JOB REFERRAL NETWORKS AND THE DETERMINATION OF EARNINGS IN LOCAL LABOR MARKETS

Ian M. Schmutte
Department of Economics
Terry College of Business
University of Georgia
Athens, GA 30602
schmutte@uga.edu

August 2013

Abstract

Despite their documented importance in the labor market, little is known about how workers use social networks to find jobs and their resulting effect on earnings. I use geographically detailed U.S. employer-employee data to infer the role of social networks in connecting workers to jobs in high-paying firms. To identify social interactions in job search, I exploit variation in social network quality within small neighborhoods. Workers are more likely to change jobs, and more likely to move to a higher-paying firm, when their neighbors are employed in high-paying firms. Furthermore, local referral networks help match high-ability workers to high-paying firms.

JEL Codes: J31, J64, R23.
Keywords: Social Interactions; Informal Hiring Networks; Wage Variation; Neighborhood Effects.

This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

1 Introduction

This paper studies spillovers across residential neighbors in earnings and place of work. I find that these spillovers are consistent with the presence of local social networks that function to help workers find employment in higher paying firms. That referral and job information networks play a major role in matching workers to jobs is widely documented (Bewley 1999; Ioannides and Loury 2004). Despite this, much less is known about why workers use social networks to find work and how networks affect labor market outcomes. A major impediment to work in this area is that the things about a person that affect her earnings also determine the people she knows, making it difficult to distinguish endogenous social interactions from sorting into social groups (Manski 1993; Moffitt 2000; Blume, Brock, Durlauf and Ioannides 2011). As a result it is hard to identify social interactions in earnings, and even harder to draw inferences about their underlying mechanisms.

I make progress by asking whether job information networks match well-connected workers to jobs in high-paying firms. In a market with search frictions, two workers can receive different pay simply because they work for different employers (Montgomery 1991; Abowd, Kramarz and Margolis 1999; Mortensen 2003). If these workers share information about pay differentials, then they can help each other find higher-paying job offers. If so, people who already have jobs in high-paying firms will be better sources of job information for a number of reasons. Workers with high-paying jobs can provide direct referrals to their present employers (Mortensen and Vishwanath 1994). Alternatively, they may pass along information about jobs they heard about but did not take (Calvo-Armengol and Jackson 2007). They might also share information about how to find 'good' jobs among their neighbors.[1]

To get at this effect empirically, I consider whether people whose neighbors work in higher-paying firms have better job search outcomes, and specifically, whether they themselves find

[1]Dolnick (2011) describes this behavior among immigrants in New York City, who share information about temp agencies that provide relatively good job opportunities.

jobs in higher-paying firms. To do so, it is necessary to measure a firm-specific contribution to pay, to observe individual job search outcomes, and to observe, for each worker, a reference group with whom she might exchange job information. I meet these requirements using matched employer-employee data from the Longitudinal Employer-Household Dynamics (LEHD) program of the U.S. Census Bureau.

I first measure the employer's contribution to pay, which I refer to as a "firm-specific wage premium", from a decomposition of earnings that separates the employer's contribution from the part of earnings associated with worker characteristics (Abowd et al. 1999). I then match each worker to her Census block of residence, restricting attention to workers who reside in one of thirty large Metropolitan Statistical Areas (MSAs).[2] In these MSAs, the Census blocks are like city blocks, and because the data are a census of private-sector employment, for any worker I observe the characteristics of all of their residential neighbors, including their firm-specific wage premia. For each worker, I compute the average employer-specific wage premium of other workers who live in the same Census block. Assuming workers exchange job information with their residential neighbors, this measures the quality of their job information network.

To identify local interaction effects, I compare the job search outcomes of workers who live in the same neighborhood but on different blocks. In my analysis, the 'neighborhood' is always defined to be the Census block group, which is contiguous group of blocks containing between 600 and 3,000 people.[3] The key identifying assumption is that, within these small neighborhoods, variation across blocks in network quality is uncorrelated with unobserved factors that affect job search. This assumption allows for the possibility that workers sort, and firms recruit, differentially across neighborhoods. It requires that there is no further sorting across blocks within neighborhoods. The economic rationale is that, because they

[2]Data from LEHD and residential address information were jointly available only for 2002–2003.
[3]See Schmutte (2012) for a complete description of these geograhpic concepts.

are constrained by the housing available when moving, people choose their neighborhood, but not their immediate neighbors. This identification strategy is closely related to that of Bayer, Ross and Topa (2008). Like them, I verify the absence of sorting within neighborhoods on the basis of observable characteristics, and show the results are not sensitive to plausible violations of the identifying assumptions.

Implementing this strategy, I find that workers who live on blocks with higher-quality networks are more likely to change jobs, and to move to jobs with higher wage premia. My estimates imply that workers receive eight to nine percent of job offers through residential job information networks, well within the range of existing estimates (Ioannides and Loury 2004). From a theoretically agnostic perspective, the results imply that a one standard deviation increase in network quality is associated with a twenty-five percent increase in the firm-specific wage premium on job change. Both interpretations confirm the presence of local interactions in earnings outcomes, consistent with workers using residential job information networks to find jobs in high-paying firms.

To motivate and frame my results, I outline a job search model. In it, workers search for jobs in high-paying firms and can get offers through formal channels or through their neighbors. This yields a non-trival prediction regarding the effect of network quality on the full distribution of accepted job offers. Conditional on job change, the marginal effect of improved network quality is stronger at high quantiles of the accepted offer distribution. Mechanically, this effect comes about because workers only change jobs when the outside offer is better than staying on the inside job, so the distribution of accepted wage offers is a truncated version of the distribution of all offers. Economically, this distributional result rules out an alternative explanation: that workers with higher-quality networks get job offers faster, but not from a better distribution.

In Section 4, I verify the assumption that workers in my sample interact with their residential neighbors to find employment. To do so, I extend the method used by Bayer et al. (2008) to document general local interactions in place of work. When changing jobs, urban workers are more likely to become coworkers of people in the same block than of people in the same neighborhood on different blocks. Because I know who was on the job first, my results suggest that referrals drive part of the local interaction effect found by Bayer et al. (2008).

I exploit this information to shed light on the nature of local referral networks. For each potential referral relationship, I know which worker would be the referral recipient and which worker would be the referral provider. The referral effect is stronger when the referral provider is employed in a higher-paying firm. Also, referrals are more likely when the receiving worker is of high ability, consistent with firms using social networks to locate high-quality job candidates. This latter finding suggests avenues for future research using these data to explore the incentives for referral use on both sides of the market.

Many studies support the premise that workers use residence-based networks as part of job search (Case and Katz 1991; Ioannides and Loury 2004; Bayer et al. 2008; Damm 2009), and indicate that some workers use them more extensively than others.[4] This suggests a natural specification check: my results should be stronger among workers who are expected to use residential networks for job search more intensively. The magnitude of the local interaction effects are nearly twice as strong for non-native workers as for natives, fifty percent stronger for young workers, and fifty percent weaker for older workers, consistent with evidence presented in Ioannides and Loury (2004).

Having verified the main result and tested the modeling assumptions, I conduct a series of exercises that explore the mechanism underlying my results. First, I ask whether the results are driven by direct referrals. By dropping workers whose new job is with a neighbor's employer, I find that direct referrals do not explain my results. Hence, job information

[4] Some argue that the Internet has eroded the importance of neighborhoods in social life. However, Mok, Wellman and Carrasco (2009) find that the emergence of the Internet been a complement for more traditional forms of local social interactions.

networks transmit job information both directly, through referral, and indirectly. I also verify that the main results hold among non-employed workers, confirming they are not driven by selection of job changers. Finally, I show that my results are not driven by networks that attract high-ability workers into high-paying agglomerations.

My results contribute to a growing literature on the role of social interactions in job search by using a clear identification strategy to provide evidence consistent with job information networks affecting earnings by directing high-paying jobs to well-connected workers. The role of job information networks in channeling access to high-paying jobs is a feature of several recent theoretical models (Mortensen and Vishwanath 1994; Cahuc and Fontaine 2009; Calvo-Armengol and Zenou 2005; Calvo-Armengol and Jackson 2007). Topa (2001) and Conley and Topa (2007) bring related models to the data, but focus on spillovers in unemployment, and use structural approaches to identification. Goel and Lang (2009) use survey evidence on network strength among Canadian immigrants to estimate social interaction effects in earnings using a much broader proxy for the network measure, and using the model structure for identification. The effect of local interactions on earnings is a feature of experimental studies on neighborhood effects (Case and Katz 1991; Katz, Kling and Liebman 2001). My analysis differs from theirs in only considering spillovers in the employer-specific component of earnings and identifying social interaction effects allowing for self-selection across neighborhoods.

This paper also joins a growing literature using matched employer-employee data to provide indirect evidence on the effects of social interactions in individual job search (Bayer et al. 2008; Laschever 2009; Hellerstein, McInerney and Neumark 2011; Cingano and Rosolia 2012; Nordström Skans and Kramarz 2011). Of these, Bayer et al. (2008) and Hellerstein et al. (2011) also study neighborhood-level interactions with respect to place-of-work. Bayer et al. (2008) consider the effects of local interactions on earnings, but do not separate the employer-specific component of pay, and can not exploit time-series variation to address whether their spillovers are associated with job search or some other mechanism. A closely-connected literature uses matched data to evaluate theories of referral use by employers (Dustmann, Glitz and Schönberg 2011; Oyer and Schaefer 2012). By considering the role of job information networks and direct referrals on job search, this paper helps link the use of job information networks by workers to the role of referral use by employers.

