

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

January 2013

Abstract

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

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 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.

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 (Bewley 1999). While the use of job information networks is well-documented, the reasons they are so widely used by workers and firms remain poorly understood.¹ One role for referral networks in job search is to help workers locate information about attractive job opportunities, and in particular, to locate higher-paying jobs. In a market with search frictions, two workers can receive different pay simply because they work in different firms (Abowd et al. 1999; Mortensen 2003). If workers share information about these pay differentials with their friends and neighbors, then who they know can affect how much they are paid.

The idea that workers use social networks to find information about relatively attractive job opportunities is not new, but little empirical work assesses how workers share information about earnings opportunities. This is because a proper analysis requires information on job information networks and the ability to measure the part of earnings associated with working for a specific employer. To make progress, I develop and empirically implement a model of job search in which employers are distinguished by idiosyncratic wage differentials and workers use social networks to find better-paying jobs. The model yields testable implications for the observed distribution of earnings 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-

¹A growing number of studies confirm the survey evidence (Ioannides and Loury 2004, see) that people help their acquaintances find work, either through direct referral to a job with their employer (Bayer et al. 2008; Hellerstein et al. 2011; Nordström Skans and Kramarz 2011) or possibly indirectly (Topa 2001; Conley and Topa 2007; Laschever 2009; Cingano and Rosolia 2012).

employee matched data for the United States from the Longitudinal Employer-Household Dynamics (LEHD) Program of the U.S. Census Bureau. These data allow me to measure the earnings differentials associated with employment in a specific firm, net of individual characteristics (Abowd et al. 2002). I then match the firm, worker, and residual components of earnings to the exact residential block for workers who lived in one of 30 large Metropolitan Statistical Areas (MSAs) in 2002–2003. Guided by the model, I measure the quality of the local job information network using information on the firm-specific wage premia held by workers living nearby.

As predicted by the model, 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 are more likely to change jobs, and conditional on changing, accept offers with higher wage premia when workers in their local referral networks are earning higher premia. Local network quality has distributional effects as well that are predicted by the job search model. A better network improves upper-tail outcomes for job changers more than lower-tail outcomes, so referral networks result in an increased dispersion of realized wage premia in addition to an increased mean. Taken at face value, my estimates indicate that workers receive about 8 to 9 percent of job offers through social interactions. This is at the lower end of existing estimates, perhaps because, by focusing on local interactions, this research misses other dimensions of interaction (Ioannides and Loury 2004).

Relative to existing literature, this paper develops a ‘pivot’ from documenting the existence of social interactions in job search to uncovering the mechanisms through which social interactions operate.² The central contribution comes from connecting a specific model of job

²A body of research has developed near the same pivot. Dustmann et al. (2011) use matched employer-employee data to estimate a model in the spirit of Simon and Warner (1992) where referrals resolve uncertainty about worker or match quality. In their study, identification is based on changes in the share of workers within a firm from the same immigration group. Goel and Lang (2009) use survey evidence on network strength among Canadian immigrants to identify social interaction effects, also in the context of a job search

information networks to tools developed for identification of local interaction effects in highly detailed matched employer-employee data. I show that the data are consistent with a model in which workers use local information to find better paying jobs. Because of its implications for the efficiency and equity of labor market outcomes, this particular role for job information networks has been the focus of many theoretical models (Mortensen and Vishwanath 1994; Cahuc and Fontaine 2002; Calvo-Armengol and Zenou 2005; Calvo-Armengol and Jackson 2007; Fontaine 2008). Topa (2001) and Conley and Topa (2007) bring explicit models of information sharing in job search models to the data, but because they use aggregate data on employment, so the interpretation of their results depends on modeling assumptions. Bayer et al. (2008), Hellerstein et al. (2011), Oyer and Schaefer (2011), and Cingano and Rosolia (2012) establish convincing research designs exploiting the detail of population-level employer-employee matched data to provide evidence that referral effects exist, but are less revealing regarding their underlying purpose. The present paper moves forward by mapping a model of local social interactions in job search to a reduced-form empirical framework with a clear source of identification. Not only does the model provide testable implications for the empirical analysis, it provides a clear interpretive framework for the results. Objections to the theoretical model can be articulated in terms of potential modifications to the empirical model, and vice-versa.

The model raises identification issues common to studies of endogenous peer effects (Manski 1993; Moffitt 2000; Blume et al. 2011). Adapting the research design of Bayer et al. (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, distinguishing neighborhood quality from network quality. My results therefore rely on the identifying assumption that workers do not sort within neighborhoods on the basis of unob-

model.

served characteristics that influence job search directly.³ The assumption is untestable, but through a formal sensitivity analysis, I show that sorting on unobservables must be much stronger than sorting on observables to explain my results. My data reveal that there is very little sorting within neighborhoods on the basis of demographic characteristics, which echoes the similar finding in Bayer et al. (2008) for the city of Boston. Furthermore, there is no sorting on the residual component of earnings. These findings suggest that my results are robust to relaxation of the identifying assumption.

Moreover, my data allow me to make advances on the original Bayer et al. (2008) design. Their data were cross-sectional, and did not reveal the identities of employers, but rather the location of employment. I observe actual employers, but more importantly, because of the longitudinal nature of my data, I can observe who held a job with an employer first. Thus, for workers who change jobs, I can see whether they begin employment with a neighbors employer. For each such potential referral relationship, I can determine which worker was referred and which worker would have been the referrer. I use this information to ascertain how much of the observed social interaction effect in earnings could possibly be due to direct referral, and how much must come through other channels. Along the way, I am able to make an additional contribution by verifying that the results in Bayer et al. (2008) hold in this rather different setting. Furthermore, I use these data to provide auxiliary evidence for the modeling assumption that workers use networks to search for jobs offering higher pay for the same set of skills.

In the empirical work, I assume that workers interact with their residential neighbors in searching for better jobs (Case and Katz 1991; Bayer et al. 2008; Damm 2009).⁴ This

³Note that this empirical design does allow for arbitrary sorting of workers across neighborhoods on the basis of search-relevant characteristics.

⁴Some argue that the Internet has eroded the importance of neighborhoods in social life. However, Mok et al. (2009) find that the emergence of the Internet has actually played a key role in strengthening local social interactions, and is a complement, rather than a substitute for more traditional relationships.

suggests a natural specification check. My results should be stronger among workers for whom this is a good proxy. I find that the magnitude of the local interaction effects are almost twice as strong for non-native workers as for natives, with similar results for young workers, and the opposite pattern for older workers, consistent with other sources. I also explore the underlying social interaction mechanism by excluding direct referral relationships from the data. By throwing out workers whose new job is with a neighbor’s employer, I find that direct referrals can explain some, but not all, of my results. Direct referrals are usually associated with firm-level screening (Kugler 2003; Dustmann et al. 2011), whereas the mechanism proposed in this paper reflects worker opportunism (Bandiera et al. 2009; Beaman and Magruder 2011). Information exchange can occur outside referrals if workers with better jobs pass on outside offers to their peers (Calvo-Armengol and Jackson 2007) or share information about how to find ‘good’ jobs among their neighbors.⁵

Since the employer wage premia are derived from a statistical decomposition, their interpretation is open to question.⁶ I document stylized features of the estimated wage premia that are broadly consistent with my augmented job search model. First, job quality is spatially correlated at the level of the Census block, but the residual component of earnings is not (Figure 1). Second, 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). These results support the subsequent analysis, and are relevant, more broadly, to those who might seek to use these or similar data to study job search or local interactions in other contexts.

⁵Dolnick (2011) describes this sort of behavior among immigrants in New York City. Immigrant workers share information about small-scale temp agencies in Chinatown that provide relatively good job opportunities.

⁶The structural interpretation of the employer and worker effects in the first stage model relies on the assumption, explicit in the theoretical model, that job mobility is exogenous to the wage residual. I discuss the implications of this assumption for my analysis at length in Section 4.3.

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.⁷

The model delivers four major predictions to be verified in the data. Proposition 1 predicts that the average outcomes of job search are better for workers with higher quality referral networks. A corollary of this proposition is that workers will also be more likely to change jobs when the members of their social network hold higher-paying jobs. Proposition 2 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.

The model also makes more basic predictions on the temporal and spatial structure of estimated employer-specific log-wage premia. Proposition 3 predicts a 'job-ladder'. Workers should move to better paying jobs on average. Proposition 4 predicts that the quality of jobs held by workers in the same referral network are positively correlated. In the empirical work, I estimate the employer-specific wage premia from a first-stage decomposition of log earnings that could be consistent with many different models. I assess the first-stage estimates against

⁷This 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 et al. 1999; 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 (Rogerson et al. 2005).

these predictions as an initial check of model validity.

