of as a combination of unemployment insurance and home production. Similarly, a vacant job produces v_j.

Shapley and Shubik (1971) show that a stable equilibrium (core allocation) includes an allocation M, and a transfer w_{ij} for each $(i,j) \in M$ which satisfies

$$\begin{aligned} \forall i' \in \mathcal{I}, j' \in \mathcal{J} &: \quad \omega_{i'} + \pi_{j'} \geq \tau_{i'j'} \\ \forall i' \in \mathcal{I} &: \quad \omega_{i'} \geq u_{i'} \\ \forall j' \in \mathcal{J} &: \quad \pi_{j'} \geq v_{j'} \end{aligned} \tag{2.1}$$

where $\omega_{i'} = \omega_{i',m(i')}$ if worker is matched and $\omega_{i'} = u_{i'}$ otherwise, and similarly $\pi_{j'} = \pi_{m^{-1}(j'),j'}$ or $v_{j'}$.

The first condition says that there is no single worker-employer combination that could deviate from their current allocation, produce together, and split the surplus in such a way that both the

[6] Pérez-Castrillo and Sotomayor (2002) show one specific mechanism that leads to the same equilibrium using subgame perfect Nash equilibrium.

employer and the worker would be better off. Note that this condition includes all possible combinations, including those that are not matched in equilibrium. The second and third conditions are participation constraints which require that every worker and employer obtain no less than their unemployment or vacancy value.[7] Shapley and Shubik (1971) shows that a stable allocation M^* is also optimal in the sense that the maximum total value is produced.

Workers' compensation in this model depends not only on the value they produce in their workplace, but also on the value they produce in other jobs. Compensation is strictly bounded by the worker and the employer's marginal contributions to the entire market (Roth and Sotomayor, 1992). Because this marginal contribution to the market is weakly smaller than the productivity at the workplace, workers are paid below their full productivity. Workers who are able to produce a similar value in more places will get a larger portion of their productivity to keep the equilibrium stable.

2.2.3 Deriving an Index for Outside Options

We next examine the role of outside options in this equilibrium. In particular, we derive the outside options index (OOI), a sufficient statistic for the impact of outside options on wages.

In any stable equilibrium, each worker must earn more in her current match than she could earn at a different employer.

$$\omega_{ij} \geq \max_{j' \neq j} \omega_{ij'} \tag{2.2}$$

This outside option $\omega_{ij'}$ is exactly what will make employer j' indifferent between their equilibrium match, and hiring i (formally, $\tau_{ij'} - \omega_{ij'} \geq \pi_{j'}$). Hence

$$\omega_{ij'} = \underbrace{\tau_{ij'}}_{\substack{\text{potential} \\ \text{value } i,j'}} - \underbrace{\pi_{j'}}_{\substack{j' \text{ equilibrium} \\ \text{compensation}}} \tag{2.3}$$

[7]Unemployment and vacancies can exist simultaneously, as long as $u_i + v_j \geq \tau_{ij}$ for every possible match of non-participants.

Combined we get a lower-bound for worker compensation[8]

$$\omega_{ij} \geq \max_{j' \neq j} \tau_{ij'} - \pi_{j'} \tag{2.4}$$

The employer decision can thus be written as the solution to a simple profit maximization problem.[9]

From these equations, we can derive an expression for worker's compensation that we can take to the data. First, define $\mathcal{X} \subseteq \mathbb{R}^{d_x}, \mathcal{Z} \subseteq \mathbb{R}^{d_z}$ to be the characteristic spaces of workers and jobs accordingly. Let X_i and Z_j denote the observed worker and job characteristics which are distributed with a density $f(X_i), f(Z_j)$ respectively.[10] We next add an assumption on the distribution of τ_{ij} based on these observables. We follow Dupuy and Galichon (2014) and assume that the value of τ_{ij} conditional on the observables is drawn from a sum of two continuous logit models, one for the workers and one for the employers. This is a generalization of the classic multinomial logit for a continuous case.

Assumption 1. *The match value τ_{ij} between a worker with observable characteristics x_i, and a job with observable characteristics z_j, can be written as*

$$\tau_{ij} = \tau(x_i, z_j) + \varepsilon_{ij}$$

where ε_{ij} has the following distribution

$$\begin{aligned} & \varepsilon_{ij} = \varepsilon_{i,z_j} + \varepsilon_{j,x_i} \\ s.t. \quad & \varepsilon_{i,z_j} \perp \varepsilon_{j,x_i} \\ & \varepsilon_{i,z_j}, \varepsilon_{j,x_i} \sim CL(\alpha) \end{aligned}$$

$CL(\alpha)$ is the continuous logit distribution, that closely resembles an extremum value type-1

[8]Equilibrium compensation ω_i must satisfy $\omega_i + \pi_{j'} \geq \tau_{ij'}$, yielding this equation. This bound will be tight as long as $\max_{j' \neq j} \tau_{ij'} - \pi_{j'} \geq u_i$. It holds with equality under an additional assumption (Assumption 1).

[9]Equation 2.2 defines the effective price that employer j needs to pay to hire worker i. In order to maximize profit, the employer needs to choose a worker that maximizes the value net of cost: $\max_{i'} \pi_{i'j} = \max_{i'} \tau_{i'j} - \omega_{i'j}$.

[10]Formally, there are measurable functions $X_i : \mathcal{I} \to \mathcal{X}, Z_j : \mathcal{J} \to \mathcal{Z}$.

distribution with scale α (Dagsvik, 1994). For details on this distribution, see Appendix 2.11.[11]

This assumption simplifies the math considerably. However, it is strong; it implies that workers have an unobserved utility or productivity in jobs with specific observed characteristics, and those unobserved shocks are uncorrelated, even between jobs with similar characteristics. Employers also have similar unobserved independent shocks based on the workers observables. Moreover, the assumption that $\varepsilon_{i,z_j} \perp \varepsilon_{j,x_i}$ implies that there are no interactions between the worker and job unobserved characteristics.

We can rewrite the latent value of outside options from Equation 2.3 as

$$\omega_{ij'} = \tau\left(x_i, z_{j'}\right) - \pi_{j'} + \varepsilon_{ij'} \tag{2.5}$$

Using $*$ to denote the best alternative offer ($\omega^*_{ij'} = \tau^*\left(x_i, z_{j'}\right) - \pi^*_{j'} + \varepsilon^*_{ij'} = \max_{j'} \omega_{ij'}$), we get a simple expression for the expected value of the best alternative offer:

$$\underbrace{E\left[\omega^*_{ij'}\right]}_{\text{Best Alternative Offer}} = \underbrace{E\left[\tau^*\left(x_i, z_{j'}\right)\right]}_{\text{Mean Value}} - \underbrace{E\left[\pi_{j'}\right]}_{\text{Employer Rents}} + \underbrace{E\left[\varepsilon^*_{ij'}\right]}_{2\alpha \cdot OOI} \tag{2.6}$$

This decomposition is the key result of our theoretical analysis.

The first component reflects the mean value the worker can produce where they typically work (without strategic sorting on ε[12]). Therefore, as in almost all labor models, workers that produce a higher value would earn a larger compensation. The second component reflects the mean employer profit, beyond costs. This could be zero, or constant if we think the employer market clears perfectly through entry, but we do not assume this is necessarily the case. This component is affected by several factors including the firm productivity, and the market price of their workers.

In this paper, we focus on the third component $E\left[\varepsilon^*_{ij'}\right]$. This expression depends on the measure of relevant options the worker has. Even though for a random match, $E\left[\varepsilon_{ij}\right]$ is constant, the

[11] Formally, every worker i draws ε_{i,z_j} shocks from a Poisson process on $\mathcal{Z} \times \mathbb{R}$. As a result, for every subset $S \subseteq \mathcal{Z}$, $\max_{z \in S}\{\varepsilon_{i,z}\} \sim EV_1\left(\alpha \log P\left(S\right) + const, \alpha\right)$ where EV_1 is extremum-value type-1 distribution, and $P\left(S\right) = \int_S f\left(z\right) dz$. A similar process exists for ε_{j,x_i}. More details in Appendix 2.11.

[12] $\tau\left(x_i, z_j\right)$ reflects the expected value a worker can produce in a random job with characteristics z_j.

expectation of $\varepsilon^*_{ij'}$ is higher because j' is positively selected, since it is the second *best* option. The more similar options in expectation worker i will have, the larger this component will be.

We derive the Outside Option Index, OOI, directly from this expectation. Specifically, we define the OOI to be the standardized expectation $\frac{1}{2\alpha}E\left[\varepsilon^*_{ij'}\right]$, where $\alpha \geq 0$ is the scale parameter that depends on the distribution of ε_{ij}. This descaling guarantees that the value of OOI is independent of the standard deviation of ε, and therefore of the units in which we define $\tau(x, z)$. α also sets the link between the OOI and wages, which we'll estimate using two distinct quasi-random sources of variation in options in Section 2.6. We assume that α is constant across all workers, implying a constant elasticity between OOI and wages. Our results on the heterogeneous impact of options on wages in Section 2.6 are consistent with this assumption.

The standard result that workers earn what they produce is a particular case of this setting. This occurs when $\alpha = 0$ (ε_{ij} is constant at 0) and entry decision of employers are optimal, such that profit is zero. This textitasizes the key distinction of this more general setting from the perfectly competitive model: heterogeneity. When $\alpha > 0$, there is no identical employer to bid wages up to the worker's full product.

Under Assumption 1, workers and employers are indifferent between matches with the same characteristics. Formally, defining f^i_j the probability density of worker i to work at job j in equilibrium we get the following lemma:

Lemma 4. *Under Assumption 1, the probability density of worker i to work at job j satisfies*

$$f^i_j = \frac{f(X_i, Z_j)}{f(X_i)\, f(Z_j)}$$

where $f(X_i, Z_j)$ is the joint density of matched worker and job observables in equilibrium.

Intuitively, this lemma implies that f^i_j is equal for all jobs with the same characteristics. Appendix 2.11 provides a full proof, as well as formal definitions for those densities.

This assumption also yields a closed-form expression for the outside options defined in Equation 2.5.

Lemma 5. *Under Assumptions 1, in equilibrium, worker i with characteristics x_i is facing a continuous logit choice between employers who are offering*

$$\max_{z_j} \omega\left(x_i, z_j\right) + \varepsilon_{i,z_j}$$

and

$$\omega\left(x_i, z_j\right) = \tau\left(x_i, z_j\right) - \pi\left(z_j\right) - \alpha \log f^i_{j'}$$

where $\pi\left(z_j\right) = E\left[\pi_{j'} | Z_j = z_j\right]$. Similarly, employers choose between

$$\max_{x_i} \pi\left(x_i, z_j\right) + \varepsilon_{j,x_i}$$

This lemma simplifies the matching procedure into two one-sided continuous logit choices. Because employers with the same characteristic z_j are willing to make the same offer, the best alternative offer ω^* equals the maximal offer the equilibrium employer is willing to make. Hence, the lower bound from Equation 2.4 can be replaced with an equality. The workers are facing a choice between the employers who are looking for workers with their observed characteristics x_i.

The market clears when the supply of workers with characteristics x_0 to jobs with characteristics z_0 equals demand. Demand is decreasing with quantity, because the marginal employer has a lower value of ε_{j,x_0}. This is why the compensation $\omega_{ij'}$ in lemma 5 depends negatively on $\alpha \log f^i_j$. However, supply increases with compensation, hence f^i_j will be increasing in $\tau\left(x_i, z_z\right) - \pi\left(z_j\right) - \alpha \log f^i_j$. Those two equalize exactly when

$$f^i_j \propto \exp \frac{1}{2\alpha} \left[\tau\left(x_i, z_j\right) - \pi\left(z_j\right)\right]$$

This result implies that we can learn about the quality of outside options based on how similar workers sort into different jobs. Workers tend to sort into jobs where their net productivity is highest, where net productivity is the difference between the mean value they can produce $\tau\left(x_i, z_j\right)$, minus the employer's expected profits $\pi\left(z_j\right)$. Therefore, the only valuable outside options for a worker are those that are taken, in equilibrium, by similar workers.

With this distributional assumption we can find an analytical expression for the OOI. The expected value of ε^*_{ij} simplifies with the following lemma

Lemma 6. *Under Assumption 1:*

$$E\left[\varepsilon^*_{ij}\right] = E\left[\varepsilon^*_{i,z_j} + \varepsilon^*_{j,x_i}\right] = -2\alpha \int f^i_j \log f^i_j \tag{2.7}$$

The last equality follows because both $\varepsilon_{i,z}, \varepsilon_{j,x}$ are drawn from a continuous logit with scale parameter α. To measure the OOI for worker i, we need to take the integral over all their potential matches. Using the definition in Equation 2.6 yields

$$OOI = \frac{1}{2\alpha} E\left[\varepsilon^*_{ij}\right] = -\int_j f^i_j \log f^i_j \tag{2.8}$$

This expression is the well-known entropy index. Entropy is frequently used to measure industry concentration. Analogously, the OOI can be thought as a concentration index across jobs. A worker with more options (a worker who is less concentrated) will have a higher OOI because their probability of being in a specific job is lower. This is concentration on all (observable) dimensions: location, occupation, industry etc. Empirically, we will estimate it based on the concentration of workers with similar observables. If similar workers tend to be concentrated in a specific region of the country, small number of occupations or industries we will estimate a lower OOI for them. We will describe this procedure in detail in Section 2.4.

The entropy index is also commonly used in measuring unpredictability. In our context, this would be the difficulty to predict the worker's job. Workers whose jobs are harder to predict, are those with more options.[13] The OOI takes values on $(-\infty, 0]$. As the measure of jobs a worker can take approaches zero, $OOI \to -\infty$; if a worker is equally likely to take any job, $OOI = 0$.

[13] It is possible that it's easier to predict the job of certain workers due to better data quality. This would imply that those workers will have a lower scaling parameter α, and therefore a lower elasticity between the OOI and wages. To test this, we estimate the OOI-wage elasticity in Section 2.6 separately by gender and education. Our results are consistent with a constant value of α for all workers.

This OOI is driven by two factors. First is worker flexibility, the ability of the worker to take jobs at different locations, use their skills in different occupations, industries etc. All of which we will measure empirically. Second is the supply of relevant jobs. More relevant jobs that the worker can take will increase the OOI directly. The OOI is only driven by relevant outside options – jobs that similar workers are actually observed taking in equilibrium ($f_j^i > 0$). Empirically, this will be options that are actually sometimes executed by workers with similar observables. Therefore, jobs that a worker could do but never would do in practice will not enter the OOI, and won't affect the equilibrium outcome.

The key advantage of using the OOI is that it does not depend on any information on a worker's alternative wages. This is useful because information on potential wages at other jobs is typically unavailable. Moreover, a worker's alternative wages also depend directly on their productivity. The OOI captures the impact of options on wages, holding productivity constant. This also implies that any link that we find between the OOI and wages is not mechanical.

2.2.4 Sufficient Statistic

The OOI is a sufficient statistic for the effect of access to more options on workers compensation, under our model assumptions. Access to options has two distinct effects on workers, both of which are captured in the OOI. First, it improves workers compensation at the same job by improving their outside options. Second, the improvement in options allows some workers to find better matches.

We first define an improvement in access to options. We define λ_x to be the measure of a random set of jobs that are accessible to workers with observables x. All jobs that are not accessible have $\tau_{ij} = -\infty$ and are therefore never chosen in equilibrium. We model an increase in access to more jobs would be an increase to this λ_x. In Appendix 2.11 we show that other definition of λ such as a linear commuting cost would yield the same results.

Theorem 7 shows that workers who get access to more outside options get an increased wage offer from their employer that equals to α times the change in their OOI.

Theorem 7. *Let j be i's equilibrium match. Access to outside options λ_{x_i} has the following effect* *on the maximum offer j is willing to make in the new equilibrium:*

$$\frac{d\omega_{i,j}}{d\lambda_{x_i}} = \alpha \frac{dOOI}{d\lambda_{x_i}}$$

The second effect of access to more options is an improvement in match quality. An improvement in outside options is only an improvement, if some workers would in practice match into those additional jobs in equilibrium. Therefore, the overall effect of access to more options is a combination of the better outside options, and the option to improve match quality. The following theorem shows that in this model, the overall effect is exactly twice the size of the effect only through outside options.

Theorem 8. *Access to options λ_{x_i} has the following overall effect on expected worker compensation in equilibrium*

$$\frac{dE\left[\omega_{i,j}\right]}{d\lambda_{x_i}} = 2\alpha \frac{dOOI}{d\lambda_{x_i}}$$

Different choices of counterfactuals could potentially lead to different results. The counterfactual we consider is giving a small group of workers access to more similar jobs. If the increase in λ affects a non-zero measure of workers, then there will be general equilibrium impacts on employer profits . For instance, mandating stable working hours in all jobs will give all women access to more jobs. This counterfactual may decrease the profits of employers that were already hiring mostly women. In such cases, the OOI would only be a sufficient statistic if the employers' market is perfectly competitive such that profit are kept constant through entry and exit. Access to better jobs (as opposed to similar jobs), in which the worker can produce greater value will also affect workers productivity, and therefore will affect compensation beyond the effect on the OOI.

2.2.5 Parametric Example

To give further intuition for the OOI, and the additional components in our key decomposition (Equation 2.6) we go over a simple parametric example.

In this simple setting, workers are characterized only by their productivity and their amount of options. Assume workers and jobs are equally dispersed across the real line $\mathbb{R}$.[14] Each worker can be described as a 3-dimensional tuple (l_i, y_i, d_i) which is her location on the real line, her productivity and the maximal distance she is able to commute. Jobs are identical other than their location l_j. The value of a match is then

$$\tau_{ij} = \begin{cases} y_i + \varepsilon_{ij} & |l_i - l_j| < d_i \\ -\infty & else \end{cases}$$

where ε_{ij} are the sum of two continuous logit distribution as before.

In this simple setting, the OOI corresponds to the log measure of options. The PDF of a worker distribution across jobs is constant at $\frac{1}{2d_i}$ for all jobs within feasible range. Therefore, the OOI is $-\log \frac{1}{2d_i} = \log 2d_i$ which is the log of the measure of jobs a worker can take. Differences in OOI are therefore the log ratio in the measure of relevant options. This result will generally hold for every pair of workers with similar distribution of jobs and different sizes of support, not only in this example. In this setting, λ is exactly $2d$, hence from Theorems 7 and 8, an infinitesimal increase in d_i leads to an increase of $\frac{\alpha}{2d_i}$ if they stay at the same job, and $\frac{\alpha}{d_i}$ overall.

The first component of Equation 2.6 (mean value) captures a worker's baseline productivity; in this case this is equal to y_i. This component represents the expected productivity in a random job that a worker could take. Equivalently, it captures productivity differences, conditional on having the same amount of options (OOI). The final component, employer rents, will be equal for all workers, as all jobs are equivalent.[15]

This example shows clearly how two workers who are on average equally productive, could still earn different wages due to differences in outside options. Assume $l_1 = l_2$, $y_1 = y_2$, and $d_1 < d_2$. Worker 2 earns a higher wage because her OOI is greater. In expectation, workers 1 and 2 are equally productive at every job in $[l - d_1, l + d_1]$. Since worker 2 has a higher price, most jobs in this range would prefer to hire worker 1. Still, as a result of heterogeneity, some employers

[14] Formally, assume each interval $[a, b]$ has a measure of $b - a$ workers and jobs. This implies an infinite measure of both workers and jobs.

[15] Its exact value would be pinned down depending on the value of unemployment, and vacant jobs.

would be willing to pay the higher price. Because worker 2 has more options than worker 1, there are enough employers who are willing to pay the higher price, so that the market clears.

2.2.6 Discussion

We summarize this section by re-examining the model assumptions and their implication on what is and what is not captured with the OOI. The primary advantage of the OOI is that it more precisely captures the size of a worker's *relevant* option set. It allows workers to use their skills in different occupations and industries. By contrast, measures such as the HHI assume that workers belong to only one industry or occupation. Similarly, the OOI accounts for heterogeneity in commuting and moving costs. Instead of assuming each worker is assigned to a specific local labor market, the OOI empirically assess the distance over which each type of worker searches for a job. Finally, the OOI accounts for variation in employer characteristics even within the same industry. For instance, if some workers are unable to work on weekends, their OOI will only be affected by employers who do not require that.

The main limitation of the OOI is that it does not account for any dynamic considerations. This is because it was derived from a static model. Dynamic considerations such as switching costs, firm-specific human capital that is acquired over time, and learning tend to limit a worker's ability to move to their outside options, but are beyond the scope of this analysis.

A second limitation is that the OOI calculates the measure of relevant jobs, not relevant employers. We assumed that employers are 1-job firms and do not account for the fact that many jobs are under the same employer. While the model will, with minor adjustments, accommodate firms, we focused on jobs due to limitations of our data (see Section 2.3.1). Therefore, the OOI will over-estimate options for workers who are more likely to work in large firms.

In contrast, some aspects of the labor market that are not explicitly modeled above could still be captured in the OOI. The most prominent one is information frictions that would generate search costs. Black (1995) has analyzed a search model where some workers have more options, and showed that in this setting as well, more options would lead to higher wages in equilibrium. Hence, it is possible that some of the effect of the OOI on wages is operating through this channel as well.

2.3 Empirical Setting and Data

We use administrative data from Germany to generate measures of individual workers' outside options. The data includes detailed information on establishment and worker characteristics, including information on a variety of amenities provided by different establishments, which allow us to estimate workers' options more accurately. Excluding some idiosyncratic features which we will now discuss, the German labor market is comparable to other low-regulated labor markets, making wages more directly affected by the market forces we want to study.

2.3.1 Data

Administrative German Employer-Employee Data

Our primary source of data is a panel of German worker employment histories known as the "LIAB Longitudinal" dataset. It is a matched employer-employee administrative data, based on a sample from the universe of German Social Security records from 1993-2014. There are four key features of the data which make it ideal for our setting. First, it is a large dataset, including about 1% representative sample of the entire German labor force. Second, there is detailed establishment-level survey information with information such as hours requirements, profitability, leave/maternity policies etc. This allows us to account for differences in outside options that may be due to differences between establishments, even within industries. Third, the panel structure of the data, allows us to track workers over long periods of time. This gives us valuable information about the workers such as their specific experience in the market, and their location before taking their job. Fourth, this data provides 4-digit occupational classification which highly improves our precision in measuring relevant options. To our knowledge, this combination of data is not available in the United States.

The data come from the Integrated Employment Biographies (IEB) dataset, which is collected by the German Institute for Employment Research (IAB). Employers are required to report daily earnings[16] (subject to a censoring limit at the maximum taxable earning level)[17], education, occupation, and demographics for each of their employees at least once per year, and at the beginning of any new employment spell. New spells can arise due to changes in job status (e.g. part-time to full-time), establishment, or occupation.

Each year the IAB selects a stratified random sample of establishments from the pool of all German establishments with at least one employee liable to Social Security. These establishments are required to complete a series of surveys on organizational structure, personnel policies, financing, and research activities. In particular, the establishments are asked for information on their annual sales, profits, establishment size and leave policies. The survey data are then merged with the complete employment histories of all individuals who worked at least one day in any of these firms between 1993 and 2014.

There are several limitations for the data, that may affect our calculations of outside options. The data do not cover civil servants or the self-employed, which comprise 18% of the German workforce. They also do not cover labor force non-participants. Therefore, we do not account for any of those options when calculating the OOI. Since the sample is done at the establishment level, we usually observe only few establishments in each industry-region combination. This is why we construct the OOI at the job level, and not the employer level.

Because our model is static, we rely on repeated cross-sections of data. For each year, we use data on employment relations on June 30th of each year. Our descriptive analysis is done for our last year in the sample, that is June 30th 2014. We use data from 1999, 2004 and 2012 in Section 2.6 to examine how quasi-random variations in the OOI effect wages.

BIBB Task Data

We supplement these data with survey information on the characteristics of occupations and industries. It includes information on the tasks completed, hours requirements and typical working conditions in these occupations/industries. These data are similar to the O*NET series, but allow

[16]Daily earnings are calculated as an average for the reported period.

[17]11% of the sample is censored. As we do not use wages to calculate OOI, it is not affected by censoring.

us to account for possible differences in the task content of occupations between the United States and Germany, as well as differences in coding.[18] The survey is conducted by the IAB and includes information on respondents' occupation, industry, in addition to responses on questions related to organizational information, job tasks, job skill requirements, health and working conditions.

2.3.2 Empirical Setting: German Labor Market

There are several distinctive features of the German labor market which are relevant for our analysis. First, there are different levels of secondary-school leaving certificates, which depend on the number of years and type of education. Our data allows us to distinguish between three categories: lower-secondary, which typically requires nine years of schooling, intermediate-secondary, which typically requires ten years of schooling, and high-secondary, which requires twelve to thirteen years of schooling, and allows the student to pursue a university degree. In our analysis we use indicators for the type of secondary education to account for years of schooling, and school quality.

Second, in addition to (or sometimes instead of) formal education, many German workers receive on-the-job training through formal apprenticeships. Individuals in apprenticeship programs complete a prescribed curriculum and obtain occupation-specific certifications (e.g. piano maker). We use this information to precisely identify the types of jobs a worker could perform.

Third, eleven percent of workers in Germany work under "fixed-term contracts" (as of 2014). These contracts expire automatically without dismissal at the end of the agreed term, at no cost to the employer. The maximal period for employment under such contract varies between 6 to 18 months over the period for which we have data. At the end of a contract, the worker and employer may choose to continue the employment relationship, but cannot use another fixed term contract to do so (Hagen, 2003).

Fourth, two percent of workers are hired through temporary work agencies. This is a triangular employment relationship, which involves the temporary work agency, a client company and a temporary worker. Historically these working relations were limited to 24 months; their duration is

[18]These data have been used in prior publications on the German task structure including Gathmann and Schönberg (2010).

no longer regulated. There are additional regulations on the pay received by workers hired through temporary agencies (in particular relating to how these workers are paid relative to other workers at the same firms) but the rules vary significantly over time. In our analysis, we distinguish between employment found via temporary work agency and work found via more traditional means (Mitlacher, 2008).

While wage setting in Germany was historically governed by strong collective bargaining agreements, employers today have considerable latitude in setting pay (Dustmann et al., 2009). While employers could always raise wages above the agreed-upon levels, it only became common for contracts to include "opening clauses" allowing employers to negotiate directly with workers to pay below-CBA wages in the 1990s. Today these clauses are very common.

2.3.3 Summary Statistics

Table 2.1 describe the characteristics of workers and jobs in our sample, for the full sample, as well as by gender.

Our sample is roughly evenly split between male and female workers. The mean age for a worker in our sample is forty-five years old and the vast majority (97%) are citizens. The workers are divided about equally between the three types of secondary education. In nineteen percent of the sample the lower- and intermediate- secondary education categories are aggregated.[19] Men and women have similar age, education and citizenship status.

On the job side, thirty one percent of the jobs in our sample are part-time. Eleven percent of jobs are on fixed contracts and only two percent are from temporary agencies. The distribution of establishment size is very skewed, with mean of 1,552 workers and a standard deviation of five times that size. The mean annual sales per worker are 163,000 Euros. Twenty-six percent of the establishments report to have females in managerial positions.

It can already be observed that men and women sort into different types of jobs. Females are much more likely to work in part-time jobs (53% compared to 13%), which is relatively high

[19] More details in data Appendix 2.12

compared to other countries.[20] They also work at smaller establishments with 827 employees on average and mean annual sales of 130,000 Euros, compared to 2,166 and 191,000 accordingly for males. Females are also more concentrated at establishments with higher share of female-management (36% compared to 17%).

2.4 Estimating Outside Options

In this section we describe how we estimate the outside options index. Our method uses the cross-sectional allocation of observably similar workers to estimate the relevant options of each worker . This allocation teaches us about the worker's ability or willingness to commute, about the set of industries or occupations that are suitable to the worker's skills, and about the worker's demand for certain workplace amenities. Section 2.4.1 states the key assumption, Section 2.4.2 describes the estimation procedure, and Section 2.4.3 describes the worker and job characteristic we use as inputs.

The OOI of a given worker requires an estimate of their probability to work in each one of the jobs observed in the data. We calculate the OOI using Equation 2.8 which shows that the OOI is only a function of the different f_j^i .This requires us to estimate N^2 distinct probabilities.

Earlier methods that were developed to estimate such densities do not work on data sets of our size. Non-parametric approaches cannot be used due to the large number of worker and firm characteristics. Parametric methods that were designed specifically for this model, work well when the number of possible combinations is around a few millions.[21] However, given the size of our data, we need to calculate the probability density of approximately 250 billion possible combinations, making these methods computationally not feasible.

To overcome this challenge, we develop a new method that is computationally feasible for large data sets, and uses a similar set of assumptions to those used in the prior literature (Dupuy and

[20] Germany is ranked 6th out of 35 OECD countries in female part-time employment (OECD, 2004).

[21] Choo and Siow (2006) develop a non-parametric method where the number of possible combinations is finite and small. Dupuy and Galichon (2014) use Iterative-Proportional-Fitting algorithm to estimate a continuous density. Their data set had about $N^2 = 10^6$ possible combinations.

Galichon, 2014). Our method relies on an equivalent representation of the probability densities as the ratio between the likelihood of a matched pair to appear in the equilibrium allocation compared to a random one. These ratios can be estimated quickly using logistic regressions. We discuss the links and differences between our method and prior methods in more detail in Appendix 2.11.

2.4.1 Assumptions

In this section we state the parametric assumptions we make to link the f_j^i densities to the data. Our data are comprised of pairs of matches between workers, and jobs (x_k, z_k), where the x_k/z_k are observed worker/job characteristics we discuss in the Section 2.4.3.

We first use the result of Lemma 4

$$f_j^i = \frac{f(X_i, Z_j)}{f(X_i)\, f(Z_j)}$$

$f(X_i, Z_j)$ is the probability of observing a match between a worker with characteristics X_i and a job with characteristics Z_j. $f(X_i)f(Z_j)$ is the product of two the marginal distributions for workers and job characteristics. This is the probability of observing a match with such observables, under a random assignment. The basic intuition for this result is that the probability of observing i matched with j depends on the frequency that workers and jobs with such observables are matched, accounting for the total measure of workers and jobs with these observables (if there are more jobs with a particular set of observables, the probability to match to a specific one is smaller). This result can be derived from weaker assumptions as well.[22]

Our second assumption parametrizes f_j^i as a function of the observables. We follow Dupuy and Galichon (2014) in assuming that the log density is linear in the interaction of worker and job characteristics.

Assumption 2. *The log of the probability density is linear in the interaction of every worker and*

[22] It is sufficient to assume If $X_i = X_{i'}$ and $Z_j = Z_{j'}$ then $f_j^i = f_{j'}^{i'}$, instead of Assumtion 1.

job characteristic:

$$\log f_j^i = X_i A Z_j + a\left(X_i\right) + b\left(Z_j\right)$$

The matrix A includes all the coefficients on each of the interactions between worker and job characteristics. The marginal distributions $f\left(x\right), f\left(z\right)$ are fully determined by $a\left(x\right)$ and $b\left(z\right)$.[23]

This assumption reduces the dimension of the problem significantly, while allowing the relationship between each pair of covariates to remain unrestricted. Dupuy and Galichon (2014) show that A is proportional to the cross-derivative of τ

$$2\alpha A = \frac{\partial^2 \tau}{\partial x \partial z} \tag{2.9}$$

where α is the scale parameter of ε we defined in Assumption 1. Intuitively, this means that if a worker characteristic and a job characteristic are complements, they will be observed more frequently in the data.

2.4.2 Empirical Procedure

Under these two assumptions, we can estimate the OOI using a simple procedure that we will now describe. The key idea of our method is to use the result of Lemma 4, that the probability density f_j^i can be written as the ratio between the probability of observing a match in the real distribution to its probability under a random assignment.

We start by expanding our data set of worker and job matches. We simulate data from a distribution $\widetilde{f}\left(x, z\right) = f\left(x\right) \cdot f\left(z\right)$, where x and z are independent. This is done by randomly sampling an observed worker and an observed job independently. We simulate a total number of random matches equal to our original data size, such that the share of real and simulated data is exactly one half. We define a binary variable Y that equals to one whenever the match is 'real' (taken from the

[23]While Dupuy Galichon are able to fit the marginal distribution precisely to their observed value in the data, we won't be able to this with our data size. Therefore, we take linear functions of all X variables and Z variables. We also include indicators for district. As we discuss in Appendix 2.11 this specification fits the first moments of the marginal distributions.

data) and zero whenever it is simulated.

We then estimate all our parameters using a logistic regression. We regress the binary variable we constructed Y_k on the matched worker and job characteristics (X_k, Z_k). Note that, as a result of Lemma 4, and a simple Bayes rule, the match probability density f_j^i is proportional to the ratio of observing this match in the real or simulated data, conditional on the observed worker and job characteristics.

$$\frac{P(Y_k = 1|x_k, z_k)}{P(Y_k = 0|x_k, z_k)} = \frac{f(x_k, z_k)}{f(x_k) f(z_k)} \frac{P(Y_k = 1)}{P(Y_k = 0)} = f_j^i \cdot const$$

Combining this result with Assumption 2 yields

$$\log \frac{P(Y_k = 1|x_k, z_k)}{P(Y_k = 0|x_k, z_k)} = x_k A z_k + a(x_k) + b(z_k) \tag{2.10}$$

We can estimate this equation using a logistic regression where we approximate $a(x), b(z)$ with linear functions. Under the assumptions this produces consistent estimates for $\widehat{A}, \widehat{a}(x_k), \widehat{b}(z_k)$. We discuss the intermediate results from this estimation procedure, in Section 2.7.2 where we analyze the underlying reasons for differences in the OOI.

We use the estimates from the logistic regression to estimate the probability density of every potential match. Specifically, we estimate the probability density of worker i to work in job j to be

$$\widehat{f_j^i} = \exp\left[x_i \widehat{A} z_j + \widehat{a}(x_i) + \widehat{b}(z_j)\right]$$

We calculate this value for all possible worker-job combination in our data set.[24] This simple functional form allows us to make this calculation directly and with minimum computational burden.

With these result in hand we can calculate the outside options index for every worker in our sample using Equation 2.8:

$$\widehat{OOI}_i = -\sum_j \widehat{f_j^i} \log \widehat{f_j^i} \tag{2.11}$$

[24] We normalize those estimated densities, such that $\sum_j \widehat{f_j^i} = 1$. So effectively, we don't use $\widehat{a}(x_i)$.

This yields a consistent estimate of the OOI, if both assumptions are correctly specified.

We verify that the OOI is robust to different choices of functional form. Instead of estimating it using the entropy index, we use the same probabilities we estimated in an HHI formula: $-\sum_j \widehat{f_j^i}^2$. We find that the results are very similar. The correlation between the two indices is .62.

In Appendix 2.11 we discuss the properties of this method, in the case where these assumptions do not hold. We show this method can be written as a GMM estimator, and discuss the moments that are being matched. We also show that if we increase the size of the simulated data, and fully saturate the functions a and b, our method becomes equivalent to Dupuy and Galichon (2014).

2.4.3 OOI Input: Job and Worker Characteristics

To estimate f_j^i using Equation 2.10 we include three groups of variables: worker characteristics (x), job characteristics (z), and the geographical distance between workers and jobs. Information on wages is intentionally not used in any of these groups, to avoid a mechanical link between the OOI and wages.

Worker Characteristics x We use x to denote the variables that describe worker demographics and worker training. The demographic variables include workers' gender, worker's level of secondary education, an indicator for whether the worker is a citizen, and a quadratic in age. For training we use the occupation in which they undertook their apprenticeship. If we do not have information on a worker's apprenticeship (e.g. if it occurred before our data begin in 1993), or if a worker did not complete an apprenticeship, we use their first occupation observed in the data, as long as this is at least ten years old.

Job Characteristics z The job characteristics z variables fall into three categories: (1) characteristics of establishments, (2) characteristics of employment contracts, and (3) characteristics of jobs. First, we take several establishment-specific variables directly from the establishment survey: size, sales and the share of females in management . We also use the first two principal components of each of the six categories of the establishment survey: business performance, investments, working hours, firm training, vocational training, and a general category. Appendix Table 2.9 shows the most weighted questions in each category.

Second, we use several variables which relate to the structure of the employment contract: whether the job is part-time, whether the contract is fixed term, and whether the position was filled by a temporary agency.

Finally, to describe the characteristic of the job we use information on the occupation and industry. Because it would not be feasible to include interactions between all of our industry and occupation codes, we use data from the BIBB to identify the characteristics associated with different industries and occupations. The BIBB survey contains modules on working hours, task type, requirements, physical conditions and mental conditions. For each 3-digit occupation and 2-digit industry, we include the first two principal components for each module. We use these to code both the occupation and industry that describe the job, and the training occupation that describes the worker. Appendix Table 2.10 shows the most weighted questions in each module. We also include occupation complexity, which codes occupations into four categories based on the type of activity they require: (1) simple , (2) technical (3) specialist and (4) complex.[25] We use a total of 18 worker characteristics and 39 job characteristics.

Geographical Distance We include the geographical distance between workers and employers. For workers we use their last place of residence before taking the job.[26] This distance could capture both the commuting, as well as the moving costs between places; empirically we cannot directly distinguish the two. Both locations are given at the district (kreis) level.

Figure 2-1 presents a map of the 402 districts in Germany. The size of the districts varies across the country and, importantly, it tends to be smaller in highly populated areas. In many cases, the major city is its own district, allowing us to separately identify the city center and the suburbs.

[25]These four categories usually reflect the type of qualification needed to perform the job, which ranges between none, vocational training, some tertiary degree and higher education. For instance, different occupations in nursing that fall under the same occupational coding (813) will be coded with different complexity, ranging between a nursing assistant, nurse, specialist nurse and general practitioner.

[26]Workers current place of residence is affected by their match. Their place of residence before taking the job better reflects the actual radius over which people are searching for jobs.

Though not perfect, this coding allows us to get a reasonable approximation of commuting and moving patterns by workers. Appendix Table 2.11 shows the mean of the distance variable by gender and education groups. We find that the mean is 15.5 miles, but there is significant variation across groups.

We allow distance to have a non-linear effect on match probability that is different for each worker type. When we estimate Equation 2.10 we use a 4th degree polynomial of the distance between a worker's lagged home district and their location of work to account for the non-linear impact of distance.[27] To account for heterogeneity in willingness to commute or move, we interact the polynomials in distance with all worker characteristics x. This allows workers to be affected differently by distance, depending on their gender, education, age, citizenship and training. As we discuss in Section 2.7.2 this turns out to be the main driver of differences in outside options.