2 Data

I use matched employer-employee data from the LEHD Program of the U.S. Census Bureau that are linked to workers' places of residence from the Statistical Administrative Records System (StARS). Here, I describe the data sources and my research sample, and discuss two measurement issues essential to my analysis: how I measure the quality of each worker's job information network, and how I identify job-to-job transitions in the data.[5]

2.1 Data Sources

The LEHD data are based on state Unemployment Insurance (UI) records. Every quarter, the UI earnings records report, for each covered job, unique identifiers for the worker and for the employer, and the total amount of UI-taxable compensation paid. These records cover approximately 98 percent of wage and salary payments in private sector non-farm jobs, and constitute the frame for LEHD. Thus, the LEHD are nearly universal in coverage. The full research file used for this project includes 660 million wage records for 190 million workers and 10 million employers across 30 states between 1990-2003.[6] The LEHD program augments the UI records with demographic and employer characteristics through links to

[5] A description of data preparation sufficient to reproduce my analysis is available in Schmutte (2012).

[6] LEHD data not available in any form before 1990. The cleaned and edited files upon which this research is based are available only to the end of 2003.

survey and administrative data.[7].

Data on place of residence come from the StARS database. StARS is a Census Bureau program that compiles administrative data from the IRS, HUD, Medicare, Indian Health Service, and the Selective Service to improve intercensal population estimates and refresh its household sampling frame. In most cases, StARS contains an exact residential address for each individual, and I use the Census block within which that address falls. For some LEHD workers, a block-level residential address is not available. In the 30 MSAs I study, over 95 percent of LEHD workers have complete data on block of residence. I restrict analysis to this group, noting that doing so does not affect basic demographic summaries.

2.2 The Estimation Samples

I use the full LEHD to estimate employer-specific wage premia, as I describe in Section 3.3. For the second-stage analysis, I use a sample of workers aged 18-70 with positive earnings in at least one quarter of 2002-2003, who resided in one of 30 large Metropolitan Statistical Areas (MSAs), and who could be matched to their Census block of residence. My sample includes data from every MSA, listed in Table I, with full geographic coverage in LEHD during 2002-2003. The restriction to to the years 2002-2003 is imposed by the joint availability of residential address information and LEHD data. I further require the recorded block of residence be in the same MSA in both years.

After imposing these restrictions, I have a master file describing all workers in the 30 MSAs who were employed at some point between 2002 and 2003. Out of these, the primary analysis centers on workers who change jobs between 2002 and 2003. I focus on job changers because (1) unemployment is not directly measured in LEHD data, and (2) the wage premium on the originating (2002) job measures the reservation value. This design choice raises

[7]See Abowd, Stephens, Vilhuber, Andersson, McKinney, Roemer and Woodcock (2009) for a thorough description of the LEHD Infrastructure File System

concerns about external validity, which I address in Section 5 by modeling selection into job change.

To estimate social interaction effects in job search, I need a measure of the quality of each worker's job information network. I construct, as a key explanatory variable, the average employer-specific wage premium among all workers in the same block who are in the full sample, but do not change jobs. This measure has a theoretical interpretation as the mean of the distribution of job offers that arrive through the local job information network. To assure sufficient precision in the estimated network quality, I restrict my estimation samples to workers who reside on blocks where there are at least 10 workers in the reference group.

To build a sample of workers who change jobs, I must identify job change events in the LEHD data. Table II presents descriptive statistics for the full sample as well as for two subsamples of job changers. The demographic characteristics in Column (1) are as expected given the focus on an employed urban workforce. Column (2) ('Annual Job Changers'), includes all in-sample workers who were employed full time in 2002 and in 2003, and who changed employers between the two years. Column (3) ('Quarterly Job Changers') uses a more restrictive definition of job-to-job transition. It only includes workers for whom I can identify the exact quarter they move from one employer to the next.[8] My preferred specifications use Annual Job Changers. Just 3.5 percent of workers are Quarterly Job Changers, significantly lower than the reported rate of job-to-job transitions in Bjelland, Fallick, Haltiwanger and McEntarfer (2011). I prefer the sample of Quarterly Job Changers to analyze the effects of selection into job change, where the precise date of job change is most relevant. The results are largely insensitive to these different ways of identifying job-to-job transitions.

[8]Since a worker may hold overlapping jobs for several quarters, I define the date of transition between dominant jobs by finding the first quarter in which earnings with the new dominant employer exceed earnings with the old dominant employer. The results are not sensitive to this definition.

3 Empirical Design

3.1 Conceptual Framework

Consider a job search framework where workers search when employed or unemployed. They are paid for their portable skills, but the price paid for skill varies from firm to firm (Montgomery 1991; Mortensen 2003). The earnings of worker i in period t are

$$y_{ijt} = p_j e_{it}, \tag{1}$$

where e_{it} measures the worker's human capital and p_j is the price paid for it by employer j. Setting $e_{it} = \exp(X_{it}\beta + \theta_i)$ so that human capital depends on observable characteristics (X_{it}) and unobserved fixed skills (θ_i), the equation for log earnings is

$$\ln y_{it} = X_{it}\beta + \theta_i + \psi_j + \varepsilon_{it} \tag{2}$$

where $\psi_j = \ln p_j$ is the employer-specific (log) wage premium. I obtain estimates of ψ and θ using the complete LEHD data.

Workers search for employment with higher-paying firms either through the formal market, or by getting job information from friends and neighbors. When a worker gets a new offer from an employer, with probability $1-\gamma$ it comes from the formal market, and with complementary probability γ it comes from the job information network. The offered premium, ψ^*, depends on the formal offer distribution and the distribution of offers received through the job information network. An employed worker changes jobs if the offered premium (ψ^*) exceeds the inside premium (ψ_0).

This basic setup yields an expression for the wage premium offer:

$$\psi_i^* = Z_i \Pi + \gamma \bar{\psi}_{b(i)} + \eta_i. \tag{3}$$

This expression allows the mean of the formal offer distribution to vary with individual observable characteristics, Z_i. The mean offer received through the job information network depends on the mean wage premium held by other workers in the network ($\bar{\psi}_{b(i)}$). The object of interest is the social interaction parameter γ. See Schmutte (2012) for details.

3.2 Testable Implications

Because workers only change jobs if the wage premium offered (ψ^*) exceeds the wage premium on the inside job (ψ_0), the distribution of wage premium offers that induce job change is a truncated version of the full wage offer distribution. Furthermore, the truncation point varies from worker to worker. As a result, the observed distribution of accepted wage premium offers is different from the (unobserved) distribution of wage offers. Increases in $\bar{\psi}_{b(i)}$ increase the mean of the overall offer distribution. Under mild regularity assumptions, described in Schmutte (2012), the combination of truncation and improvement of the offer function associated with improved network quality generate several testable implications for the observed distribution of wage premium outcomes for job changers.

First, employed workers with higher-quality job information networks are more likely to change jobs. Second, conditional on changing jobs, workers with higher-quality job information networks will move to firms with higher wage premia. Third, conditional on changing jobs, the marginal effect of a higher-quality job information network will be stronger at higher quantiles of the accepted offer distribution. That is, increasing network quality stretches the accepted offer distribution from the right. Conversely, increasing the wage premium on the inside job, from which the worker moves, compresses the accepted offer distribution from the right.

The latter distributional results are subtle, but help distinguish the social interaction mechanism. That workers with better networks are more likely to change jobs and to move to higher-paying firms might not be due to improvements in the offer distribution, but because workers with better networks get offers more quickly. This would suggest job information networks function to make search more effective, but do not necessarily channel higher-paying jobs to well-connected workers. If job information networks just speed up the offer rate, but do not move the offer distribution, then the marginal effect of improving network quality will be strongest at the lowest quantiles of the accepted offer distribution. Intuitively, this is because workers with more offers have a higher effective reservation value, so getting more offers has an effect similar to increasing the wage premium on the inside job. Verifying the distributional prediction allows me to rule out the possibility that my results are driven just by increases in the offer rate.

3.3 Primary Empirical Specification

To study wage premium spillovers, I estimate variations of an extended versions of Equation (3) on the samples of job changers described in Section 2. The basic empirical specification is

$$\psi_i = \gamma \bar{\psi}_{b(i)0} + \beta \psi_{0i} + Z_i \Pi + \bar{Z}_{b(i)} \Gamma + \kappa_{G(b(i))} + \nu_i. \tag{4}$$

The notation $b(i)$ refers to the Census block in which i resides, and $G(b(i))$ indicates the neighborhood containing $b(i)$. ψ_i is the employer wage premium on the job to which i moves in 2003, and $\bar{\psi}_{b(i)0}$ is the average wage premium across all employed workers living in the same block as i who did not change jobs. ψ_{0i} is the wage premium paid by the 2002 employer from which i transitions. Z_i is a vector of individual characteristics, and $\bar{Z}_{b(i)}$ is its mean among workers in $b(i)$. $\kappa_{G(b(i))}$ is a neighborhood effect. The neighborhood reference group in all reported results is the Census block group, so measured referral effects are among neighbors residing on the same block contrasted against the average quality of jobs found by workers living in the same block group.