2.1 Model Setup

Let 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 .⁸

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

⁸The wage function given here could arise in a matching model with worker and employer heterogeneity in production with surplus sharing when there is no wage renegotiation (Postel-Vinay and Robin 2002).

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 is the object of primary interest in the empirical analysis. It measures the strength of social interactions relative to formal channels in delivering new job offers. 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.

Note that in the model, job mobility only depends on the random arrival of job offers and their attractiveness. In the empirical work, the first stage of the estimation procedure extracts estimates of the employer log-wage premia, ψ , under the assumption that the mobility of workers across employers is exogenous to the earnings residual. This exogenous mobility assumption is common, but controversial, and I will discuss it at length when describing the empirical work.

2.2 Specification of the Referral Distribution, $h(\psi; i, t)$

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 $\Psi(t)$ is the $I \times 1$ vector of the wage premium earned by each worker and w^i is the i^{th} column of an $I \times I$ stochastic matrix, W . W_{ji} measures the probability that job information received by i originated with worker j .⁹ So the mean of the referral distribution is just the mean wage premium in one’s reference group, weighted by social proximity.

This model for the social transmission of job information is equivalent to a contagion process from epidemiology.¹⁰ Here, instead of disease, it is the wage premia on jobs held by one’s neighbors that are ‘contagious’: One is more likely to get an offer from an employer paying the same wage premium as one’s neighbors. The contagion model is consistent with workers obtaining jobs in the same firm as one of their social network contacts, as in Bayer et al. (2008). Another interpretation is that the contagion process approximates the results of secondary transmission of job offers, where workers with good wage premia pass along rejected offers to their social network contacts as in Calvo-Armengol and Jackson (2007).

Equation 2 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 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).¹¹ In this paper, I abstract from congestion effects

⁹In 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. To this end, I ultimately specify W in terms of residential proximity.

¹⁰Formally, the contagion process specifies the referral offer distribution as $h(\psi; i, t) = (w^i)^T \mathbb{1}(\Psi(t) = \psi)$ from which 2 follows.

¹¹The above-cited papers emphasize congestion effects in the transmission of job informa-

to focus on identifying the effect of local network quality on job search outcomes. This assumption is approximately correct if the congestion effect is trivially small relative to the contagion effect. Furthermore, congestion would imply that more workers desire referrals than actually receive them, suggesting that the results I find are a lower bound on the amount of underlying referral seeking behavior of workers. Future work using these data could explore the implications of allowing for congestion effects.

2.3 Implications

The model delivers simple predictions for the evolution and stationary distribution of employer wage premia, $\Psi(t)$. The first group are the key predictions of the job search model augmented with job information networks, and are the central focus in the remainder of this paper. The second group are “first-stage” results, and correspond to stylized features of the distribution of Ψ . Because the empirical analogue of Ψ is estimated from a statistical decomposition that need not be consistent with a job search model, it is important to verify that the estimated ψ behave as expected. Verifying the first-stage predictions yields more general insight into the spatio-temporal structure of earnings data, but more importantly, also helps shed light on the key identifying assumption that workers do not sort within neighborhoods on the basis of unobservable determinants of earnings.

2.3.1 Mean Effects

The following proposition shows that an increase in network quality will increase the mean of the truncated offer distribution.

tion alongside contagion effects. As Wahba and Zenou (2005) have shown, network congestion effects lead to empirically verified non-linearities in the use and effects of social contacts to find work.

Proposition 1 *If the distribution of offers received through referral, h , is log concave and $|\mathbb{E}_g(\psi|\psi > \psi_0, Z, W, \Psi^0) - \mathbb{E}_h(\psi|\psi > \psi_0, Z, W, \Psi^0)|$ is small then*

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

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

Proof. See Appendix A. ■

Proposition 1 provides conditions under which we expect the partial effect of increasing referral network quality to be associated with increased job quality, even without correcting the offer function for sample selection. The requirement that $|\mathbb{E}_g(\psi|Z, W, \Psi^0) - \mathbb{E}_h(\psi|Z, W, \Psi^0)|$ 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. In the empirical work, I estimate models that correct for sample selection formally, but also estimate the partial effect described in the proposition from the truncated offer distribution.

2.3.2 Mobility Effects

In addition to higher mean outcomes, the model also predicts that an increase in network quality will increase the probability that a worker makes a direct job-to-job transition. This result is a straightforward corollary of Proposition 1 together with Proposition 3, described below.

2.3.3 Distributional Effects

Proposition 2 provides predictions on the quantiles of the truncated offer distribution. I evaluate these in the empirical work for two purposes. First, they provide additional checks of the validity of the job search model. Furthermore, they are evidence that referral networks affect outcomes by modifying the distribution of job offers rather than through changes in search intensity. My data do not allow me to rule out the possibility that workers adjust search effort when they have better job referral prospects. However, if that was the only mechanism through which referral networks affect search, we would not expect to observe the distributional effects described in Proposition 2.

Proposition 2 *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. ■

The proposition captures the intuition that increasing ψ_0 will have a larger impact on quantiles of the offer distribution ‘close’ to the truncation point, while increases in referral network quality, which shift the mean of the offer distribution, ‘pull’ the observed offer distribution relative to its truncation point, so appear to stretch the observed job quality distribution from the right.¹²

2.3.4 First-stage implications

Workers move from lower to higher wage premium jobs.

¹²The condition of Proposition 2, that the offer distribution is log concave with a symmetric density, is satisfied by the normal distribution the uniform distribution, and the double exponential.

Proposition 3 *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 in the sense that they assume the referral offer distribution is stationary. 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. ■

Proposition 3 is true for most models of on-the-job search.¹³ 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 4 *The stationary distribution of Ψ is such that*

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

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 Calvo-Armengol and Jackson (2007), who prove an equivalent result.

¹³The 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 (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 as well as for workers aged 25–35.

3 Empirical Design

This section describes the assumptions that map the theoretical model to my preferred empirical specification for the offer function:

$$\psi_i = \gamma \bar{\psi}_{b(i)0} + Z_i \Pi + \beta \psi_{0i} + \kappa_{G(b(i))} + \bar{X}_{b(i)} \Gamma + \nu_i. \quad (3)$$

In Equation (3), ψ_i is the employer wage premium offered to i . Z_i is a vector of individual characteristics, $b(i)$ 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. $\kappa_{G(b(i))}$ is a reference group effect where the notation $G(b(i))$ indicates the reference group of contiguous blocks containing $b(i)$. The neighborhood reference group in these estimates 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. Finally, $\bar{X}_{b(i)}$ is the block-level mean of individual characteristics included in Z_i .

Evaluating the predictions of the theoretical model requires consistent estimates of γ , the social interaction parameter. Equation (3) has features in common with the linear-in-means model common in the peer effects literature, and raises similar concerns about identification. First, workers may sort into social groups on the basis of characteristics that affect job search. Second, the presence of unobservable factors affecting search, like access to transportation, or proximity to certain employers, might vary across space. As in Bayer et al. (2008), I take advantage of precise knowledge of place of residence and large sample size to pursue identification based on comparisons of workers who live on different blocks in the same narrowly defined neighborhood. Local referral network effects are identified under

the assumption that the unobserved factors affecting formal job search are identical across workers living in the same neighborhood.

Given that the variation I exploit for identification is non-experimental, I also develop a sensitivity analysis in the spirit of Altonji et al. (2005) and Krauth (2011). The sensitivity analysis describes how strong sorting and selection on unobservables must be within neighborhoods relative to sorting and selection on observables to explain the estimated effect. I provide evidence in Section 5 that there is very little within-neighborhood sorting of workers on the basis of observable characteristics, and even less sorting on unobservable characteristics, that influence earnings.

There are three additional challenges to identification: reverse causality, the reflection problem, and sample selection. The time sequencing of job mobility alleviates concerns about reverse causality and the reflection problem (Manski 1993; Brock and Durlauf 2001). Sample selection arises for the usual reason that in search models we only observe accepted job offers. As I will discuss, sample selection, while a concern, also provides an additional method for dealing with the reflection problem.¹⁴ An important, but separate identification issue, arises from the first-stage estimation of employer wage premia from a decomposition of log earnings. The key assumption, of exogenous mobility, is a feature of the theoretical model, but its validity is the subject of a debate in the literature. I defer discussion of this issue to Section 4.3

¹⁴The identification problems in this paper are related to the general problem of identifying social interactions documented by Manski (1993) and elaborated in Brock and Durlauf (2001). Blume et al. (2011) provide a thorough summary of this literature with a useful discussion of the kinds of data and models that can be used to identify social interactions.