2.5 The Empirical Distribution of Outside Options

We next turn to describing the distribution of the OOI, and the characteristics of workers with better and worse options, as measured by it. We find that the OOI is higher for men, German citizens, city residents, more educated and more experienced workers. We also find that higher skill workers tend to be more specialized in their current industry, which narrows down their outside options.

Figure 2-2 plots the raw distribution of the OOI for every worker in our data. The mean of the distribution is -4.85. We can interpret the mean by considering the share p of options a worker with this OOI would have if the probability density they worked at any given job was either $\frac{1}{p}$ or 0. A worker with an OOI of -4.85 would be found in a share $p = 0.8\%$ of jobs. The distribution is skewed, with a long left tail, indicating that there are many workers who are extremely concentrated. The standard deviation of the distribution is also quite sizable: .93. For comparison, duplicating the worker's option set by generating an additional identical job for every job option they have would increase the OOI by only .69.

[27] We find that for more than 100 miles, the effect of distance is constant. This is consistent with the idea that individuals do not commute more than 100 miles; to switch to a job that is much further away, they have to move. This moving cost may not vary significantly with distance.

We estimate the following regression to decompose the average OOI by worker characteristics

$$OOI = \beta_0 Female + \beta_1 Education + \beta_2 Citizen + \beta_3 Age + \beta_4 Age^2 + \epsilon \tag{2.12}$$

Figure 2-3 plots the results. With controls, the average OOI for women is .237 units below that of men. The average OOI for German citizen is higher by .217 units. Assuming similar distributions across jobs, this would imply that male (German citizens) have 27% (24%) more options than women (non-citizens).

Options are also better for higher-educated workers. Lower-secondary (intermediate secondary) school workers' options are on average .70 (.32) units lower than higher-secondary workers. This implies 101% (38%) more options assuming similar distribution across jobs.

Figure 2-4 shows that men have more options than women, not just on average, but across the entire distribution. The figure shows that the cumulative distribution function for men is shifted to the right. We cannot reject that the distribution for men stochastically dominates that of women.

We find an inverse U-shape relationship between the OOI and age. Figure 2-5 plots the mean OOI by age and shows that workers' options tend to improve with age before flattening off at age thirty. Older workers (over 50) see declining values of options. As we discuss in Section 2.2.6, the OOI does not capture any dynamic considerations that are particularly likely to have a differential effect across ages. Accounting for this could potentially change these results.

A large portion of the variation in options is driven by geographical variation in labor market size and density. The last category in Figure 2-3 shows a positive correlation between the district density and its OOI, controlling for other demographics. Figure 2-6 graphs the mean value of the OOI by German district. As the figure illustrates, workers in cities tend to have better options, as measured by the OOI. Workers near these cities, also appear to have better options. This result is robust for adding controls for worker demographics.

While most of our results indicate that high earning workers, such as workers with higher education or city residents, tend to have more options, this relationship is reversed at the occupation level. This is because high-skilled workers tend to have more specialized skills, which are valued by a smaller number of employers. Controlling for all other observables, workers who completed their training (apprenticeship) in higher earning occupations tend to have lower options, as measured by the OOI (raw correlation equals -.022). Figure 2-7 plots each training occupation by their (residualized) log wages and OOI.

We next look at which occupations drive the negative correlation. The upper-left corner of this figure, comprises high paying occupations with few relevant options: medical doctors, pilots, dentists. These are textbook examples of high-wage occupations with skills that cannot be easily transferred. While wages in these occupations are still high, our model predicts that, at least in partial equilibrium, if these workers were able to use their skills in more industries, their wages would have been even higher. The bottom right corner shows occupations like meter reader or car sales that have lower wages and more options. These are examples of low-wage occupations with highly general skills.

Table 2.2 presents coefficients from Equation 2.12, including additional controls for training occupation, district of residence and establishment. Controlling for training occupations (column 2) does not change the results significantly. However, adding controls for worker's district of residence as well (column 3) reduces some of the education gap, and increases the gap between German citizens to non-citizens. This suggests that higher educated, and non-citizen workers are more concentrated in large cities where there are more job options, and therefore their OOI is lower once controlling for that. Controlling for establishments (column 4) yields results that are similar to the results with controls for districts and occupation, as workers in the same establishment tend to live closely. However, we find smaller gender differences in options within establishments.

2.5.1 Mass Layoffs

We next show that the OOI is able to predict the ease with which workers recover from a job separation. These separations force workers to move to their outside option. To identify exogenous separations we follow the prior literature in focusing on mass-layoffs.

We start by constructing a sample of workers who were involved in mass-layoff between 1993 and 2014, following the approach in Jacobson et al. (1993). We define a plant in our sample as undergoing a mass layoff if it has a decline in its workforce of at least thirty percent over the year. We consider only mass layoffs that occur in establishments with at least fifty workers. We restrict our analysis to workers who had been employed at the establishment for at least three years prior to the mass layoff and who are below the age of 55. This leaves us with a final sample of 13,681 workers from 583 distinct mass-layoffs.

The outcome variable we use is relative income: the ratio between current daily income and the last daily income before the layoff. Formally we define relative income as $\widetilde{w}_t = \frac{w_t}{w_0}$, where t is months after the layoff, ranging from one to thirty-six. In case a worker is unemployed during this period their relative wage is set to zero. This choice of outcome variable takes out all productivity differences that can be captured with worker fixed effects and might be correlated with the OOI.

Appendix Figure 2-16 replicates the main result in Jacobson et al. (1993). Specifically, we look at

$$\widetilde{w}_{i,t} = \beta_t + \psi_{j(i)} \tag{2.13}$$

where $\psi_{j(i)}$ is establishment fixed effect. We plot β_t, the mean relative income of workers each month, for three years after the layoff. We find that on average workers lose 80% of their income in the month following the layoff. Their income gradually returns to its previous value over the next three years.

We next look at the differences in recovery for workers with different value of OOI. Within each establishment in our sample, we divide the laid-off workers into two groups, based on whether they are above or below the establishment median of the outside options index. Appendix Table 2.12 shows summary statistics for the two groups.

We then calculate

$$\widetilde{w}_{i,t} = \rho_t High_i + \delta_{j(i),t} \tag{2.14}$$

where $High_i$ indicates above the establishment median OOI, and $\delta_{j(i),t}$ is establishment by month fixed effects. Figure 2-8 plots ρ_t, the difference in relative income between those two groups for each month in the three years after the layoff. Our point estimates show that workers with better options, as captured by the OOI, gain an additional 8 percent of their previous income during the first year after the layoff. The groups seem to converge throughout time and after three years there are no differences.

The lower relative income is driven by both longer search time, and lower wage after search. To show this we recalculate Equation 2.14, replacing the outcome variable with $e_{i,t}$ an indicator for being employed, which we define as having a positive wage

$$e_{i,t} = \rho_t^e High_i + \delta_{j(i),t}^e \tag{2.15}$$

Figure 2-9 plots ρ_t^e, the differences in the share of employed workers on both groups. Our point estimate show that about 2% more people with higher OOI are working compared to the lower OOI. Therefore, the relative income in the new job must also be lower to explain an 8% difference between the groups.

We then repeat this analysis with a continuous measure of OOI, and a varying set of controls. We regress relative income at month t on the OOI, with fixed-effects for establishment by month $\delta_{j(i),t}$.

$$\widetilde{w}_{i,t} = \lambda_t OOI_i + X_{it} + \delta_{j(i),t} \tag{2.16}$$

We also repeat this analysis with additional worker controls X_{it} including tenure, gender, age and education. The results are reported in Table 2.3. We find virtually the same patterns we found when divided by the median, for all choices of controls.

Our findings suggest that the OOI may have an impact on additional dynamic aspects of the labor market that are not captured in our static-model. The OOI quantifies the number of options workers have. It is therefore not surprising that workers with more options (higher OOI) are able to find a new job more quickly. However, the effect on relative wage is less obvious, as the OOI is likely to already affect income before the layoff. One interpretation of our findings is that workers with higher OOI face a labor market that is closer to perfectly competitive, such that the impact of a single employer on equilibrium wages is negligible, and not affected by their exit.[28]

2.6 Effect on Wages

In this section we use two different sources of quasi-random variation in options to estimate the elasticity between the OOI and wages. The first (Section 2.6.1) focuses on the introduction of high-speed commuter rail stations in small German towns, and the second (Section 2.6.2) uses a shift-share ("Bartik") instrument. These methods yield semi-elasticities between .17-.32.

Estimating this elasticity using instruments allows us to translate differences in OOI to differences in wages, even if the model is misspecified. In our model, the effect of the OOI on wages is determined by the parameter α (Theorems 7, 8). Our model only implies that this parameter is non-negative and that in a perfectly competitive labor market, this parameter is zero. Since we do not use any of the model assumptions to estimate this parameter, the results capture the elasticity between the OOI and wages, even if the model is misspecified. This exercise can also be seen as a test for whether the OOI has an impact on wages.

Identifying α, the relationship between options and wages is challenging for two reasons. First, the OOI estimator is a function of the observed characteristics X_i. These observables are likely to also capture differences in productivity, thus creating a problem of omitted variable bias. Second, the OOI is estimated with a potentially large amount of noise, which would create an attenuation bias, especially when adding several controls. Because of those reasons, the OLS estimates of this parameter (Appendix Table 2.13), strongly depend on the set of controls we use. Adding more controls shrinks the results towards zero, and can even affect the sign of the coefficient.[29] In order to cope with both issues we use two sources of quasi-random variation in workers' options that do

[28]This intuition can be captured with the predicted effect of a plant closure on the OOI. Removing a job with low probability have only a small effect on the OOI. Since workers with higher OOI, have on average a lower probability to be in every job, their OOI is less affected from the destruction of only a few jobs. Therefore their wage will be less affected as well.

[29]Our findings in Section 2.5 suggest that the bias could go both ways, therefore the sign of the OLS coefficient could be both positive and negative. The OOI is correlated with high experience, or high education, but also with lower wage occupations, potentially because high-skill occupations tend to be more specific.

not affect their productivity.

2.6.1 High-Speed Commuter Rail

We leverage the expansion of the high-speed rail network in Germany as an exogenous shock to workers' outside options, following prior work by Heuermann and Schmieder (2018).[30] High-speed trains were first introduced in Germany in 1991-1998. During that period, stations were placed in major cities. We focus on the second wave of expansion, which began in 1999. During this wave, new stations were added in cities along *existing* routes. Cities were chosen for new stations after the routes already existed, and mostly on the basis of political considerations, and not labor market factors. As a result, towns as small as 12,000 residents were connected to the train network. Heuermann and Schmieder (2018) show that this increase in infrastructure led to an increase in commuting probability. Figure 2-10 shows the map of districts that got stations in the two waves of installation.

There are several threats to identification that should be considered. One potential concern is that the cities which received new train stations were selected on the basis of expected increase in productivity. Institutional details and prior research suggest that this is unlikely.[31] Another concern is that the new infrastructure can also be used for the transportation of goods. This would impact the workers' wages through their productivity. However, Heuermann and Schmieder showed that the introduction of these stations had no effect on the product market, as the trains were only used to transport passengers. One remaining concern is that the trains also had a similar effect on employers in those small towns, by allowing them to recruit workers from major German cities. We do not have detailed enough data on the employers to reject this possibility.

We supplement our data with train schedules in the years 1999 and 2012, before and after the installment of the second-wave stations. From this data we construct an indicator variable for every match for whether the worker can use a direct high-speed line to get to this job. We add this variable for our estimation of the match probabilities f_j^i. Therefore, we allow the match probability

[30] Daniel Heuermann and Johannes Schmieder generously provided the train data for our use in this project.
[31] See Heuermann and Schmieder (2018) for the full list of criteria that were used for location choices.

to depend on whether there's a high-speed line connecting the worker and the employer. To allow for heterogeneity between workers in their demand for trains, we interact this variable with all worker characteristic X_i. We then use these probabilities to estimate the OOI, as explained in Section 2.4.2.

The treatment group consists of all workers who, in 1999, lived in districts that got high-speed connections during the second-wave expansion (1999-2012). The control group consists of all workers who, in 1999, lived in districts that never received stations. Since the major cities were connected in the first wave, they are effectively excluded in this analysis. We then follow the same workers to the year 2012, regardless of where they live. We match workers from the treatment and the control group based on their gender, age, citizenship, education level, training occupation, state (Bundesländer) and lagged income using nearest-neighbor matching with replacement.[32] Appendix Table 2.14 presents a balance table for this match.

We estimate the following system of equations:

$$\begin{aligned}\Delta_{1999}^{2012} \log w_{im} &= \alpha \Delta_{1999}^{2012} OOI_i + \mu_m + v_{im} \\ \Delta_{1999}^{2012} OOI_{im} &= \delta Treat_i + \lambda_m + \epsilon_{im}\end{aligned}$$

where $Treat_i$ is an indicator for living in a treated district in 1999 and μ_m, λ_m are match fixed-effect. Because this is a binary instrument, $\widehat{\alpha}$ collapses to a Wald estimator

$$\widehat{\alpha} = \frac{\overline{\Delta_{1999}^{2012} \log w_{treat} - \Delta_{1999}^{2012} \log w_{control}}}{\overline{\Delta_{1999}^{2012} OOI_{treat} - \Delta_{1999}^{2012} OOI_{control}}} \tag{2.17}$$

where the average is taken over matched pairs. We develop a procedure that calculates standard errors, building on the approach of Abadie and Imbens (2006). More details in Appendix 2.11.

Table 2.4 shows the main result: an elasticity of .32 between options and wages. Column 1 shows that the OOI increased by .07 in treated districts following the introduction of the new stations. The reduced form results in column 2 suggest an increase of about 2.5% increase in

[32] We require the match to be exact on gender, education, state and 2-digit occupation.

income in the treated districts. Combining both estimates into a 2SLS estimator in column 3 yields a semi-elasticity of approximately .32 between our measure of outside options and wages. Column 4 shows that our matching process worked: there are no pre-trends. Our OLS results in column 5 show a precise zero. This is likely to be driven by an attenuation bias that is amplified substantially when using first differences of noisy variables in estimation.[33]

We next verify that our effect is driven by workers who are more likely to use the train. The high-speed commuter rail is a fairly expensive commuting option.[34] As a result, the introduction of train stations should primarily affect high-income workers. We break our sample into three education groups, which we use as a proxy for potential income. Figure 2-11 plots the first stage and the reduced form results, together with our point 2SLS estimates for each group. We find a higher first stage for workers with higher education. These are the workers we would expect to use the train the most. The reduced form is also higher for the more educated workers, though the estimate is imprecise. We cannot rule out a zero effect on the low education group. The two-stage least squares estimates are similar for all three groups, and we cannot rule out homogeneous effects by education group.

2.6.2 Shift-Share ("Bartik")

We next use a standard shift-share ("Bartik") instrument to estimate the elasticity between wages and options.[35] Though this exercise gives a lower point estimate of .17, we cannot reject that the elasticity we estimate is identical to the one estimated in the prior section.

The idea behind this strategy is to compare workers who work in the same industry, but who have different outside options, because they reside in different parts of the country with different industry mixes. Some workers happen to live near industries that are growing, while others happen

[33] Duncan and Holmlund (1983) show that this depends on the level of autocorrelation between the measurement errors, and true signal. Since worker observables tend to be constant while the measurement error might change between years, we expect attenuation bias to be much stronger in first difference.

[34] For example, a round-trip between Montabaur and Frankfurt takes 45 minutes each way and costs 60 Euros.

[35] These shocks were used in several papers including Bartik (1991); Blanchard and Katz (1992); Card (2001); Autor et al. (2013). It was used specifically for the context of a shock to outside options by Beaudry et al. (2012).

to live near industries that are contracting. Because local growth of certain industries may be due to the impact of local productivity shocks, we use national industry trends as an instrument.

The instrument is a weighted average of national industry growth, weighted by the initial share of each industry in the region. Formally, we define

$$B_r = \sum_j s_{jr}^{04} \times \widehat{g_j}$$

where s_{jr}^{04} is the share of employed workers in region r, working at industry j in the base year (2004) and $\widehat{g_j}$ is the national employment growth of industry j. Regions are defined by the administrative regions ("Regierungsbezirke") in Germany, the statistical unit which is closest to a commuting zone.[36] Industries are defined at the 3-digit level.

To estimate the national growth of different industries, controlling for region-wide shocks, we regress the change in employment in industry j in region r between 2004 and 2014 on industry and region fixed effects:[37]

$$\Delta_{04}^{14} \log E_{jr} = g_j + g_r + \varepsilon_{jr}$$

By construction, the estimator of $\widehat{g_j}$ is not driven by regional trends captured in $\widehat{g_r}$. We use the weighted average of the industry fixed effects $\widehat{g_j}$ by initial industry shares s_{jr}^{04} to calculate B_r. This construction verifies that B_r is not driven by local employment shocks in this region, or even in nearby regions.[38]

We estimate the following system of equations

$$\begin{array}{rclc}
\Delta_{04}^{14} \log w_{ijr} & = & \alpha \Delta_{04}^{14} OOI_{ijr} & + \beta \Delta_{04}^{14} X_{ijr} + Ind_j^{04} + \upsilon_{ijr} \\
\Delta_{04}^{14} OOI_{ijr} & = & \gamma B_r & + \delta \Delta_{04}^{14} X_{ijr} + Ind_j^{04} + \epsilon_{ijr}
\end{array} \tag{2.18}$$

[36] We take all 39 regions based on NUTS2 level coding of the European Union. This includes historical administrative regions that have been disbanded. Results are robust to the definition of a region and hold for the NUTS3 level (district) as well.

[37] We make a Bayesian correction of uniform prior by adding one observation in each industry, region and year combination.

[38] This is different from a leave-one-out estimate, that might still be driven by local shocks in nearby regions.

where we control for the industry of the worker in the beginning of the period (2004). We cluster standard errors at the level of the treatment, which is the region. The parameter of interest is α, the elasticity of wages with respect to options.

Table 2.5 presents the main results. Columns 1 and 2 show the first-stage and reduced form results. A 10% higher employment in other industries, which is about .1 increase in the instrument, translates to approximately 6% more relevant options, and 1% increase in wages. Combining both estimates yields a semi-elasticity of .17: a 10% increase in relevant options leads to a 1.7% increase in wages.

The identifying assumption is that growing industries are not systematically located in regions where wages are growing for other reasons (Borusyak et al., 2018). One way the assumption could be violated is if there are productivity spillovers. Workers that live near industries that are growing, may enjoy a local demand shock for their production due to the positive income effect on workers in that region. This could generate a wage increase, that is not driven by the improvement in their outside options. This is particularly a concern for workers who are producing non-tradable goods, whose productivity is set by local demand.

We address this concern by showing the results hold for workers in exporting industries, which are less likely to be affected by local demand shocks. We use information from the establishment survey to calculate the export share of each industry.[39] We divide our data into three groups based on the export share of the industry where the worker worked in 2004. Table 2.6 shows the results for each of the groups. We find a large and statistically significant elasticity between options and wages even among workers in industries with the highest exporting share. Column 1 indicates that in response to a 10% increase in OOI, workers in these industries see their wages rise by 1%. This elasticity is somewhat lower than that in our baseline results (.10 versus .17). However, we cannot reject that they are equal.[40]

[39] This is a lower bound for demand from outside the region, as it does not include sales to other regions in Germany, which we cannot see in our data. We calculate the mean industry level since we don't have the share of sales from export for all employers, but only a representative sample.

[40] Beaudry et al. (2012) find similar results when dividing the data into tradable and non-tradable industries, based on their geographical spread. They argue that non-tradable industries are geographically spread across different regions, while tradable goods could be concentrated in specific regions. They also address additional potential threats to the

We next examine heterogeneity across gender and the three education groups. We estimate Equations 2.18 separately for each group. Figure 2-12 plots the results for all groups, as well as the full population. While splitting the sample increases the size of the confidence intervals, the point estimates are quite close. This suggests that using the same value for α for all groups is a reasonable approximation.

We next use this setting to decompose the different effect of access to more options into impacts for job stayers and movers. Because the choice of whether to move is endogenous, we view this as a decomposition exercise. We interact the changes in OOI with an indicator variable for whether a worker stayed at their establishments during this period. The results are shown in Table 2.7. As our model predicts, we find that the effect on stayers is smaller. This is possibly because they only benefit through an improvement in their *outside* options. The larger effect on movers is consistent with an additional improvement in match quality.

While the elasticity we estimate in this exercise is lower from the one we estimated using the fast commuter rails, their difference is not statistically different from zero. Figure 2-13 compares our results in this section to the elasticity we estimated using the introduction of high-speed commuter rails. The fact that we found elasticities of a similar magnitude by using two distinct sources of variation suggests that this range of estimates is a reasonable benchmark for the value of α.

2.7 Implications for Wage Inequality

In this section, we combine our estimates on the distribution of the OOI, with our estimates of the OOI-wage elasticity, to assess the overall effect of options on the wage distribution. We then examine which covariates drive differences in options. We find that equalizing workers' ability to commute or move would eliminate the gender gap in OOI, and would reverse the sign of the OOI gap by education.

identification assumption and show they don't seem to have a significant effect on the result.

2.7.1 Overall Impact on the Wage Distribution

We examine what portion of between-group wage inequality can be attributed to difference in OOI. We first estimate a Mincer equation

$$\log w_i = \beta_0 X_i + \epsilon_i \tag{2.19}$$

where X_i includes indicators for each education group, a quadratic in age, gender, citizenship status, log district density and an indicator for part-time job. Since wages are top-coded we use a Tobit model to estimate $\widehat{\beta}_0$. We then add the OOI to the set of dependent variables, with a fixed coefficient.

$$\log w_i = \widehat{\alpha} OOI_i + \beta_1 X_i + \epsilon_i$$

We use $\widehat{\alpha} = .26$ which is the average of the two point estimates we derived in Section 2.6 from the two quasi-random sources of variation. $\widehat{\beta_0}$ captures the overall gaps in wages between these demographic groups, $\widehat{\beta_1}$ is the remaining gaps that are driven by factors other than the OOI, and $\widehat{\beta_0} - \widehat{\beta_1}$ is the part that can be attributed to the differences in OOI.

Figure 2-14 shows the main results. The full bars display the gaps that we estimated ($\widehat{\beta_0}$), where every bar is the wage premium for this group members. For instance, the premium for being a male (the gender gap) is .19 log units in Germany. The portion that can be attributed to the OOI $\left(\widehat{\beta_0} - \widehat{\beta_1}\right)$ is colored in red, while the remaining gap ($\widehat{\beta_1}$) is left in blue.

The OOI explains significant portions of several German wage gaps. When we add OOI to the regression, the gender gap is cut by .06 log units (30% of the overall gap). This is driven by the .23 gender gap in OOI we found in Table 2.2, multiplied by $\widehat{\alpha}$. Our results also indicate that 88% of the gap between German citizens to non-citizens (.08 log units) can be attributed to differences in options. The wage difference between high-level and intermediate-level secondary schooling is cut by .07 log units, which is about 25% of the initial gap. Our results also attribute 39% of the return to experience at age 18 in access to options. Table 2.8 shows these results numerically in columns (3) and (4), as well as results from a winsorized OLS in columns (1) and (2).

2.7.2 Explaining Differences in Outside Options

We next examine which factors impact workers' options. We find that differences in commuting costs seems to be particularly important, especially in its effect on the gender gap and return to higher education.

We start by examining the impact of different variables on the probabilities of observing a match. We analyze our results from the estimation of matrix A defined in Assumption 2, which we estimated using a logistic regression. This matrix is also the cross-derivative of match quality τ (Equation 2.9). Appendix Table 2.15 shows the top absolute values of $\widehat{A}$, when variables are standardized so the results are not affected by specific units.[41]

These results indicate that the most important factor in determining match quality is commuting and moving costs. Distance has the largest standardized coefficient in absolute terms (-4.15). While distance to a job is an important factor for all workers, it is particularly important for female workers, for less-educated workers, and for non-German citizens. Appendix Table 2.16 presents the raw coefficient on distance for different worker characteristics. This coefficient is the effect of an additional mile on the log probability of a match, at mile zero. For our baseline group, forty year old males citizens from higher secondary schools, the coefficient is -.141. The interaction with female is -.024, so women are 17% more sensitive to distance than the baseline group. Lower educated workers are significantly more sensitive to distance (coefficient -.037). Non-German citizens seem to be more sensitive than citizens (coefficient -.019). Finally, workers at first become less sensitive to distance with age, but this is a concave function that reaches its maximum at age 42.

By simulating counterfactuals from the underlying model, we can quantify the overall effect of differences in commuting and moving costs on wages through their effect on the OOI. We estimate the wage gain for every worker, if they had the minimal commuting/moving costs. Based on our estimation, these are the costs of a 40 year old, high-educated, male citizen. We generate a matrix $\widetilde{A}$ where the coefficients on distance is set to this minimum level for all workers. We then simulate

[41] Online appendix table A1 shows the full standardized results for A ; online table A2 shows the raw results.

the probabilities $\widetilde{f_j^i}$ using this matrix, calculate the $\widetilde{OOI_i}$ and translate it to $\widetilde{\log w_i}$ using $\widehat{\alpha}$. This counterfactual should be thought of as changing only a zero measure number of workers each time, and keeping all other workers and employers unchanged, so that there are no general equilibrium effects. We compare the differential gains from this exercise to assess the importance of commute in generating wage gaps.

We run a regression of the counterfactual gains in wage over basic demographics

$$\Delta \log w_i = \beta_2 X_i + \epsilon_i$$

where $\Delta \log w_i = \widetilde{\log w_i} - \log w_i$ and X_i same as in Equation 2.19. The coefficients β_2 from this regression are the part of the wage gap that would be closed, if commuting and moving costs were equalized at the lowest level, for these workers. The results of this exercise are presented at Figure 2-15. The figure plots the full gap ($\widehat{\beta_0}$, blue), the portion that can be attributed to the OOI ($\widehat{\beta_0} - \widehat{\beta_1}$, red), and the part that will be closed by equalizing commuting costs at the minimal level ($\widehat{\beta_2}$, yellow).

Differences in commuting costs seem to explain all of the gender gap that is driven by differences in options. Equalizing commuting costs would increase wages for women by about .07 log units, relative to men. This is one third of the overall gender gap. Even though we find that men and women sort into different jobs, there seems to be a similar number of jobs for males and females, therefore the only difference is the distance in which workers are searching for jobs. This does not mean that there aren't other ways to increase the OOI for women, such as increasing the supply of jobs that women typically sort into.

In contrast, equalizing commuting costs increases the wage gap between German citizens and non-citizens. These results are surprising at first glance because we found that non-citizens are more sensitive to distance. However, they can be explained by the fact that non-citizens are more concentrated in large German cities. As a result, their commuting costs are already low. German citizens are more dispersed across rural areas, and are more dependent on their ability to commute to jobs in major cities.

The education gap in OOI actually reverses once we equalize commuting costs: workers with intermediate-secondary education have more options than those with higher-secondary. Therefore, the higher-education premium drops by .15 log units (51% of the overall premium), which is more than the full effect of the OOI difference between these groups (.07 log units). This implies that, in a given area, workers with intermediate-secondary education have more relevant job options than workers with high-secondary education. It is only because higher-secondary education workers are willing to take jobs in more distant areas, that they end up with more options. This result can be explained by the fact that more educated workers tend to be more concentrated in occupations that have more industry specific skills, as shown in Figure 2-7. Additionally, intermediate-secondary workers can take both higher-skill, and lower-skill jobs in addition to staying at the same level. While high-secondary workers usually have fewer options to climb to jobs requiring even more skills.

Other than geographical distance, the most significant factor in determining a worker's match (Appendix Table 2.15) is their training occupation. Workers tend to stay in occupations similar to the ones in which they were trained. Our results in Section 2.5 show that those who undertook training in occupations with more transferable skills have more options than those who received more narrow training. Since transferable skills are more common in low-paying occupations, the OOI is reducing inequality between occupations.

2.8 Conclusion

In this paper we provide a distinctive and micro-founded approach to empirically estimate workers' outside options, and to measure the impact of outside options on the wage distribution. The starting point for our analysis was a two-sided matching model, which produced a sufficient statistic for the impact of outside options on wages, the OOI. We took the OOI to the data to identify the workers with better outside options. We then combined this result with a causal estimate of the elasticity between the OOI and wages, to assess the overall impact of options on wages. Our results suggest that differences in outside options generate lower income for females by six percent, non-citizens by eight percent and intermediate educated workers by seven percent (compared to high educated).

Our results indicate that policies that improve workers' options, including investments in transportation infrastructure or regulation of working hours, are likely to have significant general equilibrium effects. While such policies are usually analyzed only through their impact on workers that directly benefit from them, our results indicate that these policies will likely have important spillovers onto other workers through their outside options. These general equilibrium channels can be studied through their effect on the OOI.

One interesting direction for future work would be to use this framework to analyze specific industries in which outside options play a key role, and good micro-data is available. A similar analysis could also be done on the employer's side of the market, analyzing heterogeneity in the availability of options for firms, and the impact of outside options on profits. Finally, the ability to identify workers with better outside options could be useful in studying heterogeneous effects of various policies, or labor market shocks. Our analysis of the heterogeneous response to mass-layoff is one example for how this can be done.

2.9 Figures and Tables

Figure 2-1: German Districts

Note: This map illustrates the 402 districts (kreis) in Germany.

Figure 2-2: Distribution of Outside Option Index

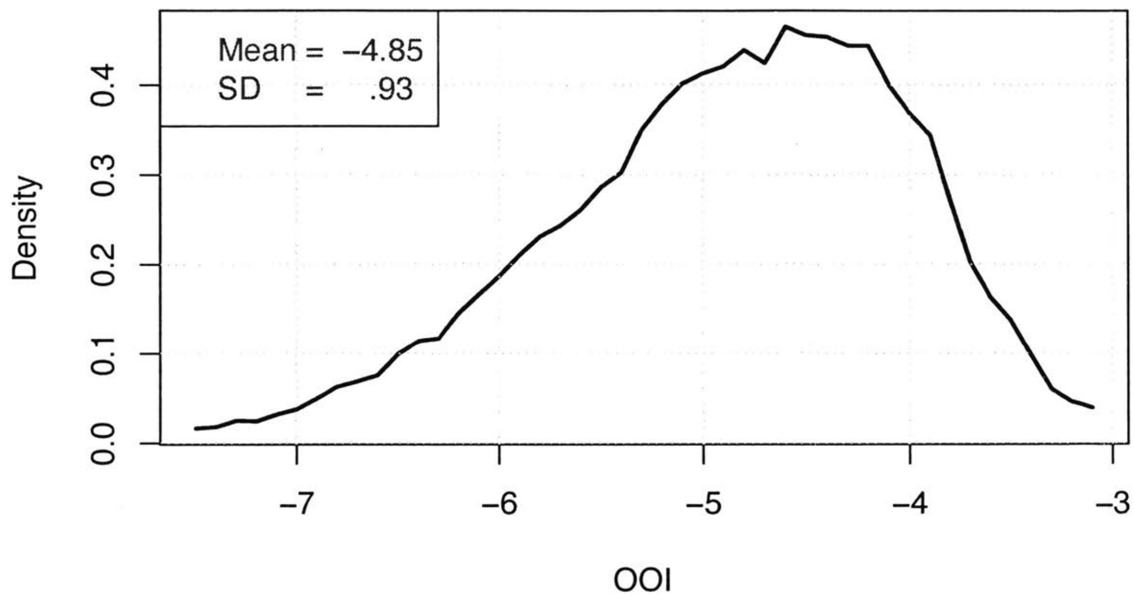

Note: This figure plots the distribution of the outside options index as calculated for the population of German workers as of June 30th, 2014. The OOI was calculated using the procedure described in Section 2.4.2. LIAB sample weights are used to make the distribution representative of the German population.

Figure 2-3: OOI by Characteristics

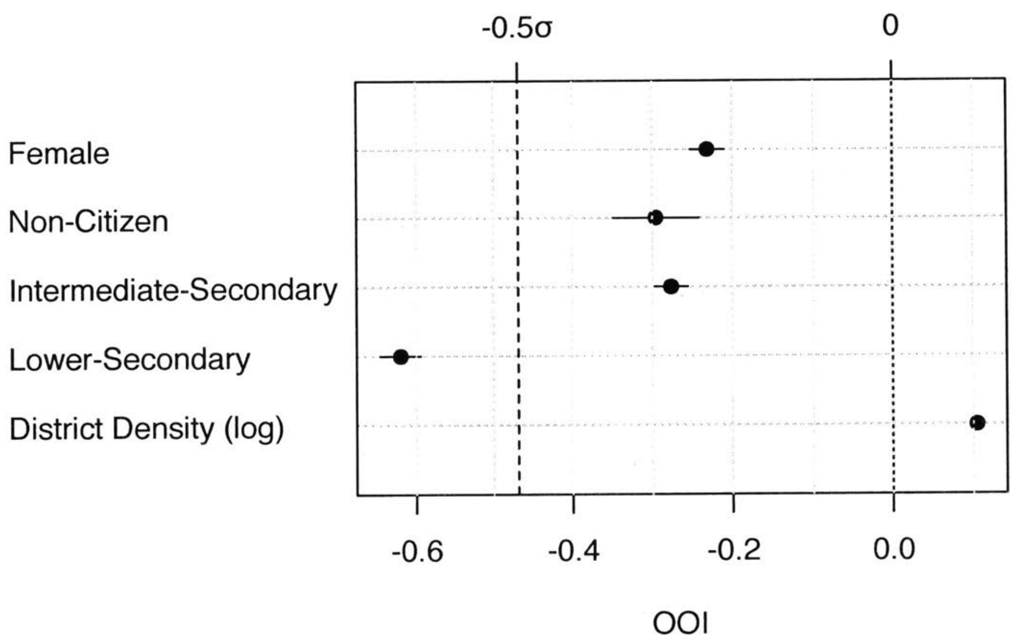

Note: This figure plots the coefficients from a regression of OOI on education, gender, citizenship and a quadratic in age. The results are also presented on column 1 of Table 2.2. Confidence intervals are plotted at the 95% level. The lower axis shows raw OOI units, while the upper axis uses standard deviation units.

Figure 2-4: Cumulative Distribution of Outside Option Index by Gender

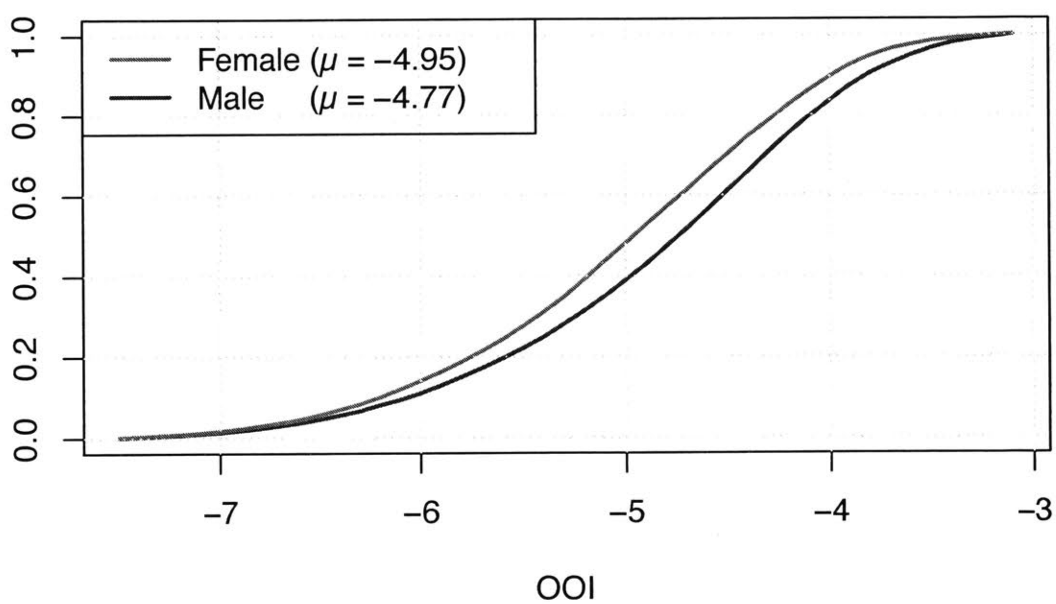

Note: This figure plots the cumulative distribution function of the outside options index by gender, as calculated for the population of German workers as of June 30th, 2014. The OOI was calculated using the procedure described in Section 2.4.2. LIAB sample weights are used to make the distribution representative of the German population.

Figure 2-5: OOI by Age

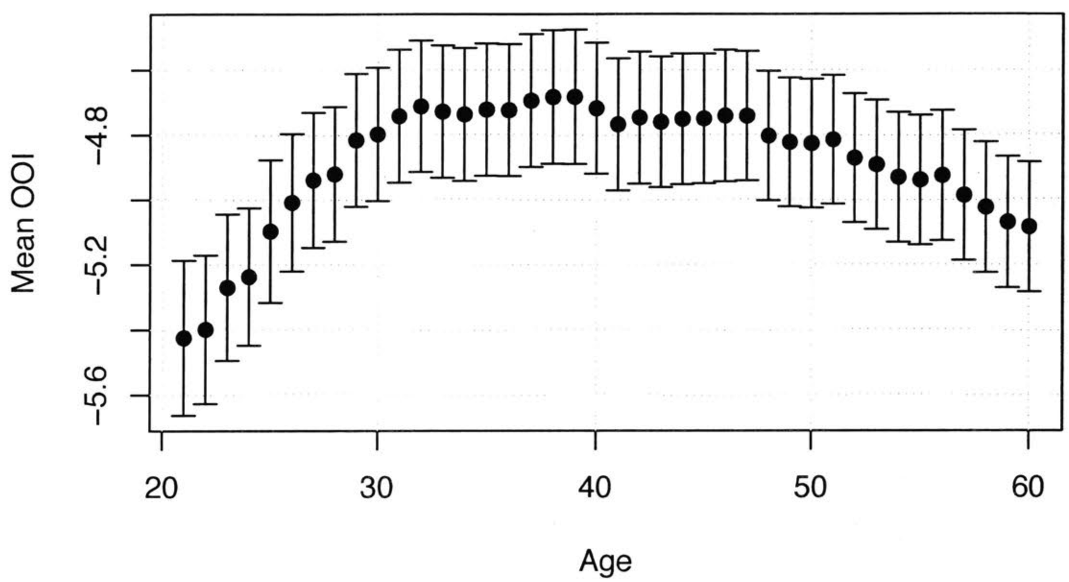

Note: This figure plots the mean OOI by age in the German population. LIAB sample weights are used to make the sample representative of the German population. Confidence intervals are plotted at the 95% level.