The employer wage premia, ψ, are measured through a first-stage AKM decomposition,[9]

$$\ln Y = X\beta + D\theta + F\psi + \varepsilon, \tag{5}$$

which is the empirical counterpart to (2). Equation (5) is estimated on the set of all LEHD work histories for workers aged 18-70. Y is a vector of annualized log earnings, and X is a design matrix of time-varying controls including a quartic in experience, year effects, and the exact within-year pattern of positive earnings, all interacted with sex. D and F are design matrices of the worker and employer effects, and ε is the statistical residual. ψ measures the employer's contribution to pay after controlling for observable and unobservable characteristics of their workforce. θ measures the part of individual earnigns associated with the return to fixed, portable, skills. In the empirical work, estimates of ψ measure the employer-specific wage premium, and estimates of θ measure worker skill.

3.3.1 Do Workers Move To Higher Paying Employers?

My interpretation of the estimated firm-effects, ψ from Equation (5), as employer-specific wage premia depends on the assumptions about job mobility and earnings embedded in the search model through Equation (2). The job search model implies that we should observe a 'job ladder', with workers moving, on average, to firms that pay them more for their skills. To build support for my interpretation of ψ, I show that the data exhibit the implied job ladder pattern.

[9]This decomposition was first introduced by Abowd et al. (1999). The estimates used in this paper are based on the exact solution method described in Abowd, Creecy and Kramarz (2002) and Abowd, Lengermann and McKinney (2003).

Table III shows the fraction of job changers that switch to a job at the same decile, or a higher decile of the empirical ψ distribution than their current job. This fraction is always strictly above 0.58, and as high as 0.94 for workers starting from jobs with log wage premia in the lowest deciles. Workers who change jobs tend to move to firms with higher ψ than the firm they left. Furthermore, the distributions of destination premia are strictly ranked, in the sense of stochastic dominance, by origin decile. These findings support the interpretation of estimated firm effects as wage premia in a job search model. Furthermore, they are the first evidence of the mobility-related structure of ψ when estimated from the AKM decomposition. Of course, some workers move to lower-premium firms – a fact that can be included in the job search framework without changing any of its implications.

3.3.2 Exogenous Mobility

To be interpreted as a wage premium, the firm effect, ψ, should measure the increment to earnings a worker receives when employed by a particular firm. This interpretation requires the assumption that job mobility and assignment are exogenous to the earnings residual in Equation (5). The validity of this exogenous mobility assumption is strongly contested. Abowd, McKinney and Schmutte (2010) formally reject exogenous mobility in LEHD data. Kinsler (2012) fails to reject it in education data. Card, Heining and Kline (2013) find in German data that estimates of plant effects are not sensitive to different methods of selecting against endogenous job mobility. In this paper, I assume exogenous mobility as part of a conventional job search framework. This framework leads to several predictions on the relationship between employer-specific wage premia and job mobility, all of which I verify. My approach is in the spirit of Combes, Duranton and Gobillon (2008) and Iranzo, Schivardi and Tosetti (2008), who use heterogeneity components from first-stage earnings decompositions to develop persuasive analyses of segregation by skill on spatial inequality and firm productivity. Progress in understanding what matched employer-employee data can reveal about labor markets will surely benefit from these approaches.

3.4 Identification

Identification of γ, the social interaction parameter, in Equation (4) is based on comparisons of the search outcomes of workers who live on different blocks in the same neighborhood. Because of thinness in the residential real estate market, workers choose a neighborhood to live in, but not a specific block. That is, they choose their neighborhood, but not their neighbors. Likewise, employers may prefer workers from certain neighborhoods, but are not likely to have strict preferences for specific blocks within neighborhoods. There is still block-by-block variation in network quality, and it drives variation across workers in job offers, and hence in observed search outcomes. The social interaction effect (γ) is measured by the correlation between destination wage premia (ψ_i) and network quality ($\bar{\psi}_{b(i)}$), deviated from their neighborhood level averages.

A strength of this identification strategy is that it allows for self-selection into neighborhoods and neighborhood-specific features that affect job search (correlated unobservables). The key assumption is that workers are not further sorted across blocks within neighborhoods and that neighborhood features that influence search affect workers in all blocks in the same way, on average. I provide evidence in support of this assumption next.

3.4.1 Is there Sorting within Neighborhoods?

Under the identifying assumption, workers do not sort within neighborhoods on the basis of factors that influence formal search outcomes. This yields the testable implication that they should not be sorted on observable characteristics either. Table IV reports on the extent of sorting within neighborhoods. The data are restricted to prime-age male workers who were employed full time in 2002, and I select one individual at random from each block. The entries in the table are R^2 from linear models that predict the characteristics of the randomly selected individual.[10] Column (1) presents the R^2 from univariate regressions predicting a given characteristic – say age, using the average characteristic – average age – of other workers in his block. Column (2) presents the within-group R^2 obtained after adding block group controls. If there is no sorting within neighborhoods, the age distribution in each block is, up to sampling noise, a copy of the neighborhood age distribution. In that case, the R^2 from the Column (1) should be eliminated after introducing block group controls. The random variation in age within block groups should not predict the age of a random individual. Likewise for other characteristics.

Supporting the identification strategy, sorting is heavily attenuated after introducing block group controls. Bayer et al. (2008) find similar evidence for Boston. Like them I find that residual sorting is not identically zero. However, there is absolutely no measurable sorting within neighborhood on the log earnings residual from the AKM decomposition. This indicates that sorting on unobservable characteristics that influence earnings is even less strong than any residual sorting on observable characteristics. In Section 5.1, I present a sensitivity analysis that suggests sorting on unobservables needs to be much stronger than sorting on observables than seems likely, given the evidence, to explain my estimated effects.

3.4.2 Other Challenges to Identification

There are three additional challenges to identification: sample selection, reverse causality, and the reflection problem. Sample selection may arise since I only observe accepted job offers. I demonstrate robustness against selection into job change in Section 5.1.3. Reverse causality is a problem if workers' residential location choices are determined by their place of work. Since I focus on workers who change jobs but do not change residence, my results

[10]Taking one individual from each block avoids the negative bias on sorting induced when the sampled workers also contribute to the mean block-level characteristic. See Bayer et al. (2008) and Guryan, Kroft and Notowidigdo (2009) for details.

can not be explained by workers learning where to live from their coworkers. This is an improvement over Bayer et al. (2008) and Hellerstein et al. (2011), which use cross-sectional matched employer-employee data to infer local interactions in place of work. The reflection problem primarily afflicts cross-sectional models (Manski 1993).[11] In my data, the average wage premium of neighbors, $\bar{\psi}_{b(i)0}$, is already determined by the time worker i changes jobs. The time-sequencing breaks the reflection problem, because $\bar{\psi}_{b(i)0}$ depends on the complete work histories of all neighbors, which can reasonably be excluded from the model. This exclusion breaks the reflection problem formally, but the data matrix can still be badly-conditioned (Conley and Udry 2010; Blume et al. 2011). In the results, I consider many alternative specifications that would be immune to reflection for various reasons, including exclusion restrictions, time-sequencing of the endogenous variables, and controls for non-random sample-selection. My main results are very robust across specifications.

4 Results: Direct Referrals

Next, I provide evidence to support the assumption that workers use local networks to share information when making direct job-to-job changes in general, and about high-paying employment opportunities in particular. In doing so, I also document that the main finding in Bayer et al. (2008) – people who live on the same block are more likely to be employed together – holds up in LEHD data. The results also suggest that referrals are more likely when the referral provider is in a high-paying firm, and the referral receipient is a high-ability (high wage) worker.

[11] The reflection problem arises when the endogenous variable – here, network quality $\bar{\psi}_{b(i)0}$, is perfectly collinear with the independent variables in the model.

4.1 Data Construction

I analyze matched pairs of workers from the sample of Annual Job Changers. A pair of workers (ℓ, m) appears when ℓ changes jobs between 2002–2003, and m resides in the same block-group, is employed, and does not change jobs. Following Bayer et al. (2008), I define variables

- $R_{\ell,m}$: equal to one if ℓ and m live on the same Census block and zero otherwise.

- $W_{\ell,m}$: equal to one if ℓ and m share the same employer in 2003 and zero otherwise.

The sample contains 1,524,733,934 matched pairs. 0.13 percent of pairs have $W_{\ell,m} = 1$ when $R_{\ell,m} = 0$. This increases by 23 percent to 0.16 among pairs with $R_{\ell,m} = 1$.

4.2 Empirical Model

I estimate referral effects using the linear probability model

$$W_{\ell,m} = \rho_{G(\ell)} + \alpha_0 R_{\ell,m} + \varepsilon_{\ell,m}. \tag{6}$$

α_0 measures the effect of living on the same block on the propensity for a job-changer to become the coworker of someone in her neighborhood. A positive estimate of α_0 is evidence of social interactions in job-finding under the identifying assumption that all variation in the propensity for neighbors to be coworkers is absorbed by the block group effect, $\rho_{G(\ell)}$. The notation intentionally corresponds to Equation (1) in Bayer et al. (2008).

I extend the specification in Equation (6) to account for characteristics of the pair:

$$W_{\ell,m} = \rho_{G(\ell)} + \beta' X_{\ell,m} + (\alpha_0 + \alpha_1' X_{\ell,m}) R_{\ell,m} + \varepsilon_{\ell,m}. \tag{7}$$

$X_{\ell,m}$ includes the employer wage premium (ψ) race, gender, ethnicity, age, and person effect

(θ) for both workers, and the interaction between those characteristics. α_1 indicates which characteristics of the pair increase the likelihood of referral.