3.1 From Theory to the Empirical Model

Let log earnings be given by

$$\ln y_{it} = \ln \gamma + \ln e_{i,t} + \ln p_{J(i,t)}. \quad (4)$$

$J(i, t) = j$ where j is the employer of i at time t . Human capital depends on observable time-varying inputs, X_{it} and 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:

$$\ln y_{it} = \alpha + X_{it}\beta + \theta_i + \psi_{J(i,t)} + \varepsilon_{it}. \quad (5)$$

The model allows heterogeneity in the formal and informal offer distributions, g and h , described in Equation 1. 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 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:

$$f(\psi|Z_i, w^i, \Psi^t) = ag(\psi|Z_i, w^i, \Psi^t) + (1 - a)h(\psi|Z_i, w^i, \Psi^t). \quad (6)$$

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.

$$\psi_{i,t}^* = a (\mathbb{E}_g(\psi|Z_i, w^i, \Psi^t) + \eta_{i,t}^g) + (1 - a)(\mathbb{E}_h(\psi|Z_i, w^i, \Psi^t) + \eta_{i,t}^h) \quad (7)$$

$$= a \mathbb{E}_g(\psi|Z_i, w^i, \Psi^t) + (1 - a) \mathbb{E}_h(\psi|Z_i, w^i, \Psi^t) + \eta_{i,t}, \quad (8)$$

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 .¹⁵

The expected offer from formal search is independent of who your neighbors are and where they work. Likewise, the mean offer received through the network does not depend on individual characteristics other than through their influence on its structure, through w^i . Imposing these conditional moment restrictions yields

$$E(\psi_{it}^*|Z_i, w^i, \Psi^t) = a E_g(\psi_{it}^*|Z_i) + (1 - a)(w^i)^T \Psi^t + a E(\eta_{i,t}^g|w^i, \Psi^t) + (1 - a) E(\eta_{i,t}^h|Z_{it}). \quad (9)$$

I make a parametric assumption that the conditional mean of the formal offer distribution is linear in observable worker characteristics so that $E_g(\psi_{it}^*|Z_i) = Z_i \tilde{\Pi}$. Accumulating all of the modeling assumptions, the offer function is given by

$$\psi_i^* = Z_i \Pi + \gamma \bar{\psi}_{b(i)0} + \eta_i, \quad (10)$$

¹⁵A more sophisticated weighting matrix would take account of other dimensions of social distance, such as age, race, ethnicity, and region of birth. For the purposes of this paper, I restrict attention to residential proximity for two reasons. First, the identification strategy is based on a fundamental intuition regarding the residential real estate market. The same intuition does not extend to demographic characteristics. Second, a specification of the social weight matrix, W , that interacts residential proximity with other characteristics will be based on additional modeling assumptions on the importance of demographic characteristics relative to proximity in social interaction. Inference about the correct metric of social distance using labor market data is a complicated topic, as Conley and Topa (2002) show. My analysis of these issues suggests that the estimated referral effects are stronger when accounting for homophily of social interactions. However, to keep the focus of this paper tight, and to maintain comparability with Bayer et al. (2008), I maintain the simple definition of W described in the text. Extensions of W to include more observable characteristics is the subject of ongoing research.

where $\gamma = (1 - a)$ and $\Pi = a\tilde{\Pi}$, and I have eliminated the time subscripts to reflect the fact that the empirical work uses data on job changers in a single year, along with lagged information on block-level characteristics. The notation $\bar{\psi}_{b(i)0}$ indicates that the mean employer wage premium among workers on the same block is measured before the worker changes jobs. The primary identification problem is embedded in the potential for correlation between the composite error term, η_i and referral network quality, $\bar{\psi}_{b(i)0}$. I turn to these issues next.

3.2 Identification Strategy and Estimation Issues

3.2.1 Sorting and Correlated Unobservables

The key identification problem arises 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.

The economic argument underlying my identification strategy 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. That is, they choose their neighborhood, but not their neighbors. Likewise, 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.

Formally, identification of Equation 10 requires the composite error term is uncorrelated with referral network quality:

$$E(\eta_i | \bar{\psi}_{b(i)0}, Z_i) = 0. \quad (11)$$

If neighborhoods differ in the types of workers they attract or access to jobs, then $E(\eta_i | \bar{\psi}_{b(i)0}, Z_i) \neq 0$. To make this concrete, specify

$$\eta_i = \omega_i + \kappa_{b(i)} + \varepsilon_i, \quad (12)$$

where ω_i captures person-specific heterogeneity and $\kappa_{b(i)}$ captures block-level effects that move the mean of the formal offer distribution.

These observations motivate the following generalization of Equation 10

$$\psi_i^* = Z_i\Pi + \gamma\bar{\psi}_{b(i)0} + \zeta_{G(b(i))} + \eta_i, \quad (13)$$

where $G(b(i))$ denotes the reference group of blocks – a neighborhood – to which the block on which i lives b belongs.¹⁶ $\zeta_{G(b(i))}$ is a reference group effect. It will be convenient to consider the model expressed as deviations from neighborhood-level means:

$$\tilde{\psi}_i^* = \tilde{Z}_i\Pi + \gamma\tilde{\psi}_{b(i)0} + \tilde{\eta}_i. \quad (14)$$

Note that $E(\tilde{\eta}_i) = \tilde{\omega}_i + \tilde{\kappa}_{b(i)}$.

My main results are based on the assumption that $\text{cov}(\tilde{\eta}_i, \tilde{\psi}_{b(i)0}) = 0$. This moment restriction is the analogue to the identification strategy pursued by Bayer et al. (2008), and requires that there is no within-neighborhood covariance between block-level deviations in referral network quality and block-level deviations in unobservable determinants of formal search. That is,

$$E(\tilde{\omega}_i\tilde{\psi}_{b(i)0}) = E(\tilde{\kappa}_{b(i)}\tilde{\psi}_{b(i)0}) = 0.$$

The effects of referral networks are thus identified from comparisons of workers who change jobs starting from different blocks within the same neighborhood.

¹⁶For 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.

3.3 Sensitivity Analysis

Section 5.1.1 presents evidence that there is very little sorting of workers across blocks on observable characteristics, and even less sorting on unobservable correlates of earnings absorbed in the AKM residual. Nevertheless, as in Bayer et al. (2008), some sorting remains, leaving open the question of whether, and to what extent, the results are driven by sorting on unobservables. In this section, I present a method, based in part on Krauth (2011), to evaluate the sensitivity of my findings to sorting on unobservables. Like Altonji et al. (2005), this method draws on the assumption that sorting on unobservables is related to sorting on observables.

Specifically, the method adopts the approach, common in empirical studies, of assessing the extent to which additional control variables move the estimated effect of interest – in this case, the estimate of the coefficient on the block-level average wage premium. When the key identifying assumption $\text{cov}(\tilde{\eta}_i, \tilde{\psi}_{b(i)0}) = 0$ does not hold, the preferred estimate and control estimate are biased. I describe conditions under which the contrast between these two biased estimates identifies the bias associated with sorting on observables, and then use this information to bound the amount of sorting on unobservables that would be required to completely explain the preferred estimate.

Consider the model of Equation (14), based on the within-transformation. The goal is to replace the assumption $\text{cov}(\tilde{\eta}_i, \tilde{\psi}_{b(i)0}) = 0$ with the arguably less strident assumption that sorting on unobservable determinants of job search is related to sorting on its unobservable determinants.

Following Krauth (2011), consider the projection of $\tilde{\eta}_i$ onto a set of control variables that vary block-by-block. Let \tilde{V} be this set of variables.

$$\tilde{\eta}_i = \tilde{V}_i \phi + \tilde{u}_i \tag{15}$$

In practice, $V_i = (Z_i, \bar{Z}_{b(i)})$, the set of all observable characteristics in the model along with their block level means. \tilde{V} captures any within-neighborhood variation in observable characteristics. Table III, discussed in Section 5.1.1, presents evidence on within-neighborhood variation in observable characteristics. Such variation is minimal, but non-zero, motivating this sensitivity analysis.

The second model for the offer function, expressed as within-neighborhood deviations, is now

$$\tilde{\psi}_i^* = \tilde{V}_i \Gamma + \gamma \tilde{\psi}_{b(i)0} + \tilde{u}_i. \quad (16)$$

Note that Γ combines ϕ and Π , and has no structural interpretation. In general Π is not separately identified without additional assumptions. Since I am only interested in identifying γ , I do not impose those assumptions. Note also that $E(\tilde{V}_i \tilde{u}_i) = 0$ by construction.

I am now prepared to define the objects that compose the sensitivity measure. Let $\hat{\gamma}_0$ be the estimator of γ based on the preferred model of Equation 14. Let $\hat{\gamma}_1$ be the estimator of γ based on the control model of Equation 16. Finally, define scalars B_V , B_u , and μ , such that

$$\text{plim } \hat{\gamma}_0 = \gamma + B_V + B_u \quad (17)$$

and $B_u = \mu B_V$. In words, B_V is the part of the bias in the preferred estimate that can be predicted by variation in the observable control variables (block-level average characteristics), and B_u is the part that cannot. μ measures how important unobservables are relative to observables.