Figure 2-6: OOI Distribution by Region

Note: This figure plots the distribution of the outside options index by district (kreis) as calculated for the population of German workers as of June 30th, 2014. The OOI was calculated using the procedure described in Section 2.4.2. The value for each district is a weighted mean of the workers in this district, using the LIAB sample weights to make the distribution representative of the population in the district.

Figure 2-7: OOI by Training Occupation

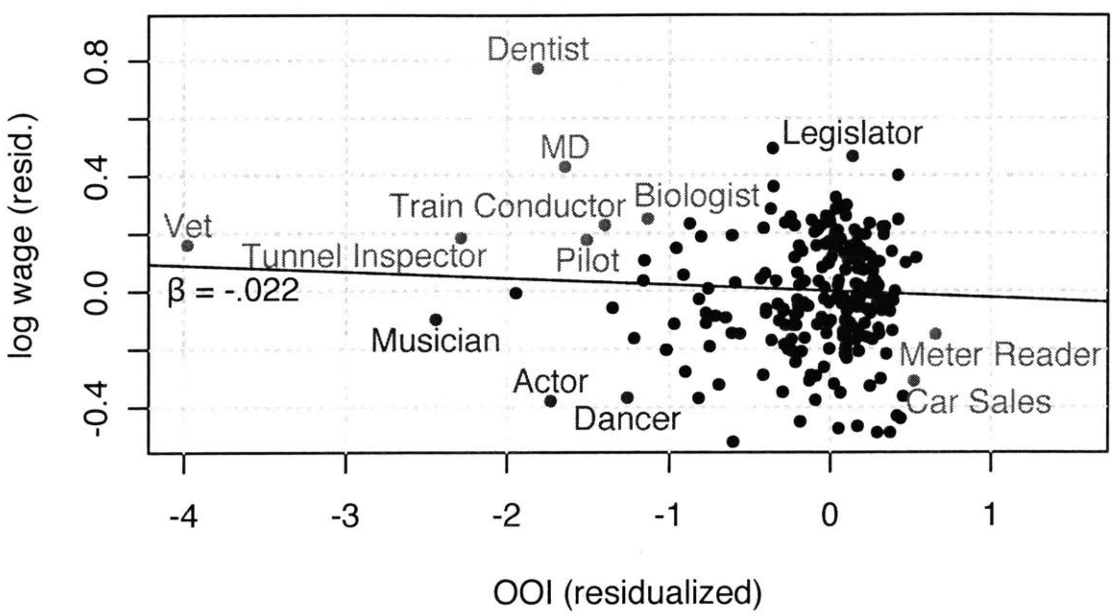

Note: This figure plots the mean residualized outside options index and log wages by training occupation as calculated for the population of German workers as of June 30th, 2014. The OOI was calculated using the procedure described in Section 2.4.2. Residuals for the OOI and log wages were taken from a regression on gender, a quadratic in age, education category, citizenship status and district of residence. Means are calculated using the LIAB sample weights to make the distribution representative of the population in the occupation. See Section 2.3.1 for exact definition of a training occupation.

Figure 2-8: Mass-Layoffs - Differences in Relative Income Between High/Low OOI Workers

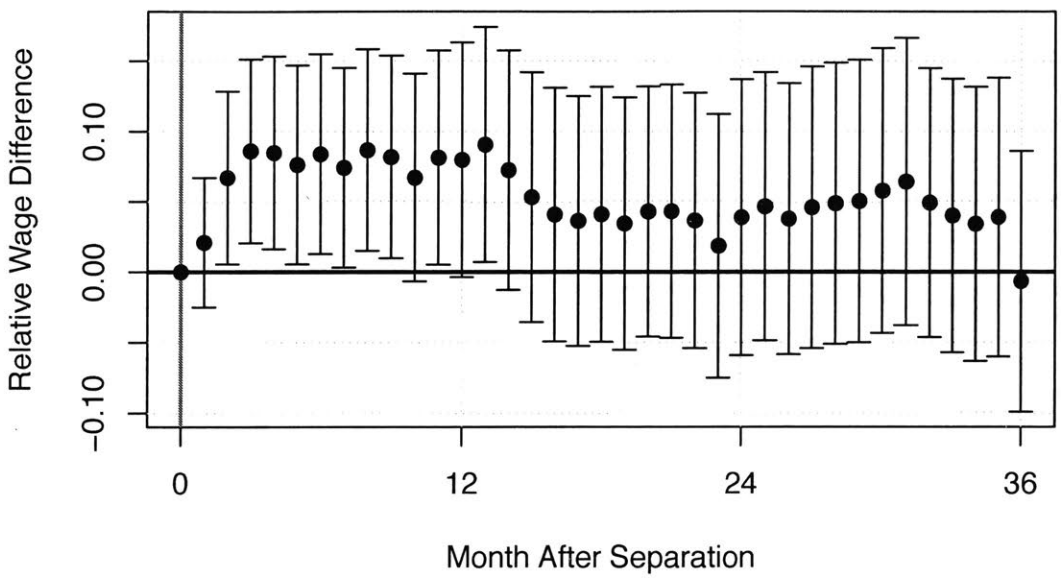

Note: This figure shows the difference in relative income for workers with OOI above and below the establishment OOI. Relative income is defined as the current daily income in that month divided by the last daily income before the layoff. Mass layoffs are defined as an establishment with at least 50 workers that reduced its workforce by at least 30% in a given year. The sample includes only workers who have worked for at least three years before the layoff and are below the age of 55. The median OOI is calculated based on the pool of laid-off workers in a given establishment and year. The coefficients are taken from a regression of relative income on an indicator for above median OOI, interacted with indicator for each month after separation (plotted), with fixed effects for establishment×month (Equation 2.14).

Figure 2-9: Mass-Layoffs - Search Time

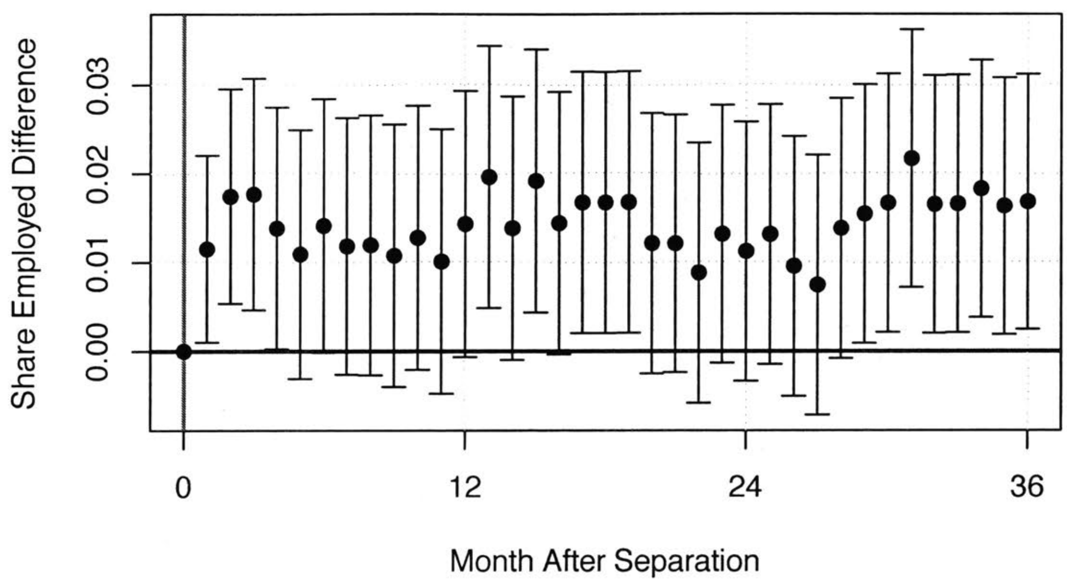

Note: This figure shows the difference in employment for workers with OOI above and below the establishment OOI. Employment is defined as any income greater than zero. Mass layoffs are defined as an establishment with at least 50 workers that reduced its workforce by at least 30% in a given year. The sample includes only workers who have worked for at least three years before the layoff and are below the age of 55. The median OOI is calculated based on the pool of laid-off workers in a given establishment and year. The difference is calculated using a regression of employment on an indicator for above median OOI, interacted with indicator for each month after separation, with fixed effects for establishment×month (Equation 2.15).

Figure 2-10: ICE Stations

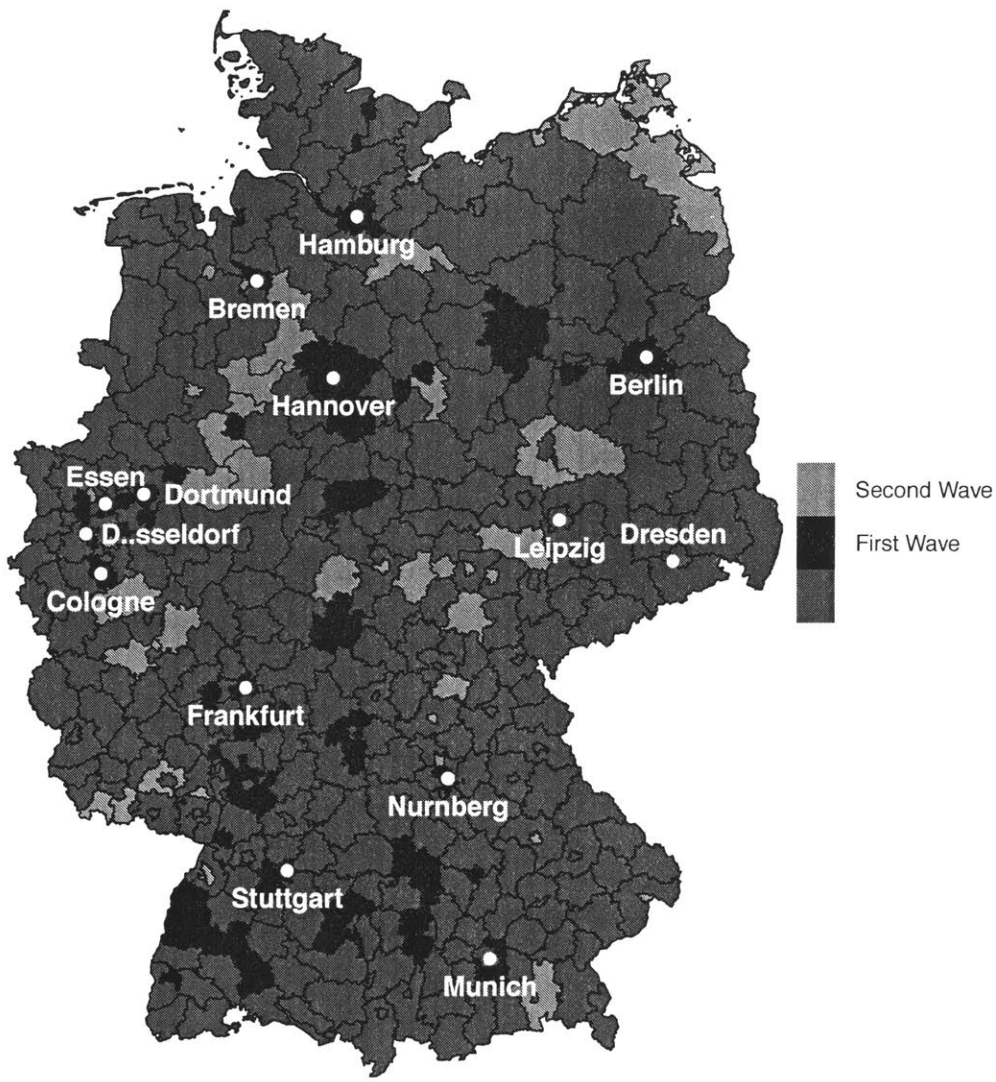

Note: This figure shows the locations of ICE train stations by districts. The first wave includes all stations that were opened pre-1999. The second wave includes all stations that were opened post-1999.

Figure 2-11: Impact of Express Trains by Schooling Level

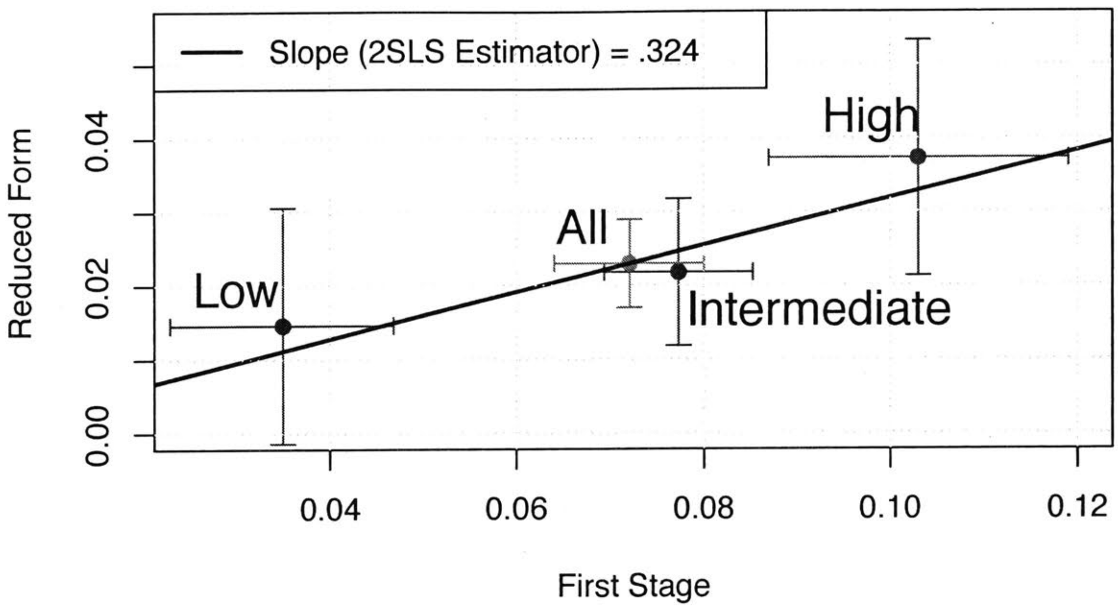

Note: This figure plots the first-stage and reduced-form results for three education groups, and their combination. First stage is the treatment effect on OOI. Reduced form is the treatment effect on log wages. Both were calculated using nearest-neighbor matching with replacement. Treatment is defined as workers that in 1999 lived in districts that got ICE stations post-1999. The control group includes workers that in 1999 lived in districts that never got ICE stations. Matching is done exactly on gender, education group, citizenship status, state and 2-digit training occupation and continuously on age, and PCA components for training occupation. Confidence intervals are at the 95% level, and are calculated based on standard errors derived from a method by Abadie and Imbens (2006). The black line represents the 2SLS point estimate for the entire sample.

Figure 2-12: Shift-Share Results by Gender and Education

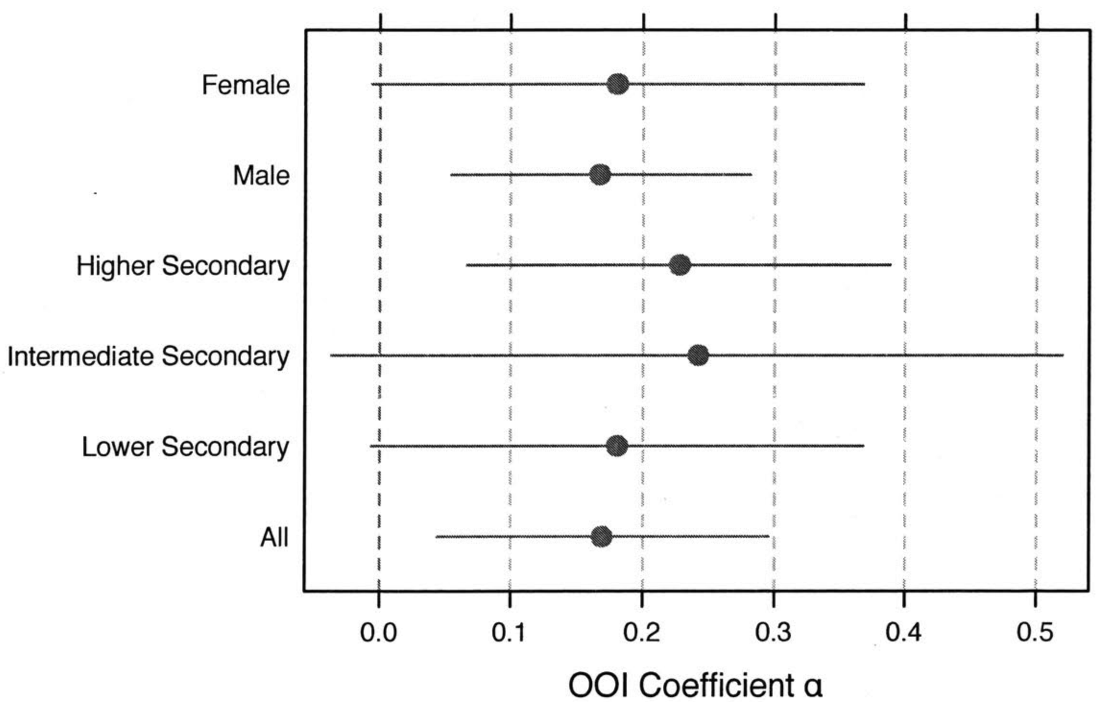

Note: Every category displays the estimate for coefficient $\widehat{\alpha}$ from Equation 2.18, ran separately for each education group or gender (blue), and for the entire population (red). This captures the effect of changes in OOI on changes in log wages between 2004-2014, when we instrument for the changes in OOI with the shift-share instrument. The instrument is constructed from an average of a 3-digit industry national employment growth weighted by the initial share of every industry in a region (see Section 2.6.2). Standard errors are clustered within the unit of treatment, which is regions. Confidence intervals are at the 95% level.

Figure 2-13: Estimates of Elasticity Between OOI and Wages

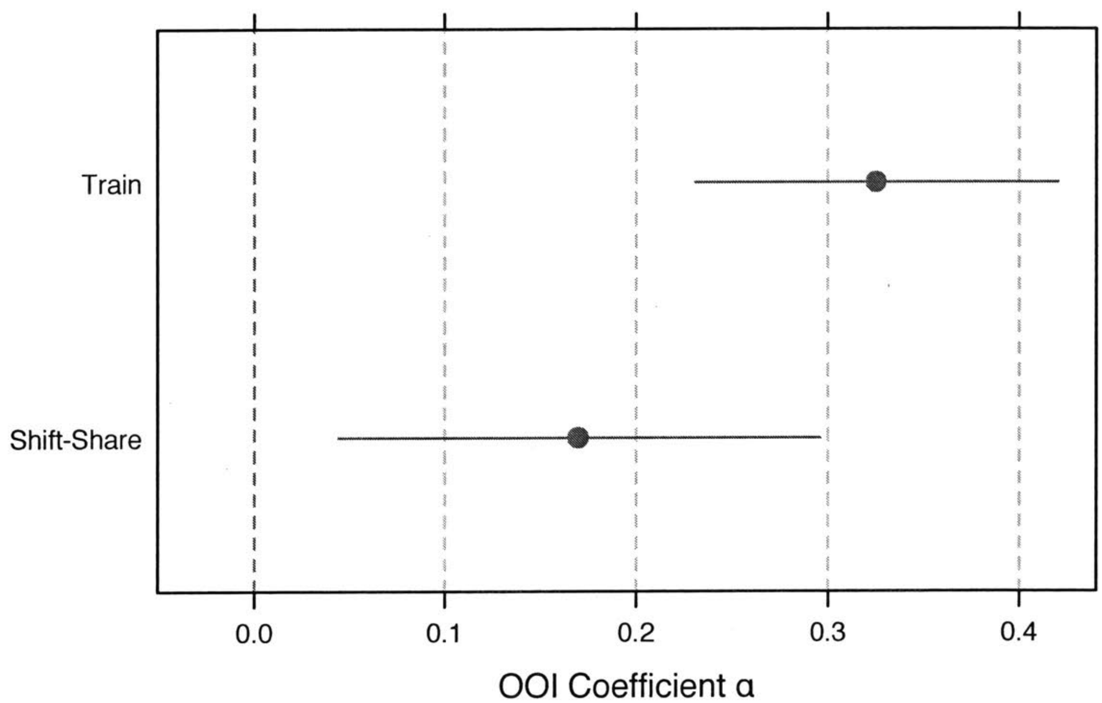

Note: This figure compares the elasticity between OOI and wages from the two different sources of quasi-random variations that we used. Train includes the results for parameter $\widehat{\alpha}$ estimated using the introduction of high-speed commuter rails (see Section 2.6.1 and notes for Figure 2-11 for more details) . Shift-Share uses an instrument based on national industry employment trends (see Section 2.6.2 and notes for Figure 2-12 for more details). Confidence intervals are at the 95% level. The difference between the point estimates is .156 (.081).

Figure 2-14: Overall Effect on Wage Inequality

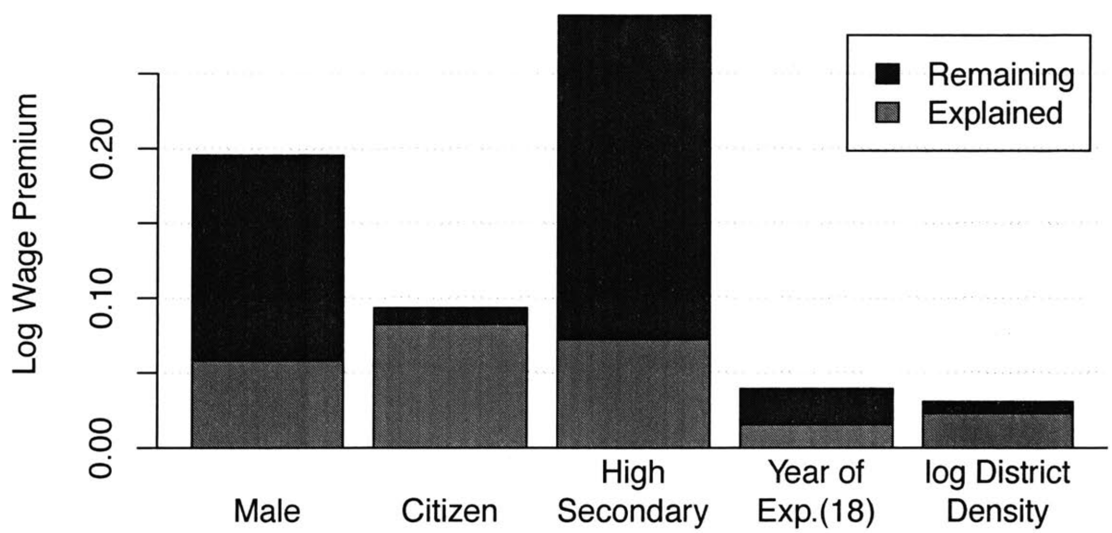

Note: Every bar in this plot is the coefficient on the corresponding category in a regression of log wages on Male, Citizen, indicator for secondary-education category, a quadratic in age, district density and an indicator for part-time job. The blue portion of the bars (remaining gap) is the coefficient from the same regression, controlling for the OOI with a coefficient fixed to .26, which was estimated with the two quasi-random variations. The part in red (explained gap) is the difference between the two coefficients. The reference workers is a female, non-citizen, with intermediate secondary education and 18 years old.

Figure 2-15: Effect of Commuting/Moving Costs

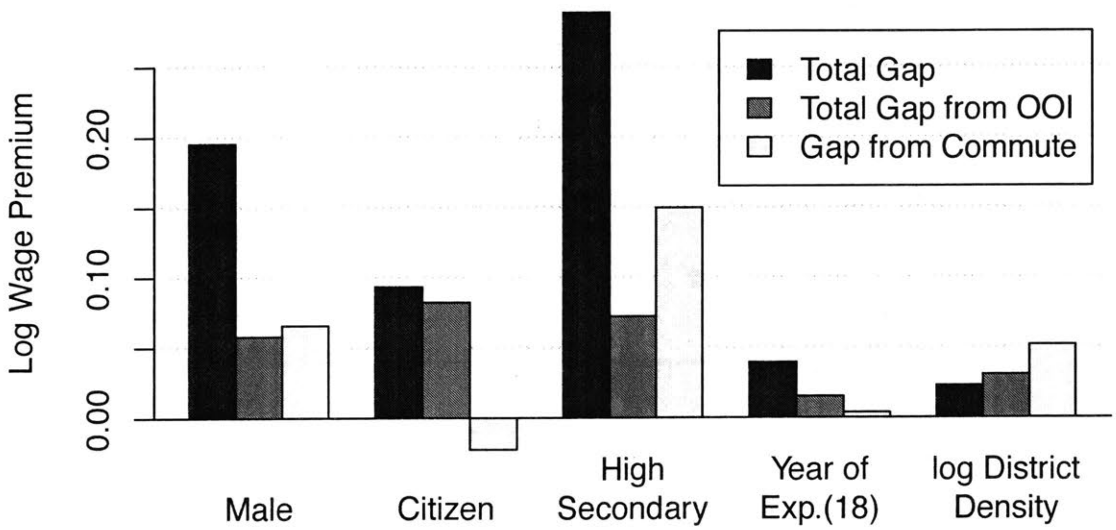

Note: Blue bars (Total gap) are derived from the coefficient on the corresponding category in a regression of log wages on Male, Citizen, indicator for secondary-education category, a quadratic in age, district density and an indicator for part-time job. The red bars (Total gap from OOI) is the difference in coefficient between the same regression and one that control for the OOI with a coefficient fixed to .26, which was estimated with the two quasi-random variations. The yellow bars (Gap from Commute) is calculate from a similar regression, replacing the dependent variable with minus the gains from reducing commuting costs to their minimal level (see Section 2.7.2 for more details). The reference workers is a female, non-citizen, with intermediate secondary education and 18 years old.

Table 2.1: Descriptive Statistics

	All	Female	Male
	(1)	(2)	(3)
Female	0.46	1	0
	(0.50)	(0.00)	(0.00)
Age	45.05	45.53	44.66
	(12.49)	(12.27)	(12.67)
German Citizen	0.97	0.98	0.96
	(0.17)	(0.15)	(0.19)
Education: Higher Secondary	0.29	0.3	0.28
	(0.45)	(0.46)	(0.45)
Education: Intermediate Secondary	0.31	0.34	0.28
	(0.46)	(0.47)	(0.45)
Education: Lower Secondary	0.22	0.16	0.26
	(0.41)	(0.37)	(0.44)
Education: Intermediate/Lower	0.19	0.2	0.18
	(0.39)	(0.40)	(0.38)
N	450917	162780	288137

	All	Female	Male
	(1)	(2)	(3)
Part-Time	0.31	0.53	0.13
	(0.46)	(0.50)	(0.34)
Fixed Contract	0.11	0.11	0.1
	(0.31)	(0.32)	(0.30)
Temporary Agency	0.02	0.01	0.02
	(0.12)	(0.08)	(0.15)
Establishment Size	1553	827	2166
	(7679)	(5014)	(9313)
Annual Sales/Worker (Euro	163286	130414	191.026
	(185651)	(163955)	(197953)
Pct Managers who are Female	0.26	0.36	0.17
	(0.31)	(0.35)	(0.24)
N	450917	162780	288137

Note: This table shows summary statistics of all workers and jobs in our sample on June 30th 2014. Sampling weights are used to make this a representative sample of the German population.

Table 2.2: OOI by Demographics

	(1)		(2)		(3)		(4)	
Female	0.237	***	-0.231	***	-0.227	***	-0.167	***
	(0.01)		(0.01)		(0.01)		(0.01)	
Education: Lower-Secondary	-0.66	***	-0.62	***	-0.596	***	-0.617	***
	(0.01)		(0.01)		(0.01)		(0.01)	
Education: Intermediate	-0.279	***	-0.277	***	-0.211	***	-0.216	***
	(0.01)		(0.01)		(0.01)		(0.01)	
Non-German Citizen	-0.307	***	-0.295	***	-0.459	***	-0.501	***
	(0.03)		(0.03)		(0.02)		(0.02)	
Age	0.099	***	0.107	***	0.107	***	0.099	***
	(0.00)		(0.00)		(0.00)		(0.00)	
Age^2	-0.001	***	-0.001	***	-0.001	***	-0.001	***
	(0.00)		(0.00)		(0.00)		(0.00)	
District Density	0.112	***	0.104	***				
	(0.00)		(0.01)					
Training Occupation FE			X		X			
District FE					X			
Establishment FE							X	
Observations	380109		380109		380109		380109	

Notes: This table shows the results of a regression of OOI on basic demographics (Equation 2.12). The sample includes all workers employed on June 30th 2014. Sampling weights are used to make this a representative sample of the German population. Training occupation fixed effects are at the 3-digit levels.

Table 2.3: Relative Income by OOI After Mass Layoff

	(1)		(2)		(3)	
3 Months	0.061	**	0.062	**	0.068	**
	(0.03)		(0.03)		(0.03)	
6 Months	0.068	**	0.069	**	0.082	**
	(0.03)		(0.03)		(0.03)	
12 Months	0.061	*	0.064	*	0.079	**
	(0.03)		(0.03)		(0.04)	
24 Months	0.033		0.039		0.064	
	(0.04)		(0.04)		(0.05)	
Mass-Layoff x Month FE	X		X		X	
Tenure			X		X	
Age					X	
Education					X	
Gender					X	
Observations	558686		558686		558686	
Workers	13707		13707		13707	
Mass-Layoff-Months	25561		25561		25561	

Note: This table shows the results of regressing relative income on OOI for workers that lost their jobs in a mass-layoff, for different times after the separation. Relative income is defined as the current daily income in that month divided by the last daily income before the layoff. Mass layoffs are defined as an establishment with at least 50 workers that reduced its workforce by at least 30% in a given year. The sample includes only workers who have worked for at least three years before the layoff and are below the age of 55. We include monthly income for the 36 months following the separation. The regression is based on Equation 2.16. Tenure includes a quadratic polynomial for days at the previous establishment. Age includes a quadratic polynomial. Education is a categorical variable for the type of secondary education (see section 2.3.2 for details).

Table 2.4: Impact of Express Trains on Options and Wages

First-Stage 1999-2012 (1)	Reduced-Form 1999-2012 (2)	2SLS 1999-2012 (3)	Reduced-Form 1993-1999 (4)	OLS 1999-2012 (5)
.073***	.024***	.324***	.002	.004
(.003)	(.004)	(.048)	(.002)	(.007)

Number of observations:	143,313
Number of treated observations:	37,695

Notes: This table shows the results of the impact of express trains on outside options, and wages. Columns 1-4 use nearest-neighbor matching with replacement. Matching is done exactly on gender, education group, citizenship status, state and 2-digit training occupation and continuously on age, and PCA components for training occupation (the third digit). The outcome variables are change in OOI 1999-2012 (column 1), change in log income 1999-2012 (columns 2,3,5) and change in log income 1993-1999 (column 4). Standard errors in matching are calculated using Abadie and Imbens (2006). 2SLS estimator is the division of the estimates in column 1 and 2 (Equation 2.17). Standard errors in column 3 are calculated using a method building on Abadie and Imbens (2006) (see Appendix 2.11 for details). OLS (column 5) estimates the regression of log wages on OOI with match fixed effects. Observations from the control group that appear in multiple matches also appear multiple times in the OLS. Standard errors are clustered for workers with the same variables we match on exactly to account for the replacement (see Appendix 2.11 for details).

Table 2.5: Effect of OOI on Wages Using Shift-Share (Bartik) Instrument

	First-Stage (1)	Reduced-Form (2)	2SLS (3)
	.622***	.106***	.170***
	(.241)	(.056)	(.064)
Industry FE	X	X	X
N	408,792	408,792	408,792
Number of Clusters	38	38	38

Notes: This table shows the results of the impact of a shift-share instrument (Bartik) on outside options, and wages (column 1 and 2). This captures the effect of changes in OOI on changes in log wages between 2004-2014, when we instrument for the changes in OOI with the shift-share instrument (column 3). The instrument is constructed from an average of a 3-digit industry national employment growth weighted by the initial share of every industry in a region (see Section 2.6.2). The outcome variables are the change in OOI 2004-2014 (column 1) and change in log daily wages (columns 2 and 3). All columns control for industry (in 2004) and age. Standard errors are clustered within the unit of treatment, which is regions.

Table 2.6: Shift-Share (Bartik) Results by Exporting Share of Sales

	Export>33% (1)	33%≥Export≥1% (2)	1%>Export (3)
OOI	.105**	.593**	.132
	(.052)	(.266)	(.141)
Industry FE	X	X	X
N	119,645	146,217	142,930

Notes: This table shows the results of the impact of OOI on wages, instrumented with a shift share instrument, calculated separately by share of export in the industry. Share of export is calculated for every 3-digit industry based on the establishment panel survey in 2014. The sample is split based on the worker industry in 2004. Outcome variable is change in log wages between 2004-2014. The dependent variable is change in OOI between 2004-2014. The instrument is constructed from an average of a 3-digit industry national employment growth weighted by the initial share of every industry in a region (see Section 2.6.2). All columns control for industry (in 2004), and age. Standard errors are clustered within the unit of treatment, which is regions.

Table 2.7: Shift-Share (Bartik) Results by Stayers and Movers

	(1)	(2)
OOI	.170***	.257***
	(.064)	(.092)
$OOI\times$		-.159***
$Stay$		(.062)
Industry FE	X	X
N	408,792	408,792

Notes: This table shows the results of the impact of OOI on wages, instrumented with a shift share instrument, interacted with whether a worker stayed at the same establishment. Outcome variable is change in log wages between 2004-2014. The dependent variable is change in OOI between 2004-2014. The instrument is constructed from an average of a 3-digit industry national employment growth weighted by the initial share of every industry in a region (see Section 2.6.2). The indicator for stay is 1 if the worker works at the same establishment on June 30th of both 2004 and 2014. All columns control for industry (in 2004), and age. Standard errors are clustered within the unit of treatment, which is regions.

Table 2.8: Mincer Equation with OOI

	OLS		OLS		Tobit		Tobit	
	(1)		(2)		(3)		(4)	
Female	-0.171	***	-0.111	***	-0.195	***	-0.137	***
	(0.01)		(0.01)		(0.01)		(0.01)	
Education: Lower-Secondary	-0.351	***	-0.174	***	-0.404	***	-0.23	***
	(0.01)		(0.01)		(0.01)		(0.01)	
Education: Intermediate	-0.245	***	-0.17	***	-0.289	***	-0.217	***
	(0.01)		(0.01)		(0.01)		(0.01)	
Non-German Citizen	0.089	***	0.007		0.093	***	0.011	
	(0.03)		(0.03)		(0.03)		(0.03)	
Age	0.057	***	0.03	***	0.061	***	0.035	***
	(0.00)		(0.00)		(0.00)		(0.00)	
Age^2	-0.573	***	-0.261	***	-0.608	***	-0.297	***
	(0.04)		(0.04)		(0.04)		(0.04)	
District Density	0.022	***	-0.008	***	0.023	***	-0.008	***
	(0.00)		(0.00)		(0.00)		(0.00)	
Part-Time	-0.913	***	-0.905	***	-0.928	***	-0.921	***
	(0.01)		(0.02)		(0.02)		(0.02)	
Observations	378776		378776		378776		378776	

Notes: This table shows the results from a regression of log wages on demographics and OOI. The coefficient of the OOI is fixed to be its point estimate from the 2SLS estimate based on the high-speed commuter rail introduction (Table 2.4). A Tobit model is used in Columns 3-4 to account for top coding of daily income at 195 Euros per day. OLS results use winsorized log income. Sampling weights are used to make this a representative sample of the German population

2.10 Appendix Figures and Tables

Figure 2-16: Relative Income Following Mass-Layoff

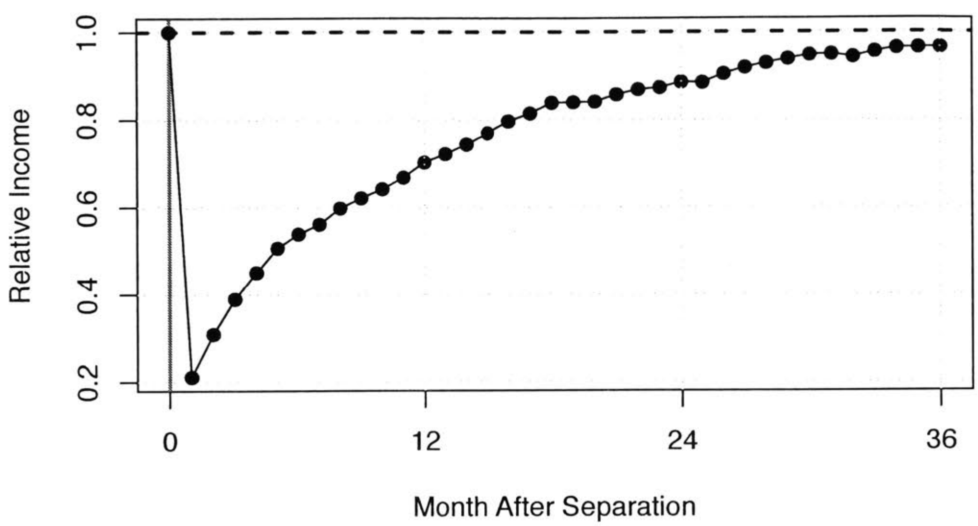

Note: This figure shows the relative income for workers who lost their jobs in mass-layoffs, for each month in the three years after the layoff. Relative income is defined as the current daily income in that month divided by the last daily income before the layoff. Mass layoffs are defined as an establishment with at least 50 workers that reduced its workforce by at least 30% in a given year. The sample includes only workers who have worked for at least three years before the layoff and are below the age of 55. The values are calculated using a regression of relative income on months after separation, with a fixed effect for every mass-layoff (Equation 2.13).

Table 2.9: Most Weighted Question in PCA - Establishment 2014 Survey

Name	N	Comp 1	Comp 2
Business Performance	8,792	Member of chamber of industry	Profit category
Investment & Innovation	8,792	IT investment	Total investment
Hours	8,792	Long leaves policy	Flextime
In-Company Training	8,792	Internal courses	Share workers in training
Vocational Training	8,792	Offer apprenticeship	Ability to fill
General	8,792	Family managed	Staff representation

This table shows the survey question that received the most weight in this principal component. We take the first two principal component from each survey category.