To address the possibility of sorting within neighborhoods, I also estimate a generalization of Equation (6). If workers who end up working together sort onto the same blocks, then α_0 is biased upward in Equations (6) and (7), and including individual effects will attenuate the estimate (Bayer et al. 2008). To this end, I estimate $W_{\ell,m} = \lambda_\ell + \lambda_m + \alpha_0 R_{\ell,m} + \varepsilon_{\ell,m}$, where λ_ℓ and λ_m are individual effects. The generalization of Equation (7) is equivalent.[12] Identification is possible since each worker appears in multiple pairs within the same block group.

I base inference on an empirical covariance matrix that is robust to heteroskedasticity and arbitrary correlation of errors within block groups. This is more conservative than Bayer et al. (2008) who use a pairs bootstrap. The size of my matched pairs dataset renders bootstrapping computationally infeasible. In unreported analysis using simple random samples, standard errors based on the pairs bootstrap are an order of magnitude smaller than the cluster robust errors. The reader should therefore take my results to represent a very conservative lower bound for precision.

4.3 Results

The results of the analysis appear in Table V. Column (1) of Panels A and B present results of the basic model in Equation (6) and Equation (4.2). The coefficient estimate on R indicates that residing on the same block is associated with a .026 percentage point increase in the probability that a mover takes a job with one of his neighbors' employers. This is an 18 percent increase over the baseline. Given the much more restricted nature of my sample and design, my results support Bayer et al. (2008), who find a 33 percent increase over

[12] The elements α_1 associated with individual characteristics included in $X_{\ell,m}$, which are not of interest here, are not identified when Equation (7) is estimated with individual fixed effects.

the baseline in their most conservative specification. Like Bayer et al. (2008), I find that controlling for individual heterogeneity has very modest positive effect on the overall social interaction effect, indicating that unobserved sorting has induced a negative bias in Column (1) of Panel A.

The results of introducing heterogeneity in the social interaction effect by estimating Equation (7) and its generalization estimated using individual effects are presented in Column (2) of Panels A and B. For this paper, interest focuses on the results involving the employer wage premium. First, there is a strong positive complementarity between θ of the worker changing jobs, which is a measure of the market return to human capital, and the wage premium held by the referral provider, ψ_{ref}. The probability of a referral increases between a high ability worker and a worker in a high-paying firm. This finding is reinforced by the results on $R \times \theta \times \theta_{ref}$, which show that referrals are also more likely between high ability workers. Together, these results are consistent with a model of referral in which high ability workers are recruited by other high ability workers into higher-paying firms, as suggested in the seminal paper of Montgomery (1991). The other results in Table V indicate that referrals are more likely between workers with the same demographic characteristics, consistent with the homophily in social interactions documented throughout the literature.

5 Results: Wage Premium Effects

5.1 Main Results

The main results use variations of Equation (4) to asses how the quality of job information networks affects job search outcomes. In the job search framework, if job information networks increase access to high-paying firms, increased network quality should make workers more likely to change jobs, and to move to firms with higher wage premia when they do. In this section, I present evidence consistent with these predictions, and show that my results are not sensitive to violations of the identifying assumptions or to sample selection effects. I then show that the effect of improving network quality is increasing through the distribution of observed outcomes, which establishes that the effect is not due solely to an increased rate of job offers.

5.1.1 Mean Effects

Table VI presents results of estimating Equation (4) with and without controls for block group heterogeneity. The key result is the contrast between the baseline specification in Column (1), which does not control for block group correlations in outcomes, and the specifications in the Columns (2) and (3) that do. Column (1) shows the raw correlation between network quality ($\bar{\psi}_{block}$) and premium on the job to which i makes a transition (ψ_i), controlling for the premium on the origin job and observable characteristics. The estimate of γ ($0.34 \pm .004$) is very large, but in this specification, γ absorbs any unobserved correlates of formal job search that are not included in the model.

The social interaction parameter, γ, is identified in the models with block group controls presented in Columns (3) and (6). The results are nearly identical across specifications. In Column (3), $\hat{\gamma} = 0.11 \pm 0.004$. Interpreted as an estimate of the offer function, this implies 11 percent of job offers arrive through the network. This is in line with Ioannides and Loury (2004), who report that 8.5 percent of employed workers use personal contacts to find work. For a theoretically agnostic interpretation of this estimate, it implies that a one standard deviation increase in network quality is associated with a 25 percent increase in the employer wage premium. Observable demographic characteristics explain little of the variation in ψ_i. This is consistent with wage premium offers being only weakly related to individual characteristics, which in turn is consistent with the notion that they are rents associated with information frictions.

Columns (2) and (5) present alternative estimates of γ based on a contrast between $\bar{\psi}_{block}$ and the mean wage premium among workers in the block group ($\bar{\psi}_{bg}$). I use this contrast to estimate the selection correction model for computational simplicity. The point estimates are nearly identical, and I conclude that $\bar{\psi}_{bg}$ controls for block group-specific correlation in outcomes.

5.1.2 Sensitivity Analysis

The evidence in Section 3.4.1 indicates very little sorting within neighborhoods. In Table VII, I investigate how sensitive the results in Table VI are to including block-level and individual controls that should pick up any sorting that occurs in violation of the identifying assumption. Column (1) of Table VII carries over the baseline result presented as Column (3) in Table VI. Columns (2)-(4) assess sensitivity by introducing block-level controls, controls for the industry of the origin job, and the interaction between MSA and initial industry.

In particular, Table VII includes the average value of θ, the person effect from the AKM decomposition, for all workers in the block. Any excess sorting within neighborhoods on the basis of individual-specific attributes correlated with earnings should show up here. If such sorting were driving the results, including $\bar{\theta}_{block}$ would severely attenuate the point estimate. The block-level controls improve fit, but attenuate the point estimate only marginally.

A formal sensitivity analysis indicates that, for my results to be driven by unobservable sorting, the bias not explained by the controls added in Table VII must be several times as strong as the bias that is.[13] Recall from Section 3.4.1 that within-neighborhood sorting

[13] Let $\hat{\gamma}_0$ be the estimator of γ based on the preferred model of Column (1) in Table VII and $\hat{\gamma}_1$ be the estimator based on the model of Column (2) in Table VII Finally, define scalars B_V, B_u, and μ, such that

$$\operatorname{plim} \hat{\gamma}_0 = \gamma + B_V + B_u \tag{8}$$

and $B_u = \mu B_V$. B_V is the part of the bias in the preferred estimate that can be predicted by block-level average characteristics, and B_u is the part that can not. μ measures how important unobservables are relative to observables. Schmutte (2012) shows that $\mu^* = \operatorname{plim} \frac{R}{1-R}$ is the value of μ required to fully explain the observed effect when the true value of γ is zero where $R = \frac{\hat{\gamma}_1}{\hat{\gamma}_0}$.

on unobservables is weaker than sorting on observables. Assuming that unobservables are no more associated with job search outcomes than observables (as in Altonji, Elder and Taber 2005), the sensitivity analysis implies that sorting on unobservables, in violation of the identifying assumption, could not explain all of the estimated social interaction effect.

5.1.3 Effects on Mobility and Sample Selection

Table VIII presents estimates from a Heckman selection-correction model, and contains two key results. First, consistent with the job search framework, the probability of making a job-to-job move is increasing in local network quality. Second, correcting for the selection of job changers does not change the estimated effect of network quality on the expected wage premium on the next job. Column (2) reports the selection equation, a first-stage probit predicting whether a worker changes jobs in 2002:Q4, using the sample of quarterly job changers, on the basis of demographic characteristics and the wage premium on the inside job (ψ_0).[14] Column (1) reports second-stage estimates the wage premium offer function, excluding ψ_0, the influence of which now enters through the Inverse Mills ratio. The exclusion of ψ_0 is consistent with the job search framework where ψ_{0i} does not affect the offer distribution.

The results in Column (2) show that ψ_0 is negatively associated with making a job-to-job move, as the search model predict. Furthermore, increases in network quality ($\bar{\psi}_{block}$) increase the probability of a job-to-job transition. For the latter effect to be explained by sorting on unobservables, living on a block with higher average wage premium must be associated with some underlying propensity to change jobs. The corrected estimate of the social interaction effect in Column (1) is $\hat{\gamma} = 0.11 \pm 0.02$, which is identical to the uncorrected estimate. It appears that there is sufficient random variation in mobility that truncation of the offer distribution has little effect on the model estimates. A caveat is warranted here,

[14] The results are not sensitive to the choice of reference quarter.

since exclusion of ψ_0 from the second-stage offer function, while internally consistent with the job search model, may nevertheless be invalid. In any event, the results for the selection equation support the job search model, and there is no evidence of sample selection bias with respect to the results estimated in Table VI.

5.1.4 Distributional Effects

Figure 1 reports results of testing the distributional predictions of Section 3.1. In each subplot, the horizontal axis reports percentiles in the distribution of wage premium outcomes. Against it are plotted estimates of the unconditional quantile partial effects (UQPE) of an explanatory variable. Figure 1a plots results for the initial wage premium (ψ_0) and Figure 1b plots results for network quality ($\bar{\psi}_{block}$).[15] The solid line in each plot connects the point estimates, and the dashed lines indicate a 95 percent confidence interval.