In Appendix B.1, I show that

$$\mu^* = \text{plim } \frac{R}{1 - R} \quad (18)$$

is the value of μ required to fully explain the observed effect when the true value of γ is zero.

The term R in Equation 18 is the ratio $R = \frac{\hat{\gamma}_1}{\hat{\gamma}_0}$.

3.4 Other Challenges to Identification

3.4.1 The Reflection Problem

Equation 16 is potentially afflicted by a reflection problem.¹⁷ Yet several features of the model and data in this research facilitate breaking the reflection problem. First, the time-sequencing of the data means that $\bar{\psi}_{b(i)0}$ is predetermined at the point in time in which a worker makes a job-to-job move. As Brock and Durlauf (2001) and Blume et al. (2011) argue, this simple time sequencing breaks the reflection problem, because $\bar{\psi}_{b(i)0}$ will depend on the complete history, of observed variables, which are reasonably excluded from the model.¹⁸ In addition, the non-random selection of realized offers provides another mechanism to break the reflection problem. The Heckman selection correction procedure described in the next section to handle non-random selection of realized offers introduces a non-linear term that breaks the reflection problem. Blume et al. (2011) aver that even though the reflection problem is in some sense easy to avoid – it is a knife’s-edge artifact of the linear model – it can still affect estimation and inference since the data matrix may be ill-conditioned. 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. The estimated effects are robust across specifications, indicating that ill-conditioning associated with a ‘near’ reflection problem is not present.

¹⁷The reflection problem (Manski 1993) arises when the endogenous variable – here, network quality $\bar{\psi}_{b(i)0}$, is perfectly collinear with the independent variables in the model. In cross-sectional data, this happens if block-level means of all of the independent variables, Z , are included as independent regressors. If it were a cross-sectional specification, γ might not be separately identified from Γ in Equation 16.

¹⁸Conley and 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.

Non-Random Selection of Job Offers

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.¹⁹ 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. Two caveats are in order, one positive and one negative. First, the selection correction model relies on an exclusion restriction, and I exclude the wage premium on the initial job, ψ_0 , which according to the model does not affect the offer distribution. Despite the internal consistency of the exclusion restriction, one might still be concerned about its validity. Second, Propositions 1 and 2 are statements about the uncorrected estimator, so for the purposes of testing these qualitative propositions, sample selection is not a concern. I discuss these matters in full when presenting results in Section 5.

Reverse Causality

Reverse causality is a problem if workers' 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 of workers who do not change residence, removing any concerns about reverse causality. Note that this represents a minor, but important, improvement over Bayer et al. (2008) and Hellerstein et al. (2011), which use cross-sectional matched employer-employee data to infer

¹⁹The estimated wage premia show workers will move to jobs with lower premia. This feature of the data is related to the finding in Nagypál (2005) that the rate of job-to-job transitions is not consistent with 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.

the presence of direct referrals by neighbors.

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 et al. 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 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 derive from 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 filesystem makes use of the unique individual and employer identifiers from this system to track workers over time as they move from employer to employer, and to identify which workers share an employer. These data are augmented with demographic and employer characteristics through links to survey and administrative records sources. For a complete description of these data, see Abowd et al. (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.

To facilitate exposition, a thorough description of the input data sources and data preparation, including construction of the main research samples are deferred to Appendix C. The discussions in this section focus on those aspects of the data of interest necessary to understand, interpret, and evaluate the analysis in the main part of the paper.

4.2 Stage 1: Estimation of Employer Wage Premia

As an empirical analogue to Equation 5, I use the Abowd-Kramarz-Margolis (AKM) decomposition:

$$\ln Y = X\beta + D\theta + F\psi + \varepsilon. \quad (19)$$

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.²⁰

²⁰This decomposition as applied to matched employer-employee data was first introduced by Abowd et al. (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 et al. (2002) and Abowd et al. (2003).

4.3 Exogenous Mobility

My analysis relies on the interpretation of the firm effects estimated from the AKM decomposition as employer-specific wage premia²¹. As wage premia, the firm effects should measure the amount any worker would be paid if they accepted employment in a particular firm in excess of the return to portable skills. Such an interpretation is equivalent to assuming that job mobility and assignment are exogenous to the wage residual or, more formally, that $E(\varepsilon|X, D, F) = 0$. The assumption of exogenous mobility is common in the literature using two-way fixed effects models to separate individual and employer heterogeneity, primarily because there are no clear, computationally attractive methods for weakening it. Here, I argue that while the exogenous mobility assumption is, like all modeling assumptions, a fiction, it is a useful fiction.²²

I have derived a simple theoretical model under the assumption of exogenous mobility that leads to several testable predictions on the relationship between employer-specific wage premia and job mobility, all of which I verify. An alternative explanation based solely on endogenous mobility bias requires assumptions that are equally, if not more, implausible. Combes et al. (2008) and Iranzo et al. (2008) use heterogeneity components from very similar 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 reveal about labor markets will surely benefit from these approaches.

²¹It is important to note that I do not need to assume that wage components from AKM have a structural interpretation as latent productivity, so the analysis is not subject to the concerns raised by, e.g. Shimer (2005) and Eeckhout and Kircher (2011)

²²Abowd et al. (2010) formally reject the null hypothesis of the exogenous mobility assumption in LEHD data. Their finding parallels those of related tests by Gruetter and Lalive (2009) for Germany, Rothstein (2010) in matched student-teacher data, and Woodcock (2008) on the importance of match-specific heterogeneity. Kinsler (2012) fails to reject the exogenous mobility assumption in education data after accounting for sample size and model misspecification.

4.4 The Estimation Sample

[Table I about here]

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. The restriction to MSAs is a modeling choice, and I selected the most populous MSAs for which all counties are included in LEHD during 2002–2003. A complete list of MSAs used is shown in Table I. The restriction to the years 2002–2003 was necessitated by availability of residential address information. Table II presents descriptive statistics for the full sample as well as for the two subsamples of workers in job-to-job transition used to estimate the main results. 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 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.

[Table II about here.]

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. I consider two different definitions of job-to-job transitions. The set of ‘Annual Job Changers’

is drawn using a less restrictive definition. I include all in-sample workers who were employed full time in 2002 and 2003, and who changed employers between the two years. The set of ‘Quarterly Job Changers’ uses a more precise definition of job-to-job transition, and only includes workers for whom I can precisely identify the exact quarter when they change from one dominant employer to the next.²³ The second, third, and fourth columns in Table II present statistics for job changers. As expected, they are younger, but otherwise similar demographically to non-movers.

The objective of using these two samples is to show that results are robust to allowing for different lag structures in the data. For the quarterly job changers, the relevant measure of network quality is the block-average wage premium in the quarter before they change jobs. For annual job changers, the relevant measure is the block-average wage premium as of the beginning of 2002 (excluding workers who change jobs). As we will see, these differences do not affect the results much. The sample of annual job changers does not impose the restriction that the worker does not experience a spell of non-employment between jobs. For the sample selection correction, and checking for a job ladder, I restrict analysis to the sample of quarterly job changers.

Among the full sample, just 3.5 percent of workers experience a transition between dominant employers, and thus belong to the set of quarterly job changers. This is significantly lower than the reported rate of job-to-job transitions in other sources (Bjelland et al. 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. Many of these workers will be picked up in the set of

²³Dominant 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.

annual job changers. 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 quarterly job changers is properly interpreted as a sample of immediate transitions from one relatively long-term job to another.

5 Results

The estimation results presented in this section fall in two categories. First, I present evidence in support of the key modeling assumptions. Specifically, in support of the identification strategy, I document the absence of substantial sorting within neighborhoods on the basis of observable and unobservable characteristics. I also show that the firm effects, ψ , estimated from the first-stage AKM decomposition are broadly consistent with the job search model outlined in Section 2.

Next, I turn to evaluation of the more specific empirical implications of the job search model. First, living on a block where workers have higher average wage premia will improve the mean outcome of job-to-job transitions. Second, living on a block where workers have higher average wage premia increases the probability of making a job-to-job transition. Third, living on a block where workers have higher average wage premia stretches the observed distribution of outcomes from job-to-job transitions. Finally, increasing the wage premium on the initial job compresses the distribution of outcomes for workers making job-to-job transitions from below.

