Job Referral Networks and the Determination of Earnings in Local Labor Markets

Ian M. Schmutte¹

December 10, 2010

¹Department of Economics, Terry College of Business, University of Georgia, schmutte@uga.edu. I have received invaluable advice from John Abowd, Larry Blume, Matt Freedman and Jon Kleinberg. Kaj Gittings, Amanda Griffith, Hope Michelson, Emily Owens, Mike Strain and Russell Toth provided extensive feedback on earlier drafts. I also thank Talia Bar, Jessica Bean, Peter Brummond, Brian Dillon, Annemie Maertens, and Liliana Sousa for valuable discussions. All errors are my own. This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the 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 and do not necessarily represent the views of the U.S. Census Bureau, its program sponsors or data providers. All results have been reviewed to ensure that no confidential information is disclosed.

Abstract

Referral networks may affect the efficiency and equity of labor market outcomes, but few studies have been able to identify earnings effects empirically. To make progress, I set up a model of on-the-job search in which referral networks channel information about high-paying jobs. I evaluate the model using employer-employee matched data for the U.S. linked to the Census block of residence for each worker. The referral effect is identified by variations in the quality of local referral networks within narrowly defined neighborhoods. I find, consistent with the model, a positive and significant role for local referral networks on the full distribution of earnings outcomes from job search.

JEL Codes: J31, J64, R23.

JEL Keywords: Social Interactions; Informal Hiring Networks; Wage Variation; Neighborhood Effects.

1 Introduction

In this paper, I study a previously unexplored connection between two features of the U.S. labor market. The first is that who you know affects where you work. The second is that where you work affects how much you are paid. Getting job information from friends and neighbors is a common strategy, and apparently a productive one: between 30 and 60 percent of new jobs are found through personal contacts (Truman F. Bewley 1999). Why is referral use so prominent? What does this prominence mean for the way the labor markets operate? These remain open questions. One role for referral networks in job search is in helping workers locate information about particularly attractive job opportunities. Two workers can receive different pay simply because they work in different firms (John M. Abowd, Francis Kramarz & David N. Margolis 1999). If workers share information about these pay differentials with their friends and neighbors, then who you know can affect how much you are paid.

The goal of this paper is to identify the role of local referral networks in the assignment of employer-specific pay differentials. In doing so, I provide the first direct evidence of referral effects and neighborhood interactions in earnings determination. I find that workers engaged in on-the-job search receive a positive and significant fraction of their job offers through local interactions. Furthermore, workers who change jobs receive offers with higher pay differentials when workers in their local referral networks are earning higher differentials. These results are robust to various specifications that attempt to address sorting and to correct for sample selection on the quality of one's current job. The magnitude of the effect is very similar to self-reported levels of referral use among employed workers found in survey data (Yannis M. Ioannides & Linda Datcher Loury 2004). My empirical approach relies on workers being likely to interact with their residential neighbors in searching for better jobs. To check the validity of my approach, I extend the model to allow for differences in the productivity of local referral networks between native and non-native workers. I find that the magnitude of the effect of local interactions on job quality are almost twice as strong for non-native workers as for natives, consistent with previous research showing that non-native workers are more likely to use referrals to find work.

My research makes several contributions to the literature on the role of local social interactions in labor markets. This is the first paper to directly identify and estimate local spillovers in earnings determination. Previous research on neighborhood effects has either focused on neighborhood-level effects in social behaviors correlated with income, or on how neighborhood characteristics affect labor market outcomes (see Ioannides & Loury (2004) for

an extensive survey). I am able to make progress because I can separate the part of earnings associated with job assignment from the part due to an individual worker's portable skills. I also have a clean strategy for identifying social interactions in job assignments and earnings by exploiting variation in network quality within neighborhoods. This strategy is derived from Patrick Bayer, Stephen L. Ross & Giorgio Topa (2008) who use it to identify local interactions in job location. In addition to its primary contributions, this paper is the first to verify that employer-specific wage premia have a mobility-related structure that is consistent with a job search model. It is also the first to document the spatial structure in the distribution of earnings components using employer-employee matched data for the United States.

To obtain my results, I develop and empirically implement a model of job search in which workers use referral networks to share information about job opportunities. Consistent with empirical evidence and a range of job search models, employers are distinguished by idiosyncratic wage differentials. In my model, these wage differentials are the only dimension of quality on which employers differ. Such wage differentials are non-market rents that accrue to the workers who find them. A worker searching for a job uses his referral network to try to find these rents. Assuming a contagion process for the social transmission of job information, the implied distribution of job offers has a simple form in which the average quality of a worker's job offer depends directly on the the average quality of his neighbors' jobs. I use the model to derive four testable implications for the observed distribution of job quality among workers making direct job-to-job transitions.

To estimate local referral effects in job quality and test the model's predictions, I use employer-employee matched data from the Longitudinal Employer-Household Dynamics (LEHD) Program. The estimation method follows two stages. In the first stage, I obtain measures of the quality of jobs held by all private sector, non-farm workers from a decomposition of log earnings into components associated with individual and employer heterogeneity (John M Abowd, Robert H Creecy & Francis Kramarz 2002). I then link these job quality estimates to the exact residential block for workers who lived in one of 30 large Metropolitan Statistical Areas (MSAs) in 2002-2003. I measure one's local network quality from the distribution of employer-specific wage premia held by workers from the same residential block.

The empirical model is similar to the conventional neighborhood-effects specification in which an individual's expected outcome from job search depends on the average outcomes of his neighbors. The model therefore raises the identification issues first addressed by Charles F. Manski (1993). Adapting the research design of Bayer, Ross & Topa (2008),

I identify the contribution to job search outcomes of the quality of local social networks from quasi-random assignment of workers to residential blocks within larger neighborhoods. This facilitates distinguishing neighborhood quality from network quality. Neighborhood quality is correlated with search outcomes for a number of reasons. For example workers may sort on the basis of characteristics that affect job search. Different neighborhoods may also have different access to transportation. Workers' residential location decisions are made on the basis of neighborhood quality, but they cannot sort perfectly by block. Thus, the variation in network quality within the neighborhood is exogenous. My model predicts the excess spatial correlation found in employer wage premia at the block-level beyond that found at the neighborhood level. Both conventional and quantile regressions confirm that the relationship between network quality and job search outcomes is significant, economically meaningful, and conforms to the predictions of my enhanced search model.

The estimation results are driven by two stylized features of the distribution of estimated employer-specific wage premia that are also predicted by the model. First, there is a 'job ladder' in the sense that workers who change jobs are more likely to move from lower- to higher-quality jobs (Figure 2). Second, job quality is spatially correlated at the level of the Census block (Figure 4). I show that these features are characteristic of the on-the-job search model with local referral networks, and then show that they hold in the estimated distribution of employer wage premia. My analysis of the spatial correlation of earnings, human capital, and employer characteristics is among the most geographically detailed of its type for U.S. cities. My results also confirm that much of the observed sorting in earnings is correlated with sorting on observable and unobservable human capital characteristics (Pierre-Philippe Combes, Gilles Duranton & Laurent Gobillon 2008, Timothy G Conley & Giorgio Topa 2002). These results are relevant to those studying residential sorting by earnings, human capital characteristics, and employer characteristics in urban labor markets.

2 A Model of Job Search with Referral Networks

I model on-the-job search with social interactions in the transmission of information about new job opportunities. Different employers offer different pay to the same worker, but workers do not know the size of the wage premium offered by any particular employer. They must engage in a process of search to collect information about new jobs. I allow for the possibility that the productivity of the search process may depend on individual characteristics, neighborhood quality, and the quality of the jobs held by people in one's

referral network.¹

The model delivers four major predictions to be verified in the data. The first two are predictions on the temporal and spatial structure of estimated employer-specific log-wage premia. Proposition 1 predicts a ‘job-ladder’. Workers should move to better paying jobs on average. Proposition 2 predicts that the quality of jobs held by workers in the same referral network are positively correlated.

The second set predict the effect of referral network quality on the distribution of outcomes of job search. Proposition 3 predicts that the average outcomes of job search are better for workers with higher quality referral networks. Proposition 4 documents the effects of referral network quality across quantiles of the outcome distribution. Specifically, increases in the quality of the origin job compress the observed job quality distribution from the left, while increases in referral network quality stretch the observed job quality distribution from the right.

2.1 Model Setup

Time evolves continuously and the observed data are snapshots taken at discrete intervals. I denote a model variable, say earnings of the i^{th} worker, evolving in continuous time, as $y(i, t)$. The data observed from this process are denoted as $y_{i1}, y_{i2}, \dots, y_{iT}$ where each $y_{i\tau}$ is an observation on this process at time $t = \tau$.