Table 2.10: Most Weighted Question in PCA - BIBB

Name	N	Comp 1	Comp 2
Hours	11,021	Sundays and public holidays	hours per week like to work
Type of Task	15,035	responsibility for other people	Cleaning, waste, recycling
Requirements	10,904	Acute pressure & deadlines	Highly specific Regulations
Physical	20,036	Oil, dirt, grease, grime	pathogens, bacteria
Mental	17,790	Support from colleagues	Often missing information

This table shows the survey question that received the most weight in this principal component. We take the first two principal component from each survey category.

Table 2.11: Commuting Distance by Gender and Education

	Distance from Job (Miles) Mean	SD
All	15.5	41.9
Female	12.1	37.1
Male	17.4	44.3
Lower-Secondary	9.4	27.9
Intermediate-Secondary	11.4	34.4
Higher-Secondary	26.2	56.1

Values are mean distance in miles between workers previous place of residence and their job.

Table 2.12: Summary Stats for Mass-Layoff Workers by Treatment Status

	Above Median OOI		Below Median OOI	
	Mean	SD	Mean	SD
Female	.36	.48	.44	.50
Age	40.0	9.7	37.2	11.3
Higher-Secondary Education	.21	.41	.14	.35
Tenure in Establishment (days)	2316.3	1272.3	2167.4	1197.8
Daily Income	63.8	43.1	57.5	42.2
N	6.839		6,887	

Note: This table shows the summary stats for workers that lost their jobs in a mass-layoff above and below the establishment median OOI. Mass layoffs are defined as an establishment with at least 50 workers that reduced its workforce by at least 30% in a given year. The sample includes only workers who have worked for at least three years before the layoff and are below the age of 55. We include monthly income for the 36 months following the separation.

Table 2.13: Correlation Between OOI and log $wage$ - OLS Results

	Dep.Var: log wage$_i$		
OOI	.107***	.027***	-.010*
	(.005)	(.005)	(.005)
Demographics		X	X
District FE			X
N		378,776	

Note: Demographics include gender, education group, a quadratic in age and citizenship status. The sample includes all workers employed on June 30th 2014. Sampling weights are used to make this a representative sample of the German population.

Table 2.14: Balance Table - High Speed Train

	Treatment		Control	
	Mean	SD	Mean	Sd
log $wage$ (1993)	3.26	2.49	3.27	2.49
log $wage$ (1999)	4.25	.62	4.27	.58
Female	.356	.479	.356	.479
Age	36.4	6.8	36.4	6.7
Citizen	.995	.073	.995	.073
Low-Secondary	.257	.437	.257	.437
Intermediate-Secondary	.508	.500	.508	.500
High-Secondary	.235	.424	.235	.424
N	37,695		26,963	

Note: This table shows the summary stats for workers used to estimate the impact of high-speed trains on OOI and log wages. Treated group includes workers who lived in districts in which a new station was introduced between 1999-2012. Control group was chosen from a pool of workers living in districts that never got a station. Control workers were chosen through nearest-neighbor matching with replacement on gender, age, citizenship, education level, training occupation, state (Bundesländer) and lagged income. We require the match to be exact on gender, education, state and 2-digit occupation.

Table 2.15: Top Standardized Values of A

Variable (X)	Variable (Z)	A_{xz}
Distance		-4.15
Train Occ - Physical Cond. 1	Occ - Physical Cond. 1	1.477
Train Occ - Task Type 2	Occ - Task Type 2	1.077
Train Occ - Task Type 2	Occ - Physical Cond. 1	-.93
Train Occ - Physical Cond. 1	Occ - Task Type 2	-.82
Lower Secondary Education	Distance	-.74
Intermediate Education	Distance	-.61
Train Occ - Contract 2	Occ - Contract 2	.56
Train Occ - Task Type 1	Occ - Task Type 1	.55
Lower/Intermediate Education	Distance	-.54

Results from logistic regression for dummy variable on real vs. simulated match, on interaction of worker and job characteristics (Equation 2.10). Results are standardized, such that each variable has standard deviation of 1.

Table 2.16: Distance Coefficient by Demographics

Baseline	-.141
Female	-.024
Non-Citizen	-.019
Lower-Secondary	-.037
Intermediate-Secondary	-.012
Age	.002
Age^2 ($\times 10^{-3}$)	-.026

Results from logistic regression for dummy variable on real vs. simulated match, on interaction of worker and job characteristics (Equation 2.10). Baseline category is a forty years old high-secondary male.

2.11 Theoretical Appendix

Continuous Logit Distribution

We follow Dagsvik (1994) in defining the continuous logit that produces ε_{i,z_j} and ε_{j,z_i}. In this section we define the distribution of ε_{i,z_j} and the distribution of ε_{j,x_i} is defined similarly.

Every worker $i \in \mathcal{I}$ draws ε_{i,z_j} shocks from a Poisson process on $\mathcal{Z} \times \mathbb{R}$ with intensity

$$f(z)\, dz \times e^{-\epsilon} d\epsilon$$

This is different from the Poisson process used in Dupuy and Galichon (2014) as the density $f(z)$ also affects the intensity, which allows this distribution to be properly defined over a larger class of functions for $\tau(x, z)$, including a constant, or simple polynomials. Denoting by P_i the infinite but countable points chosen in the process, every worker has a set

$$\left\{\varepsilon_{iz_j} = \alpha\epsilon \middle| (z, \epsilon) \in P_i\right\}$$

This process yields a distribution of ε_{i,z_j} that has several similarities to finite extremum value type-1 distribution. These similarities are all derived from one basic property of this point process.

Proposition 9. *Let $g : \mathcal{Z} \to \mathbb{R}$ be a function that satisfies*

$$\int_{\mathcal{Z}} e^{g(z)} f(z)\, dz < \infty$$

and let $S \subseteq \mathcal{Z}$ be some Borel measurable subset. Define

$$\psi_S^g = \max_{z \in S \cap P_Z} \left\{g(z) + \varepsilon_{i,z_j}\right\}$$

Then

$$\psi_S^g \sim EV_1\left(\alpha \log \int_S \exp \frac{g(z)}{\alpha} f(z)\, dz, \alpha\right)$$

and

$$S_1 \cap S_2 = \phi \iff \psi_{S_1}^{g_1} \perp \psi_{S_2}^{g_2}$$

Proof. This proposition stems from the fact that in a Poisson process, the amount of points chosen in two disjoint Borel measurable sets B_1, B_2 has an independent distribution $N(B_i) \sim Poisson(\Lambda(B_I))$ with

$$\Lambda(B_i) = \int_{B_i} \lambda(x)\, dx$$

Therefore, in our context the cumulative distribution function of ψ_S^g is

$$P\left(\psi_S^g \le x\right) = P\left(N\left((S \times \mathbb{R}) \cap \{g(z) + \alpha\epsilon > x\}\right) = 0\right)$$

From the Poisson distribution this is

$$\begin{aligned}
\log P\left(\psi_S^g < x\right) &= -\Lambda\left(S \times \left\{\epsilon > \tfrac{x-g(z)}{\alpha}\right\}\right) \\
&= -\int_S \int_{\frac{x-g(z)}{\alpha}} f(z)\, e^{-\epsilon} dz d\epsilon \\
&= -\int_S e^{-\frac{x-g(z)}{\alpha}} f(z)\, dz \\
&= -\exp\left[-\frac{x - \alpha \log \int_S \exp \frac{g(z)}{\alpha} f(z) dz}{\alpha}\right]
\end{aligned}$$

which is exactly a cumulative distribution function of $EV_1\left(\alpha \log \int_S \exp \frac{g(z)}{\alpha} f(z)\, dz, \alpha\right)$.

Since every draw of points in a Poisson process is independent, $S_1 \cap S_2 \iff \psi_{S_1}^{g_1} \perp \psi_{S_2}^{g_2}$. $\square$

This Proposition has several important implications for our context. It implies that even though ε_{i,z_j} is not defined for every $z \in \mathcal{Z}$, it is defined infinitely often for every Borel measurable subset that includes z, and the maximum for that set ψ_S^1 has an extreme-value type-1 distribution.

Since workers in equilibrium are getting a sum of a continuous function (which we mark by $\omega(x, z)$) and ε_{i,z_j} (Lemma 5), then we get that the maximum value they receive also has an EV_1 distribution, for every Borel measurable set of jobs. Moreover, the probability density to choose a particular observables z_j is similar to the finite case and its exact value is

$$f(z_j|i) = f\left(z_j = \arg\max_{z\in S\cap P_Z}\left\{g(z)+\varepsilon_{i,z_j}\right\}\right) = \frac{\exp\left[\frac{1}{\alpha}\omega(x,z_j)\right]f(z_j)}{\int_{\mathcal{Z}}\exp\left[\frac{1}{\alpha}\omega(x,z)\right]f(z)\,dz}$$

Another link to the finite multinomial logit can be drawn if we divide $\mathcal{Z}$ into a finite number of disjoint sets $\mathcal{Z} = \bigcup_{i=1}^{n} S_i, S_i \cap S_j = \phi$. Then the value of the best job for worker i in each subset $(\psi_{S_i}^{\omega})$ is EV_1 distributed. The choice of the best job characteristics $z_{m(i)}$ would be made with a finite multinomial logit, over these n options. When we increase n, the sets become smaller, and the choice becomes closer to an infinite options choice.

Note that in a standard multinomial logit, increasing the number of options to infinity will yield an infinite compensation, but this is not the case here. This is because when the number of options n grow, the mean measure of S_i decreases in a rate of $\frac{1}{n}$. Therefore the location parameter of each one of the choices, decreases in a rate of $\frac{1}{n}$ as well from the proposition.

Proofs

Proof of Lemma 4

Part 1: We will start by formally defining the densities we are using. We will use $\mathcal{I}^{\phi}, \mathcal{J}^{\phi}$ to mark the set of unmatched workers and jobs.

Definition. *Let $f(i,j) : \mathcal{I}\times\mathcal{J} \to \mathbb{R}_{\geq 0}$ be the density that satisfies for every Borel measurable subset of potential matches $B \subseteq \mathcal{I}\times\mathcal{J}$*

$$\int_B f(i,j)\,didj = \frac{\mu(B\cap M)}{\mu(M)+\mu(\mathcal{I}^{\phi})+\mu(\mathcal{J}^{\phi})}$$

where μ is the measure function.

Intuitively, this is the joint density of observing i and j matched in equilibrium. Similarly we define a density over the probability of observing worker and job with specific characteristics matched in equilibrium.

Definition. *Let $f(x,z) : \mathcal{X} \times \mathcal{Z} \to \mathbb{R}_{\geq 0}$ be the density that satisfies*

$$f(x,z) = \int_{X_i = x} \int_{Z_j = z} f(i,j)\, didj$$

From these definitions we can derive the conditional distribution of a match for a given worker.

Definition. *Let f_j^i be*

$$f_j^i = f(j|i) = \frac{f(i,j)}{f(i)} = \frac{f(i,j)}{I^{-1}}$$

Part 2: Let $i, i' \in \mathcal{I}$, $j, j' \in \mathcal{J}$ with $X_i = X_{i'}$ and $Z_j = Z_{j'}$ From Assumption 1 τ_{ij} has the same distribution as $\tau_{i'j'}$, and therefore $f(i,j) = f(i',j')$.

Hence, from Definition 2.11,

$$f(X_i, Z_j) = If(X_i)\, Jf(Z_j)\, f(i,j)$$

and from Definition 2.11

$$f_j^i = \frac{f(X_i, Z_j)}{f(X_i)\, f(Z_j)} J^{-1}$$

and we normalized $J = 1$.

Proof of Lemma 5

Let $i, i' \in \mathcal{I}$ with $X_i = X_{i'} = x_0$ and $j, j' \in \mathcal{J}$ with $Z_j = Z_{j'} = z_0$, where $m(i) = j$ and $m(i') = j'$. The sum of compensation equals the total surplus, hence

$$\omega_{ij} + \pi_{ij} = \tau(x_0, z_0) + \varepsilon_{i,z_0} + \varepsilon_{j,x_0}$$

$$\omega_{i'j'} + \pi_{i'j'} = \tau(x_0, z_0) + \varepsilon_{i',z_0} + \varepsilon_{j',x_0}$$

For stability, it must be that

$$\omega_{ij} + \pi_{i'j'} \geq \tau(x_0, z_0) + \varepsilon_{i,z_0} + \varepsilon_{j',x_0}$$

$$\omega_{i'j'} + \pi_{ij} \geq \tau(x_0, z_0) + \varepsilon_{i',z_0} + \varepsilon_{j',x_0}$$

Note that the sum of the two weak-inequalities is equal to the sum of the two equalities, therefore they must hold with equality (otherwise, the sum should hold both as an equality and strong inequality). Hence, we can rewrite

$$\omega_{ij} - \omega_{i'j'} = \varepsilon_{i,z_0} - \varepsilon_{i',z_0}$$

$$\pi_{ij} - \pi_{i'j'} = \varepsilon_{j,x_0} - \varepsilon_{j',x_0}$$

In other words, compensation for workers and employers in matches with the same characteristics is constant up to their value of ε, so we can write

$$\omega_{ij} = \omega(x_0, z_0) + \varepsilon_{i,z_0}$$

$$\pi_{ij} = \pi(x_0, z_0) + \varepsilon_{j,x_0}$$

$$\omega(x_0, z_0) + \pi(x_0, z_0) = \tau(x_0, z_0)$$

We can also pin down the alternative offers

$$\omega_{ij'} = \tau_{ij'} - \pi_{j'} = \tau(x_0, z_0) + \varepsilon_{i,z_0} + \varepsilon_{j',x_0} - \pi(x_0, z_0) - \varepsilon_{j',x_0} = \omega(x_0, z_0) + \varepsilon_{i,z_0} = \omega_{ij}$$

This implies that all employers with $Z_j = z_0$ who are matched with $X_{m^{-1}(j)} = x_0$ are willing to make the same offer. Therefore, both workers and employers are facing a continuous logit choice. Hence, we can link the values of $\omega(x_0, z_0)$ and $\pi(x_0, z_0)$ to their choice probabilities (see Appendix 2.11):

$$f(x_0|z_0) = \frac{\exp\left[\frac{1}{\alpha}\pi(x_0, z_0)\right] f(x_0)}{\int_{\mathcal{X}} \exp\left[\frac{1}{\alpha}\pi(x, z_0)\right] f(x)\, dx}$$

The denominator is the expected value π_j, which is a function of $Z_j = z_0$ so we can rewrite it as

$\pi(z_0)$. Taking logs we get

$$\alpha \log f(x_0|z_0) = \pi(x_0, z_0) + \log f(x_0) - \pi(z_0)$$

and with Lemma 4

$$\pi(x_0, z_0) = \alpha \log f_j^i + \alpha \log J + \pi(z_0)$$

therefore

$$\omega_{ij'} = \tau(x_0, z_0) - \pi(z_0) + \alpha \log f_j^i + \varepsilon_{i,z_0}$$

where J was pinned to 1.

Proof of Lemma 6

First equality is by definition, and because the first best and second best options are equivalent. We showed that $\omega_{ij} = \omega(x_i, z_j) + \varepsilon_{i,z_j}$. The expected compensation of worker i is

$$\omega(x_i) = E[\omega^*(x_i, z_j)] + E\left[\varepsilon^*_{i,z_j}\right]$$

From the continuous logit structure we know that (similar to the previous proof)

$$\omega(x_i, z_j) = \alpha \log f_j^i + \omega(x_i)$$

hence

$$\omega(x_i) = E\left[\alpha \log f_j^i + \omega(x_i)\right] + E\left[\varepsilon^*_{i,z_j}\right]$$

Therefore

$$E\left[\varepsilon^*_{i,z_j}\right] = -\alpha \int f_j^i \log f_j^i dj$$

Similarly for ε_{j,x_0} and combinedly:

$$E\left[\varepsilon^*_{i,z_j} + \varepsilon^*_{j,x_0}\right] = -\alpha \int f^i_j \log f^i_j dj$$

Proof of Theorem 5

Following the notations from the previous proofs. The ω_{ij} offer can be written as

$$\omega_{ij} = \omega(x_i, z_j) + \varepsilon_{i,z_j}$$

and ε_{i,z_j} is unaffected by λ hence

$$\frac{d\omega_{i,j}}{d\lambda_i} = \frac{d\omega(x_i, z_j)}{d\lambda_i}$$

In the previous proofs we showed that

$$\omega(x_i, z_j) = \alpha \log f^i_j + \omega(x_i)$$

$$\pi(x_i, z_j) = \alpha \log f^i_j + \pi(z_j)$$

hence

$$\omega(x_i, z_j) - \pi(x_i, z_j) = \omega(x_i) - \pi(z_j)$$

Adding $\tau(x_i, z_j)$ and dividing by 2:

$$\omega(x_i, z_j) = \frac{1}{2}\left(\tau(x_i, z_j) - \pi(z_j) + \omega(x_i)\right)$$

$\tau(x_i, z_j), \pi(z_j)$ don't change by the definition of λ. Hence the only effect is on $\omega(x_i)$.

$$\frac{d\omega_{i,j}}{d\lambda_i} = \frac{1}{2}\frac{d\omega(x_i)}{d\lambda_i}$$

We get the value for $\omega(x_i)$ from the decomposition in Equation 2.6. Since $\tau\left(x_i, z_{j'}\right), \pi\left(z_{j'}\right)$

remain constant the remaining effect is on the OOI.

$$\frac{d\omega_{i,j}}{d\lambda_i} = \alpha \frac{dOOI_i}{d\lambda_i}$$

Proof of Theorem 8

This is similar to before, only that ε_{i,z_j} is allowed to change as well. Since $E\left[\varepsilon_{i,z_j}\right] = \alpha OOI$ we get the effect from the previous lemma, in addition to the effect on the OOI.

$$\frac{d\omega_{i,j}}{d\lambda_i} = 2\alpha \frac{dOOI_i}{d\lambda_i}$$

Alternative Definitions for λ

Assume workers and equally distributed across the real line (as in Section 2.2.5). Each worker is a 3-dimensional tuple (l_i, y_i, c_i) and τ_{ij} is defined as

$$\tau_{ij} = y_i - c_i \left| l_i - l_j \right| + \varepsilon_{ij}$$

Now workers log density is a triangular function, with its peak at l_i (Laplace). Hence,

$$f_j^i = \frac{c_i}{2} \exp -c_i \left| l_i - l_j \right|$$

The OOI is (shifting l_i to 0)

$$\int_0^\infty c \exp^{-cl} \left(\log \frac{c}{2} - cl \right) dl = \log \frac{c}{2} - 1$$

The mean value $E\left[\tau\left(x, z\right)\right]$ is

$$\int_0^\infty c \exp^{-cl} \left(y_i - cl \right) = y_i - 1$$

Hence, setting the commuting cost c only affects the OOI but not net productivity and can be served as λ. This will also work for more general settings, as long as worker and job locations are not correlated with locations.

Another example is to define λ as the intensity of the Poisson process for the continuous logit process. Higher λ will mean more options on average in every subset of jobs.

$f(x,z)$ Estimation

To estimate a logistic regression following Equation 2.10, we maximize the following likelihood

$$\max_{\theta} \sum_{k} \log P\left(y_k | x_k, z_k; \theta\right)$$

where θ are the parameters defined in this equation, including matrix A. We rewrite Equation 2.10 in a more general form. Note $p_k\left(\theta\right) = P\left(Y_k = 1 | X = x_k, Z = z_k\right)$:

$$\log \frac{p_k\left(\theta\right)}{1 - p_k\left(\theta\right)} = \sum_{j=1}^{K} \beta_j h_j\left(x_k, z_k\right)$$

where K is the number of moments h_j we control for in this regression.

Then the K FOC of this maximization converge asymptotically to

$$E\left[p_k\left(\theta\right) h_j\left(x_k, z_k\right)\right] = E\left[h\left(x_k, z_k\right) | y_k = 1\right] s$$

where $s = P\left(Y = 1\right)$ is the share of real data (in our case $\frac{1}{2}$). Using $\frac{p_k(\theta)}{1-p_k(\theta)} = \frac{f(x,z)}{f(x)f(z)} \frac{s}{1-s}$ we can write

$$E\left[\frac{f\left(x,z\right)}{sf\left(x,z\right) + \left(1-s\right) f\left(x\right) f\left(z\right)} h_j\left(x,z\right)\right] = E\left[h_j\left(x,z\right) | real\right]$$

The RHS is simply the moment of $h_j\left(x,z\right)$ in the real data. The LHS is the moment of $h_j\left(x,z\right)$ in the full data (real and simulated), weighted by the probability it is real.

If the model is correctly specified and the functional form assumption on $\frac{f(x,z)}{f(x)f(z)}$ is true, θ will be estimated consistently. This is because

$$E\left[\frac{f(x,z)}{sf(x,z)+(1-s)f(x)f(z)}h_j(x,z)\right]=$$

$$\int\frac{f(x,z)}{sf(x,z)+(1-s)f(x)f(z)}h_j(x,z)\left(sf(x,z)+(1-s)f(x)f(z)\right)dxdz=$$

$$=\int h(x,z)\,f(x,z)\,dxdz=E\left[h(x,z)\,|real\right]$$

If the model is misspecified, our estimate of $\frac{f(x,z)}{sf(x,z)+(1-s)f(x)f(z)}$ will not be converging to the real density rations. Instead we will equalize moments of some other weighted average of h_j

$$E\left[w(x,z,\theta)\,h_j(x,z)\right]=E\left[h_j(x,z)\,|real\right]$$

where

$$w(x_k,z_k,\theta)=s^{-1}\frac{\exp\sum_{j=1}^{K}\beta_j h_j(x_k,z_k)}{1+\exp\sum_{j=1}^{K}\beta_j h_j(x_k,z_k)}$$

We next analyze these weights as $s\to 0$. We will mark $h_1(x,z)=1$, the offset of the regression. When $s\to 0$, $\frac{p(\theta)}{1-p(\theta)}\to 0$ as well, therefore $\exp\sum_{j=1}^{K}\beta_j h_j(x_k,z_k)\to 0$. With some abuse of notation, we will redefine β_1 as $\beta_1-\log s$. Therefore

$$\lim_{s\to 0}w(x,z,\theta)=\exp\sum_{j=1}^{K}\beta_j h_j(x_k,z_k)=\frac{\widehat{f(x,z)}}{f(x)\,f(z)}$$

The density of the full data approaches the density of the simulated data. Hence overall, we get

$$E\left[w(x,z,\theta)\,h_j(x,z)\,|sim\right]=E\left[h_j(x,z)\,|real\right]$$

In order to calculate the OOI, we simulate values from $f(x)\,f(z)$, and reweight them based on $\frac{\widehat{f(x,z)}}{f(x)f(z)}$. This is because we hold workers fixed, and simulate z values from $f(z)$. As $s\to 0$ we use weights that converge to $w(x,z,\theta)$. The above equation guarantees that we sample from a distribution with same moment value for every $h_j(x,z)$, even if the model is misspecified.

Dupuy and Galichon (2014) produce a distribution with the same second moments as the data, and same marginal distributions. Therefore, when $s \to 0$, and h include all X, Z interactions, and an indicator for every x_k, and every z_k value (that is, $h(x,z) = 1_{x=x_k}$ or $h(x,z) = 1_{z=z_k}$ for every k), we get the same distribution.

Standard Errors for a Wald Estimator with Matching

We want to estimate the standard errors of $\widehat{\alpha}$, defined in Equation 2.17. Both the nominator (reduced form), and the denominator (first stage) are standard matching estimators for average treatment effect on treated (ATET). Abadie and Imbens (2006) show how to estimate standard errors for ATET. But to estimate correctly the standard error for the Wald estimator, we also need to estimate the covariance of the first stage and reduced form. So we extend their approach for this case.

Mark the ATET on log wages (reduced form), and OOI (first stage) as:

$$\begin{aligned} \rho &= E\left[\log w(1) - \log w(0) \,|T = 1\right] \\ \gamma &= E\left[OOI(1) - OOI(0) \,|T = 1\right] \end{aligned}$$

where (1) means value when treated and (0) when not treated. T is treatment status (so this is the mean effect for treated).

The Wald estimator is then

$$\alpha = \frac{\rho}{\gamma}$$

For each match m (treated unit and one or more control unit), define

$$X_m = \begin{pmatrix} X_{1m} \\ X_{2m} \end{pmatrix} = \begin{pmatrix} \log w(1) - \overline{\log w(0)} \\ OOI(1) - \overline{OOI(0)} \end{pmatrix}$$

Our estimators are then simply

$$\begin{pmatrix} \widehat{\rho} \\ \widehat{\gamma} \end{pmatrix} = \overline{X_m}$$

Asymptotically

$$\overline{X_m} \sim N\left(\begin{pmatrix} \rho \\ \gamma \end{pmatrix}, \begin{pmatrix} \sigma_{11} & \sigma_{12} \\ \sigma_{12} & \sigma_{22} \end{pmatrix}\right)$$

With the Delta method

$$V\left(\widehat{\alpha}\right) = \frac{1}{\gamma^2}\left(\sigma_{11} - 2\frac{\rho}{\gamma}\sigma_{12} + \frac{\rho^2}{\gamma^2}\sigma_{22}\right)$$

Abadie and Imbens (2006) tells us how to find σ_{11}, σ_{22} which are $V\left(\widehat{\rho}\right), V\left(\widehat{\gamma}\right)$. We want to extend their approach to σ_{12}.

The challenge in getting the variance correctly for matching with replacement, is that the matches are not independent. Some observations from the control pool appear in more than one match. Following Abadie Imbens we write

$$V\left(\overline{X_m}\right) = \frac{1}{N_1}\sum_m \left(X_m - \overline{X_m}\right)^T \left(X_m - \overline{X_m}\right) + \frac{1}{N_1}\sum_{T=0} \left(K_i\left(K_i - 1\right)\right)\widehat{V_i}$$

with

$$\widehat{V_i} = \widehat{V}\begin{pmatrix} \log w_i \\ OOI_i \end{pmatrix}$$

where N_1 is number of treated units,K_i is the number of times observation i from the control pool was used. $\widehat{V_i}$ is a 2x2 matrix of the variance for that particular observation. The first part is a standard variance calculation. The second part corrects for the covariance between the matches.

If an observation i is used $K_i > 1$ times, then there are $K_i\left(K_i - 1\right) > 0$ pairs of matches that both use it, and so their covariance is not 0, but includes $\widehat{V_i}$.

To estimate $\widehat{V_i}$ we follow Abadie and Imbens (2006) and use nearest neighbor from the control group. So for every control observation we find a match from the control group as well and write

$$\widehat{V_i} = \frac{1}{2}\begin{pmatrix} \log w_i - \log w_{m(i)} \\ OOI_i - OOI_{m(i)} \end{pmatrix}\begin{pmatrix} \log w_i - \log w_{m(i)} \\ OOI_i - OOI_{m(i)} \end{pmatrix}^T$$

This is asymptotically unbiased.

In practice, the only difference from Abadie and Imbens (2006) is that we also have a covariance component.

$$COV\left(\widehat{\rho}, \widehat{\gamma}\right)$$

Which we estimate with

$$\frac{1}{N_1}\sum\left(\log w_m\left(1\right) - \overline{\log w_m\left(0\right)}\right)\left(OOI_m\left(1\right) - \overline{OOI_m\left(0\right)}\right) - \widehat{\rho}\widehat{\gamma} +$$

$$2 * \frac{1}{2} * \sum_{T=0} K_i\left(K_i - 1\right)\left(\log w_i - \log w_{m(i)}\right)\left(OOI_i - OOI_{m(i)}\right)$$

If our two variables were the same ($\log w = OOI$) then this would be the standard Abadie and Imbens (2006) formula for variance, as expected.

2.12 Data Appendix

LIAB

In this section we clarify the coding of some of the variables we use. Our panel data, allows us to observe some variables several times in the data, and correct for coding errors. In particular, we set German citizenship to one, if this worker was ever reported as a German citizen by her employer.

We also take the highest level of education we observed until every year. All upper secondary school certificates are coded as upper-secondary. In some years intermediate and lower secondary education are coded with the same value. In these cases, if we observe the worker in other years and can infer their schooling level we use that. Otherwise, we code these workers in a separate category for either lower or intermedia secondary education.

For training occupation, we use the occupation in which workers spent the longest time in training. The LIAB data specify whether a worker is in vocational training and their occupation. For the large majority of workers, there is only one occupation in which they perform their vocational training. In rare cases where workers have conducted training in more than one occupation, we use the occupation in which the training was longer. If the we never observe the worker during vocational training, we take the occupation in which they conducted an internship. If this is unobserved as well, we use the first occupation they were observed in, as long as at least ten years have passed since we first observed them.

We calculate distance at the district level. For each district, we calculate the district center, by taking the weighted average of the latitude and longitude coordination of each city in this district. We then calculate the distance between the districts, taking into account the concavity of the earth.

BIBB Survey

In this section we describe in more detail the BIBB survey and PCA analysis.

We use data from the 2011-2012 wave of the German Qualification and Career Survey conducted by the Federal Institute of Vocational Training (BIBB) and the Institute for Labor Market Research (IAB). The data cover 20,000 employed individuals between the ages of 16 and 65. We run PCA on this survey by questions category and aggregate the results by 2-digit industry and 3-digit occupations. We link the results to our main data. The top question in each category are shown in Table 2.10.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

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

JOINT WITH EMILY OEHLSEN[1]

"Perfect discrimination is probably rare in buying labor but imperfect discrimination may often be found. For instance there may be two types of workers (for example, men and women or men and boys) whose efficiencies are equal but whose conditions of [labor] supply are different. It may be necessary to pay the same wage within each group, but the wages of the two groups (say of men and of women) may differ.

— JOAN ROBINSON *(1933)*

[1]This paper benefitted from feedback from David Autor, David Card, Oren Danieli, Joshua Dean, Ellora Derenoncourt, Jonathan Hall, Dan Knoepfle, Elizabeth Mishkin, Suresh Naidu, Elizabeth Setren, David Silver, Kane Sweeney, Alice Wu, and Roman Andrés Zárate. This paper also benefitted from comments at the 2018 ASSA annual meeting, the MIT labor lunch, and presentations to the Boston and Houston Uber city teams. Phoebe Cai and Anran Li provided outstanding research assistance. The views expressed here are those of the authors and do not necessarily reflect those of Uber Technologies, Inc. Caldwell's work on this project was carried out under a data use agreement executed between MIT and Uber. Oehlsen is a former employee of Uber Technologies, Inc. This study is registered in the AEA RCT Registry as trial no. AEARCTR-0001656. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. 1122374 (Caldwell) and by the National Science Foundation Dissertation Improvement Grant No. 1729822 (Caldwell).

3.1 Introduction

Recent research has suggested that imperfect competition in the labor market may have a meaningful impact on wages for workers throughout the skill distribution (see, e.g. Card et al., 2013; Dube et al., 2017). 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 et al., 2010; Boal and Ransom, 1997; Bhaskar et al., 2002; Manning, 2003b). Firms have an incentive to pay higher wages to workers that are harder to recruit or retain, even if they are no more productive than other workers.

The idea that this imperfect competition could lead to a gender wage gap dates back to Joan Robinson's 1933 book, in which she coined the term monopsony.[2] 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., 2016c).[3] This can happen if women are more loyal to their employers (i.e. have higher average switching costs), have less information about their outside labor market opportunities, have different valuations for employer-provided amenities, or face smaller effective labor markets due to different commuting costs (Babcock and Laschever, 2009; Manning, 2011). However, without exogenous variation in the wages provided by a single firm it is difficult to produce credible measures of firm-specific elasticities, or to test whether these elasticities differ by gender.

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, both to individual firms and to the market. We also test whether gender-differences in firm-specific elasticities might contribute to a gender wage gap. These experiments offered random subsets of male and female drivers the opportunity to drive with higher wages. While some drivers had access to a competing ride-share company, others did not. We use data on drivers unable to drive for a competing platform to identify Frisch

[2]We thank David Card for pointing out correspondence that reveals Joan Robinson asked B.L. Hallward (a classicist) to coin the term. She credits him in her book.

[3]Similarly, search models predict that workers with lower arrival rates of job offers earn less in equilibrium (Black, 1995).

elasticities for both men and women. We identify firm substitution elasticities by comparing these Frisch elasticities to the elasticities of drivers who could drive for another ride-sharing firm. We show that, in a very simple monopsony model, these elasticities are sufficient to calculate the firm's optimal gender wage gap.

Our analysis starts with a theoretical model that allows workers to adjust both how much they work (participation and hours) and for whom they work (firm substitution). The model illustrates that when hours are flexible, the amount of monopsony power in the market 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. Most prior work on monopsony has focused on the substitution elasticity, ignoring the elasticity of workers' hours to the market; most prior work on labor supply has ignored the role of firm substitution.

We use data from a randomized experiment conducted when Uber faced little competition to provide experimental estimates of the Frisch elasticity for men and women. These elasticities serve as a baseline for our analysis of firm substitution: we can assess the degree of cross-platform shifting by contrasting these elasticities with those estimated in a market where some drivers could work for Uber's main competitor. They are also of independent interest as they are a key component of most business cycle models (King and Rebelo, 1999). These elasticities govern how labor supply (and thus output) respond to shocks to productivity.

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). This reflects the fact that it is difficult to find the type of wage variation necessary to identify Frisch elasticities: variation that is both temporary and exogenous.[4] While a few studies have exploited temporary wage variation in settings where workers can freely choose their hours, the populations in these

[4]In particular, most tax changes do not satisfy the second requirement. The tax holiday studied in Martinez et al. (2018) is a notable exception.

studies are predominantly male (Oettinger, 1999; Farber, 2005; Fehr and Goette, 2007; Farber, 2015; Stafford, 2015). Though most (more than 85%) of Uber drivers are male, we structured our experiment to include roughly equal numbers of male and female drivers (Hall and Krueger, 2015).[5]

We offered random samples of drivers the opportunity to drive for one week with 25-39% higher hourly earnings. Both the week and generosity of the offer varied from driver to driver. The offers were presented to drivers as an Uber promotion called the "Earnings Accelerator". Drivers received the experimental offers by e-mail and text message, as well as through the Uber application ("app") itself. They were required to opt-in in order to receive the wage increase.

We find that women have Frisch (market-level) elasticities double those of men. In response to a ten percent increase in wages female drivers work seven percent more hours ($\epsilon = .7$), while male drivers work only three percent more hours ($\epsilon = .3$). 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 measure part time workers shifting hours across firms or platforms.[6] Extensive margin elasticities are modest, even among our sample of marginally attached drivers. In response to a ten percent increase in wages, women are at most two percentage points more likely to drive (an elasticity of at most .18), relative to a single percentage point for men (at most .09). These elasticities are significantly smaller than those typically used to calibrate dynamic models; these models typically assume an elasticity greater than 1. The design of our experiment, which required drivers to opt-in, allows us to rule out driver inattention as a possible confounder.

To assess firm substitution, we compare these market-level Frisch elasticities to estimates from

[5]In order to ensure that we included male and female drivers with a range of different (non-treated) hours worked, we stratified active Uber drivers by their usual hours worked during the four weeks prior to sampling before selecting drivers for inclusion in the experiment.

[6]In particular, our estimates are smaller to those estimated in the Boston Earnings Accelerator experiment analyzed in Angrist et al. (2017). Part of this difference is likely attributable to city-specific factors. However, some of the difference is likely because most Boston drivers could shift hours to Lyft, if desired.

two similar experiments where a subset of the drivers cannot drive for Uber's competitor, Lyft, due to the age of their car. We find that both men and women who can drive for competing platforms are significantly more elastic. The additional trips likely come at the cost of Uber's competitor. The gaps between shifters (those who had access to both platforms) and non-shifters (those who did not) are largest for young drivers, who likely are more technologically adept. We do not see any differences between male and female drivers.

Because our experimental estimates of the firm-specific elasticity are not very precise and rely, in part, on comparing elasticities estimated in different cities, we use data from a large-scale Uber promotion we call the "Individual Driver Bonus" (IDB) to corroborate our findings. Drivers in this promotion receive offers of lump-sum bonuses in return for exceeding trip thresholds. Within the IDB sample, drivers who receive more generous ("high") bonuses are statistically indistinguishable from those given smaller incentives. We use a simple model to translate reduced form differences in opt-in rates into labor supply elasticities. We find that, just as in our experiments, those with the opportunity to drive for competing platforms are significantly more elastic. The effects are particularly pronounced for younger drivers.

We use these two sets of elasticities to compute implied firm substitution elasticities for male and female drivers. We find mean elasticities between two and four. These estimates are in line with other recent estimates of firm-specific elasticities. In particular, 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 (Panel B, Table 3) for moderate values of optimization frictions. Our low elasticities reflect the fact that, even in this setting, switching between firms is not trivial.

However, unlike most prior (primarily non-experimental) work, we do not see any significant differences between men and women (Hirsch et al., 2010b; Ransom and Oaxaca, 2010; ?).[7][8] Our

[7]Dube et al. (forthcoming) present experimental elasticities that pool men and women; only their non-experimental estimates are separately reported for men and women. They find that, in the *offline* economy, women are somewhat less elastic than men.

[8]Kline et al. (2018) use variation in wages induced by the grant of a patent to identify firm-specific elasticities and find that women are, if anything, more elastic than men.

results suggest that, even if gig economy firms wield monopsony power, as some authors suggest, they do not have any incentive to wage discriminate between men and women. We view our 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. In particular, in other contexts, women may face higher commuting costs or hours constraints, which could result in lower firm-specific elasticities. These could be fairly substantial. Our results show that, in the absence of commuting costs, women are no less strategic about switching between firms to maximize their earnings.

In addition to the papers cited above, this paper is related to a small literature on labor supply in the gig economy (see, e.g., Hall et al., 2017; Koustas, 2017). Chen and Sheldon (2015) and Angrist et al. (2017) also estimate labor supply elasticities using wage variation among Uber drivers, but do not investigate gender differences and ignore the potential for platform substitution. Our work complements recent work by Cook et al. (2018), who show that there is a gender gap in earnings on the Uber platform itself, driven by differences in driving speed, experience, and the time and location of driving. Our paper differs from Cook et al. (2018) in that it uses experimental variation in the Uber wage to comment on the sources of the non-Uber wage gap. Our results on monopsony are most relevant in settings where firms have the flexibility to wage discriminate between workers. Our results are less useful for explaining the existing gender wage gap at firms where men and women are paid via a gender-blind algorithm.