The final subsection considers alternative specifications designed to check the interpretation of my results for the offer function. I find that non-native workers and younger workers are more strongly affected in their job search outcomes by the job quality of their neighbors, which is consistent with previous research on the prevalence of referral use among those groups. Furthermore, I show that my results are partially, but not completely associated

with direct referrals, where a worker moves to a job in the same firm as one of his block-level neighbors. I defer full interpretation of these findings to the sections in which they are presented.

5.1 Evidence for the Model Assumptions

5.1.1 Sorting within Neighborhoods

My identification strategy relies on the absence of sorting of workers across blocks within neighborhoods on the basis of factors that influence their jobs search outcomes. To support that assumption, and facilitate the sensitivity analysis, this section presents two types of evidence regarding the block-by-block sorting of workers within neighborhoods.

First, I present non-parametric estimates of the spatial autocorrelation function for block- and tract-level averages of log earnings and the error components from the AKM decomposition. The main result is that there is no excess spatial correlation within neighborhoods in the residual from the AKM decomposition. A secondary result is that these estimates show evidence of spatial correlation in employer-specific wage premia, ψ , one of the stylized features of the data my model attempts to explain. Next, like Bayer et al. (2008), I present results on the basic correlation between individual characteristics and the average characteristics of neighbors living on the same block, with and without controls for block-group. I also find that there is very little sorting by demographic characteristics within block groups.

Figures 1a and 1b 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 . As in Conley and Topa (2002), each variable x_i is associated with a spatial coordinate, s_i , and the correlation between x_i and x_j depends only on the distance between s_i and s_j : $\text{Corr}(x_i, x_j) = f(\|s_i - s_j\|)$. Details on estimation of the spatial autocorrelation function, f ,

are available in Appendix B.2.

[Figures 1a and 1b about here.]

The key feature of Figures 1a and 1b is the contrast in the extent of spatial sorting in the AKM residual between the tract-level and block-level estimates. The observed sorting that appears at the tract level is completely eliminated in the block level estimates. Given that ε from the AKM decomposition is a measure of unobservables that affect earnings, if there is sorting within neighborhoods in factors affecting job search outcomes, we should also see sorting on ε . While the fact that we do not observe such sorting does not prove that the identifying assumption holds, it is consistent with that assumption. Moreover, it suggests that sorting on unobservables is, if anything, less strong than sorting on observables, a finding that lends support to the sensitivity analysis.

Table III presents complementary evidence of the extent of sorting within neighborhoods. The data for this analysis are restricted to prime-age male workers who were employed full time for the full year in 2002. From each block, I draw one worker at random. The entries in the table are R^2 from estimates of linear models that predict the characteristics of the randomly selected individual using the block-level average characteristic of other workers in the same block. The column labeled ‘Raw’ presents the R^2 from univariate regressions of characteristics of a randomly selected individual on the block-level average characteristic of other workers in the same block. The column labeled ‘Block Group Controls’ presents the within-group R^2 from the same model adding block group effects. Intuitively, if there is no sorting within neighborhoods, all of the raw correlation measured in the raw R^2 should be eliminated by the block group controls. That is, variation within block group in average characteristics should not predict the characteristic of a random individual.

Sorting is heavily attenuated after introducing block group controls. This is consistent with the identifying argument that workers sort themselves into different parts of the city,

but are not as selectively sorted within neighborhoods. Like Bayer et al. (2008), who conduct a similar exercise, I find that these within-neighborhood sorting measures are not identically zero. However, I also present a rough measure of the amount of sorting on unobservables. The amount of within neighborhood sorting on the AKM residual is zero. This indicates that sorting on unobservable characteristics that influence earnings is less strong than sorting on observable characteristics. These findings lend support to my interpretation of the sensitivity analysis. Sorting on unobservables will need to be much stronger than sorting on observables than seems likely, given the data evidence, to explain my estimated effects.

5.1.2 Evidence on Direct Referrals

In this section, I provide evidence to support the assumption that workers use local networks to share information about employment opportunities in general, and about high-paying employment opportunities in particular. I 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. I show further that the strength of the social interaction effect is increasing in the wage premium, measured by the AKM firm effect, ψ , on the destination job. Like Bayer et al. (2008), I interpret my finding as evidence of a social interaction effect in job search. Unlike Bayer et al. (2008), whose data were cross-sectional, I use longitudinal data with details on actual job changes, which admits two improvements on their original design. First, my data eliminate concerns about reverse causality: I observe workers living in the same neighborhood before they begin the same jobs. My second, and more substantial, improvement is to take advantage of information on which worker in a pair of neighbors was on the job first. I observe events in which one worker is currently employed by a firm, and then one of her neighbors begins employment with the same firm. If there is a referral relationship underlying the event, the time-sequencing in the data reveals its direction: from the referring worker to the referred worker.

I construct a sample of matched pairs of workers from the sample of yearly 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

- $R_{\ell,m}$: an indicator equal to 1 if ℓ and m live on the same Census block and zero otherwise.
- $W_{\ell,m}$: an indicator equal to 1 if ℓ and m share the same employer in 2003 and zero otherwise.

The sample of matched pairs contains 1,558,436,893 observations. The baseline probability that $W_{\ell,m} = 1$ among pairs of workers that do not reside on the same block is 0.0013, or 0.13 percent of such pairs. This number is one third the size of the baseline probability Bayer et al. (2008) find for two workers who live in the same block-group but not the same block to be employed on the same block. This probability increases to 0.0016 among pairs of workers that live on the same block.

I control for block-group specific variation in the propensity for a mover and stayer to end up employed in the same firm. The baseline empirical specification is the linear probability model

$$W_{\ell,m} = \rho_{G(\ell)} + \alpha_0 R_{\ell,m} + \varepsilon_{\ell,m}, \tag{20}$$

where the notation intentionally corresponds to Equation (1) in Bayer et al. (2008).

The effect, α_0 is the estimated effect of living on the same block on the propensity for a mover to become the coworker of a stayer, which is evidence of a social interaction in job-finding under the identifying assumption that all variation in the propensity to be employed in the same firm as a neighbor is absorbed by the block-group effect, $\rho_{G(\ell)}$.

I extend the specification in Equation (20) to account for the characteristics of the pair:

$$W_{\ell,m} = \rho_{G(\ell)} + \beta' X_{\ell,m} + (\alpha_0 + \alpha_1' X_{\ell,m}) R_{\ell,m} + \varepsilon_{\ell,m}, \quad (21)$$

I estimate this model including in $X_{\ell,m}$ the race, gender, ethnicity, age and person effect from the AKM decomposition (θ) for both the mover and stayer. I also include the employer wage premium, ψ , of the stayer. The main novelty of this model with respect to the analysis in Bayer et al. (2008) is that the directed nature of the job move allows me to infer who is doing the referring (the stayer) and who is being referred (the mover) in these matched pairs.

The results of the analysis appear in Table IV. Specification (1) corresponds to the basic model in Equation (20). The coefficient estimate says that residing on the same block is associated with a .02 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. Bayer et al. (2008) find a roughly 33 percent increase over the baseline in their most conservative specification. Given the much more restricted nature of my sample and design, these results should be interpreted as supportive of the basic Bayer et al. (2008) findings. The results from specification (2) show that the effect of geographic proximity is more positive among pairs of workers when the firm wage premium is higher. In other words, to the extent that this design measures social interaction effects, it says that the social interaction effect is stronger when the referring worker has a better job.

The results in this section suggest a range of further avenues for investigating referral relationships in longitudinal matched employer-employee data. While the presence and importance of referrals in job search are well-established at this point, the reasons referrals are used are still very poorly understood. For the purpose of this paper, I include these findings to support the modeling assumption that workers use job information networks, in part, to find better-paying jobs. Later, I will use the data developed in this section to examine how

much of the social interaction effect in my estimated wage premium offer function can be attributed to direct referrals.

5.1.3 The Job Ladder in ψ

[Table V about here.]

Table V shows the fraction of quarterly 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 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 Section 2. Workers who change jobs tend to move to firms with higher ψ than the firm they left. Figures 2b and 2a plot the cumulative and frequency distribution of destination ψ for all job transitions stratified by decile of the origin job wage premium. Figure 2b displays a clear first-order stochastic dominance relationship among the conditional distributions.²⁴ 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.

²⁴While the evidence supports the job ladder prediction, many job-to-job transitions involve moves to jobs with lower rather than higher wage premia. This is consistent with a model in which mobility is driven both by wage premia and idiosyncratic preferences by workers for employment in particular firms. It is possible to incorporate these compensating differentials into the job search model with referral networks without changing any of the model predictions. If anything, this should lead to an understatement of the effect of local referral networks on job search outcomes.