The model is populated by a finite group of workers and a continuum of employers. Let $i \in \{1 \dots I\}$ index workers and j index employers. Workers are heterogeneous in the characteristics that affect productivity and pay. Let $e(i, t)$ denote the stock of human capital characteristics held by worker i at time t . Different employers compensate workers differently. Let $p_j > 0$ be the idiosyncratic component of employer pay. The earnings function, $y(e(i, t), p_j)$ satisfies log-separability. That is,

$$\ln y(e(i, t), p_j) = \ln y_1(e(i, t)) + \psi_j,$$

where $\psi_j = \ln y_2(p_j)$. ψ_j is the log-wage premium paid by employer j .²

¹This model of wage-setting is motivated by the empirical finding that employer specific heterogeneity explains a large portion of the dispersion in log earnings (Abowd, Kramarz & Margolis 1999, Abowd, Creecy & Kramarz 2002). This is consistent with a primary theoretical result of job search models, which show that information imperfections lead labor markets to fail to eliminate all idiosyncratic differences in pay between employers (Richard Rogerson, Robert Shimer & Randall Wright 2005).

²The wage function given here could arise in a matching model with worker and employer heterogeneity

Workers are infinitely lived and can be either employed or unemployed. Unemployed workers receive new job information at Poisson rate λ_0 . Employed workers receive job information at rate λ_1 . Jobs can end due to exogenous productivity shocks that occur at rate δ . These contact and separation rates are exogenous and common across workers. Workers receive utility in unemployment equivalent to getting a job with wage premium p_b .

When a worker receives a job offer, it is sampled from an employer offering the log wage premium ψ with probability $f(\psi; i, t)$. As the notation indicates, the sampling distribution differs across workers and can change over time. This distribution is a mixture of a formal market offer distribution, denoted $g(\psi; i, t)$ and the distribution of job offers of one's social contacts, denoted $h(\psi; i, t)$. With probability a , a worker samples an offer, ψ , from the distribution of offers in the formal market, $g(\psi; i, t)$. With probability $1 - a$, he samples from the distribution of offers that come through his referral network, $h(\psi; i, t)$. Thus conditional on receiving an offer, the worker draws its type from the distribution

$$f(\psi; i, t) = ag(\psi; i, t) + (1 - a)h(\psi; i, t). \quad (1)$$

The formal offer distribution describes the availability of jobs received when applying directly to employers, answering ads or knocking on doors. The informal offer distribution describes the probability of receiving a job of a particular type conditional on the number of your social contacts who already hold that type of job.

The parameter a measures the strength of social interactions relative to formal channels in delivering new job offers. It is the object of primary interest in the empirical analysis. In setting up the model, I maintain that a is identical across workers. In the empirical work, I estimate the model under this restriction, but also allow for heterogeneity in a on observable characteristics.

2.2 The Referral Distribution, $h(\psi; i, t)$

Individuals receive information about job opportunities from their neighbors. The transmission of this job information is stylized as a contagion process from epidemiology. Here, the types of jobs held by one's neighbors are 'contagious' in the sense that their social network partners are at increased risk to get an offer for the same type of job as one they already hold.

in production with surplus sharing when there is no wage renegotiation (Fabien Postel-Vinay & Jean-Marc Robin 2002).

Let W be an $I \times I$ stochastic matrix whose $(ji)^{th}$ entry measures the probability that job information received by i through a referral originated with worker j .³ I assume the distribution of offers received through referrals satisfies

$$E_h(\psi|W, i, \Psi(t)) = (w^i)^T \Psi(t), \quad (2)$$

where w^i is the i^{th} column of W and $\Psi(t)$ is the $I \times 1$ vector of the wage premium earned by each worker. This specification is consistent with a contagion process where the probability of receiving an offer with log wage premium ψ is increasing in proximity to workers already holding jobs paying that premium:

$$h(\psi; i, t) = (w^i)^T \mathbf{1}(\Psi(t) = \psi). \quad (3)$$

This captures the intuition that referrals are used to share information about particularly attractive wage premia. In the empirical work, I identify the effects of social networks by the quasi-random variation in the residential location choices of individual workers. Since this variation facilitates the identification of local neighborhood interactions, I will specify W in terms of residential proximity.

The construction of h in Equation 3 embeds an assumption that there is no demand-side constraint that affects the distribution of offers through the referral network. This is in keeping with the partial equilibrium nature of the model. Second, and more crucial, is the assumption that the probability that i receives an offer ψ through referral, $h(\psi; i, t)$, is independent of the job search of worker k at t . In other words, i and k are not competitors for the same scarce piece of job information. This assumption is a key feature of the contagion approach, and differs from related models that focus on the routing of job information across social networks in partial or general equilibrium search and matching models (Calvo-Armengol and Jackson, 2004; Calvo-Armengol and Zenou 2005, Wahba and Zenou 2005).⁴

In this paper, I abstract from congestion effects to focus on identifying the effect of local network quality on job search outcomes. This abstraction eliminates dependencies between worker's outcomes in the instantaneous cross-section. More plainly, taking network quality

³In the empirical work, I relax the assumption that social structure is exogenous since I allow for the possibility that people sort into neighborhoods on the basis of unobservable characteristics that might be correlated with their job search outcome.

⁴The above-cited papers emphasize congestion effects in the transmission of job information alongside contagion effects. As Jackline Wahba & Yves Zenou (2005) have shown, network congestion effects lead to empirically verified non-linearities in the use and effects of social contacts to find work.

as given, the job offers received by any worker are independent of those received by any other worker. This assumption is approximately correct if the congestion effect is trivially small relative to the contagion effect. Modeling social interactions in job search as a contagion process allows independence in individual job search outcomes.

2.3 Implications

The search model just outlined yields a continuous-time Markov process over assignments of workers to types of jobs. When there are social interactions, $a \neq 1$, there are spillovers leading to correlations across individuals in the state vector $\Psi(t)$. It is conceptually straightforward to define a transition kernel, $Q(Z, W)$ for the evolution of $\Psi(t)$ from the primitives of the mobility model, λ , δ , a , h and g . The notation reflects the dependence of the kernel on a matrix of observable worker characteristics, Z , and the social distance matrix, W . The full mobility model has the form:

$$Y(t) = \ln y_1(E(t)) + \Psi(t) + \varepsilon(t) \quad (4)$$

$$\Pr(\Psi(t)|\Psi(t-\Delta)) = \exp(Q(Z, W)\Delta)\Psi(t-\Delta), \quad (5)$$

where $Y(t)$ is a vector of observed log earnings, $E(t)$ is an $I \times 1$ matrix of time-varying human capital characteristics, and $\varepsilon(t)$ is a vector of errors. The term $\exp(Q(Z^Q, D)\Delta)$ refers to the matrix exponential. The model delivers simple predictions for the evolution and stationary distribution of $\Psi(t)$. The first is that workers move from lower to higher wage premium jobs.

Proposition 1 *In the job search model described above, assume workers are expected wealth maximizers and $e(i, t)$ is independent of work history. Further, assume workers are myopic about the evolution of the offer distribution. Then employed workers will always accept an offer of a job paying a higher wage premium. In addition, unemployed workers follow a reservation strategy.*

Proof. See Appendix A. ■

The assumption that $e(i, t)$ is independent of job assignment may not hold if workers choose jobs both for their wage premia and also to optimize wage growth associated with experience in a particular sector. It is probably not a bad approximation for workers who supply labor in jobs where there is little human capital specificity, and also for workers who have already selected a career and are changing jobs within their chosen field to maximize earnings (Derek Neal 1999). My main results are based on estimates of the model for all

workers, but to acknowledge the preceding argument, I also allow for heterogeneity in the social interaction parameter a to accommodate the possibility that the model may more accurately describe certain groups of workers than others. To foreshadow the results, I find that my estimates of local interactions in job search are much stronger for non-native than for native workers.

Proposition 1 is true for most models of on-the-job search. The next result is specific to a model with on-the-job search with social transmission of job information. It simply states that the correlation in wage premia earned by socially connected workers is positive.

Proposition 2 *The stationary distribution of Ψ is such that*

$$W_{ii'} \neq 0 \implies \text{Corr}(\psi_i, \psi_{i'}) > 0.$$

That is, the presence of social interactions induces excess correlation in employer-specific wage premia.

This proposition follows from the similarity of the model to that of (Antoni Calvo-Armengol & Matthew O. Jackson 2007), who prove an equivalent result.

Once these two results are verified in the data, I check whether the relationship between referral network quality and job search outcomes exists and conforms to the predictions of the model. A job-to-job switch is an observation from a stochastic process whose mean is $E(\psi|\psi > \psi_0, Z, W, \Psi^0)$ where ψ_0 is the log wage premium on the worker's current job and Ψ^0 is the vector of log wage premia held by all workers at the time of the transition. The following proposition shows that an increase in network quality will increase the mean of the truncated offer distribution.