The rest of the paper proceeds as follows: the next section develops a conceptual framework that illustrates how monopsonistic wage discrimination can lead to a gender wage gap when workers choose both for whom and how much to work. Section 3.3 describes the empirical setting and data, and lays out the experimental variation we exploit. Section 3.4 presents market-level labor supply elasticities for men and women on the intensive and extensive margins. Section 3.5 presents estimates of platform substitution. Section 3.6 concludes.

3.2 Conceptual Framework

Our conceptual framework shows how differences in labor supply elasticities generate wage gaps when employers have monopsony power. The key difference between our framework and standard models is that we allow workers to choose both where to work, and how much to work.

3.2.1 Monopsony with Flexible Hours

Consider a simple model where a firm's potential earnings each period are a function of the hours supplied by their employees, $Y_t(H)$. Firms pick wages w_t in order to maximize their earnings, subject to the labor supply function $H(w_t)$. The firm's problem is thus

$$\max_{w_t} Y_t(H(w_t)) - w_t H(w_t)$$

and the first order condition is

$$Y_t'(H(w_t))H'(w_t) = H(w_t) + w_t H'(w_t)$$

As in a standard monopsony model, the profit maximizing wage is the marginal product of labor, marked down by the elasticity of labor supply. Suppressing time subscripts for clarity, this is

$$w^* = \frac{Y'(H(w))}{1 + 1/\epsilon}$$

where $\epsilon = \dfrac{d \log H(w)}{d \log w}$.[9] In a perfectly competitive labor market $\epsilon = \infty$ and individuals are paid their marginal product ($w^* = Y'$); as ϵ decreases, firms gain monopsony power, and the optimal wage decreases. This may occur if there are few employers in the market, if firms differ in amenities, or if there are costs (e.g. search costs) associated with finding a new job (Manning, 2003b; Card et al., 2016c).

[9] This expression is analogous to expressions used in monopoly pricing models in industrial organization, where the profit-maximizing markup depends on the inverse elasticity of demand (the "Lerner index").

Additional hours may come either from new workers or from an increase in hours worked by existing workers. Suppose that, for a given wage w_t, $N(w_t)$ individuals work for the firm, providing

$$H = \int_0^{N(w_t)} h(i, w_t) di$$

hours of labor. Hours respond to wages according to

$$\frac{dH}{dw_t} = \frac{d}{dw_t} \int_0^{N(w_t)} h(i, w_t) di = \underbrace{h(N(w_t), w_t) N'(w_t)}_{\text{new employees}} + \underbrace{\int_0^{N(w_t)} \frac{\partial}{\partial w_t} h(i, w_t) di}_{\text{increased hours}}$$

by Leibniz's rule. The first term is the change in hours that occurs because some workers join (or leave) the firm in response to the change in wages. The second term is the change in hours for workers whose firm location is unaffected by the change in wages. In elasticity terms this is

$$\frac{d \log H}{d \log w_t} = \frac{h(N(w_t), w_t) N'(w_t)}{H} w + \frac{\int_0^{N(w_t)} \frac{\partial}{\partial w_t} h(i, w_t) di}{H} w$$

For simplicity suppose that, conditional on working for the firm, workers have identical preferences, i.e. $h(i, w_t) = h(w_t)$ for all i. This is the case if individuals have identical preferences but can only work for a single firm at a time. Under this assumption, $H = N(w_t) h(w_t)$ and we can write

$$\begin{aligned} \epsilon = \frac{d \log H}{d \log w_t} &= \frac{h(w_t) N'(w_t)}{N(w_t) h(w_t)} w + \frac{N(w_t) \frac{\partial}{\partial w_t} h(w_t)}{h(w_t) N(w_t)} w \\ &= \frac{N'(w_t)}{N(w_t)} w + \frac{h'(w_t)}{h(w_t)} w \\ &= \eta + \iota \end{aligned}$$

In this case wages depend on both the 'recruiting' elasticity (η) and on the intensive margin elasticity (ι).[10]

[10]As with most monopsony models, this depends on the assumption that the firm cannot engage in perfect price

3.2.2 Monopsonistic Wage Discrimination

Suppose there are two groups of workers: men and women. The firm's problem is to pick w_m, w_w to maximize

$$\max_{w_m, w_w} Y(H_m(w_m) + H_w(w_w)) - w_m H_m(w_m) - w_w H_w(w_w)$$

A derivation similar to that in Section 3.2.1 shows that the optimal wage gap (for the monopsonist) is

$$\frac{w_m^*}{w_w^*} = \frac{1 + 1/\epsilon_w(\eta_w, \iota_w)}{1 + 1/\epsilon_m(\eta_m, \iota_m)} \tag{3.1}$$

The firm maximizes its profits by paying the less elastic group of workers less.[11]

The key difference between the wage gap in equation 3.1 and the wage gap derived from the basic monopsony model is that, in this case, the elasticity depends both on individuals' willingness to leave or join a firm (η) and on their willingness to change their hours worked in response to changes in wages (ι). Even if women are less likely to switch firms (or shift hours between firms), firms may have little incentive to price discriminate if women's overall labor supply is more responsive to wages. We can summarize the results of this section in two propositions.

Proposition 10. *If workers can flexibly choose their hours, a monopsonist would choose the wage gap:*

$$\frac{w_m^*}{w_w^*} = \frac{1 + 1/\epsilon_w(\eta_w, \iota_w)}{1 + 1/\epsilon_m(\eta_m, \iota_m)} \tag{3.2}$$

where ϵ includes intensive (hours) and extensive (firm choice) margin adjustments. If workers have identical preferences such that $h(i, w_t) = h$, this simplifies to

$$\frac{w_m^*}{w_w^*} = \frac{1 + 1/(\eta_w + \iota_w)}{1 + 1/(\eta_m + \iota_m)}$$

discrimination. This means that in order to hire more workers, the firm must also raise wages for existing workers.

[11]This is known as third degree price discrimination in the industrial organization literature (?). A monopolist who is able to price differentiate between different groups of consumers should charge lower prices to more price-elastic groups (e.g. students or senior citizens).

Proposition 11. *If workers cannot flexibly choose their hours, a monopsonist would choose the wage gap:*

$$\frac{w_m^*}{w_w^*} = \frac{1 + 1/\epsilon_w(\eta_w, \iota_w)}{1 + 1/\epsilon_m(\eta_m, \iota_m)} = \frac{1 + 1/\eta_w}{1 + 1/\eta_m} \tag{3.3}$$

where η reflects the change in the number of workers at the firm.

3.2.3 From Elasticities to Wage Gaps

We can calculate the gender wage gap implied by equations 3.2 and 3.3 using labor supply elasticities for two groups: (1) workers that are limited to a single flexible-hours employer and (2) workers that have access to multiple flexible-hours employers. In our empirical setting these correspond to elasticities for drivers who can only drive for Uber ("non-shifters"), and elasticities for drivers that also can drive for Lyft ("shifters").

Wage Gap with Flexible Hours

In a labor market with flexible hours, the optimal wage gap depends on the elasticity of hours worked to a firm's wage rate (by gender). This elasticity will reflect both true changes in hours worked, and changes in the allocation of hours across firms. We use exogenous variation in wages among "shifters" to identify this elasticity.

Wage Gap with Fixed Hours

When hours are inflexible, all that matters to a monopsonistic firm is the extent to which individuals join or leave individual firms in response to changes in relative wages (equation 3.3). We cannot directly measure this firm substitution elasticity because we do not observe hours worked at other firms. However, we can estimate firm substitution elasticities by exploiting the relationship between the market-level elasticities we estimate for non-shifters and for shifters.

Suppose that a driver shifts hours from other platforms smoothly in response to changes in relative wages. Use H to denote total hours, h to denote Uber hours and r to denote other ride-share hours. In response to a change in the Uber wage, w, the change in Uber hours will depend on both the change in total hours worked (which depends on the market elasticity) and on the change in hours worked on competing platforms.

$$\frac{dh}{dw} = \frac{dH}{dw} - \frac{dr}{dw}$$

If we rearrange this expression so total hours are on the left hand side and multiply all terms by w/H to convert this to the total market elasticity we find that

$$\begin{aligned}
\frac{dH}{dw}\frac{w}{H} &= \frac{dh}{dw}\frac{w}{H} + \frac{dr}{dw}\frac{w}{H} \\
\epsilon &= \frac{dh}{dw}\frac{w}{\phi H(1/\phi)} + \frac{dr}{dw}\frac{w}{H(1-\phi)/(1-\phi)} \\
&= \underbrace{\tau\phi}_{+} + (1-\phi)\underbrace{s}_{-}
\end{aligned} \tag{3.4}$$

where ϕ is the fraction of total hours that the driver originally worked on Uber and $s = \dfrac{d \log r}{d \log w}$ measures the elasticity of non-Uber hours to the Uber wage. The market elasticity (ϵ) is the sum of the "Uber" elasticity (τ) and firm substitution elasticity, weighted by the fraction of hours worked on and off Uber.

In order to identify the firm substitution elasticity, s, we need an estimate of ratio of hours spent on Uber to total hours worked, ϕ. For a given ϕ, s can be derived using: $s = \dfrac{\epsilon - \tau \times \phi}{1-\phi}$. Prior work reported an estimate of .93 for ϕ (Koustas, 2017).[12] We use this estimate in much of our analysis. However, we can also produce our own estimates of ϕ using data from our Earnings Accelerator experiments. Our experimental wage offers were so generous that, conditional on taking an offer, it is likely that the driver chose to shift all of her hours from Lyft to Uber.[13] Hours when treated (h_1) depend on the drivers' counterfactual Uber (h_0) and non-Uber (r) hours, the labor supply elasticity

[12]Koustas examined the value of ride-share opportunities as consumption insurance. Koustas reports that conditional on being an Uber (Lyft) driver, 93% (33%) of ride-share earnings come from Uber (Lyft).

[13]Drivers were offered wage increases of 25-39%. More details on the Earnings Accelerator are provided in the next section.

ϵ, and the size of the wage increase.

$$h_1 = (h_0 + r)(1 + \epsilon d \log w)$$

For a given treatment, the percentage change in hours worked on Uber is

$$\begin{aligned} d \log h &= \frac{1}{\phi} \epsilon d \log w + \frac{1-\phi}{\phi} && \text{for Shifters} \\ &= \epsilon d \log w && \text{for Non-Shifters} \end{aligned} \tag{3.5}$$

where ϕ is the fraction of total hours that are spent on Uber. We present estimates of both s and ϕ in section 3.5.

3.3 Empirical Setting and Data

Next, we describe the variation we use to identify the labor supply elasticities of interest. We provide background on the Uber platform and describe how drivers may work for multiple platforms (Section 3.3.1). Then we explain our two sources of empirical variation: (1) a series of experiments we conducted in Boston and Houston (Section 3.3.2), and (2) a long-running Uber promotion we refer to as the Individual Driver Bonus (IDB) program (Section 3.3.3).

3.3.1 Background on Ride-Share

Uber is a global Transportation Network Company (TNC) whose software connects drivers and riders. Uber launched its peer-to-peer operations in mid-2012 and currently has over 900,000 active drivers in the United States. In most cities in the United States there are few barriers to becoming a ride-share driver. While the exact requirements vary from city to city, drivers typically must fill out online paperwork, submit to a background check, and undergo a vehicle screening.

Uber drivers can work whenever and wherever they choose (within Uber's service region) and are paid per mile and minute for each trip they complete. These per-mile and per-minute rates increase at certain times of day and in certain locations due to Surge pricing. Throughout the course of our experiments Uber drivers paid a fixed fraction of their trip receipts to Uber in the form of the "Uber fee". This fee varied across drivers based on the city and when they joined the platform.[14]

Many drivers drive for multiple ride-share platforms. The Rideshare Guy, a popular blog aimed at TNC drivers, estimates that three quarters of drivers drive for both Uber and Lyft. The vast majority of the ride-share market is captured by these two companies.

Drivers that have signed up for multiple platforms may choose to shift between platforms at low frequency, choosing to drive for whichever app offers them the highest earnings when they start driving for the session. Alternatively, they may keep both apps on during down time, accepting the first dispatch to come in. The second strategy is known as "multi-apping" and reduces the amount of time a driver spends idle (earning no money). While it is unlikely a driver could completely eliminate idle time, a driver who cut the time he spent waiting in half would increase earnings by thirty-three percent.[15] Conversations in online forums, such as the one depicted in Figure 3-1, suggest that multi-apping requires a non-trivial amount of effort. As a result, several companies have developed third party apps to help drivers navigate between the two interfaces (e.g. Mystro, Upshift, and QuickSwitch). An advertisement from one of these companies (Figure 3-2) claims that they can help drivers increase their earnings by thirty-three percent.

Some ride-share drivers are not eligible to drive for both platforms. In some cases this is because only one platform operates in the market. For instance, between November 2016 and May 2017, Lyft did not operate in Houston. Even in cities where both platforms operate, some drivers are ineligible to work for both platforms based on the age of their car. In Boston, Lyft requires that drivers use cars model year 2004 or newer, while Uber allows vehicles as old as 2001. Similarly, San Francisco drivers need a vehicle model year 2004 or newer to drive for Lyft, but only a car

[14]In late 2017, Uber loosened the link between driver and rider pay. Now riders and drivers face distinct per-minute and per-mile rates.

[15]This calculation is based on the observed utilization rates in our Houston (pre-Lyft) experiment. The data show that drivers spend roughly 40% of their time without a passenger or active dispatch.

model year 2002 or newer to drive for Uber.[16] While we cannot identify which drivers chose to multi-app, we use the car year threshold to determine which drivers had the option of driving for Lyft. We refer to drivers that could work for both platforms as "shifters" and those that could not as "non-shifters".

3.3.2 Earnings Accelerator Experiments

Our primary source of wage variation is a series of randomized experiments we ran, known as the "Earnings Accelerator". Transportation network service companies routinely run promotions in which they change driver pay, without affecting the prices for riders. These promotions allow ride-share companies to equilibrate supply and demand without the use of surge pricing. Our experiments were modeled after such promotions. We conducted the first experiment in Boston in fall 2016.[17] We conducted two subsequent experiments in Houston: (1) in spring 2017, before Lyft entered the Houston market and (2) in fall 2017, after Lyft had gained substantial market share. Table 3.1 presents detailed timelines for each of the three experiments.

The three experiments follow a similar format. In each, we identify a set of drivers that satisfy two criteria: (1) they were active on the Uber platform (had completed at least four trips in the past month), and (2) they averaged between 5 and 25 hours per week (Boston) or 5 and 40 hours per week (Houston). So that we would have a mix of full-time and part-time drivers, we grouped drivers into one of three bins based on their usual hours per week, and randomly selected subsets of drivers from each bin. The low group consisted of drivers that averaged between 5 and 15 hours per week, the high group consisted of drivers that averaged between 15 and 25 hours per week, and the very high group consisted of drivers that averaged between 25 and 40 hours per week. Drivers in the very-high group worked more than part-time on the platform. Within each bin, we randomly selected drivers for inclusion in the experiment. For both Houston experiments, we over-sampled

[16]See https://www.lyft.com/driver-application-requirements/california and https://www.uber.com/drive/san-francisco/vehicle-requirements/. Uber has additional requirements to drive for its "premium" services, including UberXL, UberBlack, and Uber Select.

[17]The data for the Boston experiment were also analyzed by Angrist et al. (2017) who look at the value of the ride-share contract relative to taxi-style leasing.

women in each bin so that

We offered these drivers the opportunity to drive with no Uber fee for one week. Half of the drivers in each bin were offered the opportunity in week one; the other half were offered it in week two. We informed the drivers that this would result in a "X% higher payout", where X varied across drivers based on the fee they faced. Boston drivers faced either a 20% or 25% commission, depending on when they joined Uber; Houston drivers faced either a 25% or 28% commission. The offers indicated that the promotion would apply to all trips that week, including those with Surge pricing.

Drivers received the offers via e-mail and text message and through the Uber app itself. Figure 3-6 shows a sample e-mail and text message. These messages (and the in-app notification) included links to Google Forms (see Figure 3-7) like those typically used in Uber promotions. The forms were pre-filled with a driver's unique Uber identifier and included detailed information on the promotion, as well as consent language. We sent the offers one week before the promotion went live; drivers had one week to accept the offer by clicking "yes" on the Google Form. Around 60% of the drivers in each experiment accepted our offer.[18] Drivers were able to see their increased earnings reflected in-app while they were driving fee-free.

To increase our sample and generate additional wage variation, we conducted a second set of "Taxi" experiments with drivers who accepted our initial offer of fee-free driving. Treated drivers were offered random subsets of additional fee-free (or reduced fee) driving in exchange for an up-front payment, much like the lease payment a taxi driver would pay to a medallion holder. While these offers were only attractive to drivers that intended to drive enough to pay off the lease payment, these treatments allowed us to generate additional wage variation, at a much lower cost. More details on the experiment, including balance tables, messaging, and sample counts are provided in Appendix 3.10.2. Table 3.2 shows the breakdown of the sample between men and women and potential shifters (individuals who could drive for both Uber and Lyft) and non-shifters.

[18] While the offer should have been attractive to all drivers, Uber drivers receive many messages from Uber each week and many may choose to ignore some of this messaging. In addition, some drivers may not have wanted to participate in academic research.

Our baseline analysis focuses on the Boston experiment and the first Houston experiment due to implementation issues in the second experiment. Our analysis is not sensitive to this decision. See Appendix 3.10 for details.

Columns 1-4 of Table 3.3 show that male and female drivers in the Earnings Accelerator sample are similar on most dimensions. However, female drivers tend to be less experienced. In the Houston sample (columns 3 and 4), women have an average of twelve months of experience on the platform, compared to twenty months for men. The differences between male and female drivers are even larger when considering a trip-based measure of experience. By the start of the Houston experiment, the average male drivers in our sample had completed over 1700 trips, compared to 860 for the average female driver. Because differences in experience may impact a driver's responsiveness to promotions (in particular through drivers' awareness of how to multi-app), we also present results for inexperienced drivers: those with less than nine months on the platform.[19]

3.3.3 Individual Driver Bonuses

Our second source of variation comes from a promotional incentive where drivers are given lump-sum payouts if they exceed specified trip thresholds. Throughout the paper we refer to this promotion as the "Individual Driver Bonus" (IDB) program. Uber sends IDB offers twice each week, once on Monday at 4 a.m. and once on Friday at 4 a.m. The Monday offer covers all trips completed between Monday 4 a.m. and Friday 4.a.m. ("weekday") and the Friday offer covers all trips completed between Friday 4 a.m. and the following Monday 4 a.m. ("weekend"). Drivers are notified at the start of each period about offers via in-app cards, emails, and text messages, and they are able to track their progress towards trip thresholds in the app. Trip thresholds and payouts vary period-to-period and across drivers. Not all drivers receive offers each period, and some drivers receive multiple offers in a given period. Within a week, drivers with the same trip thresholds may receive different payments for exceeding the threshold. In our data there are typically two different

[19]We focus on months-based measures of experience, rather than trips-based measures, because the latter are a function of labor supply. We also present evidence that splits drivers based on the trip threshold they faced (which is also a function of labor supply).

awards for each threshold: we refer to these as "high" and "low" offers.

Our data include all Uber drivers who were included in the IDB in a single large U.S. city between July 2017 and December 2017. We limit the data to drivers who completed trips for Uber's 'peer-to-peer' services—UberX, UberPool, and UberXL. Other Uber services (e.g. Uber Eats) use different payment and promotion structures. We track total trips completed, hours worked, and total earnings for each driver-period.[20]

Table 3.4 presents summary statistics of the drivers in the IDB sample and shows that, conditional on past driving behavior, drivers that received high IDB offers are statistically indistinguishable from those that received low offers. Column 1 of this table shows that our IDB drivers complete an average of 31 trips per week and make an average of $350 per week. Sixteen percent of the drivers are female and ninety-nine percent have a car model year 2003 or newer. Column 4 of Table 3.2 shows that there are 218 female and 864 male drivers with cars that prevent them for driving for Lyft. IDB drivers also tend to be more experienced than those included in our experimental sample; the median driver has been on the platform for sixteen months, compared with only nine months in the Earnings Accelerator sample. Importantly, high and low bonus offers are as good as randomly assigned within the IDB sample. Column 3 of Table 3.4 shows that, conditional on background characteristics, the high and low offer groups are statistically indistinguishable. Column 6 shows that, conditional on the same characteristics, the dollar amount of the bonus is as good as randomly assigned.

3.4 Labor Supply to the Market

We use data from the first (pre-Lyft) Houston Earnings Accelerator to provide experimental estimates of the market labor supply elasticities for men and women. Because our experiment involved short-run, anticipated wage increases, we interpret all of our estimates as Frisch elasticities.

[20]Not all Uber trips count towards IDB's thresholds (e.g. trips completed in another city). For simplicity, we focus on total trips completed; the vast majority of trips qualify.

3.4.1 Intensive Margin Frisch Elasticities

We estimate intensive margin elasticities by regressing log hours and log log wages. We use treatment offers as an instrumental variable and estimate

$$\log h_{it} = \varepsilon \log w_{it} + \beta X_{it} + \eta_{it} \tag{3.6}$$

$$\log w_{it} = \gamma Z_{it} + \lambda X_{it} + v_{it}, \tag{3.7}$$

where X_{it} includes dummies indicating the strata used for random assignment, the number of months a driver has been on the Uber platform, one lag of log earnings, an indicator for whether the driver drove at all in the prior week, and an indicator for whether a driver uses Uber's "vehicle solutions" leasing program. The parameter of interest is ϵ. Because program take-up is endogenous and impacts driver hourly earnings, we instrument log wages with treatment offers, Z_{it}.

We present estimates for just-identified models where the instrument is an indicator for whether an individual was offered treatment (either fee-free driving or a taxi offer) and for over-identified models where we use separate instruments for each hours group, commission, and week. The additional instruments in the over-identified model allow us to better account for natural differences in hourly earnings across different groups of drivers and for differences in take-up rates.[21] We use a stacked model to test whether women and men have the same elasticities. To ensure that our test has enough power, we require that men and women have the same covariates. Standard errors are clustered by driver.

Table 3.5 presents estimates of ϵ for men and women and shows that, across a variety of samples and specifications, women are about twice as elastic as men. The fee-free week data reveal that in response to a ten percent increase in wages, women spend seven percent more time driving and men between two and four percent more time driving. Because the over-identified model suffers from weak instruments (in columns 3 and 4) we also present results produced with limited

[21] Appendix 3.11.2 shows that the first stage is a function of both the experimentally induced change in the Uber fee and in the take-up rate of the offer. Given prior research on the gender wage gap on Uber, it is especially important to include separate treatment indicators for each gender (Cook et al., 2018).

information maximum likelihood. The elasticities in columns 3 and 4, which use data from the second "Taxi" phase of the experiment, are larger, likely reflecting the fact that the Taxi compliers are a particularly elastic subset of drivers.

Table 3.6 shows that these results are not an artifact of the particular sample we use. This table presents estimates from a stacked model where, in order to boost power in small samples, the coefficients on some covariates (months since signup, vehicle year, and one lag of log hours worked) are constrained to be equal for men and women. Columns 1 through 3 show results separately for each of the hours bins we used for random assignment. Moving across columns we see that elasticities are largest in the low hours group and smallest in the very high hours group. This is consistent with recent evidence in Chen and Sheldon (2015) and Mas and Pallais (2018) on the value of non-work time. Columns 4 and 5 split the sample by median months on the Uber platform (nine months) and show that the results are largely driven by the experienced drivers. We cannot reject that inexperienced men and women are equally elastic.

One alternative explanation for our findings is that women are more elastic because they are less likely to hold outside employment. To address this concern we present estimates for the subset of drivers who were observed working more than forty hours per week in the period before sample selection. The results in Column 6 of Table 3.6 show that, even among full-time drivers, women are twice as elastic on the intensive margin.

3.4.2 Extensive Margin Frisch Elasticities

We next turn to examining how drivers' decision to drive in a given week responded to the offer of higher wages. We present estimates of the reduced form equation

$$\text{Drive}_{it} = \eta^F Z_{it} \times \text{Female}_i + \eta^M Z_{it} \times \text{Male}_i + \beta X_{it} + \epsilon_{it} \tag{3.8}$$

where X_{it} includes dummies indicating the strata used for random assignment, driver gender, the number of months a driver has been on the Uber platform, and indicators for whether a driver uses Uber's "vehicle solutions" leasing program. Drive_{it} is an indicator for whether the driver was active on the Uber platform that week. Z_{it} indicates the percentage increase in wages offered to the driver. This is clearly defined based on the structure of the experiment: each driver is told what percentage increase in wages they will see if they opt in to the treatment. For control drivers it is equal to zero. The sex-specific parameter η measures how driver participation decisions respond to percentage changes in the offered wage. Standard errors are clustered by driver.

To estimate these elasticities we use data from the first two weeks of the experiment, which did not require drivers to make an up-front payment in order to get higher wages. Because Taxi offers were only attractive to drivers who planned to drive at least a minimum number of hours, it is unlikely that they had a large impact on whether drivers chose to drive. Taxi offers have no impact on whether drivers choose to drive.

Table 3.7 shows that, across all groups, women are more responsive to the offer of higher wages than men are. The results in column 1 reveal an average participation elasticity of .12 for women and .04 for men: in response to a 10% increase in the offered wage, women are 1.2 percentage points more likely to drive, compared with only 0.4 percentage points for men. The next three columns break out the results by hours bin and show that the effects are largest in the low hours group, which contains the drivers that are least attached to the platform, and smallest in the very high hours group. The remaining columns divide drivers by median months on platform (nine months) and by age. The results suggest that (1) less experienced drivers are more responsive to the promotion and (2) there aren't significant differences between older and younger drivers.[22]

The reduced form estimates do not correct for driver inattention. If drivers do not start driving because they did not see the Earnings Accelerator offer, our estimates of η will be biased downward.

[22] The experienced and inexperienced groups each contain roughly equal numbers of drivers in the low, high, and very high bandwidths. This is largely because when we selected drivers for the experiment, we stratified on both commission and hours group. Drivers with a 20% commission are necessarily more experienced drivers, because they had to join the platform before the commission changed.

Panels B and C of Table 3.7 present two-stage least squares estimates of

$$\text{Drive}_{it} = \eta D_{it} + \beta X_{it} + \epsilon_{it} \tag{3.9}$$

$$D_{it} = \gamma Z_{it} + \lambda X_{it} + \upsilon_{it}, \tag{3.10}$$

where D_{it} is an indicator for whether the driver accepted an Earnings Accelerator offer in week t. The instrument in the just-identified model (Z_{it}) is the same as before: the offered percentage increase in wages. The over-identified model uses indicators for whether a driver was offered fee-free driving interacted with week-of-offer and driver commission.

Column 1 of Table 3.7 shows that, once we scale by the participation rate, we obtain extensive margin elasticities of .16 and .07 for women and men, respectively (Panel C). These elasticities are significantly larger than the reduced form estimates in Table 3.7, but still significantly smaller than most estimates in the literature.[23] The estimated elasticities are largest among low hours drivers—whose baseline participation rates are lowest—and inexperienced or young drivers.

Interpretation Of course drivers participating in the Earnings Accelerator may differ from those that did not participate. The econometric issue is that there are two types of never-takers: (1) inelastic never-takers and (2) consent/inattention never-takers. Drivers in the first group do not accept the offer because the offer is not generous enough to induce them to drive; drivers in the second group do not accept the offer because they do not want to participate in academic research or because they did not see the messaging. While the two-stage least squares estimates identify the effect on compliers, the true extensive margin elasticity combines the impact on compliers with the impact on inelastic never-takers. Without information on the relative proportions of these two groups, it is impossible for us to identify the true extensive margin elasticity. The reduced form and two-stage least squares estimates give us lower and upper bounds, respectively.

The primary concern with interpreting our extensive margin results as extensive margin Frisch elasticities is that drivers in our sample may hold second jobs in the non-gig economy. However,

[23] Chetty et al. (2013) report a mean extensive margin elasticity of .28 among the fifteen studies in their meta-analysis.

this is not a concern for our analysis as long as long as the worker cannot adjust their hours with less than a week's notice. Our elasticities are within the range of recent quasi-experimental results (Martinez et al., 2018). In general we expect our results to be an upper-bound on the 'true' extensive margin elasticity since adjustment costs are minimal in this setting.

3.5 Firm Substitution

We use data from repeated Earnings Accelerator experiments and a large-scale Uber promotion to look at how drivers shift hours between platforms. Labor supply elasticities for drivers that can work for multiple platforms ("shifters") combine the market-level elasticities we estimated in the previous section with firm-specific shifting. We use the formulas derived in Section 3.2.2 to convert these elasticities into implied firm substitution elasticities.

3.5.1 Evidence from the Earnings Accelerator

We stack data from the three rounds of the Earnings Accelerator in order to identify firm- and market-elasticities. The market labor supply elasticity—the increase in total hours worked in response to a wage change—is identified by the responses of two groups: (1) Houston drivers in the first Houston experiment and (2) Boston drivers that were ineligible for Lyft.

The opportunity to drive for other platforms makes drivers appear more elastic. Panel A of Table 3.8 presents separate estimates of equation 3.6 for shifters and non-shifters. Column 1 shows that, on average, a non-shifter will increase hours worked by 8% in response to a 10% increase in hours. A shifter will increase hours by much more - 12.8% vs. 8%. The gap between shifters and non-shifters is most pronounced among young drivers. This result is consistent with younger drivers being more technologically adept, since more technologically adept drivers find it easier to shift platforms.

Panel B breaks out the results by driver gender and shows that men and women respond equally to the opportunity to multi-app. Column 1 shows that, across the three Earnings Accelerator experiments, male drivers that cannot shift to competing platforms drive 6 percent more when confronted with a 10% increase in hourly wages. However, male drivers that can shift drive nearly 12 percent more. These additional hours likely come from Lyft, and therefore do not reflect real increases in labor supply. Female drivers are generally more responsive to increases in wages; both female shifters and non-shifters are more elastic than their male counterparts. However, the gaps between shifters and non-shifters are roughly the same size. The remaining columns of Table 3.8 show that the same pattern emerges across different groups of drivers defined by experience and age.[24]

We can look for additional evidence of multi-apping by examining the utilization rates (the fraction of the time a driver's app is on that he/she is actively on a trip) of shifters and non-shifters, and by looking at the impact of the treatment on utilization rates for each group. Because drivers who multi-app spend less time waiting for dispatches, we should see higher utilization rates among these drivers. Appendix Section 3.9.2 presents additional analysis showing that utilization rates are in fact higher among shifters. This is important because only shifters can use multi-apping as a way to increase their earnings; non-shifters can only work for Uber.

3.5.2 Evidence From Individual Driver Bonuses

Because our experimental elasticities in Section 3.5.1 are imprecise, we use data from a large-scale Uber promotion we call the "Individual Driver Bonus" (IDB) program to corroborate our findings. This promotion has two main advantages. First, unlike the Earnings Accelerator, the data come from a single large city. Second, due to the structure of the promotion, we are able to examine high earnings/hours drivers who we were unable to include in our experiment. It is possible that a gender gap in shifting could emerge among these drivers.

As discussed in Section 3.3.3, drivers in this program were offered lump-sum payouts for exceeding pre-specified trip thresholds. Figure 3-9 shows how the IDB incentive affects a driver's pay. The black line denotes the normal relationship between trips and total earnings. The red line

[24]Note that unlike in Section 3.4.1, we do not stratify by hours bin when examining shifting behavior. Because our bandwidths are based on only Uber hours, drivers who can shift across platforms but have high hours on Uber, are less likely to be taking advantage of their option to shift.

shows how this changes with the IDB incentive. A comparison of the solid and dashed red lines reveals the difference between the low and high groups. The two groups face the same earnings until the trip threshold, but there is a larger discontinuity in the high group. The incentive structure is most attractive to drivers who, in the absence of treatment, would be close to the trip threshold. For these drivers, the implied increase in wages due to the incentive (the bonus spread across the additional trips they would need to complete) is largest. Whether a driver completes more trips in response to the incentive depends on three factors: (1) the size of the bonus, (2) the number of additional trips a driver needs to take to cross the threshold, and (3) the curvature of the driver's utility function.

Individuals that are offered the high bonuses are more likely to exceed the pre-specified trip thresholds and complete more trips. Figure 3-3 plots kernel density estimates of trips completed for drivers who faced a 40 trip threshold (indicated by a solid red line). While the densities of both groups of drivers have a mass at exactly 40 trips, there is a larger spike for drivers in the high group. Figure 3-4 plots similar kernel densities for the high group, splitting drivers by sex and whether their car made them eligible for Lyft. The figure reveals that there is a significantly larger spike among the male shifters than the male non-shifters. We present similar, regression-based results, in Appendix Section 3.9.3.

Estimation Strategy

We can derive estimates of drivers' labor supply elasticities by assuming a parametric form for trips completed without the incentive. Because all drivers face a fifty percent chance of obtaining the high and low offer each period, there is no income effect. Use t_{i0} to denote the number of trips driver i completes without the promotion and t_{i1} to denote the number of trips driver i completes when given an offer. Individuals will receive the bonus if their treated trips exceed the trip threshold. Use B to denote the lump sum bonus and T to denote the trip threshold. Individuals will exceed the trip threshold if:

$$
\begin{aligned}
t_{i0} &\geq T && (3.11)\\
t_{i0}(1+\epsilon \log \frac{B/T}{w}) &\geq T \quad t_{i0} < T && (3.12)
\end{aligned}
$$

The first line simply states that, if a driver would have exceeded the threshold without the incentive, they will with the incentive. The second line describes the conditions under which a driver who would not have crossed the threshold without the incentive, crosses the threshold with the incentive. This depends on the driver's untreated trips (t_{i0}), the amount by which the incentive changes the wage $(B/T)/w$, and the labor supply elasticity (ϵ). A larger trip elasticity leads more drivers to cross the trip threshold.[25]

We can estimate driver labor supply elasticities with or without assuming assumptions about the distribution of trips completed. First, suppose trips are log-normally distributed (perhaps conditional on some covariates). We can take logs of expression 3.11 and use the approximation that $\log(1+x) \approx x$ for small x to re-write the expression as

$$
\begin{aligned}
\log t_{i0} - \epsilon \log T + \epsilon \log B - \epsilon \log w &\geq \log T\\
\log t_{i0} &\geq -\epsilon \log B + \epsilon \log w + (1+\epsilon) \log T
\end{aligned}
$$

Assuming a mean of μ and a variance of σ^2, the probability a driver exceeds the trip threshold is

$$1 - \Phi \left[-\frac{1}{\sigma}\epsilon \log B + \frac{1+\epsilon}{\sigma} \log T + \frac{1}{\sigma}\epsilon \log w - \frac{\mu}{\sigma} \right] \tag{3.13}$$

We can estimate this model using a probit where the dependent variable is whether a driver crossed the trip threshold, B is the lump sum bonus and T is the trip threshold. The final term is a function of average per-trip earnings. Because these may vary over time, we include period fixed effects. We use the relationship between the coefficients on $\log B$ and $\log T$ to estimate ϵ.

[25]Note that here the elasticity is in terms of trips, rather than hours.

Figure 3-15 shows that the log-normal assumption is reasonable. First, we regress log trips on date fixed effects and the hours bins from Table 3.4. Then, we plot the residuals, along with a fitted normal curve. The figure on the left plots the residuals for the full sample. The data roughly follow a normal distribution, but there is a spike to the right of the mean. This is likely driven by bunching at the trip threshold. While the parametric assumption applies to the control distribution (in the absence of IDB offers), we only observe the treated distribution. Because the treatment is only likely to affect the distribution of trips completed in a neighborhood of the trip threshold, we present a similar histogram, omitting observations for drivers within a two-trip band of the trip threshold, in Panel B. This distribution looks very similar to a normal distribution. The residual variance in the four groups defined by sex and Lyft eligibility is nearly constant, ranging from .76 (male non-shifters) to .81 (all other groups).

Alternatively, we can derive estimates of drivers' labor supply elasticity without assuming a parametric distribution for trips completed. Use $p_{B,T}$ to denote the fraction of drivers in the treatment group and F_0 to denote the distribution of trips for the control group. We can re-write the opt-in equation as

$$F_0^{-1}\left[1 - p_{B,T}\right] = \frac{T}{1 - \epsilon \frac{B/T}{w}} \tag{3.14}$$

The left hand side of this equation is the quantile of the trip distribution corresponding to the fraction of drivers in the high bonus group who exceeded the trip threshold. We estimate equation 3.14 using non-linear least squares. See Appendix 3.11.3 for a complete derivation and for more details on the estimation.

IDB Elasticities

Table 3.10 presents labor supply elasticities for four different groups: (1) male non-shifters, (2) male shifters, (3) female non-shifters, and (4) female shifters. We calculate these elasticities using the structural relationship between the coefficients in the probit model described in equation

3.13.[26] The probit coefficients are reported in Appendix Table 3.20. The first two columns present estimates from the baseline model in equation 3.11. The third and fourth columns present results from a similar model where we re-weight the sample so that male and female drivers are equally distributed across treatments.[27] Because female drivers drive fewer trips on average, they are more concentrated in 'low' treatment groups. Re-weighting the sample allows us to account for the fact that drivers with different (untreated) driving patterns may have different elasticities.

Our preferred specification is the instrumental variables specification presented in column 2. In response to a ten percent increase in wages, male drivers that cannot drive for competing platforms increase their labor supply by ten percent; male drivers that can drive for competing platforms increase their labor supply by almost fourteen percent. We expect these additional hours came at the expense of Uber's competitor, Lyft. We see a similar pattern among women: female drivers that are limited to a single platform drive only eight percent more in response to a ten percent wage increase, compared with nearly twelve percent. For both male and female drivers, we reject the hypothesis that shifters and non-shifters have the same elasticity. These results indicate that drivers shift between platforms in response to changes in relative wages.

The theory described earlier says that firms have an incentive to pay lower wages to workers who are less likely to leave for another firm. Prior, non-experimental, work has suggested that women are less likely to leave. We find no evidence of that here. While the male shifters are more elastic than the female shifters, the gap between shifters and non-shifters is roughly equal for the two groups. In the next section we show that the firm specific elasticities for men and women are statistically indistinguishable.