5.2 Tests of the Model Predictions

5.2.1 Mean Effects

The main results are estimates of linear and unconditional quantile regression models of the form

$$\psi_i = \gamma \bar{\psi}_{b(i)0} + Z_i \Pi + \beta \psi_{0i} + \kappa_{G(b(i))} + \bar{X}_{b(i)} \Gamma + \nu_i, \quad (22)$$

which was introduced as Equation (3) in Section 3. Primary interest lies with estimates of the parameter γ , which, under the identifying assumption, measures the effect of local interactions on job offers. In this section, I test the prediction that local referral network quality is positively correlated with job search outcomes. I also show that increases in network quality are associated with an increased probability of making a job-to-job transition in the first place. The former outcome could be the result of a spurious correlation in wage premia among neighboring workers. The latter outcome is harder to explain in terms of sorting on existing wage premia.

[Table VI about here.]

The key result in Table VI is in the contrast between the baseline specification in column (1), which does not control for reference group correlations in outcomes, and the specifications in the remaining columns that do. Inference in the conditional mean regressions is based on heteroscedasticity-corrected standard errors that have been clustered at the MSA level.²⁵ The baseline model presented in the first column of Table VI 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.

²⁵This specification is conservative. Under the empirical model, clustering at the county or tract level would be appropriate. As Cameron et al. (2008) point out, asymptotic tests based on data with around 30 or fewer clusters may over-reject. Even with standard 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.

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 columns (3) and (4) of Table VI. Column (3) is restricted to quarterly job changers, for whom we know there is a direct job-to-job transition. Column (4) uses annual job changers. The results are nearly identical in both cases, and I focus the remaining discussion on yearly job changers. The point estimate of interest is $\hat{\gamma} = 0.11 \pm 0.01$, and is statistically significant. Interpreted in terms of the job search model of Section 2, this implies that 11 percent of job offers arrive through referrals. This is in line with the analysis in Ioannides and 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 personal contacts to search for work. Alternatively, the point estimate indicates that a one standard deviation increase in network quality is associated with a 25 percent increase in the employer wage premium.

In all models, 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.

Sensitivity to Sorting

The other important contrast in Table VI is between columns (4), (5), and (6). Column (5) investigates the influence of sorting on observable demographic characteristics. As discussed in Section 3.3, this provides information on the sensitivity of the key result to sorting within neighborhoods on unobservables, following the arguments in Krauth (2011). Let $\hat{\gamma}_{(4)}$ be the estimate from column (4) and $\hat{\gamma}_{(5)}$ the estimate from column (5). Then based on the argument in Section 3.3, we can calculate how strong the bias from within-neighborhood sorting on unobservables must be relative to sorting explained by the controls. Using the results, I find

$$\hat{\mu}^* = \frac{R}{1 - R} = 4.6 \quad (23)$$

That is, the bias that is not explained by the controls must be 4.6 times as strong as the bias that is explained by the controls. Whether this is high or low depends, of course, on the context, and quantitatively there is no clear benchmark. However, given the evidence in Section 5.1.1 that within-neighborhood sorting on unobservables is weaker than sorting on observables suggests a qualitative benchmark. Assuming, as in Altonji et al. (2005), that unobservables are no more highly associated with job search outcomes than observables, a natural benchmark is $\hat{\mu}^* = 1$. The sensitivity analysis is, therefore, supportive of the main conclusion regarding the presence of local interaction effects in employer-specific wage premia.

Note that the model in Column (5) of Table VI includes as a control the average value of θ for all workers in the block. Recall that θ is the person-effect from the AKM decomposition (Equation 19) and captures all worker-specific characteristics that do not change over time associated with earnings. If there is excess sorting within neighborhoods on the basis of any individual-specific attributes correlated with earnings, it should show up here. If such sorting were driving the results, the inclusion of $\bar{\theta}$ might severely attenuate the point estimate. That

it does not support my interpretation of the sensitivity analysis.

Columns (4) and (6) in Table VI 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.

5.2.2 Mobility Effects

As a further test of robustness of the preceding results, as well as to provide additional evidence in support of the theoretical model, I present estimates from a selection-correction model that corresponds to the full job search model described in Section 3. Two key results appear in Table VII. First, I find, consistent with the model, that the probability of making a job-to-job move is increasing in referral network quality. Second, I find that the selection correction has no effect on the estimated effect of referral network quality on the wage premium on the next job. After presenting the estimation results, I discuss concerns about the validity and identification of the selection correction model.

[Table VII about here.]

In the job search model model, only workers who receive sufficiently attractive offers change jobs, leading to truncation of the observed offer distribution. For employed workers, the attractiveness of a job offer depends on the wage premium of one's current job, ψ_0 . I implement a standard Heckman selection model with normal errors to formally model this selection process. The theoretical model supports exclusion of ψ_0 from the offer function. I estimate on the sample of quarterly job changers. The first stage estimates the probability of changing dominant employers in 2002:Q4. The results are not sensitive to the choice of

reference quarter.

The first-stage estimates a probit with the observed outcome being whether the worker changed jobs. As the search model predicts, and consistent with the ‘job ladder’ evidence presented earlier, the selection equation shows that the log wage premium on the worker’s initial job ψ_0 is negatively associated with making a job-to-job move. Even more relevant for this paper, the results show that workers living on blocks with better than average network quality for their neighborhood are more likely to make a job-to-job transition. For this effect to be explained by sorting on unobservables, it must be the case that living on a block with higher average wage premium is also associated with being more likely to change jobs.

The point estimate on the social interaction parameter in the second stage is $\hat{\gamma} = 0.11 \pm 0.02$, which is essentially identical in magnitude and significance 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. As we saw in the discussion of the job ladder prediction, there is a fair amount of variation in job mobility consistent with the presence of idiosyncratic compensating differentials.

From an economic perspective, the first-stage results are more compelling than the second-stage results. The reduced-form probit estimated in the first-stage supports the prediction in Proposition 3 without invoking additional exclusion restrictions. The second-stage relies on the internally consistent, but perhaps dubious, exclusion of ψ_0 from the offer function. The selection correction model is not required to test Propositions 1 and 2. I offer the second stage result of the selection correction model in the interest of completeness, and as additional evidence that basic robustness checks do not alter the effects estimated in Table VI.

5.2.3 Distributional Effects

The job search model makes predictions on the full distribution of job search outcomes, as described in Proposition 2. Specifically, an increase in ψ_0 will compress the observed job quality distribution from the left, while an increase in network quality, $\bar{\psi}_{b(i)0}$, should ‘stretch’ the observed job quality distribution from the right. Both of these effects are driven in part by the truncation of the observed offer distribution with respect to the wage premium on the initial job, ψ_0 . I test these predictions by evaluating how changes in both variables affect quantiles of the observed outcome distribution using the unconditional quantile regression approach proposed in Firpo et al. (2009). The unconditional quantile regression permits estimation of the partial effect of a variable of interest on the marginal outcome distribution. By contrast, the better known conditional quantile regression approach only measures the effect of changing the variable on interest on quantiles of the residual distribution.²⁶

[Figure 3 about here.]

Figure 3 displays estimates of the unconditional quantile partial effects (UQPE) of the employer wage premium on the origin job (ψ_0) and block-level referral network quality ($\bar{\psi}_{b(i)0}$) at each quartile of the outcome distribution for yearly job changers. The results largely support the theoretical predictions. First, the model predicts that the effect of an increase in ψ_0 should be decreasing across quantiles. Figure 3a displays a pattern which is largely consistent with this prediction. The exceptions are at fifth percentile, which appears too low given the model, and above the 75th percentile.

²⁶Firpo et al. (2009) show that the UQPE is directly related to the recentered influence function (RIF) of the outcome variable. My estimates use RIF-OLS. The specification for the RIF at each quantile is identical to the specification in column (5) of Table VI, though using specification (4) or (6) does not have a meaningful affect on the results. Details of the estimation procedure are available upon request. The qualitative results are essentially the same when estimated using conditional quantile regression.

The results in Figure 3a show the effect of increasing referral network quality across the distribution of observed wage premia on new jobs. The model predicts that these partial effects should be increasing across quantiles, which is precisely what the data show. A closer look at the results shows that most of the increases occur beyond the 75th percentile of the outcome distribution. This pattern is not predicted, nor is it ruled out, but the job search model.

5.2.4 Robustness: Referral effects by Demographic Subgroups

As a robustness check, I estimate the model allowing the effect of local referral network quality to differ for different demographic groups. Table VIII presents results for native workers versus non-native workers, younger workers, and older workers. Previous research indicates that immigrant workers and younger workers are considerably 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 referral network quality and job search outcomes and these demographic groups.

[Table VIII about here.]

Column (1) in Table VIII reproduces the model in column (6) from Table VI. Column (2) presents the results of allowing the local interaction effect through $\bar{\psi}_{b(i)0}$ vary with nativity. Non-native workers have $\hat{\gamma} = 0.16 \pm 0.02$, which is twice the magnitude of the pooled estimate.