Proposition 3 *If the distribution of offers received through referral, h , is log concave and $|E_g(\psi|\psi > \psi_0, Z, W, \Psi^0) - E_h(\psi|\psi > \psi_0, Z, W, \Psi^0)|$ is small then*

$$\frac{\partial \mu_{f^*}(\psi_0)}{\partial \mu_h} > 0$$

where $\mu_{f^}(\psi_0) = E(\psi|\psi > \psi_0, Z, W, \Psi^0)$ and $\mu_h = E_h(\psi|Z, W, \Psi^0)$*

Proof. See Appendix A. ■

The requirement that $|E_g(\psi|Z, W, \Psi^0) - E_h(\psi|Z, W, \Psi^0)|$ is small means that the distribution of acceptable offers from referrals is not too different from the distribution of acceptable

offers from formal search. The jobs available through the referral network should generally be fairly close to the distribution of offers that workers would receive through formal search, including those features of job search productivity that are correlated across individuals.

The job search model also yields predictions on the quantiles of the truncated offer distribution. I evaluate these in the empirical work as additional checks of the validity of the job search model.

Proposition 4 *If the cumulative distribution function of the wage premium offer distribution, $F(\psi)$, is log concave, twice continuously differentiable, and its density function symmetric, then (i) an increase in ψ_0 has a monotonically decreasing effect on quantiles of the ψ distribution, and (ii) increases in referral network quality have an increasing effect on quantiles of the ψ distribution.*

Proof. See Appendix A. ■

In the search model, increasing ψ_0 affects outcomes by increasing the reservation offer that triggers mobility. Intuitively, increasing ψ_0 will have a larger impact on the distribution of acceptable offers close to the truncation point than those further away. The condition of Proposition 4, that the offer distribution is log concave with a symmetric density, is satisfied by the normal distribution, the uniform distribution, and the double exponential.

3 The Determination of Job Search Outcomes: Econometric Framework

To bring the model to the data, I must fully specify the offer function. Recall that observed earnings are denoted by y_{it} , and specify the earnings determination process so that

$$y_{it} = \gamma e_{i,t} p_{J(i,t)} \quad (6)$$

$$\ln y_{it} = \ln \gamma + \ln e_{i,t} + \ln p_{J(i,t)}. \quad (7)$$

$J(i, t) = j$ where j is the employer of i at time t . Human capital depends on observable time-varying inputs, X_{it} and observable and potentially unobservable correlates of ability, θ_i , so that

$$e_{it} = \exp(X_{it}\beta + \theta_i).$$

Since $\psi_j = \ln p_j$ the final expression for log earnings is:

$$\ln y_{it} = \alpha + X_{it}\beta + \theta_i + \psi_{J(i,t)} + \varepsilon_{it}. \quad (8)$$

The model allows arbitrary heterogeneity in the formal and informal offer distributions:

$$f_{it}(\psi) = ag_{it}(\psi) + (1 - a)h_{it}(\psi). \quad (9)$$

I assume that this heterogeneity is fully captured by observable worker characteristics, Z_i , the vector describing i 's referral network, w^i , and the log wage premia held by workers at the time of the transition. The latter quantity is the data analogue to $\Psi(t)$, denoted Ψ^t , where the i^{th} entry is $\psi_{J(i,t)}$, the log wage premium paid by employer $j = J(i, t)$. The offer distribution is:

$$f(\psi|Z_i, w^i, \Psi^t) = ag(\psi|Z_i, w^i, \Psi^t) + (1 - a)h(\psi|Z_i, w^i, \Psi^t). \quad (10)$$

It is a simple formality to express a realized offer, $\psi_{i,t}^*$, in terms of the means of the formal and informal distributions, g and h , and deviations from those means.

$$\psi_{i,t}^* = a (E_g(\psi|Z_i, w^i, \Psi^t) + \eta_{i,t}^g) + (1 - a)(E_h(\psi|Z_i, w^i, \Psi^t) + \eta_{i,t}^h) \quad (11)$$

$$= a E_g(\psi|Z_i, w^i, \Psi^t) + (1 - a) E_h(\psi|Z_i, w^i, \Psi^t) + \eta_{i,t}, \quad (12)$$

where $\eta_{i,t} = a\eta_{i,t}^g + (1 - a)\eta_{i,t}^h$. Restrictions on the sources of observable variation and the error processes clarify the essential identification problem and provide a template for implementing the model empirically. The model specifies the mean of the informal offer distribution in Equation (2), which is implemented empirically as:

$$E_h(\psi|w^i, \Psi^t) = (w^i)^T \Psi^t.$$

In the empirical work, W puts equal weight on all workers residing in the same Census block, and no weight elsewhere. That is, $(w^i)^T \Psi^t = \bar{\psi}_{b(i)t}$ where $b(i)$ indicates the block of residence for worker i , and $\bar{\psi}_{b(i)t}$ is the average wage premium in jobs held by workers at time t .

The conceptual separation of the formal and informal distributions implies that the expected offer from formal search is independent of who your neighbors are and where they work after conditioning on observable characteristics that might correlate with the productivity of formal job search. Likewise, the mean offer received through the referral network

does not depend on individual characteristics when conditioning on the quality of the referral network. Imposing these conditional moment restrictions yields

$$E(\psi_{it}^*|Z_i, w^i, \Psi^t) = a E_g(\psi_{it}^*|Z_i) + (1-a)(w^i)^T \Psi^t + a E(\eta_{i,t}^g|w^i, \Psi^t) + (1-a) E(\eta_{i,t}^h|Z_{it}). \quad (13)$$

I make a parametric assumption that the conditional mean of the formal offer distribution is linear in observable worker characteristics.

$$E_g(\psi_{it}^*|Z_i) = Z_i \tilde{\Pi}. \quad (14)$$

Accumulating all of the modeling assumptions, the offer function is given by

$$\psi_{i,t}^* = Z_i \Pi + \gamma \bar{\psi}_{b(i)t} + \eta_{i,t}, \quad (15)$$

where $\gamma = (1-a)$ and $\Pi = a\tilde{\Pi}$. The primary identification problem is embedded in the potential for correlation between the composite error term, $\eta_{i,t}$ and referral network quality, $\bar{\psi}_{b(i)t}$. I turn to these issues next.

3.1 Identification

There are four major challenges to the identification of the effect of referral network quality on earnings: reverse causality, reflection, sorting and correlated unobservables, and sample selection. The time sequencing of job mobility alleviates concerns about reverse causality and the reflection problem (Manski 1993). More crucially, the large sample size and fine detail of the residential address information mean that self-selection and correlated effects affecting formal job search outcomes for workers with the same referral network can be separately identified from referral network quality under mild assumptions. Sample selection occurs because job changes are only observed when the offer is sufficiently attractive. In what follows, I articulate the aspects of the data and assumptions required for identification.⁵

Reverse causality would suggest that worker's residential location choices are determined by their job quality. I use the time dimension of the LEHD data to measure the quality of jobs in a worker's referral network prior to changing jobs. Moreover, I focus on job changes

⁵The identification problems in this paper are related to the general problem of identifying social interactions documented by Manski (1993) and elaborated in William A. Brock & Steven N. Durlauf (2001). Lawrence E. Blume & Steven N. Durlauf (2005) provide a concise introduction to this literature with a useful discussion of the kinds of data and models that can be used to identify social interactions.

of workers who do not change residence, removing any concerns about reverse causality.

The reflection problem (Manski 1993) occurs when network quality is perfectly collinear with independent variables in the model. In cross-sectional data, this happens if block-level means of all of the independent variables, Z , are included. One solution is to exclude these ‘contextual effects’. However, the time sequencing of the data make it possible to include contextual effects. Since the quality of the network is predetermined at the point in time at which a worker makes a job-to-job move, breaking the reflection problem only requires excluding lagged values of block-level means of independent variables (lags of the contextual effects). As a practical matter, the results presented exclude all contextual effects. Including contemporaneous contextual effects does not substantially alter the results.⁶

The most substantial identification problem comes from the possibility that offers received through formal job search might be spatially correlated, either because of worker sorting, or because of features of the urban landscape that differentially facilitate or impede formal job search. In terms of Equation 15, identification requires the composite error term is uncorrelated with referral network quality. Formally,

$$a E(\eta_{i,t}^g | w^i, \Psi^t) + (1 - a) E(\eta_{i,t}^h | Z_{it}) = 0 \quad (16)$$

The assumption that $E(\eta_{i,t}^h | Z_i) = 0$ can be justified as follows. If the social interaction process has been properly specified, then the influence of one’s own characteristics on the arrival of offers through referral should already be included through w^i .