[26] Each elasticity is calculated using the ratio of the coefficients on $\log B$ and $\log T$. We use the fact that: $\frac{\beta_{\log B}}{\beta_{\log B}+\beta_{\log T}} = \frac{-\epsilon/\sigma}{-\epsilon/\sigma + (1+\epsilon/)/\sigma} = \frac{-\epsilon}{-\epsilon+1+\epsilon} = -\epsilon$. Table 3.10 presents estimates of ϵ.

[27] For each group g we assign male drivers weights of $\frac{p(g|m)}{p(g|f)}$ where $p(g|f)$ is the probability that a driver is in group g, conditional on being female.

3.5.3 Firm Substitution Elasticities

We can use the the formulas in Section 3.2.3 to convert our labor supply elasticities into implied firm substitution elasticities. We can also use these formulas to calculate the fraction of time spend on other platforms.

We use data from the Earnings Accelerator experiments to estimate the fraction of time male and female drivers spend on Uber, relative to ride-share as a whole. While these are not firm-substitution elasticities, these provide information about how aggressively each group optimizes their earnings. Because multi-apping is likely to always be a profitable strategy, we should see lower fractions for men if they make more strategic labor supply decisions. Table 3.9 presents the main results.

The first column estimates that men spend about half of their total ride-share/gig time on Uber, though the standard errors can't rule out relatively large fractions. Female drivers appear to spend less total time on Uber, but the standard errors are again large and the effects are insignificant. The experienced drivers appear to spend more time on competing platforms, but, again, the results are imprecise.

With these fractions in hand, we compute firm-substitution elasticities using our IDB estimates from the previous section and equation 3.4. These substitution elasticities measure the extent to which drivers move hours onto Uber in response to changes in the Uber wage.

Column 1 of Table 3.12 shows estimated elasticities of between 2-4 for both male and female drivers. These estimates are surprisingly similar to recent work by Dube et al. (2017).[28] We do not see any significant differences between male and female drivers. If anything, women appear to be more elastic. The remaining columns show that significant differences do not emerge in different subgroups. The fact that we do not see gender differences in these firm-substitution elasticities indicates that gig economy firms have little incentive to pay equally productive men and women different amounts.

[28]The authors use a bunching estimator to estimate firm-substitution elasticities from administrative wage data and the CPS and from Amazon mTurk. Our estimates are larger than those reported for the online gig economy (mTurk) in that paper, but are very similar to those reported for the offline economy.

3.6 Conclusion

We provided new evidence on the potential for gender differences in labor supply to explain the gender wage gap. Firms with market power in the labor market have an incentive to pay lower wages to workers who are less elastic to the firm: workers who are less willing to leave in search of better wages elsewhere. We illustrated that once workers can choose their hours freely, the optimal monopsonistic markdown depends on both the intensive margin elasticity and on the firm substitution elasticity.

We then used experimentally induced variation to estimate intensive and extensive margin Frisch elasticities for men and women. We found that women have Frisch elasticities roughly double those of men. In response to a ten percent increase in wages, women are nearly two percentage points more likely to drive at all, compared to one percentage point for men. Conditional on driving, women drive eight percent more hours, compared to four percent more for men. These elasticities—in particular the extensive margin elasticities—are modest relative to those usually used to calibrate macro models.

We found that drivers shift hours between platforms (firms) when given the opportunity to do so and that women are not significantly less likely to do so than men. To our knowledge, we are the first to experimentally estimate separate firm-specific elasticities for men and women. Taken as a whole, these results suggest that, at least in the gig economy, firms do not have a strong incentive to wage discriminate between their male and female employees (or independent contractors). To the extent that women may be particularly drawn to gig economy employers due to a desire for flexible work arrangements, this is encouraging.

3.7 Figures and Tables

Figure 3-1: Multi-Apping Discussion on UberPeople

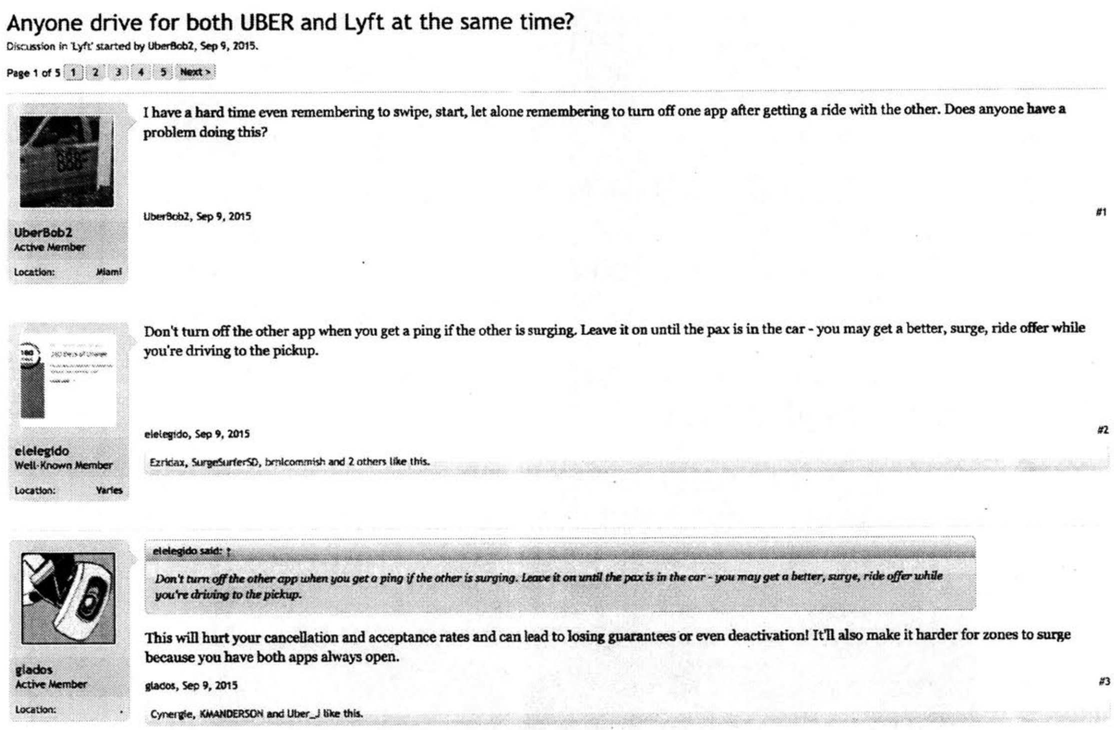

Note: The above picture is a screenshot from "Uber People", an online forum and discussion board where drivers discuss ride-share related topics. The forum is not affiliated with Uber Technologies, Inc. or any other ride-share company. The conversation highlights that drivers are interested in multi-apping but find it requires a non-trivial amount of effort.

Figure 3-2: Example of a Third Party Multi-Apping Application


News Video Events Crunchbase

Mystro is an app aiming to bring in more bacon for Uber and Lyft drivers

Posted Aug 2, 2017 by Sarah Buhr (@sarahbuhr)

Mystro hopes to put 33 percent more money per year in the pockets of Uber and Lyft drivers through an app.

Note: This is a screenshot of a TechCrunch article discussing a third party app, Mystro, which helps drivers quickly switch between competing ride-share platforms.

Figure 3-3: Example of Bunching Around IDB Threshold

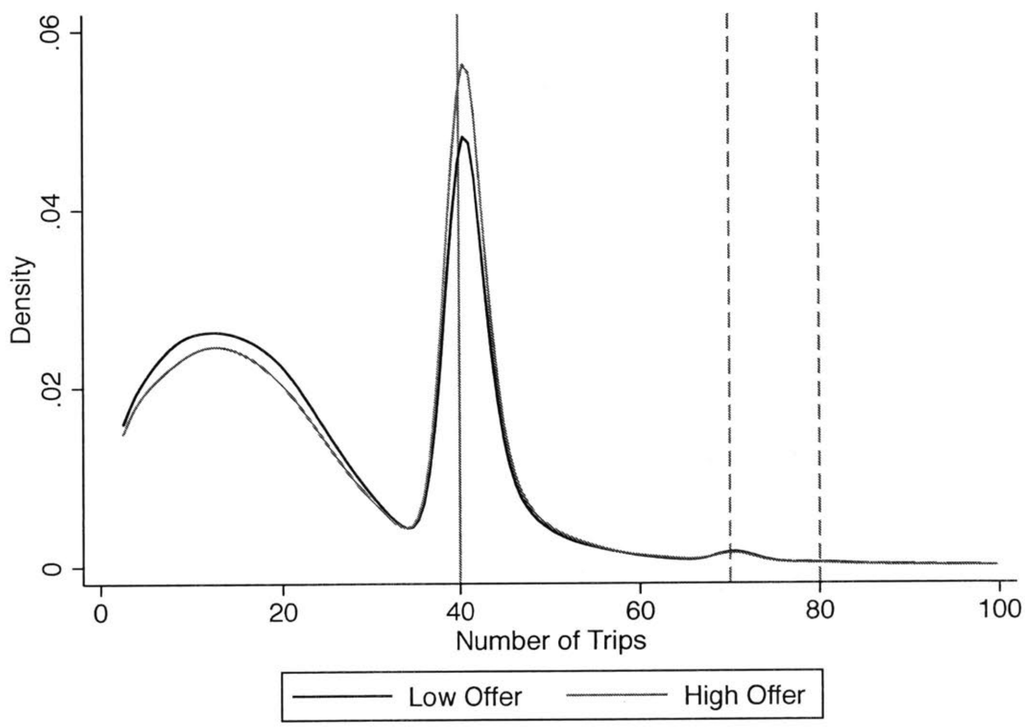

Note: This figure plots Gaussian kernel density estimates of the distribution of trips completed for drivers in the high and low bonus groups with a 40 trip threshold using a bandwidth of 2. We selected this trip threshold because it contains the largest number of female non-shifters (and the second largest number of drivers overall). The dashed red lines denote additional trip thresholds. Drivers were offered up to two incentives per period.

Figure 3-4: Example of Differences in Bunching by Subgroup

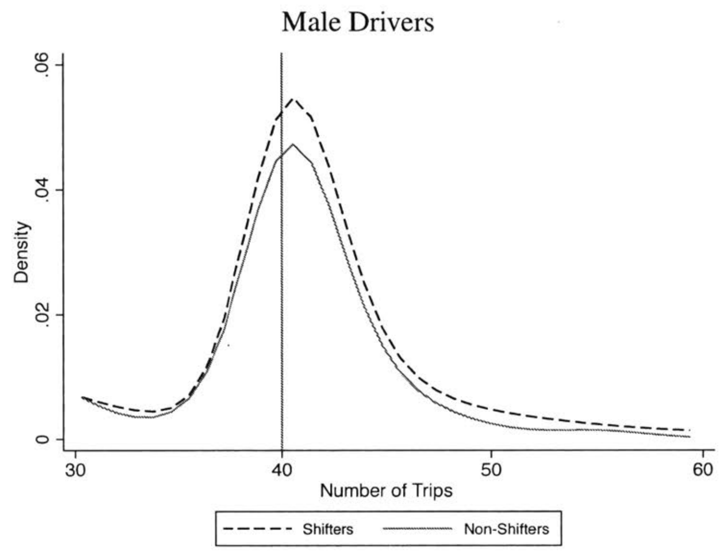

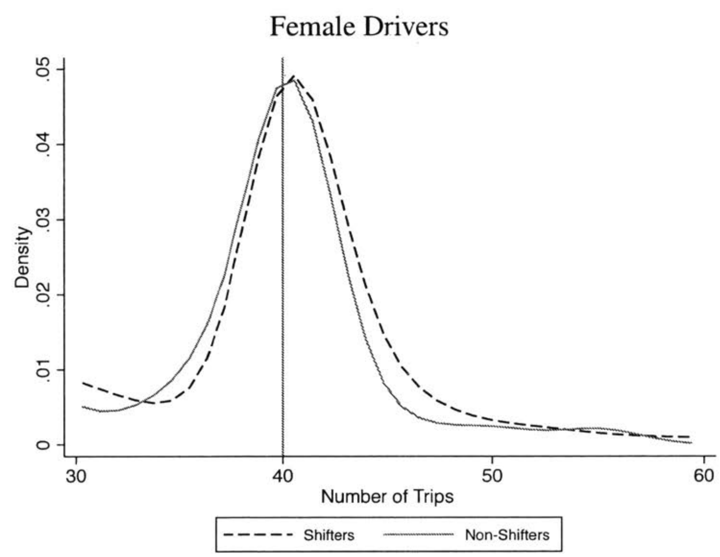

Note: This figure plots Gaussian kernel density estimates of the distribution of trips completed for drivers with a 40 trip threshold using a bandwidth of 2. We calculate the density separately for four groups of drivers, based on sex and car year, all of whom were in the more generous (high) treatment. We selected this cutoff as it was associated with the largest number of female non-shifters (the smallest group). Regression results that pool all strata are presented in Appendix Table 3.9.3.

Table 3.1: Timeline

City	Week Beginning	Action
Boston	August 15, 2016	Sample selection for Boston experiment
	August 22, 2016	Wave 1 Notifications and Opt-In
	August 29, 2016	Wave 1 Opt-Ins Drive Fee Free; Wave 2 Notifications and Opt-In
	September 5, 2016	Wave 2 Opt-Ins Drive Fee Free
	September 12, 2016	Taxi 1 Offers and Opt-In
	September 19, 2016	Taxi 1 Live
	September 26, 2016	
	October 3, 2016	
	October 10, 2016	Taxi 2 Offers and Opt-In
	October 17, 2016	Taxi 2 Live
Houston	March 27, 2017	Sample selection for round 1 of Houston
	April 3, 2017	Wave 1 Notifications and Opt-In
	April 10, 2017	Wave 1 Opt-Ins Drive Fee-Free; Wave 2 Notifications and Opt -In
	April 17, 2017	Wave 2 Opt-Ins Drive Fee-Free
	April 24, 2017	
	May 1, 2017	Taxi 1 Offers and Opt-In
	May 8, 2017	Taxi 1 Live
	May 15, 2017	Taxi 2 Offers and Opt-In
	May 22, 2017	Taxi 2 Live
	May 29, 2017	Lyft Enters Houston
	September 11, 2017	Sample selection for round 2 of Houston
	September 18, 2017	Wave 1 Notifications and Opt-In
	September 25, 2017	Wave 1 Opt-Ins Drive Fee-Free; Wave 2 Notifications and Opt -In
	October 2, 2017	Wave 2 Opt-Ins Drive Fee-Free
	October 9, 2017	
	October 16, 2017	Taxi 1 Offers and Opt-In
	October 23, 2017	Taxi 1 Live

Note: This table presents the timeline of the three Earnings Accelerator experiments. Each experiment unfolded in three stages. First, we defined the eligible sample, based on drivers' trips and hours over the prior four weeks and randomly selected a subset of drivers to participate in the experiment. We picked half of the drivers to receive the offer of fee-free driving in one week (wave 1); the second half received offers the following week (wave 2). Finally, we offered random subsets of drivers who opted in to fee-free driving the opportunity to buy additional weeks of fee-free or reduced-fee driving, for an upfront payment. We conducted two weeks of Taxi treatments in Boston and in the first Houston experiment. We were only able to conduct one week in the second Houston experiment, as a result of changes in the Uber app.

Table 3.2: Sample Counts

	Boston (1)	Houston 1 (2)	Houston 2 (3)	Individual Driver Bonus (4)
Male Drivers	1431	972	1283	48527
Shifters	230	0	1283	47958
Non-Shifters	1201	972	0	569
Female Drivers	232	1048	817	10923
Shifters	28	0	817	10794
Non-Shifters	204	1048	0	129

Note: This table gives the sample counts of male and female drivers in the three Earnings Accelerator experiments and in the IDB sample. We call drivers who could drive for Lyft shifters, and those could not non-shifters. Boston drivers are considered non-shifters if they have a car model year 2003 or older. No drivers in Houston 1 are considered shifters because Lyft was not present in the market at that time. All drivers in Houston 2 are considered shifters because Lyft had already re-entered the market. IDB drivers are considered shifters if their car model year is 2003 or older.

Table 3.3: Characteristics of Male and Female Drivers

	Earnings Accelerator				Individual Driver Bonus			
	Full Sample		Houston 1		Non-Shifters		Shifters	
	Male	Female	Male	Female	Male	Female	Male	Female
	(1)	(2)	(3)	(4)	(5)	(6)	(5)	(6)
Age	42.8	43.6	44.0	44.3	41.8	42.6	41.0	40.3
	(12.5)	(11.9)	(12.5)	(11.9)	(13.6)	(13.6)	(12.0)	(11.9)
	[42.0]	[43.0]	[43.0]	[44.0]	[40.0]	[42.2]	[39.3]	[38.9]
Months on Platform	14.2	8.1	20.1	12.2	10.5	8.6	17.5	12.1
	(10.1)	(8.5)	(9.5)	(9.3)	(8.9)	(7.4)	(13.7)	(10.9)
	[12.0]	[4.2]	[21.5]	[9.3]	[8.5]	[6.4]	[14.3]	[9.4]
Trips Completed	1413.3	581.6	1707.3	861.2				
	1653.4	893.4	1820.3	1111.2				
	797.0	250.0	1047.0	474.0				
Vehicle Solutions	0.1	0.1	0.1	0.1	0.0	0.0	0.0	0.0
	(0.3)	(0.3)	(0.3)	(0.4)	(0.0)	(0.0)	(0.0)	(0.0)
	[0]	[0]	[0]	[0]	[0]	[0]	[0]	[0]
Model Year	2012.4	2013.3	2013.4	2013.7	2002.6	2002.6	2013.7	2013.7
	(3.8)	(3.0)	(2.5)	(2.5)	(0.6)	(0.5)	(3.2)	(3.1)
	[2013.0]	[2014.0]	[2014.0]	[2014.0]	[2003.0]	[2003.0]	[2015.0]	[2015.0]
Average Hours/Week	19.4	18.0	20.7	18.6	9.5	7.7	12.3	8.9
	(12.1)	(11.7)	(14.5)	(12.7)	(6.3)	(4.3)	(8.3)	(6.9)
	[17.0]	[15.6]	[17.7]	[16.0]	[8.6]	[7.0]	[10.6]	[7.5]
Observations	3623	2097	972	1048	569	129	47958	10794

Note: This table includes all experimental drivers and IDB drivers included in our analysis. The first two columns compare male and female drivers included in any of the three Earnings Accelerator experiments. The third and fourth columns compare male and female drivers in the first Houston experiment (pre-Lyft). These are the drivers included in the analysis in Section 3.4. The remaining four columns compare male and female drivers in the Individual Driver Bonus sample.

Table 3.4: Individual Driver Bonus Balance

	IDB Low Mean	High - Low	p-value	Scaled High - Low	p-value
	(1)	(2)	(3)	(4)	(5)
Months on Platform	18.77	0.004	0.884	0.000	0.944
		(0.026)		(0.001)	
Female	0.14	0.000	0.664	0.000	0.920
		(0.001)		(0.000)	
Vehicle Year	2013.78	0.003	0.670	0.000	0.839
		(0.006)		(0.000)	
Vehicle Year>=2003	1.0	0.000	0.708	0.000	0.397
		(0.000)		(0.000)	
Qualifying Trips in Prior Period	30.28	0.038	0.417	-0.002	0.447
		(0.047)		(0.002)	
Trips in Prior Period	31.34	0.036	0.454	-0.002	0.538
		(0.048)		(0.002)	
Earnings in Prior Period	350.04	0.045	0.938	-0.003	0.921
		(0.571)		(0.030)	
High in Prior Period	0.50	0.000	0.850	0.000	0.554
		(0.001)		(0.000)	
Observations	519122	1047998		1047998	

Note: Column 1 shows the mean for drivers in the "low" bonus treatment. Column 2 shows the adjusted difference between the high and low bonus treatments. We control for date fixed effects and for indicators for eight hours groups based on an individual's driving behavior in the prior four weeks. More information is in appendix 3.10. Column 3 shows the p-value for the treatment effect estimated in column 2. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.5: Frisch Elasticities

	Free Week		Taxi		Stacked	
	Male	Female	Male	Female	Male	Female
	(1)	(2)	(3)	(4)	(5)	(6)
			A. Just-Identified 2SLS			
Log Wages	0.43*	0.75***	0.61**	1.31***	0.52***	1.00***
	(0.23)	(0.16)	(0.26)	(0.24)	(0.18)	(0.14)
			B. Over-Identified 2SLS			
Log Wages	0.22	0.69***	0.64**	1.10***	0.44**	0.90***
	(0.23)	(0.16)	(0.27)	(0.22)	(0.18)	(0.14)
First Stage F Statistic	13.4	14.6	6.5	7.5	10.5	11.2
p-value from test of equality	0.042		0.051		0.012	
			C. LIML			
Log Wages	0.26	0.73***	0.71**	1.17***	0.53**	0.97***
	(0.30)	(0.17)	(0.30)	(0.24)	(0.22)	(0.15)
p-value from test of equality	0.045		0.062		0.015	
Drivers	714	766	479	565	752	805
Observations	1341	1425	868	1016	2209	2441

Note: All models control for the strata used for random assignment, date fixed effects, the number of months a driver has been on the platform, whether a driver uses the "vehicle solutions" program, and one lag of log earnings. We also include a dummy variable for whether the driver drove at all in the prior week; for drivers that did not drive, we recode their lag earnings to 0. The p-values for the 2SLS and LIML models come from stacked models where the coefficient on each covariate is restricted to be equal for men and women. The over-identified model includes 12 binary instruments for free week and 12 binary instruments for Taxi. Within each treatment type there is a binary instrument for each combination of: commission group (2), treatment week (2), and hours group (3). Table 3.19 presents analogous results without baseline covariates. Standard errors are clustered by driver. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.6: Frisch Elasticities by Subgroup

	By Hours Group			By Months on Platform		By Age		Usual Hours		
	Low (1)	High (2)	Very High (3)	Experienced (4)	Inexperienced (5)	35 or Younger (6)	Older than 35 (7)	40+ Hours (8)	Weekday Afternoon (9)	No Late Nights (10)
Log Wages * Male	0.40	0.46**	0.40	0.29	0.38	0.40*	0.73**	0.33	0.35*	0.35
	(0.31)	(0.20)	(0.25)	(0.19)	(0.28)	(0.21)	(0.31)	(0.25)	(0.20)	(0.33)
Log Wages * Female	1.13***	0.75***	1.06***	1.12***	0.56**	0.96***	0.75***	0.78***	0.71***	0.73***
	(0.26)	(0.15)	(0.27)	(0.20)	(0.25)	(0.15)	(0.17)	(0.25)	(0.15)	(0.22)
p-value for equality	0.059	0.225	0.063	0.001	0.617	0.024	0.944	0.182	0.140	0.302
Drivers	406	1151	604	1018	414	1140	539	414	843	478
Observations	1115	3535	1938	3062	1182	3458	1588	1337	2766	1303

Note: Columns 1-3 stratify by the hours groups used for random assignment. Columns 4 and 5 present results for experienced drivers (those who have been on the platform for more than 9 months) and inexperienced drivers. Columns 6-7 split the sample by driver age. Columns 8-10 present estimates for different subgroups, based on usual (non-treatment week) hours worked. Column 8 includes all drivers who worked more than 40 hours in at least one of the four weeks we used for sample selection. Column 9 includes drivers who, during the course of the experiment, were observed working for at least ten minutes between 3 p.m. and 7 p.m. on at least ten distinct week-days (out of forty maximum in our sample). Column 10 includes drivers that never work after 11 P.M. or before 4 A.M. All models control for the strata used for random assignment, date fixed effects, the number of months a driver has been on the platform, whether a driver uses the "vehicle solutions" program, and one lag of log earnings. We also include a dummy variable for whether the driver drove at all in the prior week; for drivers who did not drive, we recode their lag of earnings to 0. The p-value comes from a stacked model where the coefficients on each covariate is restricted to be equal for men and women. The over-identified models include 24 binary instruments for each gender; these reflect the 4 treatment weeks, 3 hours groups, and 2 commission groups. Standard errors are clustered by driver. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.7: Extensive Margin Elasticities

		By Hours Group			By Months on Platform		By Age	
	All	Low	High	Very High	Experienced	Inexperienced	35 or Younger	Older than 35
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				A. Reduced Form				
Male	0.04	0.04	0.05	0.00	0.01	0.12*	-0.01	0.06
	(0.03)	(0.06)	(0.04)	(0.04)	(0.04)	(0.07)	(0.06)	(0.04)
Female	0.12***	0.20***	0.08**	-0.05	0.09*	0.14***	0.13**	0.12***
	(0.03)	(0.07)	(0.04)	(0.04)	(0.05)	(0.04)	(0.07)	(0.04)
p-value	0.10	0.07	0.48	0.40	0.19	0.83	0.10	0.35
				B. Just-Identified 2SLS				
Male	0.09	0.05	0.11*	-0.01	0.05	0.20*	0.02	0.11
	(0.06)	(0.11)	(0.07)	(0.07)	(0.07)	(0.11)	(0.10)	(0.07)
Female	0.18***	0.32***	0.12**	-0.05	0.15**	0.20***	0.20*	0.17***
	(0.04)	(0.10)	(0.05)	(0.06)	(0.06)	(0.06)	(0.10)	(0.05)
p-value	0.23	0.07	0.93	0.68	0.30	0.98	0.22	0.50
				C. Over-Identified 2SLS				
Male	0.06	0.06	0.07	0.00	0.02	0.23**	0.00	0.10
	(0.06)	(0.11)	(0.06)	(0.07)	(0.06)	(0.11)	(0.10)	(0.07)
Female	0.16***	0.31***	0.11**	-0.06	0.13**	0.19***	0.22**	0.15***
	(0.04)	(0.10)	(0.05)	(0.05)	(0.06)	(0.06)	(0.10)	(0.05)
p-value	0.15	0.09	0.61	0.53	0.21	0.77	0.12	0.51
Observations	4040	1334	2706	1352	2566	1474	1186	2860

Note: This table presents estimates of the extensive margin labor supply elasticity estimated based on models in equation 3.8 and 3.9. All models control for the strata used for random assignment, date fixed effects, the number of months a driver has been on the platform, whether a driver uses the "vehicle solutions" program, and driver gender. The over-identified model includes 12 binary instruments for each gender, one for each combination of: commission group (2), treatment week (2), and hours group (3). The just-identified model includes a single treatment indicator for each gender. Standard errors are clustered by driver. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.8: Earnings Accelerator Elasticities for Shifters and Non-Shifters

		By Months on Platform		By Age		Including
	Baseline	Experienced	Inexperienced	35 or Younger	Older than 35	Houston 2
	(1)	(2)	(3)	(4)	(5)	(6)
A. Pooled						
Non-Shifters	0.84***	0.89***	0.85***	0.53***	0.87***	0.84***
	(0.11)	(0.14)	(0.15)	(0.21)	(0.12)	(0.11)
Shifters	1.28***	1.18***	1.21***	1.39***	1.11***	0.95***
	(0.12)	(0.17)	(0.18)	(0.20)	(0.15)	(0.09)
p-value for equality	0.009	0.200	0.121	0.003	0.212	0.439
B. By Gender						
Male Non-Shifters	0.63***	0.61***	1.04***	0.45*	0.78***	0.76***
	(0.16)	(0.19)	(0.26)	(0.25)	(0.18)	(0.13)
Male Shifters	1.18***	1.02***	1.22***	1.28***	0.97***	0.92***
	(0.13)	(0.18)	(0.18)	(0.19)	(0.17)	(0.10)
p-value for equality	0.007	0.135	0.536	0.005	0.439	0.292
Female Non-Shifters	0.88***	1.04***	0.75***	0.47**	0.86***	0.82***
	(0.12)	(0.19)	(0.16)	(0.22)	(0.14)	(0.11)
Female Shifters	1.39***	1.36***	0.98***	1.28***	1.03***	0.95***
	(0.20)	(0.28)	(0.26)	(0.29)	(0.24)	(0.13)
p-value for equality	0.008	0.186	0.383	0.004	0.452	0.369
p-value: Female = Male Non-Shifters	0.169	0.092	0.276	0.959	0.722	0.631
p-value: Female = Male Shifters	0.179	0.090	0.313	0.999	0.744	0.767
Observations	9061	5364	3697	2723	6301	12113

Note: All models control for the strata used for random assignment, date fixed effects, the number of months a driver has been on the platform, whether a driver uses the "vehicle solutions" program, and one lag of log earnings. Standard errors are clustered by driver. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.9: Estimates of the Fraction of Ride-Share Time Spent on Uber

		By Months on Platform		By Age	
	Baseline	Experienced	Inexperienced	35 or Younger	Older than 35
	(1)	(2)	(3)	(4)	(5)
Male Drivers	0.54***	0.61***	0.86***	0.35*	0.80***
	(0.15)	(0.22)	(0.22)	(0.20)	(0.23)
Female Drivers	0.64***	0.76***	0.76***	0.37**	0.83***
	(0.10)	(0.14)	(0.22)	(0.16)	(0.19)
p-value for equality	0.20	0.15	0.19	0.94	0.71
Observations	9061	5364	3697	2723	6301

Note: This table uses the elasticities in Table 3.8 to estimate ϕ using the formula in equation 3.5. All models control for the strata used for random assignment, date fixed effects, the number of months a driver has been on the platform, whether a driver uses the "vehicle solutions" program, and one lag of log earnings. Standard errors are clustered by driver. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.10: Individual Driver Bonuses: Elasticities

	Parametric				
	Unweighted		Weighted		Non-
	Probit	IV-Probit	Probit	IV-Probit	Parametric
	(1)	(2)	(3)	(4)	(5)
Male Non-Shifters	0.881***	1.080***	0.894***	1.184***	1.143***
	(0.014)	(0.035)	(0.016)	(0.041)	(0.214)
Male Shifters	1.357***	1.367***	1.379***	1.358***	1.564***
	(0.008)	(0.008)	(0.010)	(0.010)	(0.027)
p-value: Male Shifters = Non-Shifters	0.000	0.124	0.000	0.386	0.050
Female Non-Shifters	0.708***	0.821***	0.672***	0.824***	1.509***
	(0.052)	(0.102)	(0.051)	(0.102)	(0.302)
Female Shifters	1.163***	1.183***	1.185***	1.225***	1.176***
	(0.005)	(0.006)	(0.006)	(0.008)	(0.053)
p-value: Female Shifters = Non-Shifters	0.049	0.263	0.027	0.217	0.284
p-value: Female = Male Non-Shifters	0.494	0.479	0.382	0.335	0.336
p-value: Female = Male Shifters	0.000	0.006	0.000	0.066	0.000
Observations	1047998	1047998	1047998	1047998	1047994

Note: This table presents elasticities for male and female non-shifters, estimated using data from the IDB program. The parametric and non-parametric models are discussed in Section 3.5.2 in the text and in Appendix Section 3.11.3. The standard errors for the parametric model are clustered by driver. Observations in columns 3 and 4 are re-weighted so that male drivers have the same distribution across treatment groups as female drivers. The IV-probit model uses treatment (high bonus) indicators, interacted with gender and "shift" status as instruments for the size of the bonus the driver was offered. Appendix Table 3.20 presents the raw probit coefficients. The standard errors in column 5 are computed using 500 bootstrap replications. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.11: Individual Driver Bonuses: Elasticities by Subgroup

	By Trip Group		By Age		By Experience	
	Very High	Low	Younger than 35	Older than 35	Above Median	Below Median
	(1)	(2)	(3)	(4)	(5)	(6)
Male Non-Shifters	0.390***	1.626***	1.865***	1.217***	1.452***	1.421***
	(0.146)	(0.420)	(0.537)	(0.314)	(0.413)	(0.381)
Male Shifters	0.571***	1.869***	1.873***	1.656***	1.648***	1.701***
	(0.073)	(0.421)	(0.396)	(0.255)	(0.250)	(0.323)
p-value: Male Shifters = Non-Shifters	0.206	0.565	0.987	0.175	0.623	0.453
Female Non-Shifters	0.366	1.05**	0.96**	0.802	1.168	0.89**
	(0.412)	(0.483)	(0.415)	(0.510)	(1.550)	(0.377)
Female Shifters	0.498***	1.620***	1.429***	1.364***	1.268***	1.549***
	(0.061)	(0.303)	(0.248)	(0.169)	(0.162)	(0.246)
p-value: Female Shifters = Non-Shifters	0.715	0.222	0.338	0.212	0.930	0.103
p-value: Female = Male Non-Shifters	0.957	0.332	0.163	0.478	0.859	0.294
p-value: Female = Male Shifters	0.099	0.174	0.042	0.040	0.013	0.330
Observations	437622	610376	346945	701053	525201	522797

Note: This table presents elasticity estimates based on the parametric model discussed in Section 3.5.2. The elasticities come from a probit model where we instrument the size of the offered bonus with treatment indicators, interacted with gender and and indicator for whether the driver is a "shifter". Standard errors are clustered by driver. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.12: Firm-Specific Elasticities

		By Experience		By Age	
	All Drivers	Above Median	Below Median	35 or Younger	Older than 35
	(1)	(2)	(3)	(4)	(5)
Male	2.736	0.096	3.477	3.477	1.842
	(2.647)	(3.497)	(3.264)	(3.264)	(3.673)
Female	3.983	4.178	5.023	5.023	5.892
	(4.602)	(5.317)	(6.064)	(6.064)	(4.743)
p-value for equality	0.813	0.516	0.822	0.822	0.494
Observations	1047998	437622	610376	346945	701053

Note: This table uses the formulas in Section 3.2.3 to translate the elasticities presented in Table 3.10 into implied firm-specific elasticities. Levels of significance: *10%, ** 5%, and *** 1%.

3.8 Appendix Tables and Figures

Figure 3-5: Multi-Apping

How to run the apps at the same time

Once you have everything set up you will need to actually open and run both apps at the same time. This is pretty easy for the most part, but there are a few tricks that will help you out.

First, close out any apps you have open. Driving for Lyft and Uber at the same time requires you to keep both apps open until you get a request, which puts quite a bit of strain on your phone, sucks the battery dry, and uses a ton of data in the process. So bring a charger, stay on task, and be prepared to pull down a lot of data. And I mean, a TON.

Second, try to minimize surfing the web or being active on social media. Keep in mind you have two apps that are constantly talking to their respective platforms, so the chance of your phone crashing or glitching out becomes much higher than usual. The last thing you want is to get a ride request and have your phone freeze up or crash, costing you both time and money.

It is also worth noting that when you have both of the apps open, you will need to have Uber open on the main screen, with Lyft running in the background behind it. Uber will automatically close out after a minute or two if it is running in the background, but Lyft will stay open.

Accepting ride requests

When you get a ride request, make sure to accept one and immediately log right out of the other. It won't take long to get a ride request, and on a busy night you may get two at the same time. If this happens, pick the once closest to you and decline the other. As every driver knows, the best way to make money is in volume, and the less dead miles you have, the better.

Note: This is a screenshot from rideshareapps.com's guide on how to drive for Lyft and Uber at the same time. This is intended to illustrate that multiple non-Uber/Lyft affiliated forums provide information to drivers on how to maximize earnings via multi-apping.

Figure 3-6: Earnings Accelerator Messaging

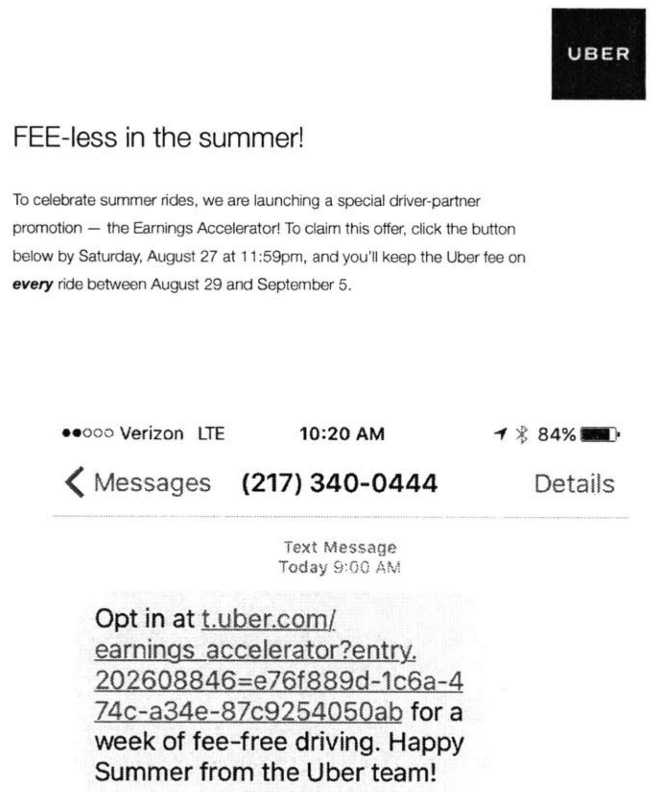

Note: This figure shows the e-mails and text messages that were sent to drivers selected for the Earnings Accelerator. The link in the text message and the link in the email (not included in the picture) directed the driver to a more detailed opt-in page with information on how the incentive worked and with consent language. This opt-in form is depicted in Figure 3-7.

Figure 3-7: Earnings Accelerator Opt-In Form

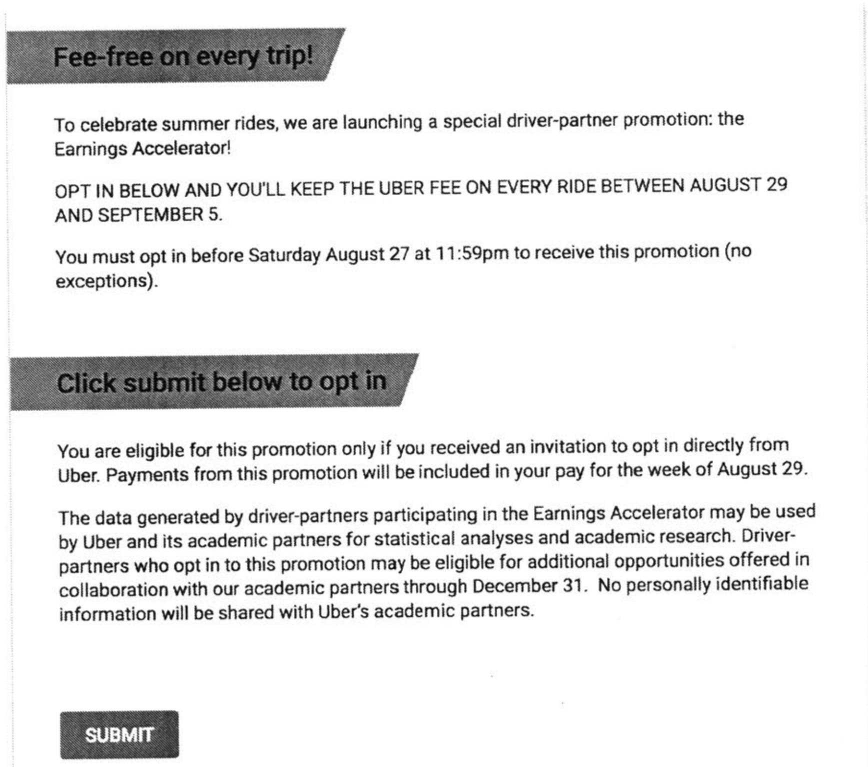

Fee-free on every trip!