The positive slope and negative intercept in Figure 1a support the prediction that an increase in ψ_0 compresses the outcome distribution. Exceptions appear at fifth percentile, and above the 75th. The positive slope and intercept in Figure 1b support the prediction that the effects of increasing network stretch the observed outcome distribution from the right.[16] The latter prediction is non-trivial, since if the effect of job information networks is to increase search effort rather than improve the offer distribution, the effect of higher network quality will be to compress the distribution of realized outcomes from the right, and we would expect a negative slope in Figure 1b. I therefore rule out the possibility that this alternate social interaction mechanism can fully explain my results.

[15] Unconditional quantile regression estimates the partial effect of a variable of interest on the marginal outcome distribution. Firpo, Fortin and Lemieux (2009) show how to estimate the UQPE from the recentered influence function (RIF). I use RIF-OLS adopting the specification in column (3) of Table VI. Using specification (1) or (2) does not affect on the results.

[16] For details of these predictions and their proof, see Proposition 2 in Schmutte (2012).

5.2 Alternative Specifications

5.2.1 Robustness: Referral effects by Demographic Subgroups

My results rely on the assumption that workers interact with their residential neighbors to find work. Previous research indicates that immigrant workers and younger workers are more likely to use referrals in job search, and to participate in local social networks. I therefore expect to see a much stronger relationship between network quality and job search outcomes for these demographic groups.

Table IX reports results of allowing for demographic heterogeneity in the influence of network quality on job search outcomes. Column (1) reproduces estimates from the benchmark model in Column (3) of Table VI. Column (2) allows the effect of $\bar{\psi}_{block}$ to vary with nativity. Non-native workers have $\hat{\gamma} = 0.16 \pm 0.006$, which is twice the magnitude of the pooled estimate. This result is consistent with work finding that immigrants are more likely to use personal contacts to find work, to reside in ethnic enclaves, and to find jobs by referral than their native counterparts (Andersson, Burgess and Lane 2008; Damm 2009; Beaman and Magruder 2012). Elliott (2001) finds that non-native workers in large cities are roughly twice as likely to have been hired to a recent job through referral.

Columns (3) and (4) allow the effect of $\bar{\psi}_{block}$ to vary by age group. Column (3) contrasts 'younger workers', defined as those between 25-35, with the rest of the population. Column (4) does the same for 'older workers', defined as those between 45-55. The local interaction effects is stronger for younger workers, and weaker for older workers, consistent with evidence cited by Ioannides and Loury (2004). In unreported results, allowing for heterogeneity by racial categories does not produce significant differentials, consistent with Holzer (1988). If the main result were driven by sorting across blocks within neighborhoods, interactions with race would alter the measured effect.

5.2.2 Robustness: The Relationship to Direct Referrals

Given the evidence on local referrals in Section 4, and the fact that ψ is a fixed characteristic of firms, there is a concern that the local interaction effects documented in this section arise mechanically by workers getting jobs in their neighbor's firms. Here I show that the local interaction effect is not driven entirely by direct referrals. This result also implies that measured spillovers in wage premia reflect the transfer of indirect information about better paying jobs through local networks.

The analysis is presented in Table X. Column (1) carries over results from Column (3) of Table VII for comparison. Columns (2)-(4) show the effect of allowing separate effects of $\bar{\psi}_{block}$ for workers who, on changing jobs, move to a firm that already employs one of their block-level neighbors. If the spillovers in wage premia are driven by direct referrals, then the estimate of $\bar{\psi}_{block}$ for the other workers should be eliminated. I find that workers who move to the same employer as a neighbor have worse overall outcomes ($-.07 \pm .002$). The estimate of γ in Column (4) ($0.08 \pm .004$) takes out the influence of workers who move to a job with the employer of a block-level neighbor. It is identical to the baseline estimate in Column (1). The interaction effect in Column (4) ($0.18 \pm .009$) shows a much stronger, possibly mechanical, correlation between $\bar{\psi}_{block}$ and job search outcomes among workers moving to a neighbor's firm. These findings complement the results in Table V, which suggest local referrals are associated with high-ψ employers. However, they also reveal that not all of the estimated social interactions are driven by direct referral, consistent with other, indirect forms of information sharing in job information networks.

5.2.3 Robustness: Analysis for the Non-Employed

Up to this point, the analysis has focused entirely on workers who change jobs. In the job search framework, conditional on becoming employed, unemployed workers should find jobs in higher-paying firms when their neighbors are already in higher-paying firms. Table XI reports results of estimating the specifications in Table VII for workers who are not employed in 2002. Since LEHD data do not record search behavior, I restrict the analysis to prime-age males to limit the influence of non-participation. For the non-employed, ψ_0 is not available, and I report only the estimate for ψ_{block}.

The results support the job search prediction, but the effects are smaller for the non-employed. The estimate of γ is 0.06 ± 0.010 in the preferred specification that includes block group controls, and attenuated (0.04 ± 0.010) by introducing block mean characteristics. Column (4) shows that the estimate of γ is wiped out (0.01 ± 0.006) after controlling for the sector of the destination job. This indicates that for non-employed workers, job information networks help locate employment in higher-paying sectors. As we will see in the next section, employed workers find employment in higher-paying firms within a given sector.

Section 4 highlights the importance of unobserved worker ability in direct referrals. The literature on scarring suggests that long-term unemployed workers are perceived to be of lower ability. Hence, local job information and referral networks may not operate as strongly for the unemployed. These results suggest job information networks also help the unemployed find higher-paying jobs, but perhaps in different ways than for the employed.

5.2.4 Alternative Mechanism: Agglomeration Economies

The presence of urban agglomeration economies suggests an alternative mechanism for local interaction effect. Firms in agglomerated locations may pay more because of human capital spillovers (Moretti 2004), learning (Glaeser and Mare 2001), job-hopping (Freedman 2008), or compensating differentials (Fu and Ross forthcoming). If agglomeration drives the results, then the estimated effect of local interactions on the wage premium is really the effect of being drawn toward employment in the agglomerated location.

Table XII reports extensions that control for characteristics of the destination employer that are associated with agglomeration. Column (1) carries over the baseline specification from Column (3) of Table VII. Column (2) controls for a cubic in the size of the destination firm ($0.08 \pm .004$), Column (3) controls for the destination sector ($0.07 \pm .003$), and Column (4) controls for the interaction between the destination sector and the MSA, which should allow for city-specific industrial agglomerations ($0.06 \pm .003$).[17] These controls attenuate, but do not eliminate, the estimate of γ. The presence of agglomeration economies may explain some, but not all, of the social interaction effect.

6 Conclusion

The purpose of this paper is to understand how workers use local social networks to find high-paying jobs. I show that workers, when changing jobs, are disproportionately likely to become coworkers of a neighbor, and are more likely to become employed in a high-wage firm when their neighbors already work in high-wage firms. My results indicate the presence of local social interactions in job search, and are consistent with workers using local networks to find higher-paying jobs. As a result, there are modest, but economically meaningful local interactions in earnings. I am able to shed light on these mechanisms by applying a clear identification strategy to study the implications of a basic job search framework. I also provide evidence consistent with direct referrals driving local spillovers in employment location. In doing so, I extend recent research that documents the relevance of local social interactions in job search (Topa 2001; Bayer et al. 2008; Hellerstein et al. 2011).

This paper also speaks to the closely related problem of how employers use referrals to recruit workers. The implication that workers use job information networks to get high-paying jobs raises the question of where these wage premia come from. Montgomery (1991)

[17] Note that Column (3) of Table XII and Column (4) of Table XI report the same model for employed and non-employed workers. Unlike unemployed workers, network quality helps employed workers find higher-paying jobs within a sector.

shows that firms use referrals to find high ability workers, and when they do, they pay a premium that dissipates the profit. Many results of this paper are consistent with such a mechanism. They suggest that direct referrals are more likely when the referral recipient has high-ability and the referral provider is in a high-wage firm. Furthermore, it appears that job information networks operate differently for non-employed workers than for employed workers. These results are consistent with employers using information networks to screen and attract high-ability applicants. My results are also partially consistent with job information networks acting to channel workers to higher paying jobs in agglomerated sectors, where firms are more likely to pay well and also to recruit through referral to find better workers. These findings suggest further avenues for research using matched data to understand the interaction between agglomeration economies and the incentives on both sides of the labor market to use job information networks.

References

Abowd, John M, Robert H Creecy, and Francis Kramarz. 2002. "Computing person and firm effects using linked longitudinal employer-employee data." Tech. Rep. TP-2002-06, LEHD, U.S. Census Bureau.

Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High wage workers and high wage firms." *Econometrica* 67 (2):251–333.

Abowd, John M., Paul Lengermann, and Kevin L. McKinney. 2003. "The measurement of human capital in the U.S. economy." Tech. Rep. TP-2002-09, LEHD U.S. Census Bureau.

Abowd, John M., Kevin L. McKinney, and Ian M. Schmutte. 2010. "How important is endogenous mobility for measuring employer and employee heterogeneity?" Unpublished Manuscript. Cornell University School of Industrial and Labor Relations.

Abowd, John M., Bryce E. Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock. 2009. "The LEHD infrastructure files and the creation of the Quarterly Workforce Indicators." In *Producer Dynamics: New Evidence from Micro Data*, *Studies in Income and Wealth*, vol. 68, edited by Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, chap. 5. The University of Chicago Press, 149–234.

Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber. 2005. "Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools." *Journal of Political Economy* 113 (1):151–184.

Andersson, Fredrik, Simon Burgess, and Julia Lane. 2008. "Recent immigrants, social networks and labor markets." Unpublished manuscript.