The assumption that formal search outcomes are uncorrelated with referral network quality, $E(\eta_{i,t}^g | w^i, \Psi^t) = 0$, may be too strict. Many economic processes generate spatially correlated outcomes from formal search. One particularly problematic process is the residential sorting of workers in terms of latent characteristics that affect job search. Also, neighborhoods differ in their proximity to jobs with particular characteristics so that workers in those areas have correlated search outcomes simply due to proximity.

Following the identification argument developed by Bayer, Ross & Topa (2008), I assume all economic processes generating spatial correlation in search outcomes are homogeneous within pre-defined reference groups of geographically contiguous Census blocks – neighborhoods. Referral network effects are then identified by block-level variation in network quality within neighborhoods. Formally, the identifying assumption is $E(\eta_{i,t}^g | w^i, \Psi^t, G(b(i))) = 0$,

⁶Timothy G. Conley & Christopher R. Udry (2010) also use the time sequencing of information transmission to identify the effect of social learning by farmers in Ghana about new agricultural practices.

where $G(b(i))$ denotes the neighborhood in which i resides.

Workers living in the same neighborhood confront minor variations in the quality of their referral. Within a neighborhood, spatially correlated factors affecting search are identical across workers. Therefore, the effects of referral networks are identified by controlling for unobserved effects driving search outcomes at the neighborhood level. The network effect is identified from the within-neighborhood variation in network quality.

The economic rationale underlying identification is as follows: Thinness of the residential real estate market means that workers can choose the neighborhood in which they live, but generally not a specific block. Similarly, employers may prefer to hire workers from a certain part of the city, but it is unlikely that they have strict preferences for workers from specific blocks within the same neighborhood. Finally, in urban areas, transportation access is similar for workers residing in the same neighborhood. Bayer, Ross & Topa (2008) show that this assumption is largely valid in their study of the Boston MSA. I present additional evidence in support of this assumption in Section 5.2.

For the empirical work, I use Census block groups as the reference group. Block groups are a convenient choice for several reasons. They are the lowest level of geography above the block for which the Census Bureau releases data, and are structured to collect relatively homogeneous, geographically contiguous blocks that do not cross tract boundaries.

The final specification of the offer function is:

$$\psi_i^* = Z_i\Pi + \gamma\bar{\psi}_{b(i)0} + \zeta_{G(b(i))} + \eta_i, \quad (17)$$

where $\bar{\psi}_{b(i)0}$ is the within-block average wage premium across all employed workers whose jobs were already in progress before the quarter in which i makes a transition, and that remained in progress in the quarter after. $\zeta_{G(b(i))}$ is a reference group effect where $G(b(i))$ denotes the reference group within which b belongs.

3.2 Sample Selection

With the offer function identified, the problem of sample selection remains. An already employed worker changes jobs only for a job with a higher premium. Hence, the observed wage premium distribution is truncated. I address this through a standard selection correction procedure.⁷

⁷The estimated wage premia show workers will move to jobs with lower premia. This feature of the data is related to the finding in Éva Nagypál (2005) that the rate of job-to-job transitions is not consistent with

The offer function is given in Equation 17. Mobility depends on the net utility difference between the new offer and their current premium, ψ_{io}^*

$$v_i^* = \psi_i^* - \psi_{io}^*,$$

where $\psi_{io}^* = \psi_{io} - \phi_i$ measures the worker's idiosyncratic preference for his current job. An indicator for whether the move occurs is $I_i = \mathbf{1}(v_i > 0)$. The conditional expectation of the observed wage premium distribution for job changers is:

$$E(\psi_i | \psi_{io}, Z_i, i) = Z_i \Pi + \gamma \bar{\psi}_{b(i)0} + \zeta_{G(b(i))} + E(\eta | I_i = 1, \psi_{io}, Z_i, i).$$

Unbiased estimation of the offer function requires a correction for selection by ψ_0 , which involves predicting $E(\eta | I_i = 1, \psi_{io}, Z_i, i) = E(\eta | \psi_i^* + \phi_i > \psi_{io})$. In the empirical work, I estimate the selection correction model under the theoretically justified restriction that ψ_0 is excluded from the offer function. I also estimate models that simply control for ψ_{io} through a linear term:

$$E(\psi_i | \psi_{io}, Z_i, i) = Z_i \Pi + \gamma \bar{\psi}_{b(i)0} + \zeta_{G(b(i))} + \beta \psi_{io}.$$

I find that the selection correction procedure has a very minor, statistically insignificant effect on the estimate of γ . In spite of the job ladder behavior of workers there is sufficient randomness in worker mobility so that truncation of the observed offer distribution induces very little bias.

4 Data and Estimation Procedure

I analyze the model on work histories drawn from the Longitudinal Employer Household Dynamics (LEHD) Program of the U.S. Census Bureau linked to data on workers' Census blocks of residence from the Statistical Administrative Records System (StARS). I follow a two-stage estimation procedure. In the first stage, estimates of employer-specific wage premia are generated by applying the Abowd-Kramarz-Margolis (AKM) log earnings decomposition (Abowd, Kramarz & Margolis 1999) to the complete universe of LEHD data. In the second stage, I focus on workers who make direct job-to-job moves. For each job, I merge the estimated employer-specific wage premium, ψ . For each Census block, in every quarter, I

the strong job ladder model. They are consistent with a modified on-the-job search model where workers have idiosyncratic non-pecuniary preferences for particular jobs.

measure referral network quality as the mean wage premium of workers on that block. All of the empirical work centers on identifying and estimating the effect of these network quality measures on the size of the wage premium a worker receives when making a direct job-to-job transition.

4.1 Data Sources

The LEHD data are built around the longitudinal employer-employee links represented by state Unemployment Insurance (UI) wage records which constitute the job frame. UI records cover approximately 98 percent of wage and salary payments in private sector non-farm jobs. The LEHD infrastructure makes use of the unique individual and employer identifiers from this system to track workers over time as they move from job to job, and to identify which workers share an employer. These data are augmented with demographic characteristics through administrative record and statistical links as well as to employer characteristics, including employer size, industry, and ownership type. For a complete description of these data, see John M. Abowd, Bryce E. Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer & Simon Woodcock (2009).

Data on place of residence come from the StARS database. StARS is a Census Bureau program originally designed to improve intercensal population estimates as well as refresh its household sampling frame. It incorporates administrative data from the IRS, HUD, Medicare, Indian Health Service and the Selective Service to update information on residential geography and other variables once a year. Geocodes of Census block precision are available for at least 90 percent of all LEHD workers who appear in one of the 30 sample MSAs during 2002-2003.

4.2 Stage 1: Estimation of Employer Wage Premia

As an empirical analogue to Equation 8, I use the Abowd-Kramarz-Margolis (AKM) decomposition:

$$\ln Y = X\beta + D\theta + F\psi + \varepsilon. \quad (18)$$

This model is estimated on the set of all LEHD work histories for workers aged 18-70. These data cover 30 states between 1990-2003, and include 660 million wage records for 190 million workers and 10 million employers. Y is a vector of annualized earnings on the dominant job, and ε is a statistical residual. D and F are design matrices of the worker and employer

effects. X is a matrix of time-varying controls consisting of a quartic in experience, year effects, and the exact within-year pattern of positive earnings. All of these measures are interacted with sex.⁸

4.3 The Estimation Sample

Table 1: Descriptive Statistics

Variable	Full Sample	Job Changers	Job Changers*
White	0.6572	0.6280	0.6220
Black	0.1151	0.1220	0.1205
Hispanic Origin	0.1167	0.1369	0.1400
Male	0.5098	0.4985	0.4979
Born in U.S.	0.8098	0.8098	0.8026
Age in 2002	40.5456	35.05848	34.9561
Any job transition	0.3116	1	1
Transition to new job	0.0351	1	1
Transition out of sample	0.2634	0	0
N	25,689,739	899,147	816,138

Summary statistics for a sample of workers with reported UI earnings in one of 30 large MSAs between 2002 and 2003. The sample is restricted to workers who did not move MSAs during 2002-2003, were at least 14 years of age in 2002, and had valid data for block of residence in 2002 and 2003. The summaries in column 3 are for job changers who lived on blocks where at least 10 other workers contribute data to compute the block-level average ψ .