To celebrate summer rides, we are launching a special driver-partner promotion: the Earnings Accelerator!

OPT IN BELOW AND YOU'LL KEEP THE UBER FEE ON EVERY RIDE BETWEEN AUGUST 29 AND SEPTEMBER 5.

You must opt in before Saturday August 27 at 11:59pm to receive this promotion (no exceptions).

Click submit below to opt in

You are eligible for this promotion only if you received an invitation to opt in directly from Uber. Payments from this promotion will be included in your pay for the week of August 29.

The data generated by driver-partners participating in the Earnings Accelerator may be used by Uber and its academic partners for statistical analyses and academic research. Driver-partners who opt in to this promotion may be eligible for additional opportunities offered in collaboration with our academic partners through December 31. No personally identifiable information will be shared with Uber's academic partners.

SUBMIT

Note: This is a screenshot from the opt-in form sent to drivers included in the first Earnings Accelerator experiment (in Boston). Drivers were sent to this page via the in-app notification and via text messages and e-mails they received throughout opt-in week. The form pre-filled with their unique Uber identifier (not included in screenshot). To opt in, the driver only needed to scroll to the bottom of the page and click submit.

Figure 3-8: Sample Weekly Pay Statement

Weekly Earnings

Trip Earnings	**$19.34**
Fare	$24.12
Uber Fee	- $6.03
Toll	+ $1.25
Estimated Payout	**$19.34**

Note: This figure shows that, when we ran the experiment, drivers' weekly pay statements listed (1) how much they collected in trip receipts, (2) how much of this went to Uber in the form of the Uber fee, and (3) what, if any, reimbursements they received for tolls. Their estimated payout was the sum of these three items. The structure of drivers' weekly earning statements has changed since we ran the experiment.

Figure 3-9: Individual Driver Bonus: Budget Set

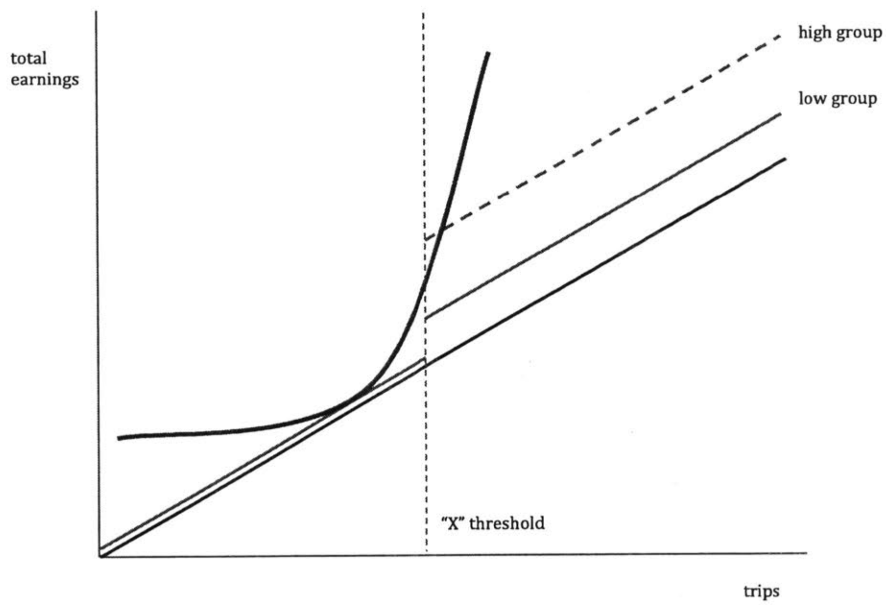

Note: This figure shows a stylized budget set for the IDB incentive studied in Section 3.5.2 of this paper. The x-axis denotes trips and the y-axis denotes total take-home earnings. For a given trip threshold X, drivers in the IDB were told that they would receive a lump-sum bonus for exceeding the trip threshold. The amount of the bonus and the trip threshold varied across drivers and weeks. The blue curve shows a utility curve of a driver who, in the absence of the incentive would not exceed the trip threshold. The budget set for the high bonus cuts through her indifference curve, though the budget set for the low bonus does not. Therefore this driver will only exceed the trip threshold (labeled X) if offered the high bonus.

Figure 3-10: Sample Trip Receipt

Note: Individual trip receipts showed both the fare and the amount of the fee, if applicable. If a driver was driving fee-free, the fare would be equal to the estimated payout.

Figure 3-11: Earnings Accelerator Lease Calculator

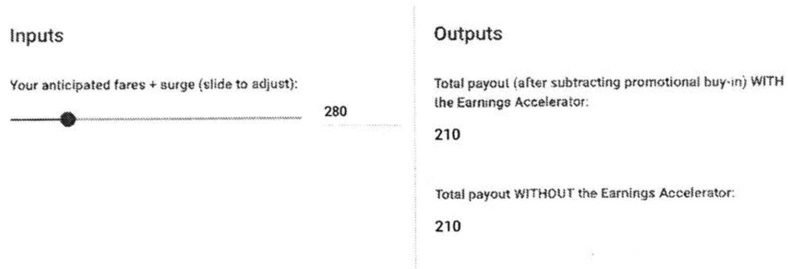

Note: Each driver who was offered a Taxi offer was sent a slider that allowed them to compare the earnings they would receive if they accepted the offer (net of the lease) to the earnings they would normally receive. The slider was set to load at the breakeven (the place where treated and untreated earnings would be identical).

Figure 3-12: Earnings Accelerator Buy-In

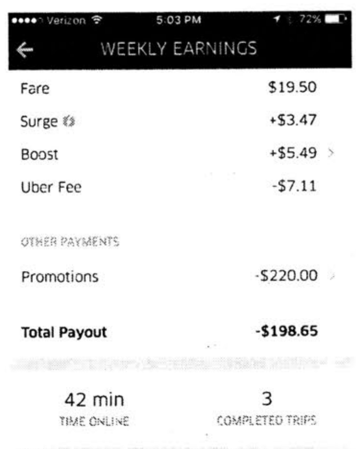

Note: This shows a sample trip receipt of a driver who accepted one of the Taxi offers. The lease amount is broken out from the trip receipts, and is identified as a promotional payment. This is not a screenshot from a driver in our experiment; we did not offer leases that cost $220.

Figure 3-13: Hours Worked by Male and Female Drivers

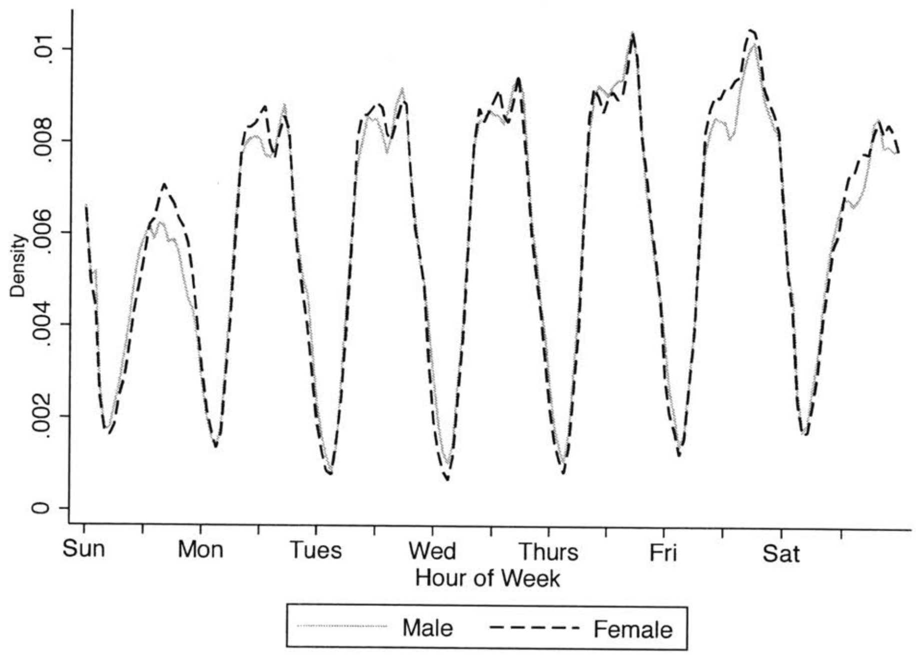

Note: This figure uses data from the first Houston experiment to plot the density of hours worked (by male and female drivers) over the course of the week. The data include only non-treatment weeks.

Figure 3-14: Age and Experience Distributions by Subgroup: IDB Sample

Panel A: Age

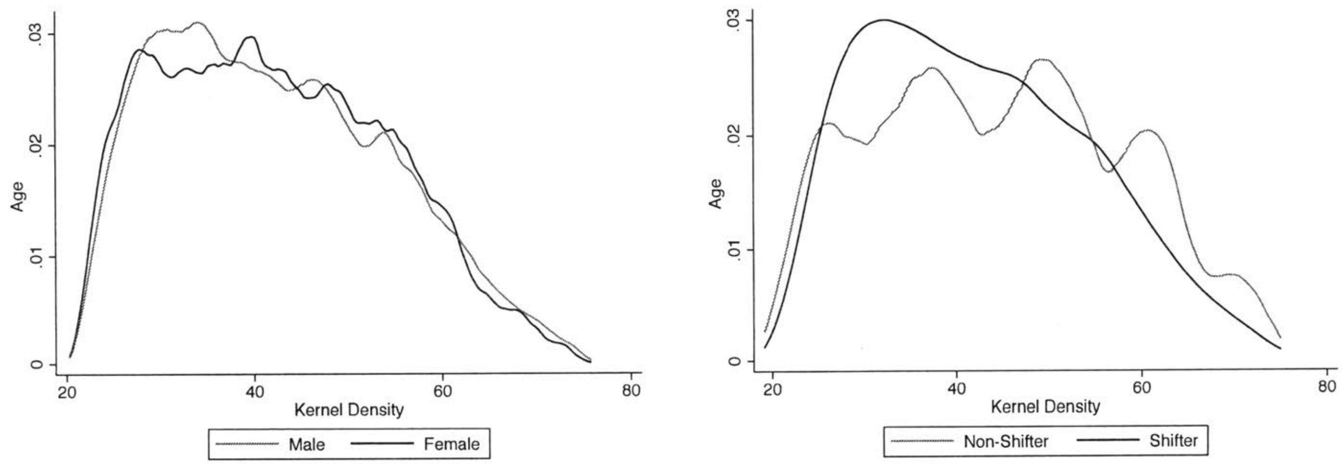

Panel B: Months on Platform

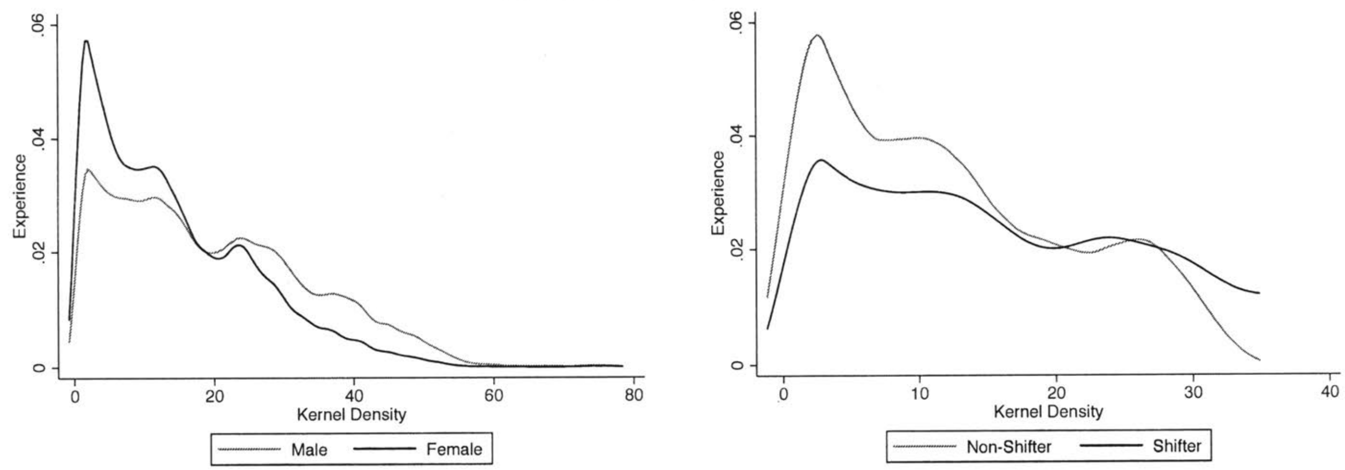

Note: This figure presents kernel densities of driver age and experience (months on platform) in the IDB sample.

Figure 3-15: Residual Density of Log Trips Distribution

Panel A: All Drivers

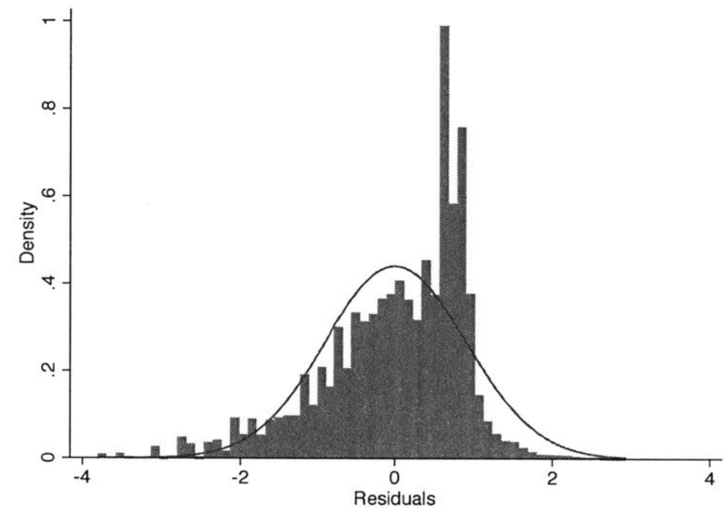

Panel B: Omitting Drivers Near Threshold

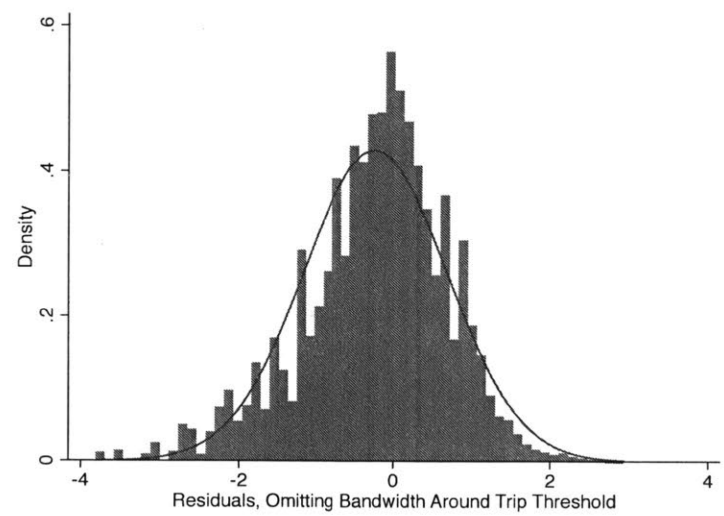

Note: Panels A and B present histograms of the residual log trip distribution for the IDB sample. The residuals are computed by first regressing log trips on date dummies and on the strata used to balance the high and low groups. Panel A plots the distribution of the full set of residuals. Panel B plots the same residuals but omits observations from drivers whose trips fell within a 2-trip bandwidth of their assigned trip threshold.

Table 3.13: Earnings Accelerator Balance: Boston

	Eligible Drivers (1)	Experimental Difference (2)	Week 1 - Week 2 Difference (3)	Taxi Week 1 Treated - Control Difference (4)	Taxi Week 2 Treated - Control Difference (5)
Female	0.14	0.00	-0.02	-0.01	0.01
		(0.010)	(0.02)	(0.02)	(0.02)
Age	41.28	0.19	1.05	1.08	-0.08
		(0.362)	(0.66)	(0.76)	(0.76)
Vehicle Solutions	0.08	0.00	0.00	0.03	0.02
		(0.008)	(0.02)	(0.02)	(0.02)
Months Since Signup	14.26	-0.26*	0.00	-0.19	-0.35
		(0.138)	(0.25)	(0.29)	(0.29)
Hours Week Prior to Offer	15.14	-0.08	-0.74	0.29	0.83
		(0.259)	(0.51)	(0.85)	(0.77)
Earnings Week Prior to Offer	20.06	0.25	-10.76	2.85	11.39
		(0.189)	(11.46)	(22.40)	(23.28)
F-statistic		1.11	1.240	0.915	1.577
p-value		0.35	0.283	0.483	0.150
Observations		8685	1600	1031	1031

Note: Column 1 presents the mean value of the indicated characteristic for Boston drivers who were eligible inclusion in the first Earnings Accelerator experiment. These are drivers who completed at least four trips in the prior month and whose average hours per week (conditional on driving) are between 5 and 25 hours. Column 2 presents the strata-adjusted difference between drivers selected for the experiment and all eligible drivers. Column 3 presents the strata-adjusted difference between the 800 drivers offered free week in the first week and the 800 drivers offered free week in the second week. Columns 4 and 5 present the strata-adjusted difference between drivers offered a Taxi contract and drivers not offered a contract. Only the 1031 drivers who accepted the free week offer were included in this phase of the experiment. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.14: Earnings Accelerator Balance: Houston 1

	Eligible Drivers (1)	Experimental Difference (2)	Week 1 - Week 2 Difference (3)	Taxi Week 1 Treated - Control Difference (4)	Taxi Week 2 Treated - Control Difference (5)
Female	0.20		0.02	0.01	0.00
			(0.02)	(0.02)	(0.02)
Vehicle Solutions	0.10	0.00	0.01	-0.02	0.00
		(0.01)	(0.01)	(0.02)	(0.02)
Months Since Signup	10.48	0.24	0.30	-0.42	0.55**
		(0.16)	(0.21)	(0.26)	(0.26)
Hours Week Prior to Offer	16.84	-0.59	-0.37	-0.88	-0.16
		(0.39)	(0.57)	(0.82)	(0.93)
Earnings Week Prior to Offer	262.74	-4.11	-6.88	-15.06	-4.29
		(7.32)	(8.97)	(10.65)	(14.36)
F-statistic		1.49	0.84	1.18	0.94
p-value		0.20	0.52	0.32	0.45
Observations		10641	2020	1355	1355

Note: Column 1 presents the mean value of the indicated characteristic for Houston drivers who were eligible inclusion in the second Earnings Accelerator experiment. These are drivers who completed at least four trips in the prior month and whose average hours per week (conditional on driving) are between 5 and 40 hours. Column 2 presents the strata-adjusted difference between drivers selected for the experiment and all eligible drivers. Drivers were selected within strata based on hours bandwidths, commissions, and gender. Column 3 presents the strata-adjusted difference between the drivers offered free week in the first week and the drivers offered free week in the second week. Columns 4 and 5 present the strata-adjusted difference between drivers offered a Taxi contract and drivers not offered a contract. The Taxi randomization was conducted within hours bandwidth by commission groups. Only the 1355 drivers who accepted the free week offer were included in this phase of the experiment. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.15: Earnings Accelerator Balance: Houston 2

	Eligible Drivers (1)	Experimental Difference (2)	Week 1 - Week 2 Difference (3)	Taxi Week Treated - Control Difference (4)
Female	0.09		0.01	0.01
			(0.02)	(0.03)
Vehicle Solutions	0.07	0.01	0.01	-0.02
		(0.01)	(0.01)	(0.02)
Months Since Signup	9.62	-0.09	-0.19	-0.33
		(0.18)	(0.25)	(0.33)
Hours Week Prior to Offer	20.02	-0.50	0.71	0.75
		(0.41)	(0.62)	(0.81)
Earnings Week Prior to Offer	264.46	-3.08	3.75	7.44
		(6.01)	(8.44)	(12.59)
F-statistic		1.18	0.96	0.69
p-value		0.32	0.44	0.63
Observations		9124	2100	1270

Note: Column 1 presents the mean value of the indicated characteristic for Houston drivers who were eligible inclusion in the third Earnings Accelerator experiment. These are drivers who completed at least four trips in the prior month and whose average hours per week (conditional on driving) are between 5 and 40 hours. Column 2 presents the strata-adjusted difference between drivers selected for the experiment and all eligible drivers. Drivers were selected within strata based on hours bandwidths, commissions, and gender. Column 3 presents the strata-adjusted difference between the drivers offered free week in the first week and the drivers offered free week in the second week. Column 4 presents the strata-adjusted difference between drivers offered a Taxi contract and drivers not offered a contract. The Taxi randomization was conducted within hours bandwidth by commission groups. Only the 1270 drivers who accepted the free week offer were included in this phase of the experiment. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.16: Taxi Treatments: Boston

Boston Taxi Treatments						
	20% Fee Class			25% Fee Class		
Bandwidth (1)	Lease (2)	New Fee (3)	Treatment Fraction (4)	Lease (5)	New Fee (6)	Treatment Fraction (7)
			Week 1			
High	$110	0	40%	$165	-0.125	20%
	$110	0	40%	$165	-0.125	20%
Low	$45	0	40%	$75	-0.125	20%
	$45	0	40%	$75	-0.125	20%
			Week 2			
High	$110	0	40%	$165	-0.125	20%
	$55	0	30%	$35	0.125	30%
Low	$40	0	30%	$15	0.10	30%
	$35	0	30%	$15	0.125	30%

Note: This table presents the taxi treatments offered to drivers included in the Boston Earnings Accelerator experiment. Only the 1031 drivers who accepted fee-free driving were included in this phase of the experiment and the treatment fraction refers to the fraction of consented drivers in each hours bandwidth and commission who were offered a given taxi offer.

Table 3.17: Taxi Treatments: Houston 1

Houston Taxi Treatments						
	20% Fee Class			28% Fee Class		
Bandwidth	Lease	New Fee	Treatment Fraction	Lease	New Fee	Treatment Fraction
(1)	(3)	(4)	(5)	(3)	(4)	(5)
			Week 1			
Very High	$100	0	60%	$120	0	60%
High	$40	0	60%	$50	0	60%
Low	$15	0	60%	$15	0	60%
			Week 2			
Very High	$65	0	60%	$90	0	60%
High	$35	0	60%	$35	0	60%
Low	$10	0	60%	$10	0	60%

Note: This table presents the taxi treatments offered to drivers included in the taxi phase of the first Houston experiment. Only drivers who accepted fee-free driving were included in this experiment and the treatment fraction refers to the fraction of consented drivers in each hours bandwidth and commission who were offered a given taxi offer.

Table 3.18: Opt-In Rates

		By Hours Group			By Months on Platform		By Age	
	All	Low	High	Very High	Experienced	Inexperienced	35 or Younger	Older than 35
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				A. Male Drivers				
Opt-In	0.61***	0.58***	0.63***	0.65***	0.60***	0.63***	0.59***	0.63***
	(0.02)	(0.03)	(0.02)	(0.03)	(0.02)	(0.03)	(0.03)	(0.02)
Observations	3130	1004	2126	1076	2488	642	998	2144
				B. Female Drivers				
Opt-In	0.73***	0.67***	0.75***	0.78***	0.74***	0.71***	0.64***	0.76***
	(0.01)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)
Observations	2096	698	1398	698	1016	1080	556	1540

Note: This table presents opt-in rates for fee-free driving in the first Houston experiment. The opt-in rates are adjusted for the strata used for random assignment. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.19: Frisch Elasticities Without Covariates

	Free Week		Taxi		Stacked	
	Male	Female	Male	Female	Male	Female
	(1)	(2)	(3)	(4)	(5)	(6)
			A. Just-Identified 2SLS			
Log Wages	0.47*	0.63***	0.71**	1.38***	0.58***	0.96***
	(0.24)	(0.16)	(0.29)	(0.27)	(0.19)	(0.14)
			B. Over-Identified 2SLS			
Log Wages	0.34	0.56***	0.79***	1.23***	0.58***	0.90***
	(0.25)	(0.16)	(0.30)	(0.24)	(0.20)	(0.14)
First Stage F Statistic	16.5	22.1	5.8	7.2	10.1	11.8
p-value from test of equality	0.136		0.060		0.023	
			C. LIML			
Log Wages	0.34	0.57***	0.81***	1.26***	0.59***	0.93***
	(0.26)	(0.17)	(0.31)	(0.25)	(0.21)	(0.15)
p-value from test of equality	0.136		0.062		0.024	
Drivers	785	852	504	604	804	871
Observations	1450	1548	928	1105	2378	2653

Note: All models control for the strata used for random assignment and for date fixed effects, both of which are interacted with gender. The p-values for the 2SLS and LIML models come from stacked models where the coefficient on each covariate is allowed to vary by sex. Table 3.5 presents analogous results, controlling for baseline covariates. Standard errors are clustered by driver. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.20: Probit Coefficients

	Probit Coefficients			
	Unweighted		Weighted	
	Probit	IV-Probit	Probit	IV-Probit
	(1)	(2)	(3)	(4)
Log(Bonus)				
Male Non-Shifters	0.633***	0.779***	0.614***	0.787***
	(0.063)	(0.108)	(0.062)	(0.103)
Male Shifters	0.672***	0.703***	0.634***	0.621***
	(0.007)	(0.014)	(0.008)	(0.015)
Female Non-Shifters	0.550***	0.650***	0.504***	0.612***
	(0.132)	(0.199)	(0.131)	(0.191)
Female Shifters	0.679***	0.713***	0.644***	0.660***
	(0.015)	(0.028)	(0.015)	(0.027)
Log(Threshold)				
Male Non-Shifters	-1.352***	-1.500***	-1.301***	-1.451***
	(0.091)	(0.122)	(0.088)	(0.118)
Male Shifters	-1.167***	-1.218***	-1.093***	-1.078***
	(0.034)	(0.039)	(0.034)	(0.041)
Female Non-Shifters	-1.327***	-1.442***	-1.254***	-1.355***
	(0.181)	(0.218)	(0.177)	(0.218)
Female Shifters	-1.263***	-1.316***	-1.187***	-1.198***
	(0.039)	(0.047)	(0.038)	(0.048)
Observations	1047998	1047998	1047998	1047998

Note: This table presents probit coefficients from equation 3.13. The corresponding elasticities are presented in Table 3.10. Standard errors are clustered by driver. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.21: Mean Utilization Rates

		By Months on Platform		By Age	
	All	Experienced	Inexperienced	35 or Younger	Older than 35
	(1)	(2)	(3)	(4)	(5)
Male Non-Shifter	0.592	0.589	0.601	0.580	0.597
	(0.149)	(0.139)	(0.175)	(0.166)	(0.140)
Male Shifter	0.639	0.644	0.632	0.662	0.662
	(0.177)	(0.172)	(0.183)	(0.162)	(0.165)
Female Non-Shifter	0.590	0.599	0.582	0.596	0.588
	(0.157)	(0.146)	(0.167)	(0.162)	(0.156)
Female Shifter	0.618	0.680	0.612	0.616	0.660
	(0.194)	(0.142)	(0.197)	(0.198)	(0.170)
Observations	5229	2558	2671	1049	2261

Note: This table presents mean weekly utilization rates (minutes en route or on trip / minutes online) for male and female shifters calculated using data from all three Earnings Accelerator experiments. The data include all opt-in and non-treatment weeks between the start of the experiment and the end of the last taxi week. The results are discussed in Appendix Section 3.9.2.

Table 3.22: Mean Utilization Rates by City

		By Months on Platform		By Age	
	All	Experienced	Inexperienced	35 or Younger	Older than 35
	(1)	(2)	(3)	(4)	(5)
			A. Boston		
Male Non-Shifter	0.639	0.628	0.649	0.633	0.640
	(0.194)	(0.191)	(0.199)	(0.197)	(0.194)
Male Shifter	0.662	0.680	0.641	0.662	0.662
	(0.164)	(0.145)	(0.181)	(0.162)	(0.165)
Female Non-Shifter	0.610	0.610	0.611	0.654	0.599
	(0.210)	(0.256)	(0.194)	(0.152)	(0.225)
Female Shifter	0.643	0.681	0.623	0.616	0.660
	(0.183)	(0.147)	(0.197)	(0.198)	(0.170)
Observations	1495	744	751	537	947
			B. Houston		
Male Non-Shifter	0.583	0.584	0.577	0.569	0.590
	(0.137)	(0.132)	(0.158)	(0.158)	(0.126)
Male Shifter	0.617	0.611	0.624		
	(0.186)	(0.187)	(0.184)		
Female Non-Shifter	0.590	0.599	0.581	0.595	0.588
	(0.156)	(0.144)	(0.166)	(0.162)	(0.154)
Female Shifter	0.611	0.675	0.610		
	(0.196)	(0.100)	(0.197)		
Observations	3734	1814	1920	512	1314

Note: This table presents mean weekly utilization rates (minutes worked / minutes active) for male and female shifters calculated using data from all three Earnings Accelerator experiments. The data include all opt-in and non-treatment weeks between the start of the experiment and the end of the last taxi week. The results are discussed in Appendix Section 3.9.2. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.23: Treatment Effects on Utilization

	Effect on Utilization for Shifters and Non-Shifters				
		By Months on Platform		By Age	
	All	Experienced	Inexperienced	35 or Younger	Older than 35
	(1)	(2)	(3)	(4)	(5)
Non-Shifters	-0.01**	-0.01	-0.01**	-0.01	-0.02**
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Shifters	-0.02***	-0.03***	-0.02***	-0.01	-0.02***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Observations	12686	4384	8302	3155	6699

Note: This table presents estimates of the impact of treatment offers on utilization rates. The results are discussed in Appendix Section 3.9.2. Levels of significance: *10%, ** 5%, and *** 1%.

Table 3.24: IDB: Treatment Effects

	Trips			Above Threshold			IDB Payout		
	All	Male	Female	All	Male	Female	All	Male	Female
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Low Bonus Mean	29.587	30.654	23.166	0.378	0.387	0.323	49.496	51.553	37.156
Non-Shifter	1.199***	1.486***	1.118	0.050***	0.050***	0.06**	11.990***	12.269***	11.437***
	(0.403)	(0.391)	(0.789)	(0.011)	(0.011)	(0.028)	(1.413)	(1.554)	(3.050)
Shifter	1.524***	1.636***	1.282***	0.054***	0.055***	0.051***	14.552***	15.018***	11.924***
	(0.048)	(0.045)	(0.100)	(0.001)	(0.001)	(0.003)	(0.158)	(0.170)	(0.383)
Observations	845080	731483	113597	845080	731483	113597	845080	731483	113597

Note: The first row shows the mean outcome among all observations where the driver was offered the "low" bonus. Subsequent rows show the impact of the high bonus treatment on outcomes for non-shifters and shifters. All models control for date fixed effects, the strata used in Table 3.4, and one lag of trips. Standard errors are clustered by driver. The results are discussed in Appendix Section 3.9.3.

3.9 Supplementary Results

3.9.1 Robustness of Intensive Margin Elasticities

We consider two alternative explanations for our finding that women are more elastic than men on the intensive and extensive margins: differences in family responsibilities and differences in the availability of alternative jobs. We do not find any evidence for either of these explanations.

One concern is that changes in hours worked on Uber represent shifts in hours from alternative jobs. While we are able to rule out shifting from Lyft, Uber's main competitor in the marketplace, it is possible that Uber drivers have other sources of employment. In order to bias our estimates, however, the hours at these alternative jobs would need to be changeable at relatively high frequency (within a week) and differ by gender. We are not aware of any evidence that female Uber drivers are more engaged in other types of flexible work than their male counterparts. However, to address this concern we re-estimate the elasticities within a group of drivers whom we observe working more than forty hours per week in any of the four weeks of data we used to sample drivers. It is unlikely that this group of workers has substantial outside (non-Uber) employment. Column 7 of Table 3.6 shows that the gap between male and female elasticities shrinks, but that there is still a large difference between the responsiveness of male and female Uber drivers in this group.

Another possibility is that even conditional on the number of hours worked, women may be less able to adjust their hours due to family responsibilities. These constraints on women's hours would lead us to under-estimate their labor supply elasticities, and under-state the differences between their behavior and that of men. To address this concern, we identify a sample of women who typically drive during hours when we think family responsibilities may loom large: women who regularly work between 3 and 7 P.M. on week-days. We then re-estimate equation 3.6 within this subsample. Column 8 of Table 3.6 shows that the gap in male-female elasticities is qualitatively similar in this subsample.

3.9.2 Utilization

We can look for suggestive evidence of platform shifting by comparing the utilization rates of shifters and non-shifters. While switching between the Lyft and Uber apps at low frequency will not impact utilization rates, multi-apping (running both apps at the same time) will allow drivers to achieve a higher Uber utilization rate (and as a result higher hourly earnings when considering earnings from multiple platforms) by spending less time idle between trips.

Table 3.21 presents suggestive evidence of higher utilization rates among shifters, consistent with multi-apping. The gap is slightly larger for men than for women. The fact that young drivers (those under 35) seem to have larger gaps in utilization is consistent with young people being more adept with technology.

This pattern is not driven by city differences in utilization rates. Table 3.22 breaks down the results for Boston and Houston. While the Houston shifters and non-shifters necessarily come from different weeks, which could have different utilization rates due to customer demand, the Boston results show that there are still gaps between shifters and non-shifters when we compare utilization rates within a single city and week.

We can look for more evidence of multi-apping by comparing the impact of the treatment on utilization rates of shifters and non-shifters. Assuming drivers first pick the "best" hours in a day to drive, and slowly move down the utilization curve (perhaps conditional on personal hours constraints), treated compliers should see a decrease in utilization rates, relative to untreated compliers. Because shifters move a greater number of hours to the Uber platform they should see a larger decrease in utilization rates.

Table 3.23 shows that this is indeed the case. Specifically, it presents estimates of β from:

$$\text{utilization}_{it} = \beta \text{Offer} + \gamma_t + X_{it} + \epsilon_{it}$$

where Offer is the experimental percentage increase in wages. The sample is a pooled sample that includes drivers in all three Earnings Accelerator experiments. Column 1 shows that both shifters and non-shifters have lower utilization rates when they are treated. This is consistent with the fact that, as drivers work more hours, they start to work less valuable hours (with lower utilization rates). The fact that the impact is larger for shifters is consistent with these drivers having artificially high utilization rates pre-treatment, due to multi-apping. When they receive the Earnings Accelerator wage increase, they decide to shift all (or most) of their hours to Uber, even though this means lower utilization rates on both marginal and infra-marginal hours. We see a similar pattern across subgroups of drivers defined by driver experience or age. It is somewhat hard to interpret the magnitudes, because, depending on the usual number of hours worked shifters and non-shifters (or men and women) may be on parts of the utilization curve with very different slopes.

3.9.3 IDB Treatment Effects

Table 3.24 presents treatment effects for four groups of drivers, defined by sex and whether the driver is eligible to drive for Lyft (based on the age of their car). Specifically, we estimate

$$y_{it} = \beta^S \text{Treated}_{it} \times \text{Shift}_i + \beta^N \text{Treated}_{it} \times \text{Non} - \text{Shift}_i + X_{it} + \epsilon_{it} \qquad (3.15)$$

where Treated_{it} is an indicator for whether individual i is in the high bonus group in period t, Shift_i is an indicator for whether the driver has a car that allows them to drive for Lyft, γ_t is a full set of date fixed effects, and X_{it} includes one lag of trips, and the strata used for random assignment. Standard errors are clustered by driver. Because high and low offers are, conditional on strata, as good as randomly assigned, this specification allows us to measure the impact of a high offer on drivers' labor supply, relative to the impact of a low offer.

The first three columns present estimates of β^N and β^S when we use total trips completed as the outcome variable. Column 1 shows that shifters who receive the high bonus complete 1.5 additional trips; non-shifters who receive the high bonus complete 1.2 additional trips. Columns 2 and 3 show that among both male and female drivers, shifters increase their labor supply more than non-shifters in response to the high offer. While the difference between shifters and non-shifters is not statistically significant, this is not surprising, given the structure of the promotion. The promotion only incentivizes drivers to complete more trips if they believe they are able to cross the threshold. Furthermore, there is no incentive to drive more once the driver has crossed the threshold.

Columns 4-6 present estimates of equation 3.15 where $1\{\text{cross threshold}\}$ is the outcome variable. Shifters who receive the high bonus are 5.5 percentage points more likely to cross the threshold, relative to 5.0 percentage points for non-shifters. However, the difference is not statistically significant. This largely reflects the fact that we have not exploited week to week and driver to driver variation in the strength of the incentive. We start to see a significant difference between shifters and non-shifters when we look at the amount of the bonus drivers receive.[29]

[29] The fact that male shifters and non-shifters in the high bonus group appear to benefit more (in dollar terms) than their female counterparts reflects both differences in labor supply, and the fact that the male drivers are more concentrated in the lucrative, high-threshold groups.

3.10 Empirical Appendix

3.10.1 Construction of Hours and Earnings

Hours A driver is considered to be working whenever their Uber app is on and they have indicated that they are available for a dispatch. This includes three distinct periods: time waiting for a trip, time traveling to a pickup, and time on a trip. The utilization rate is the fraction of hours spent in the second two periods.

Earnings Uber distinguishes between gross earnings– which include promotional incentives—and net earnings—which subtract the amount the driver paid in Uber fees. Both of these measures are not net of costs the driver may incur, including gas or depreciation to the driver's vehicle. We focus on gross earnings.

3.10.2 The Earnings Accelerator

This section provides more detail on the implementation of the three Earnings Accelerator experiments. The first experiment was conducted in fall 2016 in Boston and was analyzed in Angrist et al. (2017). The second and third experiments were conducted in Houston and have not been used in other work.