Bayer, Patrick, Stephen L. Ross, and Giorgio Topa. 2008. "Place of work and place of residence: Informal hiring networks and labor market outcomes." *Journal of Political Economy* 116:1150–1194.

Beaman, Lori and Jeremy Magruder. 2012. "Social networks and the dynamics of labor market outcomes: Evidence from refugees resettled in the U.S." *Review of Economic Studies* :128–161.

Bewley, Truman F. 1999. *Why wages don't fall during a recession*. Harvard University Press.

Bjelland, Melissa, Bruce Fallick, John Haltiwanger, and Erika McEntarfer. 2011. "Employer-to-employer flows in the United States: Estimates using linked employer-employee data." *Journal of Business & Economic Statistics* 29 (4):493–505.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Yannis M. Ioannides. 2011. "Identification of social interactions." In *Handbook of Social Economics*, vol. 1, edited by Alberto Bisin Jess Benhabib and Matthew O. Jackson. North-Holland, 853–964.

Cahuc, Pierre and François Fontaine. 2009. "On the efficiency of job search with social networks." *Journal of Public Economic Theory* 11 (3):411–439.

Calvo-Armengol, Antoni and Matthew O. Jackson. 2007. "Networks in labor markets: Wage and employment dynamics and inequality." *Journal of Economic Theory* 132 (1):27–46.

Calvo-Armengol, Antoni and Yves Zenou. 2005. "Job-matching, social networks and word-of-mouth communication." *Journal of Urban Economics* 57 (3):500–522.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace heterogeneity and the rise of West German wage inequality." .

Case, Anne C. and Lawrence F. Katz. 1991. "The company you keep: The effects of family and neighborhood on disadvantaged youths." Tech. rep., National Bureau of Economic Research.

Cingano, Federico and Alfonso Rosolia. 2012. "People I know: job search and social networks." *Journal of Labor Economics* 30 (2):291 – 332.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon. 2008. "Spatial wage disparities: Sorting matters!" *Journal of Urban Economics* 63:723–742.

Conley, Timothy G. and Giorgio Topa. 2007. "Estimating dynamic local interactions models." *Journal of Econometrics* 140:282–303.

Conley, Timothy G. and Christopher R. Udry. 2010. "Learning about a new technology: Pineapple in Ghana." *American Economic Review* 100:35–69.

Damm, Anna Piil. 2009. "Ethnic enclaves and immigrant labor market outcomes: Quasi-experimental evidence." *Journal of Labor Economics* 27 (2):pp. 281–314.

Dolnick, Sam. 2011. "Many immmigrant's job search starts in Chinatown." *New York Times* :A19.

Dustmann, Christian, Albrecht Glitz, and Uta Schönberg. 2011. "Referral-based job search networks." IZA Discussion Papers 5777, Institute for the Study of Labor (IZA).

Elliott, James R. 2001. "Referral hiring and ethnically homogeneous Jobs: How prevalent is the connection and for whom?" *Social Science Research* 30 (3):401 – 425.

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux. 2009. "Unconditional quantile regressions." *Econometrica* 77 (3):953–973.

Freedman, Matthew L. 2008. "Job hopping, earnings dynamics, and industrial agglomeration in the software publishing industry." *Journal of Urban Economics* 64 (3):590–600.

Fu, Shihe and Stephen Ross. forthcoming. "Wage premia in employment clusters: How important is worker heterogeneity?" *Journal of Labor Economics* .

Glaeser, Edward L. and David C. Mare. 2001. "Cities and skills." *Journal of Labor Economics* 19 (2):316–42.

Goel, Deepti and Kevin Lang. 2009. "Social ties and the job search of recent immigrants." Working Paper 15186, National Bureau of Economic Research.

Guryan, Jonathan, Kory Kroft, and Matthew J. Notowidigdo. 2009. "Peer effects in the workplace: Evidence from random groupings in professional golf tournaments." *American Economic Journal: Applied Economics* 1 (4):34–68.

Hellerstein, Judith K., Melissa McInerney, and David Neumark. 2011. "Neighbors and coworkers: The importance of residential labor market networks." *Journal of Labor Economics* 29 (4):659 – 695.

Holzer, Harry. 1988. "Search method use by unemployed youth." *Journal of Labor Economics* 6:1–20.

Ioannides, Yannis M. and Linda Datcher Loury. 2004. "Job information networks, neighborhood effects, and inequality." *Journal of Economic Literature* 42:1056–1093.

Iranzo, Susana, Fabiano Schivardi, and Elisa Tosetti. 2008. "Skill dispersion and firm productivity: An analysis with matched employer-employee data." *Journal of Labor Economics* 26:247–285.

Katz, Lawrence F, Jeffrey R Kling, and Jeffrey B Liebman. 2001. "Moving to opportunity in Boston: Early results of a randomized mobility experiment." *The Quarterly Journal of Economics* 116 (2):607–654.

Kinsler, Josh. 2012. "Assessing Rothstein's critique of teacher value-added models." *Quantitative Economics* 3 (2):333–362.

Laschever, Ron. 2009. "The Doughboys network: Social interactions and the employment of World War I veterans." Purdue University Department of Economics.

Manski, Charles F. 1993. "Identification of endogenous social effects: The reflection problem." *Review of Economic Studies* 60:531–542.

Moffitt, Robert A. 2000. "Policy interventions, low-level equilibria and social interactions." In *Social Dynamics*. MIT Press, 45–82.

Mok, D., B. Wellman, and J. A. Carrasco. 2009. "Does distance still matter in the age of the Internet?" *Urban Studies* 46:2747–2783.

Montgomery, James D. 1991. "Social networks and labor market outcomes: Toward an economic analysis." *American Economic Review* 81:1408.

Moretti, Enrico. 2004. "Human capital externalities in cities." In *Cities and Geography, Handbook of Regional and Urban Economics*, vol. 4, edited by J. Vernon Henderson and Jacques-Franois Thisse. Elsevier, 2243 – 2291.

Mortensen, Dale and Tara Vishwanath. 1994. "Personal contacts and earnings: It is who you know!" *Labour Economics* 1:187–201.

Mortensen, Dale T. 2003. *Wage dispersion: Why are similar workers paid differently.* Zeuthen lecture book series. MIT Press.

Nordström Skans, Oskar and Francis Kramarz. 2011. "When strong ties are strong – networks and youth labor market entry." Working Paper Series 2011:18, IFAU - Institute for Labour Market Policy Evaluation.

Oyer, Paul and Scott Schaefer. 2012. "Firm/employee matching: An industry study of american lawyers." Working Paper 18620, National Bureau of Economic Research.

Schmutte, Ian. 2012. "Job referral networks and the determination of earnings in local labor markets." Unpublished Manuscript. University of Georgia Department of Economics.

Topa, Giorgio. 2001. "Social interactions, local spillovers and unemployment." *Review of Economic Studies* 68:261–295.

7 Tables and Figures

Table I: List of Metropolitan Statistical Areas Used

Austin-Round Rock, TX	Philadelphia-Camden-Wilmington, PA-NJ-DE-MD
Baltimore-Towson, MD	Pittsburgh, PA
Charlotte-Gastonia-Concord, NC-SC	Portland-Vancouver-Beaverton, OR-WA
Chicago-Naperville-Joliet, IL-IN-WI	Richmond, VA
Dallas-Fort Worth-Arlington, TX	Riverside-San Bernardino-Ontario, CA
Houston-Sugar Land-Baytown, TX	Sacramento–Arden-Arcade–Roseville, CA
Indianapolis-Carmel, IN	San Antonio, TX
Jacksonville, FL	San Diego-Carlsbad-San Marcos, CA
Kansas City, MO-KS	San Francisco-Oakland-Fremont, CA
Los Angeles-Long Beach-Santa Ana, CA	San Jose-Sunnyvale-Santa Clara, CA
Louisville-Jefferson County, KY-IN	Seattle-Tacoma-Bellevue, WA
Miami-Fort Lauderdale-Miami Beach, FL	St. Louis, MO-IL
Milwaukee-Waukesha-West Allis, WI	Tampa-St. Petersburg-Clearwater, FL
Minneapolis-St. Paul-Bloomington, MN-WI	Virginia Beach-Norfolk-Newport News, VA-NC
Oklahoma City, OK	
Orlando-Kissimmee, FL	

NOTE.– List of Metropolitan Statistical Areas used in the analysis with population summaries based on publicly available Census data. All observations used in the analysis were for workers whose Census block of residence in 2002 and 2003 fell in one of these 30 MSAs.

Table II: Descriptive Statistics

Variable	Full Sample (1)	Quarterly Job Changers (2)	Annual Job Changers (3)
White	0.6572	0.6220	0.6495
Black	0.1151	0.1205	0.1129
Hispanic Origin	0.1167	0.1400	0.1274
Male	0.5098	0.4979	0.5886
Born in U.S.	0.8098	0.8026	0.8145
Age in 2002	40.5456	34.9561	37.0943
N	$25,689,739$	$815,899$	$2,189,659$

NOTE.–Summary statistics for a sample of workers with reported UI earnings in one of 30 large MSAs between 2002 and 2003. Except 'Age in 2002', each variable is an indicator, and the summaries in the table represent the fraction of the sample with the associated characteristic. The sample is restricted to workers who did not move MSAs during 2002-2003, were at least 14 years of age in 2002, had valid data for block of residence in 2002 and 2003, and who lived on a block where at least 10 other workers contribute data to compute the block-level average ψ. 'Quarterly Job Changers' are from a sample that identifies the exact quarter of a direct job-to-job transition. 'Annual Job Changers' use a less restrictive definition of job-to-job transition and captures any change in dominant job from 2002–2003 as long as the worker is employed all year in both years.