The final analysis sample includes workers aged 18-70 who resided in one of 30 large Metropolitan Statistical Areas (MSAs) during 2002-2003 with information on the wage premia for any job they held in that two year period. A complete list of MSAs used is shown in Appendix B, Table 7. Table 1 presents descriptive statistics for the full sample as well as for the subsample of workers involved in a job-to-job transition. An observation in the sample is a worker from the LEHD infrastructure with positive earnings in at least one quarter of 2002-2003 who could be matched to a consistent block of residence in 2002-2003. For the

⁸This decomposition as applied to matched employer-employee data was first introduced by Abowd, Kramarz & Margolis (1999) as a means of correcting biases in the estimation of industry and other more aggregated types of wage premia. The estimates used in this paper were conducted as part of the Human Capital Estimates Project within LEHD according to the estimation procedure described in Abowd, Creecy & Kramarz (2002) and John M. Abowd, Paul Lengermann & Kevin L. McKinney (2003).

urban workers that are the focus of the paper, this selection rule has little effect: over 95 percent of workers have consistent data on block of residence in both years. I require the recorded block of residence be in the same MSA in both years; that is, this analysis is for the group of workers who do not move between MSAs during the sample period. The demographic characteristics of this sample of urban workers are consistent with other published sources on labor force characteristics. The sample of movers is marginally less white, more Hispanic, and substantially younger.

Model testing focuses on workers who make job-to-job transitions. The job history information for workers includes information on transitions between dominant jobs. A dominant job in a given year is the one on which the worker had the most earnings in that year.⁹ The second column in Table 1 presents statistics for the sub-sample of workers who experienced a transition from one employer to another employer without an observed intervening spell out of the sample. As expected, they are younger, but otherwise similar demographically to non-movers.

Just 3.5 percent of workers experience a transition between dominant employers. This is significantly lower than the reported rate of job-to-job transitions in other sources (Melissa Bjelland, Bruce Fallick, John Haltiwanger & Erika McInterfer 2008). However, many cases where a worker holds a short-term job between dominant employers will not be picked up. Bjelland et al. (2008), using a different definition of ‘main job’, find that roughly 31 percent of all transitions from jobs with tenure greater than one year are to jobs that last only 2-3 quarters. So, as many as 12 percent of workers who appear to make a transition out of sample are actually transitioning into temporary jobs. Thus, my sample of job-to-job transitions is properly interpreted as a sample of immediate transitions from one relatively long-term job to another. Given the objective of the study, this is the correct set of transitions to focus on. A worker who takes a stop-gap job in between long-term employers is perhaps more likely to have separated from the previous employer for other reasons or is adopting a different kind of search strategy.

⁹Dominant job to dominant job transitions occur at most once per year. 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.

5 Stylized Facts

In this section, I show that workers are more likely to move to jobs better than the one they currently have. Furthermore, the distributions of wage premia on destination jobs conditional on the premium in the origin job are strictly ranked in the sense of stochastic dominance. This is the first evidence that there is any mobility-related structure to ψ when estimated from the AKM decomposition.

I next show evidence of spatial correlation in the wage-premia held by workers. I compute non-parametric estimates of the spatial autocorrelation function for block- and tract-level averages of log earnings and the components from the AKM decomposition. These reveal positive spatial autocorrelation in estimated wage premia.

These results should allay concern about a potential objection to the model. The empirical method is a two-stage estimation procedure. The first stage consists of estimating the empirical wage premia, ψ , from the AKM decomposition, and the second stage the estimation of the realized offer distribution from data on workers making direct dominant job transitions. This procedure is consistent under the assumption that the errors in the earnings equation are not correlated with errors in the job mobility process. This exogenous mobility assumption is a feature of the extended on-the-job search model developed in this paper. It is nevertheless a strict one. That the estimated wage premia conform to the stylized predictions of the model means that the assumption may not be too strong.

5.1 Evidence of a job ladder

Table 2: Unconditional Transition Probabilities

Origin ψ -decile, ψ_0^d	1	2	3	4	5	6	7	8	9	10
$\Pr(\psi_1^d \geq \psi_0^d)$	1	0.94	0.85	0.79	0.73	0.73	0.69	0.61	0.58	0.59

Probability that the decile of the log wage premium on the destination job is greater than or equal to the decile of the origin job.

Table 2 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 probability is always strictly above 0.58, and significantly higher for workers starting from jobs with log wage premia in the lowest deciles. This evidence is consistent with the job ladder prediction of the basic search model developed above. Additional details on the nature of the job

ladder evidence in these data appear in the corresponding Figures 1 and 2. Figure 1 plots the cumulative frequency of destination wage premia for all job transitions stratified by decile of the origin job wage premium. The plots show decile-to-decile transitions, but the same results hold when looking at more detailed quantiles. First, note that there is a clear first-order stochastic dominance relationship among the conditional distributions. Workers starting from jobs with higher wage premia are more likely to move to jobs with better premia. Second, for each conditional distribution, the probability of moving to a job with the same or a higher premium is always strictly higher than the probability of moving to a job with a lower premium.

Figure 1: Cumulative probability of transition to each decile of the wage premium (ψ) distribution, by decile of origin

Figure 2: Probability of transition to each decile of the wage premium (ψ) distribution, by decile of origin

Figure 2 plots the transition matrix between deciles of the wage premium distribution. Each ribbon shows, for job changers whose initial wage premium fell in a certain decile, the fraction of transitions to jobs in each decile of the wage premium distribution. The saddle shape indicates that workers tend to move to jobs similar to, or better than the jobs they already have. The conditional densities are all peaked at the origin decile. This suggests a relationship between the current job and the offer distribution. Such a pattern will arise if workers tend to move among jobs within the same industry that all offer roughly the same wage premium. Such preference heterogeneity is consistent with the model as long as worker preferences are uncorrelated with earnings residuals.

5.2 Evidence of spatial correlation in ψ

By Proposition 2, the data must also exhibit correlation in wage premia between workers in the same referral network. To evaluate this implication of the model, I compute the spatial autocorrelation function for each of the components of earnings from the AKM decomposition, both as tract-level and block-level means. To my knowledge, these are the first estimates of their kind using matched employer-employee data for the U.S. Furthermore, the spatial autocorrelation estimates are the first of their kind to be estimated on earnings data at high spatial resolution.

5.2.1 Estimates of the spatial autocorrelation function

Figures 3 and 4 plot averages of the estimated spatial autocorrelation function in each MSA for tract- and block-level means of log earnings, the estimated person effect θ , the estimated wage premium ψ , and the residual from the AKM decomposition, ε . The discussion in this section closely follows Conley & Topa (2002) from which the method for this analysis was derived. The core statistical model is one in which random variable x_i is associated with a spatial coordinate, s_i . The spatial process generating the data is one in which the correlation between x_i and x_j depends only on the distance between s_i and s_j .

$$\text{Corr}(x_i, x_j) = f(\|s_i - s_j\|) \quad (19)$$

Figure 3: Spatial Autocorrelation Function: tract-level means

Figure 4: Spatial Autocorrelation Function: block-level means

This assumes that the correlations do not depend either on the precise locations in space of these random variables, nor the direction of the vector between them. I estimate the

spatial autocovariance function at distance δ , $f(\delta)$, non-parametrically by

$$\hat{f}(\delta) = \sum_{i=1}^N \sum_{i'=1}^N \phi \left[\frac{|\delta - A_{ii'}|}{\sigma} \right] (X_i - \bar{X}) (X_{i'} - \bar{X}) \quad (20)$$

where $A_{ii'}$ is the distance between i and i' . $\phi(\cdot)$ denotes the standard normal kernel. The spatial autocovariance function is estimated as the kernel-weighted average of the products of demeaned observations. To convert this to the spatial autocorrelation, one must divide the resulting estimate by relevant product of standard deviations. With the normal kernel, this is just the sample variance.

I implement this estimator for tract-level and block-level means of all earnings and the components of the AKM decomposition. I compute $\hat{f}(\delta)$ at distances from 0 to 5 miles at half-mile gridpoints. $A_{ii'}$ is measured as the great-circle distance between internal points of the block or tract. For the block-level estimates, the bandwidth parameter, σ , is set to 0.5. For the tract level estimates, it is set at 0.7. Since the computation scales in the square of the number of observations, for the block-level calculation some simplification is required. I randomly sample block pairs at the rate of 1/100. For a hypothetical MSA with 5,000 blocks, which would be a fairly small one for this study, this means the spatial autocorrelation function is estimated from approximately 125,000 unique data points. To satisfy the disclosure avoidance restrictions required to publish these results, each point in the figures represents the unweighted average of the estimated $\hat{f}(\delta)$ across 30 MSAs. There is some variation between the MSA-level estimates, but not enough to change the qualitative features of the plot. These plots are representative of most of the individual MSAs.

Both figures clearly show positive spatial autocorrelation in the tract- and block-level means of earnings, θ , and ψ . The main point to take away is the spatial correlation in estimated wage premia of workers in nearby blocks. This is consistent with the social interactions model of this paper. To be clear, there are also many other models that could generate these correlation patterns. The key challenge given the stylized fact is to identify the effect of social interactions in wage premia separately from other spatially correlated influences that could produce the result.