This finding is consistent with other 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 et al. 2008; Damm 2009; Beaman and Magruder 2012). As a point of reference, 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) present results of allowing the local interaction effect through $\bar{\psi}_{b(i)0}$ 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).

While they are not presented here, results of estimating the same model allowing for heterogeneity by racial categories (‘white’, ‘black’) do not produce statistically significant differentials. This finding is consistent with evidence reported by Holzer (1988) that, at least for young workers, there are minimal differences in the use of informal contacts to find work. These findings provide additional robustness to the main result. If the main result were driven by sorting across blocks within neighborhoods, interactions with race would likely magnify the measured effect. If anything, the relevant point estimates are negative, so they weaken the measured effect.

An interesting extension of the results presented would incorporate demographic characteristics directly into the definition of the reference group, so that the block-level average wage premium is computed, for instance, just for workers of the same race. Preliminary analysis of such models indicates that this specification strengthens the estimated referral effect, as expected. However, the interpretation of such specifications is less clear, and there is no clear prior information to use in specifying the reference groups. See Footnote 15 for additional discussion on this point.

5.2.5 Robustness: The Relationship to Direct Referrals

In this section, I address the extent to which the local interaction effect is associated with direct referrals. A direct referral occurs when one worker gets another worker a job with his current employer. The existing work on social interactions in labor markets using matched employer-employee data studies direct referral (Bayer et al. 2008; Hellerstein et al. 2011;

Dustmann et al. 2011; Nordström Skans and Kramarz 2011). One objective of the present research is to ask whether the data are informative about more general mechanisms of informal job search. As discussed in the introduction, the ‘contagion’ process described in the theoretical model might capture several different mechanisms, of which true referrals are just one. The key intuition is that workers use their (local) contacts to find better paying jobs, and this could occur for a number of reasons.

To address this, using the sample of yearly job changers, I identify those workers who, on changing jobs, move to an employer that already hires one of their block-level neighbors. Not all such moves are direct referrals – some may occur by chance. By including them, I obtain an upper bound on how much of the local interaction effect in employer wage premia could be explained by direct referral. If all such interactions are driven by direct referrals, then there should be no effect of $\bar{\psi}_{b(i)0}$ on other workers when changing jobs.

Table IX presents the effect of controlling for direct referrals as well as for the interaction between direct referral and $\bar{\psi}_{b(i)0}$. Workers who move to the same employer as a neighbor have slightly worse outcomes than average. However, controlling for direct referrals does not affect the estimated effect of $\bar{\psi}_{b(i)0}$ on job search outcomes overall. The point estimate on γ is $0.08 \pm .01$ in Column (3), and is not attenuated at all in column (4) after taking out the part of the effect of $\bar{\psi}_{b(i)0}$ associated with workers who move to a job with the same employer as a block-level neighbor. The interaction effect in Column (4) shows that workers moving to a neighbors firm have a much stronger correlation between $\bar{\psi}_{b(i)0}$ and job search outcomes. This latter effect may be partially mechanical. These findings complement our earlier finding in Table IV that suggested local referrals are associated with high- ψ employers. However, they also reveal that not all of the correlation picked up by the neighborhood effect parameter, γ , is driven by direct referral, consistent with the use of local social interactions for indirect exchange of job information.

6 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 et al. (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 8 to 10 percent of a worker's job offers come from local job information networks. 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.

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 job search model makes a number of predictions on the quality of job search outcomes for workers making job-to-job moves, and on the probability that a worker will change jobs. All of these are verified in the data. The identifying assumption, that workers randomly sort across blocks within neighborhoods, is untestable, but auxiliary evidence suggests it is approximately valid. Moreover, my results survive a sensitivity analysis that includes controls for sorting on the basis of unobservable worker characteristics that are associated with earnings.

The research I have presented relies not only on indirect measures social proximity, motivating sensitivity and robustness checks, but also relies on indirect measures of the employer wage premia. I therefore provide auxiliary evidence that the estimated firm-specific component of earnings exhibits certain stylized features. These stylized features are consistent with my model, but also with many other models. I show that workers who switch jobs tend

to move into jobs with higher wage premia than their current job. 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 studies of residential sorting by earnings, human capital, and employer characteristics, in urban labor markets. Even for the most skeptical reader, these results are of interest, independent of the referral network interpretation.

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.

A Appendix: Model Details

A.1 Proof of Proposition 3

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

$$rV(p) = p + \lambda_1 \int_0^\infty [\max\{V(p'), V(p)\} - V(p)] d\tilde{F}(p') + \delta [U - V(p)], \quad (24)$$

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

$$rU = p_b + \lambda_0 \int_0^\infty [\max\{V(p') - U\}] d\tilde{F}(p'). \quad (25)$$

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

$$p_R = p_b + (\lambda_0 - \lambda_1) \int_{p_R}^\infty \left[\frac{1 - \tilde{F}(p)}{r + \delta + \lambda_1 [1 - \tilde{F}(p)]} dp \right]. \quad (26)$$

■

A.2 Proof of Proposition 1

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 2

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.²⁷

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,

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

which by the assumption above implies

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

²⁷For details on log concave functions and their application to search models, see Bagnoli and Bergstrom (2005) and Flinn and Heckman (1983).

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 Appendix: Estimation Details

B.1 Sensitivity Analysis

This appendix contains details of the sensitivity measure introduced in Section 3.3. The preferred model of Equation 14 is

$$\tilde{\psi}_i^* = \tilde{Z}_i \Pi + \gamma \tilde{\psi}_{b(i)0} + \tilde{\eta}_i. \quad (27)$$

To simplify exposition, consider partitioned regression. First, project variables into the null space of \tilde{Z} yields the univariate regression

$$\tilde{\psi}_i^Z = \gamma \tilde{\psi}_{b(i)0}^Z + \tilde{\eta}_i^Z. \quad (28)$$

where the superscripted Z indicates that the variable is expressed as the residual from projection onto \tilde{Z} . The OLS estimator, $\hat{\gamma}_0$ has the property that

$$\text{plim } \hat{\gamma}_0 = \gamma + \frac{\text{cov}(\tilde{\eta}_i^Z, \tilde{\psi}_{b(i)0}^Z)}{\text{var}(\tilde{\psi}_{b(i)0}^Z)}. \quad (29)$$

As in the text, I consider the linear projection

$$\tilde{\eta}_i^Z = \tilde{V}_i^Z \phi + \tilde{u}_i, \quad (30)$$

noting that $\tilde{u}_i^Z = \tilde{u}_i$, so

$$\text{plim } \hat{\gamma}_0 = \gamma + \frac{\text{cov}(\tilde{V}_i^Z \phi, \tilde{\psi}_{b(i)0}^Z)}{\text{var}(\tilde{\psi}_{b(i)0}^Z)} + \frac{\text{cov}(\tilde{u}_i, \tilde{\psi}_{b(i)0}^Z)}{\text{var}(\tilde{\psi}_{b(i)0}^Z)}, \quad (31)$$

$$= \gamma + B_V + B_u. \quad (32)$$

Next, introduce the relative bias term, μ , that captures how strong bias from unobservables is relative to observables: $B_u = \mu B_V$. In practice, I am only concerned with cases where $\mu > 0$. Next, consider an estimate of the control model from Equation 16. Repeating the ideas from above, consider the model expressed as residuals from projection onto V

$$\tilde{\psi}_i^V = \gamma \tilde{\psi}_{b(i)0}^V + \tilde{\eta}_i^V, \quad (33)$$

$$= \gamma \tilde{\psi}_{b(i)0}^V + \tilde{u}_i. \quad (34)$$

The OLS estimator, $\hat{\gamma}_1$ satisfies $\text{plim } \hat{\gamma}_1 = \gamma + B_{u'}$, where

$$B_{u'} = \frac{\text{cov}(\tilde{u}_i, \tilde{\psi}_{b(i)0}^V)}{\text{var}(\tilde{\psi}_{b(i)0}^V)}. \quad (35)$$

Define another scalar k such that $B_{u'} = kB_u$, and a statistic, R , such that $R = \frac{\hat{\gamma}_1}{\hat{\gamma}_0}$. Let μ^* measure how strong bias from unobservables must be relative to bias from observables to explain the entire estimated effect. Setting the true value $\gamma = 0$ yields $\text{plim } R = \frac{\mu^* k}{1 + \mu^*}$. It follows that

$$\mu^* = \text{plim } \frac{R}{k - R} \quad (36)$$

for given k . The case $k = 1$ yields the condition presented in Section 3.3. In this paper, $0 < R < 1$ always, and $k > R$ must obtain. Therefore, the setting $k = 1$ gives the most

conservative estimate of μ^* .²⁸

B.2 Estimation of the Spatial Autocovariance

To produce Figures 1a and 1b, I estimate the spatial auto-covariance function at distance δ , $f(\delta)$, non-parametrically by

$$\hat{f}(\delta) = \sum_{i=1}^N \sum_{i'=1}^N \phi \left[\frac{|\delta - A_{ii'}|}{\sigma} \right] (X_i - \bar{X}) (X_{i'} - \bar{X}) \quad (37)$$

where $A_{ii'}$ is the distance between i and i' . $\phi(\cdot)$ denotes the standard normal kernel. The spatial auto-covariance 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

²⁸ k reflects the amount bias associated with unobservables that is removed through the control variables. The case $k = 0$ occurs when the control variables eliminate all of the bias. The case $k = 1$ occurs when the control variables do not predict any of the within-neighborhood variation in $\bar{\psi}_{b(i)0}$. In practice, I will only be concerned with cases where $0 < k \leq 1$.