Boston: Fall 2016

In August-October 2016, we conducted the first of our three Earnings Accelerator experiments. As in both subsequent experiments, there were three phases: (1) the selection of eligible drivers, (2) "fee-free" offers, and (3) taxi offers. Table 3.1 lists the timeline for the three phases.

Drivers were eligible for inclusion in the Boston experiment if they had completed at least 4 trips in July 2016 (were "active" drivers) and if their average hours per week, conditional on driving, were between 5 and 25 hours per week. We excluded very high hours drivers to reduce the cost of the experiment. We grouped drivers into two bandwidths based on their average hours per week. "Low-hours" drivers drove an average of 5-15 hours/week and "high-hours" drivers drove an average of 15-25 hours/week. Roughly 45% of Boston drivers were eligible for inclusion in the experiment.

We randomly selected 1600 eligible drivers for inclusion in the experiment within strata defined by average hours driven in July, driver fee class (commission rate), and vehicle model year. All of these drivers were offered one week of fee-free driving. Half were offered fee-free driving one week (wave 1); half were offered it in the next week (wave 2). Column 3 of Table 3.13 shows that drivers offered fee-free driving in wave 1 were statistically indistinguishable from those offered fee-free driving in wave 2. Drivers were notified about the Earnings Accelerator offer via e-mails, text message, and in-app notification. The in-app notification stayed at the top of each driver's Uber app for the entire opt-in period, and drivers received reminder e-mails and text messages throughout the week. Figure 3-7 shows sample messaging. Each message contained a link to a Google Form, which provided more information on the incentive. In particular, this form indicated the exact time the incentive would be active (Monday 4 A.M. for one week, following the standard Uber week) and informed the drivers that if they opted in to the Earnings Accelerator, their data would be used by academic researchers. One thousand and thirty-one of the 1600 drivers chose to opt-in to the promotion.

At the time we ran the experiment, drivers' trip receipts typically showed three things: the amount collected from the rider, the amount collected by Uber (due to the proportional fee), and the amount they were paid (the difference between the two). Drivers who accepted the offer of fee-free driving were able to see in-app that their fees were zero (see Figure 3-10 for a sample trip receipt). They also received e-mail, text message, and in-app reminders throughout the week that the "Earnings Accelerator [was] on" and that they were earning more on every trip. The messaging was crafted so as to mimic that used for standard Uber promotions.

Drivers who opted in to fee-free driving were included in the third phase of the experiment, the Taxi treatments. In each of two weeks, we offered drivers the opportunity to buy additional weeks of fee-free driving for a pre-specified cost. Table 3.16 shows the Taxi contracts offered in each of the two weeks, along with the probability of selection and the percentage of drivers who accepted our offers. Columns 4 and 5 of Table 3.13 shows that the taxi treatment and control groups were balanced during both weeks of treatment.

In each week we offered two types of taxi contracts, where the lease varied by hours bandwidth and commission. In the first week some drivers received the opportunity to buy an additional week of fee-free driving; others received the opportunity to buy a week of negative fee (-.125%) driving. Offers in the second week were less generous, but were priced accordingly. The e-mails and opt-in forms contained information on the "break-even" a driver must exceed (in terms of gross earnings) to make buying the contract worthwhile. They also contained links to online calculators that allowed the drivers to calculate their earnings with and without the Earnings Accelerator. A screenshot from one of these calculators is shown in Figure 3-11. Driver who accepted the offer had the lease payment subtracted from their opt-in week earnings statement, as shown in Figure 3-12. They saw their increased earnings in-app, just as they did during fee-free week, with one exception: drivers who bought a "negative fee" contract during the first Boston Taxi week. It was not possible to implement a negative fee using the Uber platform. These drivers saw no fee in-app and received additional text and email reminders that they would receive an additional 12.5% on their weekly pay statement.

Houston: Spring 2017

In spring 2017 we conducted a second round of the Earnings Accelerator in Houston, Texas. Uber launched operations in Houston in July 2013 and by the spring of 2017 had over 15,000 active drivers (drivers who had completed at least four trips in the previous month). Lyft entered the market in February 2014 but suspended operations in August 2017 after the Houston City Council passed new TNC regulations which mandated a stricter background check for drivers. They had fully withdrawn from the Houston market by November 2017. Uber remained operational in Houston despite the new regulation.

Relative to the Boston experiment, there were two key modifications. First, we included a third hours group, including drivers who drove between 25 and 40 hours per week on average in the month prior to the experiment. These drivers had been omitted from the Boston experiment due to budgetary considerations. Their inclusion allowed us to examine the responsiveness of drivers who were working more than part-time on the Uber platform. Second, we over-sampled female drivers so that we could explore gender differences in labor supply elasticities.

Drivers were eligible for inclusion in the first Houston experiment if they completed at least 4 trips in the prior month (were "active" drivers) and if their average hours per week, conditional on driving, were between 5 and 40 hours per week in the month before the experiment. Within this sample of eligible drivers, we randomly selected 2020 drivers for inclusion in the experiment within six strata defined by the interaction of hours bandwidth and gender. The messaging and notifications mirrored that of the Boston experiment. As before, we offered half of the drivers the opportunity to drive fee-free in one week, and the other half of the drivers the same opportunity the next week. Column 3 of Table 3.14 shows that drivers offered fee-free driving in wave 1 were statistically indistinguishable from those offered fee-free driving in wave 2. One thousand, three hundred and fifty-five drivers accepted our offer and were included in the third phase of the experiment.

Table 3.17 shows the taxi contracts offered in each of the two Taxi weeks, along with the probability of selection and the percentage of drivers who accepted our offers. Due to the addition of the third hours bandwidth and logistical constraints on the number of treatments we could implement at a single time, we eliminated the negative and half fee treatments and only offered drivers the opportunity to buy weeks of fee-free driving. Columns 4 and 5 of Table 3.14 shows that the taxi treatment and control groups were balanced during both weeks of treatment.

Houston: Fall 2017

We conducted a third round of the Earnings Accelerator several months after Lyft re-entered the Houston market. In May 17, 2017, the Texas State Legislature passed bill, H.B. 100, with a supermajority in the Senate (21-9), and on May 29, the Governor signed it, immediately removing mandatory fingerprinting. Lyft announced its intention to resume operations and re-entered Houston at 2 p.m. C.T. on May 31, 2017.

The second Houston experiment was hampered by a number of implementation issues, which complicate the analysis. First, the experiment took place only a few weeks following Hurricane Harvey, which flooded much of Houston. While we made every attempt to recruit drivers who were not affected by the hurricane, we saw significantly lower opt-in rates than we saw in either Boston or in our first Houston experiment, suggesting that some drivers may not have fully recovered. Second, changes to the Uber app made it impossible to implement the increased wages in the usual manner, except during the first week of fee-free driving.

Drivers were eligible for inclusion in the third iteration of the Earnings Accelerator if – as before – they had completed at least four trips in the prior month (were "active" drivers), if their average hours per week, conditional on driving, were between 5 and 40 hours per week, and if they had completed a trip in Houston after Uber re-started operations following Hurricane Harvey.[30] The messaging, notifications, and timeline were similar to the first Houston experiment. The one key change was reference to Uber's fee: Uber changed its policy in June 2017 to loosen the link between rider fares and driver earnings—this was called "up front pricing"—and removed the concept of the Uber "fee". What drivers earned per trips did not change; it remained a function of a base fare plus a per-mile rate and a per-minute rate. As a result, we did not mention the "Uber fee" in the second Houston experiment. Instead we focused our messaging on the proportional increase in earnings. Column 3 of Table 3.15 shows that drivers were balanced across waves one and two.

We included 2100 drivers in the second Houston experiment. In the first treatment, drivers received one of four multipliers on total earnings at no cost (the equivalent of fee-free driving): 1.2x, 1.3x, 1.4x, and 1.5x. The first wave were offered the multiplier in the week of September 18 and saw the treatment in-app (as a proportional increase in the base fare, per-mile rate, and per-minute rate). Due to technical constraints, the second wave—who were offered the multiplier in the week of September 25—did not see the treatment in-app and instead received a lump-sum bonus at

[30]The total volume of trips had rebounded to pre-hurricane levels by the time we conducted the experiment, but we did not want to include drivers who had stopped driving because they had been personally impacted by the hurricane.

the end of the week. A total of 1270 drivers accepted our "fee-free" offer and were included in the third ("taxi") phase of the experiment, a single week of Taxi offers. We do not include data from the second phase in our analysis as it is not comparable to earlier experiments due to significant changes in the driver app.

3.10.3 Individual Driver Bonuses

The complexity of the algorithm Uber uses to assign in the IDB program generates random variation in assignment to high and low bonus offers, conditional on prior driving behavior. We group drivers into eight hours bins based on their driving behavior in the prior four weeks. In Table 3.4 We show that conditional on these eight strata, there is no statistical difference between drivers in the high and low offer groups.

3.11 Theoretical Appendix

3.11.1 Firm- and Market- Labor Supply Elasticities

Consider a simple intertemporal labor supply model where individuals can work for two jobs. Hours at the first job are denoted h_t and hours at the second job are denoted r_t. At time t, individuals choose consumption, c_t, and hours $\{h_t, r_t\}$ to maximize the present discounted value of future utility. Their instantaneous utility function, $u(c_t, l_t)$ depends on consumption and leisure where $l_t = T - h_t - r_t$. Utility is increasing in both consumption and leisure and, as is standard, as $c_t \rightarrow 0, u_c \rightarrow \infty$.

Individuals earn an exogenous income stream y_t and face a constant (within period) wage rate w_t at their main job. At their second job, r_t hours nets them $L(r_t)$ in earnings where $L(\cdot)$ is concave. Individuals can borrow and save, and face no borrowing constraints. Assets in period t are denoted A_t. As in standard labor supply models, since utility is additive we can write the problem recursively as

$$V_t(A_t) = \max_{c_t, h_t, r_t} u(c_t, T - h_t - r_t) + \beta E_t[V_{t+1}(A_{t+1})]$$

subject to

$$\begin{aligned} h_t, r_t &\geq 0 \\ A_{t+1} &= (1 + R_t)(A_t + y_t + w_t h_t + L(r_t) - c_t - \kappa 1\{r_t > 0\}) \\ A_T &= 0 \end{aligned}$$

where κ is the psychic cost of working multiple jobs.

The model yields simple predictions for the responsiveness of hours at both the main and second job to changes in the main job's wage and in the cost of working on two jobs. The intratemporal conditions are:

$$\begin{aligned} w_t &= \frac{u_l + \mu_t^h}{u_c} \\ L'(r_t) &= \frac{u_l + \mu_t^r}{u_c} \end{aligned}$$

where the μ_t are the Lagrange multipliers on the constraint that hours in both jobs must be greater than or equal to zero. Assuming an interior solution for hours at the main job, hours are chosen to equate the ratio of the marginal utilities of leisure and consumption (taking into account jobs at the second job) to the wage. This simple set-up yields three intuitive predictions.

Proposition 12. *Conditional on working a second job, hours in the second job are decreasing in* w_t.

Proof. We can rearrange the first order conditions to see that $L'(r_t) - w_t = 0$. By the implicit function theorem

$$\frac{\partial r_t}{\partial w_t} = \frac{1}{L''(r_t)}$$

and since L is concave, r_t is decreasing in w_t. □

Proposition 13. *If the second job's hours are not flexible, the response of total hours worked is the same as the response of hours worked at the main job.*

Proof. If the second job's hours are fixed, $dr/dw = 0$. □

This proposition allows us to ignore non-gig employment when estimating labor supply elasticities in markets where Lyft is unavailable. While the Uber drivers in our data may work traditional jobs (just like taxi drivers or stadium vendors), as long as they are not able to change their hours at these jobs at high frequency (within a week), our estimates of market labor supply elasticities will reflect real increases in hours worked.

Proposition 14. *The elasticity of hours worked at the main job with respect to the wage is greater than the elasticity of total hours worked.*

Proof. This follows from propositions 12 and 13. Temporarily drop the time subscripts and define $H = h + r$ and $\phi = \dfrac{h}{H}$ (fraction of total hours spent at the primary job). We can write

$$
\begin{aligned}
\frac{dH}{dw}\frac{w}{H} &= \frac{dh}{dw}\frac{w}{H} + \frac{dr}{dw}\frac{w}{H} \\
&= \frac{dh}{dw}\frac{w}{\phi H(1/\phi)} + \frac{dr}{dw}\frac{w}{H(1-\phi)/(1-\phi)} \\
\epsilon &= \underbrace{\tau\phi}_{+} + (1-\phi)\underbrace{s}_{-}
\end{aligned}
$$

The first term is positive. The second term is negative by Proposition 12 □

Propositions 12 and 14 show that if individuals can shift hours between employers easily, estimates of the labor supply elasticity using a single platform will conflate changes in hours supplied to the market and changes in the allocation of hours across firms .

3.11.2 Derivation of the First Stage

This section goes through a derivation of the first stage. A similar derivation is provided in the main text of Angrist et al. (2017).

The first stage effect of offers on log wages depends on: (1) the experimental participation rate, and (2) the magnitude of experimentally-induced fee changes. Use w_{it}^0 to denote a driver's potential average hourly earnings in the absence of treatment and t_0 to denote the driver's Uber fee. Then, using the potential outcomes framework, the hourly earnings we observe satisfy:

$$
\begin{aligned}
w_{it} &= w_{it}^0(1-t_0)(1-D_{it}) + w_{it}^0(1-t_1)D_{it} \\
&= w_{it}^0(1-t_0) + w_{it}^0(t_0-t_1)D_{it}.
\end{aligned}
$$

where D_{it} is a binary indicator for whether the driver is driving fee-free. Because offers, Z_{it} are independent of w^0_{it} , the first stage effect of offers on wages is

$$E[w_{it}|Z_{it}=1,t_0,t_1]-E[w_{it}|Z_{it}=0,t_0,t_1]$$
$$=(t_0-t_1)E[w^0_{it}|D_{it}=1]\times P[D_{it}=1|Z_{it}=1]. \tag{3.16}$$

The first stage impact on wages is just the experimental fee change multiplied by the opt-in rate and wages for non-participants.[31]

The *proportional* change in wages is obtained by dividing (3.16) by the wages of the controls, $E[w_{it}|Z_{it}=0]=E[w^0_{it}](1-t_0)$. The proportional wage increase is:

$$\frac{E[w_{it}|Z_{it}=1,t_0,t_1]-E[w_{it}|Z_{it}=0,t_0,t_1]}{E[w_{it}|Z_{it}=0,t_0,t_1]}$$
$$=\frac{(t_0-t_1)}{1-t_0}P[D_{it}=1|Z_{it}=1]. \tag{3.17}$$

In other words, the first stage for log wages is the change in fee divided by the baseline take-home rate (1-fee), multiplied by the treatment take-up rate. For example, with a take-up rate of 2/3, the proportional first stage for an experiment that eliminates a 25% fee is roughly $\frac{.25}{.75}.66=.22$.[32]

3.11.3 Individual Driver Bonuses and Labor Supply Elasticities

Use t_{i0} to denote the number of trips driver i completes when untreated (given the "low") offer and t_{i1} to denote the number of trips driver i completes when treated. Given the structure of the treatment, there are three possible cases:

1. If $t_{i0}\geq T$, the driver already exceeds the trip threshold. Assuming no income effects, his/her labor supply is unaffected and $t_{i1}=t_{i0}$.

2. If $t_{i0}<T$ and $t_{i0}(1+\epsilon\frac{B/T}{w})<T$, $t_{i1}=t_{i0}$ where $\epsilon=\frac{d\log t}{d\log w}$ is the elasticity of trips with respect to the per trip wage. Assuming a reasonably constant number of trips per hour,

[31]The derivation here uses the fact that $D_{it}=1$ implies $Z_{it}=1$, which in turn yields $E[w^0_{it}|D_{it}=1,Z_{it}=1]=E[w^0_{it}|D_{it}=1]$.

[32]The first stage in logs is $\ln\frac{1-t_1}{1-t_0}\times P[D_{it}=1|Z_{it}=1]$, but $\ln\frac{1-t_1}{1-t_0}\approx\frac{(t_0-t_1)}{1-t_0}$.

ϵ is also equivalent to the elasticity of hours worked to the wage. The driver is either too far below the trip threshold or not elastic enough to reach the threshold. His labor supply is unaffected.

3. If $t_{i0} < T$ and $t_{i0}(1+\epsilon\frac{B/T}{w}) \geq T$, $t_{i1} = T$. The driver is close enough to the trip threshold and elastic enough to reach the trip threshold.

This can be summarized by the following:

$$t_{i1} = \begin{cases} t_{i0} & \text{if } t_{i0} \geq T \\ t_{i0} & \text{if } t_{i0}(1+\epsilon\frac{B/T}{w}) < T \\ T & \text{otherwise} \end{cases}$$

and

$$t_{i1} \geq T \iff t_{i0}(1+\epsilon\frac{B/T}{w}) \geq T$$

If $p_{B,T}$ is the opt-in rate among the high-bonus group we can rewrite this as

$$\begin{aligned} p_{B,T} &= 1 - F_0\left[T/(1+\epsilon(B/T)/w)\right] \\ 1 - p_{B,T} &= F_0\left[T/(1+\epsilon(B/T)/w)\right] \\ F_0^{-1}\left[1 - p_{B,T}\right] &= T/(1+\epsilon(B/T)/w) \end{aligned}$$

Estimation Procedure

We estimate ϵ using the following procedure: we group drivers into strata based on: sex, shifter/non-shifter, and date.

1. We calculate the number of drivers in the high bonus group that exceeded the trip threshold. We denote this $p_{B,T}$

2. We find the $1 - p_{B,T}$ quantile of the corresponding low bonus group

3. We fit equation 3.14 by non-linear least squares. We allow the elasticity to vary across the four groups.

4. We bootstrap steps 1-3 500 times to obtain standard errors.

Bibliography

Abadie, Alberto and Guido W Imbens, "Large sample properties of matching estimators for average treatment effects," *Econometrica*, 2006, *74* (1), 235–267.

Abowd, John A and Thomas Lemieux, "The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada," *The Quarterly Journal of Economics*, 1993, *108* (4), 983–1014.

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

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *arXiv preprint arXiv:1806.07928*, 2018.

Altonji, Joseph G, Todd E Elder, and Christopher R Taber, "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal of political economy*, 2005, *113* (1), 151–184.

Andersen, Torben M and Michael Svarer, "Flexicurity—Labour Market Performance in Denmark," *CESifo Economic Studies*, 2007, *53* (3), 389–429.

Angrist, Joshua D., Sydnee Caldwell, and Jonathan V. Hall, "Uber Vs. Taxi: A Driver's Eye View," Working Paper 23891, National Bureau of Economic Research September 2017.

Ashenfelter, Orley and Alan B Krueger, "Theory and Evidence on Employer Collusion in the Franchise Sector," Working Paper 24831, National Bureau of Economic Research 2018. Princeton University Working Paper.

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

Autor, David, David Dorn, and Gordon H Hanson, "The China syndrome: Local labor market effects of import competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

Azar, José, Ioana Marinescu, and Marshall I Steinbaum, "Labor Market Concentration," Technical Report, National Bureau of Economic Research 2017.

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

Bagger, Jesper, Bent Jesper Christensen, and Dale T Mortensen, "Productivity and Wage Dispersion: Heterogeneity or Misallocation," Technical Report 2014.

_ , Fran Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin, "Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics," *American Economic Review*, 2014, *104* (6), 1551–96.

Bailey, Mike, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel, "The Economic Effects of Social Networks: Evidence from the Housing Market," *Journal of Political Economy*, forthcoming.

Barth, Erling, Alex Bryson, James C Davis, and Richard Freeman, "It's Where You Work: Increases in Earnings Dispersion Across Establishments and Individuals in the US," *Journal of Labor Economics*, 2016, *34* (S2), S67–S97.

Bartik, Timothy J, *Who benefits from state and local economic development policies?*, Kalamazoo, MI: W.E. Upjohn Institute for Employment Research, 1991.

Bayer, Patrick, Stephen L Ross, and Giorgio Topa, "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes," *Journal of Political Economy*, 2008, *116* (6), 1150–1196.

Beaman, Lori, "Social Networks and the Labor Market," in "The Oxford Handbook of the Economics of Networks" 2016.

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

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

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–890.

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

Bidner, Chris and Ben Sand, "Job Prospects and Pay Gaps: Theory and Evidence on the Gender Gap from US Cities," 2016.

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

Blanchard, Olivier and Lawrence F Katz, "Regional Evolutions," *Economic Studies Program, The Brookings Institution*, 1992, *23* (1), 76.

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

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-experimental Shift-share Research Designs," *arXiv preprint arXiv:1806.01221*, 2018.

Botero, Juan C, Simeon Djankov, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer, "The Regulation of Labor," *The Quarterly Journal of Economics*, 2004, *119* (4), 1339–1382.

Bowles, Hannah Riley, Linda Babcock, and Lei Lai, "Social incentives for gender differences in the propensity to initiate negotiations: Sometimes it does hurt to ask," *Organizational Behavior and human decision Processes*, 2007, *103* (1), 84–103.

Burdett, Kenneth and Dale T Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 1998, pp. 257–273.

Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin, "Wage Bargaining with On-The-Job Search: Theory and Evidence," *Econometrica*, 2006, *74* (2), 323–364.

Caldwell, Sydnee and Nikolaj Harmon, "Outside Options, Bargaining, and Wages: Evidence from Coworker Networks," Working Paper, Massachusetts Institute of Technology 2018.

_ and Oren Danieli, "Outside Options in the Labor Market," Technical Report, Harvard University Working Paper 2018.

Card, David, "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration," *Journal of Labor Economics*, 2001, *19* (1), 22–64.

_ , Ana Rute Cardoso, and Patrick Kline, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *Quarterly Journal of Economics*, 2016, *131* (2), 633–686.

_ , _ , and _ , "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *Quarterly Journal of Economics*, 2016, *131* (2), 633–686.

_ , _ , Joerg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, Jan 2018, *36* (S1), S13–S70.

_, _, Jörg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2016, *53* (1), S13–S70.

_, Francesco Devicienti, and Agata Maida, "Rent-Sharing, Holdup, and Wages: Evidence from Matched Panel Data," *Review of Economic Studies*, 2014, *81* (1), 84–111.

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

Chen, M Keith and Michael Sheldon, "Dynamic Pricing in a Labor Market: Surge Pricing and the Supply of Uber Driver-Partners," Technical Report, University of California Los Angeles Working Paper 2015.

Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Does indivisible labor explain the difference between micro and macro elasticities? A meta-analysis of extensive margin elasticities," *NBER macroeconomics Annual*, 2013, *27* (1), 1–56.

Choo, Eugene and Aloysius Siow, "Who Marries Whom and Why," *Journal of Political Economy*, 2006, *114* (1), 175–201.

Conley, Timothy G, "GMM estimation with cross sectional dependence," *Journal of Econometrics*, 1999, *92* (1), 1–45.

_ and Giorgio Topa, "Identification of Local Interaction Models with Imperfect Location Data," *Journal of Applied Econometrics*, 2003, *18* (5), 605–618.

Cook, Cody, Rebecca Diamond, Jonathan V. Hall, John A. List, and Paul Oyer, "The Gender Earnings Gap in the Gig Economy: Evidence from over a Million Uber Drivers," Technical Report, Stanford University Working Paper 2018.

Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg, "Peer effects in the workplace," *American Economic Review*, 2017, *107* (2), 425–56.

Council of Economic Advisors, "Labor Market Monopsony: Trends, Consequences and Policy Responses," 2016.

Dagsvik, John K, "Discrete and continuous choice, max-stable processes, and independence from irrelevant attributes," *Econometrica*, 1994, pp. 1179–1205.

Dahl, Christian M, Daniel Le Maire, and Jakob R Munch, "Wage Dispersion and Decentralization of Wage Bargaining," *Journal of Labor Economics*, 2013, *31* (3), 501–533.

Doniger, Cynthia L, "Wage Dispersion with Heterogeneous Wage Contracts," 2015.

Dube, Arindrajit, Alan Manning, and Suresh Naidu, "Monopsony and Employer Mis-Optimization Account for Round Number Bunching in the Wage Distribution," Technical Report, Working Paper 2017.

_ , **Jeff Jacobs, Suresh Naidu, and Siddharth Suri**, "Monopsony in Online Labor Markets," *AER: Insights*, forthcoming.

Duncan, Greg J and Bertil Holmlund, "Was Adam Smith right after all? Another test of the theory of compensating wage differentials," *Journal of Labor Economics*, 1983, *1* (4), 366–379.

Dupuy, Arnaud and Alfred Galichon, "Personality Traits and the Marriage Market," *Journal of Political Economy*, 2014, *122* (6), 1271–1319.

Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker, "Referral-Based Job Search Networks," *The Review of Economic Studies*, 2015, *83* (2), 514–546.

_ , **Johannes Ludsteck, and Uta Schönberg**, "Revisiting the German Wage Structure," *The Quarterly Journal of Economics*, 2009, *124* (2), 843–881.

Eliason, Marcus, Lena Hensvik, Francis Kramarz, and Oskar Skans, "When Connections Become Hires: Firm-Level Responses to Supply Shocks Within Their Workers' Social Networks," 2017.

Farber, Henry S, "The Analysis of Union Behavior," *Handbook of Labor economics*, 1986, *2*, 1039–1089.

Farber, Henry S., "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers," *Journal of Political Economy*, 2005, *113* (1), 46–82.

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

Fehr, Ernst and Lorenz Goette, "Do Workers Work More If Wages Are High? Evidence from a Randomized Field Experiment," *The American Economic Review*, 2007, *97* (1), 298–317.

Flinn, Christopher and Joseph Mullins, "Firms' Choice of Wage-Setting Protocols in the Presence of Minimum Wages," 2017.

Fortin, Nicole M and Thomas Lemieux, "Changes in Wage Inequality in Canada: An Interprovincial Perspective," *Canadian Journal of Economics/Revue Canadienne D'économique*, 2015, *48* (2), 682–713.

Garin, Andrew and Filipe Silverio, "How Does Firm Performance Affect Wages? Evidence from Idiosyncratic Export Shocks," 2018. Harvard University Working Paper.

Gathmann, Christina and Uta Schönberg, "How General Is Human Capital? a Task-Based Approach," *Journal of Labor Economics*, 2010, *28* (1), 1–49.

Gee, Laura K, Jason Jones, and Moira Burke, "Social Networks and Labor Markets: How Strong Ties Relate to Job Finding on Facebook's Social Network," *Journal of Labor Economics*, 2017, *35* (2), 485–518.

Gertler, Mark, Christopher Huckfeldt, and Antonella Trigari, "Unemployment Fluctuations, Match Quality, and the Wage Cyclicality of New Hires," Technical Report, National Bureau of Economic Research 2016.

Giorgi, Giacomo De, Anders Frederiksen, and Luigi Pistaferri, "Consumption Network Effects," Working Paper 22357, National Bureau of Economic Research June 2016.

Glitz, Albrecht, "Coworker Networks in the Labour Market," 2013.

_ and Rune Vejlin, "Learning Through Coworker Referrals," 2018.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," Technical Report, National Bureau of Economic Research 2018.

Gourieroux, Christian, Alain Monfort, and Eric Renault, "Indirect inference," *Journal of applied econometrics*, 1993, *8* (S1), S85–S118.

Granovetter, Mark, *Getting a Job: A Study of Contacts and Careers*, University of Chicago Press, 1995.

Granovetter, Mark S, "The Strength of Weak Ties," *American Journal of Sociology*, 1973, *78* (6), 1360–1380.

Green, David A, "Chasing After "good Jobs." Do They Exist and Does It Matter If They Do?," *Canadian Journal of Economics/Revue canadienne d'économique*, 2015, *48* (4), 1215–1265.

Groes, Fane, Philipp Kircher, and Iourii Manovskii, "The U-Shapes of Occupational Mobility," *The Review of Economic Studies*, 2014, *82* (2), 659–692.

Hagedorn, Marcus and Iourii Manovskii, "Job Selection and Wages over the Business Cycle," *The American Economic Review*, 2013, *103* (2), 771–803.

_ , Tzuo Hann Law, and Iourii Manovskii, "Identifying Equilibrium Models of Labor Market Sorting," *Econometrica*, 2017, *85* (1), 29–65.

Hagen, Von Tobias, "Does Fixed-term Contract Employment Raise Firms' Adjustment Speed," *Jahrbücher f. Nationalökonomie u. Statistik*, 2003, *223* (4).

Hall, Jonathan V and Alan B Krueger, "An Analysis of the Labor Market for Uber's Driver-Partners in the United States," *Princeton University Industrial Relations Section Working Paper*, 2015, *587*.

_ , John J Horton, and Daniel T Knoepfle, "Labor Market Equilibration: Evidence from Uber," Technical Report, Working Paper, 1–42 2017.

Hall, Robert E and Alan B Krueger, "Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-The-Job Search," *American Economic Journal: Macroeconomics*, 2012, *4* (4), 56–67.

_ **and Edward P Lazear**, "The Excess Sensitivity of Layoffs and Quits to Demand," *Journal of Labor Economics*, 1984, *2* (2), 233–257.

Hazell, Jonathon and Bledi Taska, "Posted Wage Rigidity," Technical Report 2018.

Hellerstein, Judith, David Neumark, and Melissa McInerney, "Changes in Workplace Segregation in the United States Between 1990 and 2000: Evidence from Matched Employer-Employee Data," in "The Analysis of Firms and Employees: Quantitative and Qualitative Approaches," University of Chicago Press, 2008, pp. 163–195.

Hensvik, Lena and Oskar Nordström Skans, "Social Networks, Employee Selection, and Labor Market Outcomes," *Journal of Labor Economics*, 2016, *34* (4), 825–867.

Herkenhoff, Kyle, Jeremy Lise, Guido Menzio, and Gordon M Phillips, "Production and Learning in Teams," Technical Report, National Bureau of Economic Research 2018.

Heuermann, Daniel F and Johannes F Schmieder, "The Effect of Infrastructure on Worker Mobility: Evidence From High-speed Rail Expansion In Germany," *Journal of Economic Geography*, 2018, p. lby019.

Hirsch, B, T Schank, and C Schnabel, "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*, 2010, *28* (2), 291–330.

Hirsch, Boris, Thorsten Schank, and Claus Schnabel, "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*, 2010, *28* (2), 291–330.

Hornstein, Andreas, Per Krusell, and Giovanni L Violante, "Frictional Wage Dispersion in Search Models: A Quantitative Assessment," *American Economic Review*, 2011, *101* (7), 2873–98.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow, "The Allocation of Talent and U.S. Economic Growth," Working Paper 18693, National Bureau of Economic Research January 2013.

Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang, "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," *American Economic Review*, 2014, *104* (6), 1597–1629.

Ioannides, Yannis M. and Linda Datcher Loury, "Job Information Networks, Neighborhood Effects, and Inequality," *Journal of Economic Literature*, 2004, *42* (4), 1056–1093.

Jacobson, Louis S, Robert J. Lalonde, and Daniel G Sullivan, "Earnings Losses of Displaced Workers," *American Economic Review*, feb 1993, *83* (4), 685–709.

Jaeger, David A, Joakim Ruist, and Jan Stuhler, "Shift-share instruments and the impact of immigration," Technical Report, National Bureau of Economic Research 2018.

Jäger, Simon, Benjamin Schoefer, Sammy Young, and Josef Zweimuller, "Wages, Bargaining, and the Nonemployment Option: Evidence from Unemployment Insurance," 2018.

Jarosch, Gregor, "Searching for Job Security and the Consequences of Job Loss," *manuscript, Stanford University*, 2015.

_ , Ezra Oberfeld, and Esteban Rossi-Hansberg, "Learning from Coworkers," Technical Report, Princeton University Working Paper 2018.

Killingsworth, Mark R and James J Heckman, "Female Labor Supply: A Survey," *Handbook of Labor Economics*, 1986, *1*, 103–204.

King, Robert G and Sergio T Rebelo, "Resuscitating real business cycles," *Handbook of macroeconomics*, 1999, *1*, 927–1007.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar, "Who Profits from Patents? Rent-Sharing at Innovative Firms," Technical Report, NBER Working Paper 25245 2018.

Koustas, Dmitri, "Labor Supply and Consumption Insurance: Evidence from Rideshare Drivers," Working Paper, UC Berkeley 2017.

Kramarz, Francis and Oskar Nordström Skans, "When Strong Ties Are Strong: Networks and Youth Labour Market Entry," *Review of Economic Studies*, 2014, *81* (3), 1164–1200.

Kreps, David M and Jose A Scheinkman, "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *The Bell Journal of Economics*, 1983, *14* (2), 326–337.

Kristoffersen, Mark Strøm, "Geographical Job Mobility and Wage Flexibility in Denmark," Technical Report, Danmarks Nationalbank Working Papers 2016.

Krueger, Alan B, "The Rigged Labor Market," 2017.

Lamadon, Thibaut, "Productivity Shocks, Dynamic Contracts, and Income Uncertainty," *manuscript, University of Chicago*, 2014.

Lise, Jeremy and Jean-Marc Robin, "The Macrodynamics of Sorting Between Workers and Firms," *American Economic Review*, 2017, *107* (4), 1104–35.

_ , Costas Meghir, and Jean-Marc Robin, "Matching, Sorting and Wages," *Review of Economic Dynamics*, 2016, *19*, 63–87.

Manning, Alan, "Monopsony in Motion," *Monopsony in Motion*, 2003, pp. 3–28.

_ , *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton University Press, 2003.

_ , "Imperfect Competition in the Labor Market," *Handbook of Labor Economics*, 2011, *4*, 973–1041.

_ **and Barbara Petrongolo**, "How Local Are Labor Markets? Evidence from a Spatial Job Search Model," *American Economic Review*, 2017, *107* (10), 2877–2907.

Martinez, Isabel Z, Emmanuel Saez, and Michael Siegenthaler, "Intertemporal Labor Supply Substitution? Evidence from the Swiss Income Tax Holidays," Technical Report, National Bureau of Economic Research 2018.

Mas, Alexandre and Amanda Pallais, "Labor Supply and the Value of Non-Work Time: Experimental Estimates from the Field," *American Economic Review: Insights*, 2018.

McClelland, Robert and Shannon Mok, "A Review of Recent Research on Labor Supply Elasticities," 2012.

Mitlacher, Lars W., "Job quality and temporary agency work: Challenges for human resource management in triangular employment relations in Germany," *The International Journal of Human Resource Management*, 2008, *19* (3), 446–460.

Mogstad, Magne, Bradley Setzler, Thibaut Lamadon et al., "Earnings Dynamics, Mobility Costs, and Transmission of Market-Level Shocks," in "2017 Meeting Papers" number 1483 Society for Economic Dynamics 2017.

Myers and G.P. Shultz, *The Dynamics of a Labor Market*, Prentice-Hall: NY, 1951.

Naidu, Suresh, "Recruitment Restrictions and Labor Markets: Evidence from the Postbellum US South," *Journal of Labor Economics*, 2010, *28* (2), 413–445.

_ **, Yaw Nyarko, and Shing-Yi Wang**, "Monopsony Power in Migrant Labor Markets: Evidence from the United Arab Emirates," *Journal of Political Economy*, 2016, *124* (6), 1735–1792.

Nimczik, Jan Sebastian, "Job Mobility Networks and Endogenous Labor Markets," 2017.

OECD, "Female Labour Force Participation: Past Trends and Main Determinants in OECD Countries," 2004.

Oettinger, Gerald S, "An Empirical Analysis of the Daily Labor Supply of Stadium Vendors," *Journal of political Economy*, 1999, *107* (2), 360–392.

Pérez-Castrillo, David and Marilda Sotomayor, "A Simple Selling and Buying Procedure," *Journal of Economic Theory*, 2002, *2* (103), 461–474.

Pissarides, CA, *Equilibrium Unemployment Theory*, MIT Press, 2000.

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

_ and _ , "To Match or Not to Match?: Optimal Wage Policy with Endogenous Worker Search Intensity," *Review of Economic Dynamics*, 2004, *7* (2), 297–330.

Pouliakas, Konstantinos, "Multiple Job-Holding: Career Pathway or Dire Straits?," *IZA World of Labor*, 2017.

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

_ and Ronald L Oaxaca, "New Market Power Models and Sex Differences in Pay," *Journal of Labor Economics*, 2010, *28* (2), 267–289.

Reenen, John Van, "The Creation and Capture of Rents: Wages and Innovation in a Panel of U. K. Companies," *The Quarterly Journal of Economics*, 1996, *111* (1), 195–226.

Rees, A. and G.P. Shultz, *Workers in an Urban Labor Market*, University of Chicago Press: Chicago, 1970.

Rees, Albert, "Information Networks in Labor Markets," *The American Economic Review*, 1966, *56* (1/2), 559–566.

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

Roth, Alvin E and Marilda Sotomayor, "Two-sided matching," *Handbook of Game Theory with Economic Applications*, 1992, *1*, 485–541.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica: Journal of the Econometric Society*, 1982, pp. 97–109.

Saygin, Perihan, Andrea Weber, and Michèle Weynandt, "Coworkers, Networks, and Job Search Outcomes," 2014.

Schmutte, Ian M, "Job Referral Networks and the Determination of Earnings in Local Labor Markets," *Journal of Labor Economics*, 2014, *33* (1), 1–32.

Shapley, L. S. and M. Shubik, "the Assignment Game I: The Core," *International Journal of Game Theory*, 1971, *1* (1), 111–130.

Shue, Kelly, "Executive Networks and Firm Policies: Evidence from the Random Assignment of MBa Peers," *The Review of Financial Studies*, 2013, *26* (6), 1401–1442.

Simon, Curtis J and John T Warner, "Matchmaker, Matchmaker: The Effect of Old Boy Networks on Job Match Quality, Earnings, and Tenure," *Journal of Labor Economics*, 1992, *10* (3), 306–330.

Sorkin, Isaac, "Ranking Firms Using Revealed Preference," *Quarterly Journal of Economics*, August 2018, *133* (1), 1331–1393.

Stafford, Tess M, "What Do Fishermen Tell Us That Taxi Drivers Do Not? an Empirical Investigation of Labor Supply," *Journal of Labor Economics*, 2015, *33* (3), 683–710.

Staiger, D, Joanne Spetz, and C Phibbs, "Is There Monopsony in the Labor Market? Evidence from a Natural Experiment," *Journal of Labor Economics*, 2010, *28* (2).