These results contain a wealth of interesting information beyond the analysis in this paper. I mention just two points briefly. First, the block-level estimates show no spatial autocorrelation in the block-level average residual. This is consistent with the key identifying assumption that workers are not systematically sorted within neighborhoods in terms

of productive characteristics. Whatever process puts people in a particular block is not correlated with the earnings residual. Second, these plots give evidence on the relationship between the spatial correlation in earnings and sorting on unobservables. The spatial correlation in earnings is mirrored almost exactly by the spatial correlation in estimated person effect, which captures the effect on earnings of unobserved and observed non-time-varying characteristics. These results confirm the findings of Combes, Duranton & Gobillon (2008) that sorting on observable and unobservable human capital characteristics explain a large amount of the spatial wage distribution in cities.

6 Estimation Results

Having established that the estimated wage premia, ψ , are consistent with the broad predictions of the model, I now use these data to estimate the influence of local referral network quality on job search outcomes. The main results are estimates of linear and quantile regression models of the form

$$\psi_i = Z_i\Pi + \beta\psi_{0i} + \gamma\bar{\psi}_{b(i)0} + \varphi\bar{\psi}_{G(b(i))} + \zeta_{G(b(i))} + \nu_i. \quad (21)$$

Primary interest lies with estimates of the parameter γ , which measures the effect of local interactions on job offers. Z_i is a vector of individual characteristics including age and its square, indicators for whether the worker is white or not, Hispanic or not, and male or not, as well as the estimated person effect, θ from the first-stage. The notation $b(i)$ indicates the Census block in which i resides. $\bar{\psi}_{b(i)0}$ is the within-block average wage premium across all employed workers whose jobs were already in progress before the quarter in which i makes a transition, and that remained in progress in the quarter after. ψ_{0i} is the wage premium of the employer from which i transitions. $\zeta_{G(b(i))}$ is a reference group effect where the notation $G(b(i))$ indicates the reference group of contiguous blocks containing $b(i)$. The reference group in these estimates is the Census block group.

The key result in Table 3 is in the contrast between the baseline specification, which does not control for reference group correlations in outcomes, and the two specifications that do. Inference in the conditional mean regressions is based on heteroscedasticity-corrected standard errors that have been clustered at the MSA level¹⁰. The baseline model presented

¹⁰This specification is very conservative. Under the empirical model, clustering at the county or tract level would be appropriate. As (Colin A. Cameron, Jonah B. Gelbach & Douglas L. Miller 2008) point out, asymptotic tests based on data with around 30 or fewer clusters may over-reject. Even with standard

Table 3: Offer Function Estimates

Premium on next job, ψ	Baseline	Bl.Grp. Means	Tract Means	Bl.Grp. Controls
Initial ψ : ψ_0 (β)	0.46* (.008)	0.45* (.008)	0.45* (.008)	0.45* (.008)
Avg. ψ in block: $\bar{\psi}_{block}$ (γ)	0.33* (.016)	0.10* (.011)	0.10* (.011)	0.10* (0.012)
Avg. ψ in block group: $\bar{\psi}_{bg}$ (ϕ)		0.34* (.024)	0.20* (.017)	
Avg. ψ in tract: $\bar{\psi}_{tract}$			0.15* (0.025)	
white	0.001 (.001)	0.001 (.002)	0.002 (.002)	-0.002 (.001)
Hispanic Origin	-0.02* (0.005)	-0.01 (0.005)	-0.014* (.005)	-0.017* (.004)
male	0.03* (0.004)	0.03* (0.004)	0.04* (.004)	0.04* (0.004)
age in 2002	0.01* (.001)	0.01* (.001)	0.01* (.001)	0.01* (0.001)
Square of age in 2002	-0.00* (.000)	-0.00* (.000)	-0.00* (.000)	-0.00* (.000)
Born in U.S.	0.00* (.003)	0.01 (.002)	0.01 (.002)	0.01* (.003)
θ from wage eqn.	-0.00 (.009)	-0.01 (.010)	-0.01 (.010)	-0.01 (.010)
block group controls	no	no	no	yes
N	815,899	815,889	815,889	815,899
R^2	0.3149	0.3175	0.3176	0.2711

Estimates of the log wage premium, ψ , for job changers. Standard errors are clustered on 30 MSAs. * entries have p-value < 0.025 .

in the first column of Table 3 shows the raw correlation between $\bar{\psi}_{b(i)0}$ and ψ_i , the premium on the job to which i makes a transition, controlling for the premium on the origin job and observable characteristics that may influence formal search. The point estimate on γ in the baseline model of 0.33 is on the same order of magnitude as the point estimate of β . In this specification, though, γ is absorbing any unobserved correlates of formal job search that aren't included in the model.

The social interaction parameter, γ , is identified in the model with reference group controls presented in the fourth column of Table 3. The point estimate is $\hat{\gamma} = 0.10 \pm 0.01$, and is statistically significant. To interpret the point estimate in terms of the model, this means that 10 percent of job offers arrive through referrals. This is in line with the analysis in Ioannides & Loury (2004) of referral use by workers in the Panel Study of Income Dynamics, which shows that 8.5 percent of employed workers report using referrals to search for work.

Observable demographic characteristics explain relatively little of the variation in the data. The signs on the coefficients associated with demographic and human capital characteristics have the same sign as would be expected in a Mincerian wage regression, but with only marginal significance in most cases. All of these estimates are an order of magnitude smaller than the point estimates of the social interaction parameter γ , and the effect associated with the initial job type. These findings are consistent with the arrival of information about wage premia being only weakly related to individual ability, which is in turn consistent with the notion that they are non-economic rents associated with information frictions in the labor market.

The other columns in Table 3 present alternative estimates of γ based on a contrast between $\bar{\psi}_{b(i)0}$ and $\bar{\psi}_{G(b(i))}$. The point estimates are nearly identical, and I conclude that the coefficient on the group-level average log wage premium has absorbed all of the unobserved correlation in outcomes. Because of its computational simplicity, I use this contrast to estimate the selection correction model as well as the quantile regressions.

6.1 Robustness Checks

As a check for robustness of my estimate of the social interaction parameter, γ , I estimate the full econometric model of job-to-job mobility described in Section 3. It allows for sample selection driven by the fact that only workers who receive sufficiently attractive offers change jobs. The attractiveness of a job offer depends on the wage premium of one's current job, ψ_0 .

errors clustered on 30 MSAs, the point estimates of interest are significantly different from zero in all cases. Clustering on county or tract does not alter the qualitative results.

Following the theoretical model, ψ_0 is excluded from the offer function, but does appear in the selection equation. I estimate the selection correction model using data on all employed workers at risk to change jobs in 2002, Quarter 4. The results are presented in Table 4.

Table 4: Selection Correction Model Estimates

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

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. * entries have p-value < 0.025. Both models include all controls from Table 3. ρ is the estimated correlation between the errors in the selection equation and the offer function.

As expected, in the selection equation, the log wage premium on the worker's initial job ψ_0 has a strong negative effect on the probability of a job-to-job move. Workers living on blocks with better than average network quality for their neighborhood are also more likely to make a job-to-job transition. The point estimate on the social interaction parameter in the selection correction model is $\hat{\gamma} = 0.11 \pm 0.02$.

To check whether my results are sensitive to heterogeneity in the effect of referrals, I estimate the model with block group controls and allowing for the use of referrals to be different for native workers and non-native workers. One objection to my research design is that the local referral interactions I model are most relevant to certain kinds of jobs, and are more likely to be used by certain groups of workers. Previous research indicates that the use and efficiency of referrals differ considerably by demographic group. Furthermore, the kinds of jobs that are shared among residential neighbors are more likely to be jobs

Table 5: Heterogeneous Referral Effects

Premium on next job, ψ		
Mean premium in block: $\bar{\psi}_{block}$	0.10*	0.17*
	(.012)	(.023)
Born in U.S. $\times \bar{\psi}_{block}$		-0.09*
		(.021)
block group controls	yes	yes
N	815,899	815,889
R^2	0.2711	0.2712

Standard errors are clustered on 30 MSAs. * entries have p-value < 0.025. Models include controls from Table 3.

with relatively little specific skill requirements. The results, reported in Table 5, show that non-native workers have $\hat{\gamma} = 0.17 \pm 0.02$, which is a 70 percent increase over the pooled estimate. This finding is consistent with other work finding that immigrants are more likely to find jobs by referral than their native counterparts.

6.2 Distributional Effects

The job search model has distributional implications as well, which are captured in Proposition 4. Specifically, increases in ψ_0 compress the observed job quality distribution from the left, while increases in network quality, $\bar{\psi}_{b(i)0}$, stretch the observed job quality distribution from the right.

Table 6 presents estimates of conditional quantile specifications for the 10th, 25th, 50th, 75th and 90th percentiles of the destination- ψ distribution for job changers. The key result is the pattern in the coefficient estimates associated with $\bar{\psi}_{b(i)0}$, $\bar{\psi}_{G(b(i))0}$ and ψ_0 .