Table III: Job Ladder Transition Probabilities

	Destination ψ-Decile									
Origin Decile	1	2	3	4	5	6	7	8	9	10
1	0.17	0.16	0.15	0.12	0.08	0.09	0.08	0.07	0.05	0.03
2	0.06	0.18	0.18	0.13	0.09	0.10	0.10	0.08	0.06	0.02
3	0.04	0.11	0.17	0.15	0.11	0.11	0.11	0.10	0.07	0.03
4	0.03	0.07	0.12	0.16	0.14	0.13	0.13	0.12	0.09	0.04
5	0.02	0.04	0.08	0.13	0.14	0.17	0.16	0.13	0.09	0.04
6	0.01	0.03	0.05	0.08	0.10	0.16	0.20	0.23	0.11	0.04
7	0.01	0.02	0.04	0.05	0.07	0.12	0.23	0.22	0.19	0.05
8	0.01	0.02	0.03	0.04	0.05	0.09	0.17	0.28	0.25	0.08
9	0.00	0.01	0.02	0.03	0.03	0.05	0.09	0.20	0.37	0.21
10	0.00	0.01	0.01	0.02	0.02	0.03	0.04	0.08	0.20	0.59

NOTE.–Frequency of transitions to an employer in a decile of the ψ distribution ('Destination Decile') conditional on the decile of the origin job ('Origin Decile'). The row frequencies sum to one. On-the-job search implies that the mass should be on and above the diagonal, indicating a 'job ladder' as workers move from lower to higher-paying firms.

Table IV: Sorting within neighborhoods, R^2 method.

Variable	Raw (1)	Block Group Controls (2)
White	.2915	.0132
Hispanic	.2859	.0125
Born U.S.	.2245	.0114
Age	.0301	.0067
ε	.0038	.0002
N	$394,305$	

NOTE.–Measures of sorting within Census block groups. The input dataset contains one individual-level observation per block and the fraction of people (not including the individual) in the block who share the listed characteristic, or its average. Each entry is the R-squared from a regression of the individuals characteristic on the block-level average. The second column controls for block group specific effects. The sample is restricted to blocks with more than six individuals.

Table V: Direct Referral Effects: With and Without Covariates

		Panel A: Block Group Controls			
		No Covariates (1)		With Covariates (2)	
	Variable	Coeff.	t-Stat.	Coeff.	t-Stat.
Reside on same block	R	0.026	23.85	0.146	5.56
AKM Wage effects	$R \times \theta \times \theta_{ref}$			0.040	3.28
Wage premium	$R \times \psi \times \psi_{ref}$			$-.006$	$-.32$
Interaction	$R \times \theta \times \psi_{ref}$			0.036	1.64
Both white	$R \times white \times white_{ref}$			0.014	4.39
Both male	$R \times male \times male_{ref}$			$-.018$	-5.56
Both native-born	$R \times nat \times nat_{ref}$			0.076	5.58
Both hispanic	$R \times hisp \times hisp_{ref}$			0.019	4.69
Age	$R \times age \times age_{ref}$			0.000	3.04
		Panel B: Individual Controls			
		No Covariates (1)		With Covariates (2)	
		Coeff.	t-Stat.	Coeff.	t-Stat.
Reside on same block	R	0.029	38.37	0.108	11.30
AKM Wage effects	$R \times \theta \times \theta_{ref}$			0.028	4.25
Wage premium	$R \times \psi \times \psi_{ref}$			$-.012$	-1.42
Interaction	$R \times \theta \times \psi_{ref}$			0.029	2.51
Both white	$R \times white \times white_{ref}$			0.015	6.45
Both male	$R \times male \times male_{ref}$			$-.018$	-8.19
Both native-born	$R \times nat \times nat_{ref}$			0.052	9.63
Both hispanic	$R \times hisp \times hisp_{ref}$			0.022	7.29
Age	$R \times age \times age_{ref}$			0.000	3.85
Num. Obs.		$1,524,733,934$			

NOTE.–The table reports linear probability models where an observation is a pair of workers who live in the same Census block group. The first worker in the pair (changer) changes jobs between 2002–2003. The second worker in the pair (non-changer) did not. The dependent variable ,W, indicates whether the pair become coworkers. The coefficients have been multiplied by 100 to reflect percentage point changes. Models in Panel A include block group fixed effects. Models in Panel B include individual fixed effects. The "With Covariates" models also control for level effects of the reported variables, as well as the interaction between R and the individual covariates. In Panel A, the model also includes controls for the level of individual covariates, which are not identified in the individual fixed effects model of Panel B. Standard errors are robust and clustered at the block group level, and t-statistics are reported. The results are more precisely estimated when standard errors are estimated by pairwise bootstraps.

Table VI: Wage Premium Estimates: Baseline

	Premium on next job, ψ					
	Annual Job Changers			Quarterly Job Changers		
	(1)	(2)	(3)	(4)	(5)	(6)
Initial ψ: ψ_0 (β)	0.46***	0.45***	0.45***	0.46***	0.45***	0.45***
	(.001)	(.001)	(.001)	(.002)	(.008)	(.008)
Avg. ψ in block:	0.34***	0.11***	0.11***	0.33***	0.10***	0.10***
$\bar{\psi}_{block}$ (γ)	(.003)	(.004)	(.004)	(.004)	(.011)	(.012)
Avg. ψ in block group:		0.34***			0.34***	
$\bar{\psi}_{bg}$ (ϕ)		(.005)			(.024)	
white	−.00***	−.00***	−.01***	0.00	0.00	−.00
	(.001)	(.001)	(.001)	(.001)	(.002)	(.001)
Hispanic Origin	−.02***	−.02***	−.03***	−.02***	−.01**	−.02***
	(.001)	(.001)	(.001)	(.001)	(.005)	(.004)
male	0.03***	0.03***	0.03***	0.03***	0.03***	0.04***
	(.000)	(.000)	(.000)	(.001)	(.004)	(.004)
age in 2002	0.010***	0.01***	0.01***	0.01***	0.01***	0.01***
	(.000)	(.000)	(.000)	(.000)	(.001)	(.001)
Square of age in 2002	−.00***	−.00***	−.00***	−.00***	−.00***	−.00***
	(.000)	(.000)	(.000)	(.000)	(.000)	(.000)
Born in U.S.	0.00***	0.01***	0.01***	0.00**	0.01*	0.01***
	(.001)	(.001)	(.001)	(.001)	(.002)	(.003)
θ from wage eqn.	−.00	−.00***	−.01***	−.00***	−.01	−.01
	(.001)	(.001)	(.001)	(.001)	(.010)	(.010)
block group controls	no	no	yes	no	no	yes
N		2,198,659			815,899	
R^2	.275	.305	.278	.315	.318	.271

NOTE.–Estimates of the log wage premium, ψ, on the destination job for quarterly and annual job changers. Models (3) and (6) are the preferred specifications. Cluster robust standard errors in parentheses. In columns (1)–(3) and (4), standard errors are clustered at the block group level, and at the MSA level in models (5) and (6). (*), (**), or (***) indicate the coefficient is statistically different from zero at the 10, 5, and 2.5 percent level.

Table VII: Wage Premium Estimates: Sensitivity Analysis

	Premium on next job, ψ			
	Baseline	Alternative Specifications		
	(1)	(2)	(3)	(4)
Initial ψ: ψ_0 (β)	0.45^{***}	0.45^{***}	0.37^{***}	0.36^{***}
	(.001)	(.001)	(.001)	(.001)
Avg. ψ in block:	0.11^{***}	0.09^{***}	0.08^{***}	0.08^{***}
$\bar{\psi}_{block}$ (γ)	(.004)	(.004)	(.004)	(.004)
frac. white on block		0.02^{***}	0.02^{***}	0.01^{***}
		(.003)	(.003)	(.003)
frac. Hispanic on block		$-.03^{***}$	$-.03^{***}$	$-.02^{***}$
		(.004)	(.004)	(.004)
frac. male on block		$-.02^{***}$	$-.02^{***}$	$-.01^{***}$
		(.003)	(.003)	(.003)
frac. native born on block		$-.01^{***}$	$-.01^{**}$	$-.01^{*}$
		(.003)	(.003)	(.003)
avg. age on block		0.00^{***}	0.00^{***}	0.00^{***}
		(.000)	(.000)	(.000)
avg. θ on block		0.06^{***}	0.06^{***}	0.06^{***}
		(.002)	(.002)	(.002)
block group controls	yes	yes	yes	yes
Industry of origin job controls	no	no	yes	yes
MSA$\times$ initial industry controls	no	no	no	yes
N		2,198,659		
R^2	.278	.305	.323	.334

NOTE.–Estimates of the log wage premium, ψ, on the destination job for quarterly and annual job changers. To facilitate comparison, Column (1) repeats the benchmark specification reported as Column (3) of Table VI. Column (2) adds block-level mean characteristics to the benchmark specification. Column (3) adds controls for the sector of initial employment. Column (4) adds controls for the interaction of MSA and sector of initial employment. All models include controls for race, gender, quadratic in age, Hispanic origin, nativity, and the person-effect from the AKM decomposition, θ. Robust standard errors clustered at the block group are reported in parentheses. (*), (**), or (***) indicate the coefficient is statistically different from zero at the 10, 5, and 2.5 percent level.