Let $\beta(q)$ be the coefficient associated with ψ_0 in the conditional regression of the q th quantile, and define $\gamma(q)$ as the effect of network quality on the q th quantile. Proposition 4 predict that for $q < q'$, $\beta(q) > \beta(q')$ and $\gamma(q) < \gamma(q')$. In the search model, increasing ψ_0 affects outcomes by increasing the reservation offer that triggers mobility. The estimates clearly show $\beta(0.1) > \beta(0.5) > \beta(0.9)$. However, $\gamma(0.1) > \gamma(0.25) = \gamma(0.5) < \gamma(0.75) < \gamma(0.9)$. The predicted pattern appears if one considers the estimates associated with the neighborhood (block-group) level mean, $\bar{\psi}_{G(b(i))0}$. It is also impossible to reject the hypothesis that $\gamma(0.1) = \gamma(0.25) = \gamma(0.5)$, so the data weakly support the proposition. In short, the data

Table 6: Quantile Regression Estimates

Premium on next job, ψ	Quantile				
	$q(0.1)$	$q(0.25)$	$q(0.5)$	$q(0.75)$	$q(0.9)$
Initial premium: ψ_0 (β)	0.59*	0.59*	0.50*	0.39*	0.32*
	(.002)	(.001)	(.001)	(.001)	(.002)
Mean premium in block: $\bar{\psi}_{block}$ (γ)	0.09*	0.08*	0.08*	0.11*	0.13*
	(.010)	(.006)	(.005)	(.005)	(.007)
Mean prem. in block group: $\bar{\psi}_{bg}$ (ϕ)	0.17*	0.22*	0.30*	0.40	0.57*
	(.012)	(.007)	(.006)	(.007)	(.010)
N	815,899	815,889	815,889	815,899	815,899
pseudo- R^2	0.1908	0.2187	0.2002	0.2711	0.1741

Quantile regression model estimates of the log wage premium, ψ , for job changers. * entries have p-value < 0.025. Models include all controls from Table 3.

support all of the main predictions of the job search model.

7 Related Literature

The research in this paper is most closely related to theoretical and empirical literature on referral networks and their effects on labor market outcomes. Dale Mortensen & Tara Vishwanath (1994) develop a general equilibrium labor market model with on-the-job search in which workers can either sample an offer from the formal offer distribution, or sample directly from the distribution of realized job offers. My model extends this to allow for heterogeneity in the underlying referral network. I thus combine their approach with that of Antoni Calvo-Armengol & Matthew O Jackson (2004) which allows for a more general specification of transmission of job information, but in a partial equilibrium setting. François Fontaine (2007) and Pierre Cahuc & François Fontaine (2002) also study the transmission of job offers through referral networks in a general equilibrium matching model to study the implications of referral networks for macroeconomic efficiency.

Empirical work related to these models has primarily studied employment outcomes. Giorgio Topa (2001) and Timothy G. Conley & Giorgio Topa (2007) use a formal job search model to estimate local interactions in unemployment. Bayer, Ross & Topa (2008) and Judith K. Hellerstein, Melissa McInerney & David Neumark (2008) both use cross-sectional employer-employee matched data to show local social networks affect where people find jobs.

My paper advances on these studies by considering the direct effect of referral networks on earnings, while formally combining the search model with a clean identification strategy.

8 Conclusion

I find evidence of local social interactions in the transmission of information about employer-specific wage premia. Workers whose neighbors have jobs paying higher wage premia are more likely to experience a job transition, and when they do, are more likely to move to a job with a better premium. I apply and extend the identification strategy of Bayer, Ross & Topa (2008), using variation in local network quality among workers who reside in the same Census block group. The best estimate from the model indicates that 10 percent of a worker's job offers come from referrals. This is consistent with figures reported by other authors on the extent of referral use. These are the first results on direct local interactions in earnings outcomes in the context of a job search model. They complement existing work on local interactions in employment status and hours of work (Topa 2001, Bruce A. Weinberg, Patricia B. Reagan & Jeffrey J. Yankow 2004).

To motivate and structure the empirical work, I construct a model of job search augmented to allow for transmission of job information through referral networks. I show that the distribution of wage premia received by job movers responds to variation in referral network quality in a manner consistent with this model. The model also predicts that workers who switch jobs tend to move into jobs with higher wage premia than their current job, and that there will be correlation in the wage premia held by workers who are socially connected to each other. I show that the log wage premia estimated from matched employer-employee data exhibit both of these properties. This is the first evidence of mobility-related structure in employer wage premia estimated from matched employer-employee data. I also estimate the spatial correlation structure of earnings, employer-specific wage premia, and worker ability. The block-level analysis in this paper is among the most geographically detailed studies of sorting by earnings, human capital, and employer characteristics in U.S. cities and is relevant to those interested in residential sorting by earnings, human capital characteristics, and employer characteristics in urban labor markets.

My findings add to a growing body of evidence on the importance of social interactions for job search and labor market outcomes. The data support a model in which referral networks facilitate the exchange of information about particularly attractive job opportunities. This has implications for the distribution of earnings, and also for the efficiency of labor market

matching. The details of these distributional and efficiency impacts are important areas for future research.

References

- 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*. Vol. 68 of *Studies in Income and Wealth*, , ed. Timothy Dunne, J. Bradford Jensen and Mark J. Roberts, Chapter 5, 149–234. The University of Chicago Press.
- 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.” LEHD U.S. Census Bureau TP-2002-09.
- Abowd, John M., Robert H Creecy, and Francis Kramarz.** 2002. “Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data.” LEHD, U.S. Census Bureau TP-2002-06.
- Bagnoli, Mark, and Ted Bergstrom.** 2005. “Log-concave probability and its applications.” *Economic Theory*, 26(2): 445–469.
- 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.
- Bewley, Truman F.** 1999. *Why Wages Don’t Fall During a Recession*. Harvard University Press.
- Bjelland, Melissa, Bruce Fallick, John Haltiwanger, and Erika McInterfer.** 2008. “Employer-to-employer flows in the United States: Estimates using linked employer-employee data.” National Bureau of Economic Research Working Paper 13867.
- Blume, Lawrence E., and Steven N. Durlauf.** 2005. “Identifying Social Interactions: A Review.” July 22, 2005.

- Brock, William A., and Steven N. Durlauf.** 2001. "Interactions-based models." In *Handbook of Econometrics*. Elsevier.
- Cahuc, Pierre, and Francois Fontaine.** 2002. "On the Efficiency of Job Search with Social Networks." IZA 583.
- Calvo-Armengol, Antoni, and Matthew O Jackson.** 2004. "The Effects of Social Networks on Employment and Inequality." *American Economic Review*, 94: 426–454.
- 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.
- Cameron, Colin A., Jonah B. Gelbach, and Douglas L. Miller.** 2008. "Bootstrap-based improvements for inference with clustered errors." *Review of Economics and Statistics*, 90: 414–427.
- 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 Christopher R. Udry.** 2010. "Learning About a New Technology: Pineapple in Ghana." *American Economic Review*, 100: 35–69.
- Conley, Timothy G, and Giorgio Topa.** 2002. "Socio-Economic Distance and Spatial Patterns in Unemployment." *Journal of Applied Econometrics*, 17: 303–327.
- Conley, Timothy G., and Giorgio Topa.** 2007. "Estimating Dynamic Local Interactions Models." *Journal of Econometrics*, 140: 282–303.
- Flinn, Christopher J., and James J. Heckman.** 1983. "Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?" *Journal of Labor Economics*, 1(1): 28–42.
- Fontaine, François.** 2007. "A job search model with social networks: The better match hypothesis." September 19, 2007.
- Hellerstein, Judith K., Melissa McInerney, and David Neumark.** 2008. "Measuring the Importance of Labor Market Networks." National Bureau of Economic Research.

- Ioannides, Yannis M., and Linda Datcher Loury.** 2004. “Job Information Networks, Neighborhood Effects, and Inequality.” *Journal of Economic Literature*, 42: 1056–1093. Priority: normal.
- Manski, Charles F.** 1993. “Identification of Endogenous Social Effects: The Reflection Problem.” *Review of Economic Studies*, 60: 531–542.
- Mortensen, Dale, and Tara Vishwanath.** 1994. “Personal contacts and earnings: It is who you know!” *Labour Economics*, 1: 187–201.
- Nagypál, Éva.** 2005. “On the extent of job-to-job transitions.” September 2005.
- Neal, Derek.** 1999. “The Complexity of Job Mobility among Young Men.” *Journal of Labor Economics*, 17(2): 237–61.
- Postel-Vinay, Fabien, and Jean-Marc Robin.** 2002. “Wage Dispersion with Worker and Employer Heterogeneity.” *Econometrica*, 70: 2295–350.
- Rogerson, Richard, Robert Shimer, and Randall Wright.** 2005. “Search-Theoretic Models of the Labor Market: A Survey.” *Journal of Economic Literature*, 43(4): 959–988.
- Topa, Giorgio.** 2001. “Social Interactions, Local Spillovers and Unemployment.” *Review of Economic Studies*, 68: 261–295.
- Wahba, Jackline, and Yves Zenou.** 2005. “Density, social networks and job search methods: Theory and application to Egypt.” *Journal of Development Economics*, 78: 443–473.
- Weinberg, Bruce A., Patricia B. Reagan, and Jeffrey J. Yankow.** 2004. “Do Neighborhoods Affect Hours Worked? Evidence from Longitudinal Data.” *Journal of Labor Economics*, 22: 891–924.