Table VIII: Selection Correction Model Estimates

	ψ Offer Function (1)	Selection Equation (2)
Initial premium		-0.58^{***}
ψ_0		(.017)
Mean premium in block:	0.11^{***}	0.10^{***}
$\bar{\psi}_{block}$	(.023)	(.020)
Mean premium in block group:	0.64^{***}	0.32^{***}
$\bar{\psi}_{bg}$	(.060)	(.069)
λ (Inv. Mills)	0.48^{***}	
	(.058)	
ρ	0.79	
σ	0.61	
N	1,330,475	
$\chi^2_{(9)}$	683.23	

NOTE.–Heckman selection correction model for the log wage premium offer function. Selection on whether a job-to-job move was observed across all employed workers in 2002:Q4. Bootstrapped standard errors clustered on 30 MSAs. (*), (**), or (***) indicate the coefficient is statistically different from zero at the 10, 5, and 2.5 percent level. Both models include all controls from Table VI. ρ is the estimated correlation between the errors in the selection equation and the offer function.

Table IX: Wage Premium Estimates: Demographic Heterogeneity in the Local Interaction Effect

	Premium on next job, ψ			
	Baseline	Native Workers	Younger Workers	Older Workers
	(1)	(2)	(3)	(4)
Initial ψ: ψ_0 (β)	0.37***	0.37***	.37***	0.37***
	(.001)	(.001)	(.001)	(.032)
Avg. ψ in block: $\bar{\psi}_{block}$ (γ)	0.08***	0.16***	.07***	0.09***
	(.004)	(.006)	(.004)	(.004)
Born in U.S. $\times\bar{\psi}_{block}$		−.09***		
		(.005)		
Younger Worker $\times\bar{\psi}_{block}$			.04***	
			(.004)	
Older Worker $\times\bar{\psi}_{block}$				−.04***
				(.005)
block group controls	yes	yes	yes	yes
block mean characteristics	yes	yes	yes	yes
Industry of origin job	yes	yes	yes	yes
N		2,198,659		
R^2	.323	.323	.323	.323

NOTE.– Estimates of the log wage premium, ψ, on the destination job for workers in the sample of annual job changers. 'Younger workers' are those between 25–35 years of age in 2002. 'Older workers' are those between 45–55 years of age in 2002. Column (1) carries over the baseline specification from Column (3) of Table VII. All models include controls for race, gender, quadratic in age, Hispanic origin, nativity, and the person-effect from the AKM decomposition, θ. All models also include these characteristics at their block-level mean, block-group controls, and controls for the major NAICS sector of the origin job. Robust standard errors clustered at the block group are reported in parentheses. (*), (**), or (***) indicate the coefficient is statistically different from zero at the 10, 5, and 2.5 percent level.

Table X: Wage Premium Estimates: The Influence of Direct Referrals

	Premium on next job, ψ			
	Baseline (1)	(2)	(3)	(4)
Initial ψ: ψ_0 (β)	0.37***	0.37***	.37***	0.37***
	(.024)	(.001)	(.001)	(.001)
Avg. ψ in block: $\bar{\psi}_{block}$ (γ)	0.08***		.08***	0.08***
	(.004)		(.004)	(.004)
Move to same job (as a block-neighbor)		$-.03^{***}$	$-.03^{***}$	$-.07^{***}$
		(.001)	(.001)	(.002)
Move to same job $\times\bar{\psi}_{block}$				.18***
				(.009)
block group controls	yes	yes	yes	yes
block mean characteristics	yes	yes	yes	yes
Industry of origin job	yes	yes	yes	yes
N		2,198,659		
R^2	.323	.323	.323	.323

NOTE.– Estimates of the log wage premium, ψ, on the destination job for workers in the sample of annual job changers. Column (1) carries over the baseline specification from Column (3) of Table VII. All models include controls for race, gender, quadratic in age, Hispanic origin, nativity, and the person-effect from the AKM decomposition, θ. All models also include these characteristics at their block-level mean, block-group controls, and controls for the major NAICS sector of the origin job. Robust standard errors clustered at the block group are reported in parentheses. (*), (**), or (***) indicate the coefficient is statistically different from zero at the 10, 5, and 2.5 percent level.

Table XI: Wage Premium Estimates: Non-employed Workers

	Premium on next job, ψ			
	(1)	(2)	(3)	(4)
Avg. ψ in block:	0.23***	0.06***	0.04***	0.01
$\bar{\psi}_{block}$ (γ)	(.007)	(.010)	(.010)	(.006)
block group controls	no	yes	yes	yes
control for block mean characteristics	no	no	yes	yes
control for industry of destination job	no	no	no	yes
N		223,159		
R^2	.278	.305	.323	.334

NOTE.– Estimates of the log wage premium, ψ, on the destination job for prime-age male workers who are not employed in 2002. All models include controls for race, gender, quadratic in age, Hispanic origin, nativity, and the person-effect from the AKM decomposition, θ. Robust standard errors clustered at the block group are reported in parentheses. (*), (**), or (***) indicate the coefficient is statistically different from zero at the 10, 5, and 2.5 percent level.

Table XII: Wage Premium Estimates: Alternative Explanations

	Premium on next job, ψ			
	(1)	(2)	(3)	(4)
Initial ψ: ψ_0 (β)	0.37***	0.38***	0.33***	0.32***
	(.001)	(.001)	(.001)	(.001)
Avg. ψ in block:	0.08***	0.08***	0.07***	0.06***
$\bar{\psi}_{block}$ (γ)	(.004)	(.004)	(.003)	(.003)
block group controls	yes	yes	yes	yes
block mean characteristics	yes	yes	yes	yes
Industry of origin job	yes	yes	yes	yes
Cubic in destination firm size	no	yes	no	no
Industry of destination job	no	no	yes	yes
MSA×Industry of desitination job	no	no	no	yes
N		2,198,659		
R^2	.323	.325	.430	.453

NOTE.– Estimates of the log wage premium, ψ, on the destination job for workers in the sample of annual job changers. Column (1) carries over the baseline specification from Column (3) of Table VII. All models include controls for race, gender, quadratic in age, Hispanic origin, nativity, and the person-effect from the AKM decomposition, θ. All models also include these characteristics at their block-level mean, block-group controls, and controls for the major NAICS sector of the origin job. Robust standard errors clustered at the block group are reported in parentheses. (*), (**), or (***) indicate the coefficient is statistically different from zero at the 10, 5, and 2.5 percent level.

Figure 1: Unconditional Quantile Regression Estimates

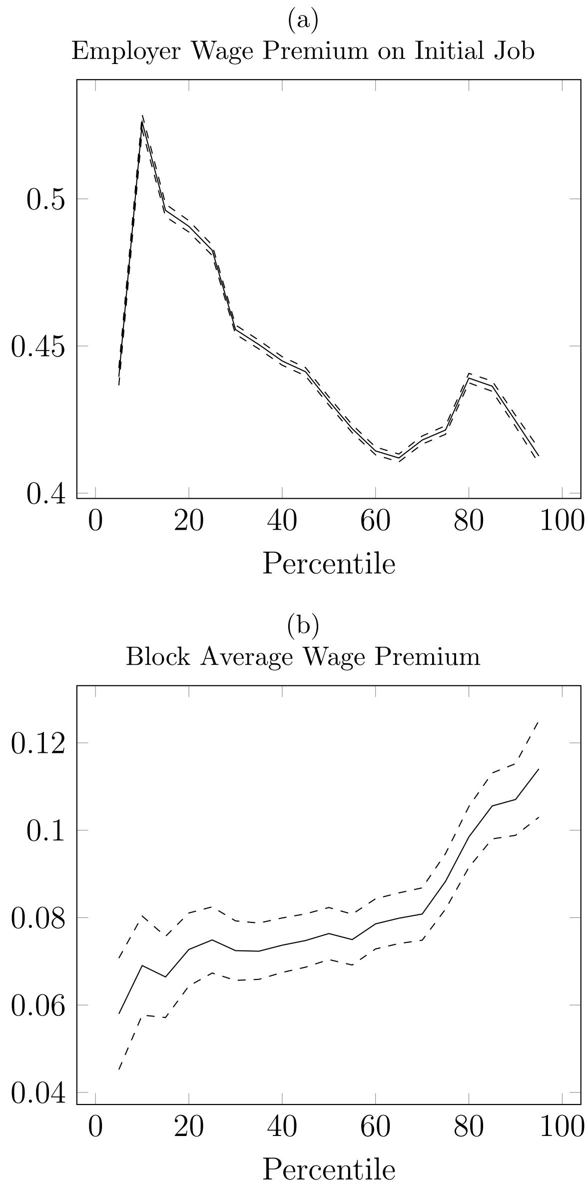

Figures (a) and (b) plot unconditional quantile partial effects (UQPE) at each percentile of the wage premium distribution for job-changers plus or minus two standard errors (dashed lines). Figure (a) reports the UQPE of the employer wage premium on the initial job, and shows that increasing the wage premium paid by the inside firm compresses the distribution of accepted offers from the left. Figure (b) reports the UQPE of the block average wage premium (network quality), and shows that increasing network quality stretches the distribution of accepted offers from the right. The models used to estimate the UQPE include the same variables as column (3) of Table VI.