A Model Details

A.1 Proof of Proposition 1

Proof. Since workers are wealth maximizers, and the evolution of portable skills e_{it} is unrelated to $p_{J(i,t)}$, we can model search over wage premia, p , and ignore e . Since workers

are myopic about the evolution of the referral network, the decision environment is stationary so the value of holding a job with wage premium p is given by the Bellman equation

$$rV(p) = p + \lambda_1 \int_0^\infty [\max\{V(p'), V(p)\} - V(p)] d\tilde{F}(p') + \delta [U - V(p)], \quad (22)$$

where r is the discount rate, U is the value of becoming unemployed, and $\tilde{F}(p)$ is the cumulative distribution of offers, p , appropriately transformed from $F(\psi)$. The myopia assumption means workers behave as if $\tilde{F}$ is fixed. The corresponding Bellman equation for the value of being unemployed is

$$rU = p_b + \lambda_0 \int_0^\infty [\max\{V(p') - U\}] d\tilde{F}(p'). \quad (23)$$

It is clear that $V(p)$ is increasing in p and that U is constant. Therefore, employed workers will adopt a strategy where they exit unemployment whenever $p > p_R$ for some constant p_R and switch jobs whenever they receive an offer with $p' > p$. The reservation premium, p_R will satisfy

$$p_R = p_b + (\lambda_0 - \lambda_1) \int_{p_R}^\infty \left[\frac{1 - \tilde{F}(p)}{r + \delta + \lambda_1 [1 - \tilde{F}(p)]} dp \right]. \quad (24)$$

■

A.2 Proof of Proposition 3

Proof. For the proof, I suppress dependence on Z, W and Ψ . Stars added to a distribution indicate that they are the truncated versions of the unstarred distribution. For instance, $g^*(\psi) = g(\psi|\psi > \psi_0)$. The truncated mean is a mixture:

$$\begin{aligned} E_{f^*}(\psi) &= a^* E_{g^*}(\psi) + (1 - a^*) E_{h^*}(\psi) \\ &= a^* \mu_{g^*} + (1 - a^*) \mu_{h^*}, \end{aligned}$$

where $a^* = \frac{a(1-G(\psi_0))}{1-aG(\psi_0)-(1-a)H(\psi_0)}$. Taking derivatives,

$$\frac{\partial \mu_{f^*}}{\partial \mu_h} = \frac{\partial a^*}{\partial \mu_h} \mu_{g^*} + a^* \frac{\partial \mu_{g^*}}{\partial \mu_h} - \frac{\partial a^*}{\partial \mu_h} \mu_{h^*} + (1 - a^*) \frac{\partial \mu_{h^*}}{\partial \mu_h}.$$

Eliminating $\frac{\partial \mu_{g^*}}{\partial \mu_h}$ and rearranging:

$$\frac{\partial \mu_{f^*}}{\partial \mu_h} = \frac{\partial a^*}{\partial \mu_h} (\mu_{g^*} - \mu_{h^*}) + (1 - a^*) \frac{\partial \mu_{h^*}}{\partial \mu_h}.$$

Log concavity of h ensures $\frac{\partial \mu_{h^*}}{\partial \mu_h} > 0$. Furthermore it is clear that $\frac{\partial a^*}{\partial \mu_h} > 0$. Thus, as long as

$$|(\mu_{g^*} - \mu_{h^*})| < \frac{(1 - a^*) \frac{\partial \mu_{h^*}}{\partial \mu_h}}{\frac{\partial a^*}{\partial \mu_h}}$$

, we have $\frac{\partial \mu_{f^*}}{\partial \mu_h} > 0$ ■

A.3 Proof of Proposition 4

Proof. The q th quantile of the distribution of observed offers, ψ^q is defined implicitly by

$$\int_{-\infty}^{\psi^q} f(\psi | \psi > \psi_0) d\psi = \int_{\psi_0}^{\psi^q} \frac{f(\psi)}{1 - F(\psi_0)} d\psi = q.$$

Which gives

$$F(\psi^q) = q + (1 - q) F(\psi_0).$$

Renormalize the offer distribution in terms of deviations from its mean, μ :

$$F(\psi^q - \mu) = q + (1 - q) F(\psi_0 - \mu).$$

First, consider the effect of a shift in the initial offer on the q th quantile of observed jobs

$$F'(\psi^q - \mu) \frac{\partial \psi^q}{\partial \psi_0} = (1 - q) F'(\psi_0 - \mu).$$

This establishes that a shift in the initial offer is expected to have a positive effect on all quantiles of the observed offer distribution. The goal is to assess how the magnitude of this effect varies with respect to the quantile q . Hence, we want to establish the sign of

$$\frac{\partial^2 \psi^q}{\partial \psi_0 \partial q}.$$

Note

$$\frac{\partial \psi^q}{\partial q} = \frac{1 - F(\psi_0 - \mu)}{F'(\psi^q - \mu)}.$$

Differentiating this with respect to ψ_0 ,

$$\frac{\partial^2 \psi^q}{\partial \psi_0 \partial q} = \frac{-F'(\tilde{\psi}_0) - F''(\tilde{\psi}^q) \frac{\partial \psi^q}{\partial q} \frac{\partial \psi^q}{\partial \psi_0}}{F'(\tilde{\psi}^q)},$$

where I have replaced $\psi - \mu = \tilde{\psi}$ for simplicity.

$$= -\frac{F'(\tilde{\psi}_0)}{F'(\tilde{\psi}^q)} - \frac{F''(\tilde{\psi}^q)(1-q)(1-F(\tilde{\psi}_0))}{F'(\tilde{\psi}^q)^3}.$$

When $F''(\tilde{\psi}^q) > 0$, this is negative. Suppose $F''(\tilde{\psi}^q) < 0$. I will show that $\frac{\partial^2 \psi^q}{\partial \psi_0 \partial q} > 0$ is impossible as long as

$$\frac{F'(\tilde{\psi}^q)^2}{|F''(\tilde{\psi}^q)|} \geq 1 - F(\tilde{\psi}^q).$$

This condition simply places limits on the amount of curvature in the density function. Note that in the case described in the statement of the proposition, where F' is a symmetric density function and F is log concave, we have

$$\frac{F'(\psi^q - \mu)^2}{|F''(\psi^q - \mu)|} = \frac{F'(\mu - \psi^q)^2}{F''(\mu - \psi^q)} \geq F(\mu - \psi^q) = 1 - F(\psi^q - \mu),$$

where the first and last equalities follow by symmetry of the density function, the inequality follows from log concavity.¹¹

Continuing with the proof, suppose $F''(\tilde{\psi}^q) < 0$ and $\frac{\partial^2 \psi^q}{\partial \psi_0 \partial q} > 0$. Then

$$\frac{-F''(\tilde{\psi}^q)(1-q)(1-F(\tilde{\psi}_0))}{F'(\tilde{\psi}^q)^2} > 1,$$

that is,

$$\frac{F'(\tilde{\psi}^q)^2}{|F''(\tilde{\psi}^q)|} < (1-q)(1-F(\tilde{\psi}_0)),$$

¹¹For details on log concave functions and their application to search models, see Mark Bagnoli & Ted Bergstrom (2005) and Christopher J. Flinn & James J. Heckman (1983).

which by the assumption above implies

$$\begin{aligned}
1 - F(\tilde{\psi}^q) &< (1 - q)(1 - F(\tilde{\psi}_0)) \\
1 - \left(q + (1 - q) F(\tilde{\psi}_0)\right) &< (1 - q)(1 - F(\tilde{\psi}_0)) \\
(1 - q)(1 - F(\tilde{\psi}_0)) &< (1 - q)(1 - F(\tilde{\psi}_0)),
\end{aligned}$$

a contradiction. It follows that $\frac{\partial^2 \psi^q}{\partial \psi_0 \partial q} < 0$. The proof that $\frac{\partial^2 \psi^q}{\partial \mu \partial q} > 0$ is analogous. ■

B MSAs Used

Table 7: 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	

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