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.

C Appendix: Data Sources and Construction

The analysis conducted in this paper is based on various subsamples of a master research database generated based on data from the Census Bureau’s LEHD program. Since the proposed analysis is at the individual level, I begin with the core LEHD data to produce a person-frame containing all individuals who ever appear in the 2004 Snapshot²⁹ Specifically, I extracted all records from the Individual Characteristics File (ICF). In the LEHD filesystem, each worker is uniquely identified by a Protected Identification Key (PIK), which may for most purposes be thought of as a scrambled Social Security Number. This frame includes the core individual characteristics (sex, date of birth, place of birth, race, and ethnicity), the sources of which are various administrative records.³⁰

C.1 Data on Residential Geography

The main novelty of the analysis in this paper is the inclusion of detailed residential address information. I also augmented the person frame with detailed data on residential geography

²⁹The 2004 Snapshot is an internally consistent, stable, extract of the core LEHD Infrastructure files in place as of 2004. Until 2011, this was the main source of LEHD data for both external research projects and internal research and development not directly associated with the production systems behind LEHD’s core public-use products. LEHD subsequently released a 2008 Snapshot, also available for external research projects through the Census Research Data Center network. Visit <http://www.census.gov/ces/> for details regarding access.

³⁰Education is included in the ICF for the 2004 Snapshot, but is imputed for almost all records, so I chose to leave it out of the analysis.

from an extract of the LEHD copy of the Composite Person Record which is itself an extract of the Statistical Administrative Records System (StARS). These records provide the Census geography down to the block of residence for most of the observations in the sample. The programs used to draw this information from the LEHD infrastructure files were also used to produce OnTheMap version 3(OTM3).³¹

The LEHD program receives an annual extract every year, beginning in 2002, of data on residential geography from the Statistical Administrative Records System (StARS). These data are delivered in a file called the Composite Person Record (CPR). The primary purpose for which LEHD obtains the CPR is to construct its LEHD Origin-Destination Employment Statistics (LODES), which are the microdata that feed OnTheMap, one of LEHDs two flagship public-use data products. ³² LEHD also uses residential address information from the CPR in making various improvements to the infrastructure filesystem, most recently to improve the quality of models used to impute race, ethnicity, and education, when missing.

The CPR exists for the years 1999–2011, with the exception of 2000. The 2001 CPR is based on proxy responses, and had many errors when compared to other sources. Linking the CPR to the LEHD files is straightforward. The CPR extract received by LEHD is a

³¹At the time the research for this paper was conducted, these residential address data were not available for external research projects. As of 2011, external researchers can apply to work with the 2008 Snapshot of the LEHD infrastructure filesystem, which includes data on place of residence as part of the Individual Characteristics File (ICF). This data appendix describes the place of residence data both to provide comprehensive detail on the analysis contained in this paper, and hopefully will be of use to researchers in developing proposals for approved external research projects using LEHD.

³²LODES contains information on the place of residence and place of employment for every job that appears in the LEHD infrastructure. As described above, this nearly covers the universe of all UI taxable employment. OnTheMap is a web-based mapping and reporting tool that allows users to exploit the information on commuting and labor demand patterns in LODES in a variety of applications, including regional labor market analysis, emergency management, and transportation planning. OnTheMap is available to the public at <http://lehdmap.did.census.gov/>. A publicly-available version of the underlying microdata are also available for research and analysis from <http://lehd.did.census.gov/led/onthemap/>.

PIK-level file with two key variables: PIK and HUID, with two associated data quality flags. Linking these files to the LEHD infrastructure is as simple as merging them to the Individual Characteristics File by PIK.

HUID is the StARS systems housing unit identifier and encodes residential address through one of several different Census geographic information systems. The HUID contains information, when it is available, down to the precise latitude and longitude of the place of residence for that PIK. Processing the HUID into a useable geocode is a complex task. For this research, I was able to borrow code, and some internal production files used as part of OnTheMap processing, to facilitate the conversion of HUID to conventional Census geographic units (block, block group, tract, county, CBSA).

Quality Concerns

Quality concerns are muted for the analysis conducted in this paper. The CPR data are quite likely to have block-level geocodes for employed workers in urban areas, and for these geocodes to be accurate when validated against external sources. The StARS collects eight types of administrative record from various government agencies to produce the annual place of residence data that appear in the CPR extracts.³³ These include

- IRS 1040,
- IRS W2/1099,
- HUD Tenant Rental Assistance Certification System (TRACS),
- Center for Medicare and Medicaid Services Medicare Enrollment Database (MEDB),
- Indian Health Services (IHS) Patient Registration System File,

³³StARS receives these records once a year and processes them to generate the CPR. As part of their process, for individuals that appear with multiple addresses, they select a ‘best’ address through a simple statistical procedure.

- Selective Service System (SSS) Registration File.

In this study, I focus largely on workers who appear in jobs in all, or most quarters of 2002–2003. These highly-attached workers likely also pay taxes, and consequently appear in one of the IRS records. Indeed, after linking and processing, I obtain a geocode that identifies at least the county for 95 percent of the LEHD workers. I use this county geography to identify those workers who live in one of the thirty in-sample MSAs. Of those workers who appear in one of my sample CBSAs, 95 percent have a block-level geocode. There were no major differences in demographic characteristics between LEHD workers in the sample CBSAs who did and did not have block-level geocodes. The only published evaluation of the quality of the administrative record data is Bye and Judson (2004), who report on the so-called Administrative Records Experiment (AREX) wherein Census evaluated the feasibility of replicating Census enumeration with data from StARS.

C.2 Job Information

I also extract a frame of annualized job histories developed for the Human Capital Estimates Project (HCEP). The job frame records one dominant employer-employee relationship (job) per year. The dominant employer is defined as the employer that reported the most earnings for a given worker over the calendar year. To the dominant job frame, I have included information on employer and job characteristics using the research files from the human capital estimates as a source. These variables include annualized earnings, NAICS 2002 major sector of the employer, imputed annual hours of work (full-time status impute), and the variance components of earnings estimated from the Abowd et al. (1999) human capital decomposition using the grouping and conjugate gradients algorithms described in Abowd et al. (2002).

C.3 Construction of the Analysis Samples

The Master Sample for this paper draws from LEHD all PIKS that meet the following criteria:

1. age in 2002 between 14 and 70 years old;
2. has a valid geocode identifying Census block in 2002;
3. lived in one of the 30 sample MSAs in 2002;
4. was employed in 2002.

The MS thus constructed has 25,689,739 PIK-level observations.

Research File for Annual Job Changers

To construct the research file for annual job changers, I first define the reference population, which is the set of workers that influence any individuals' job search outcomes. I then define the population of interest, which is the set of workers at risk to be observed with a different dominant employer in 2003 than they have in 2002.

Reference Population To define the population of workers in a given job-changer's block-level job information network, I restrict the MS to workers who did not change dominant employer between 2002 and 2003 (this sample contains 20,015,032 PIK-level observations). Importantly, this is the sample from which the local (block-level) referral network quality measure for each worker is computed. Its definition effectively excludes the job changers from construction of the mean of the reference distribution.

Annual Job Holders and Job Changers To construct a sample of annual job changers, I restrict the reference population to those workers employed full-time in both 2002 and 2003. These are the workers at risk for a job-to-job transition from 2002 to 2003. (16,562,870 PIK-level observations). The set of annual job changers is the previous sample restricted to

workers who change dominant jobs between 2002 and 2003. (2,436,970 PIK-level observations)

In the empirical work, I restrict attention to workers who live on blocks with a reference population that includes at least 10 workers who could contribute to estimation of its block-level mean characteristics. This cuts the sample of annual job holders to 15,098,786 workers, and the sample of annual job changers to 2,206,421. As can be seen in the results, an additional handful of observations are lost due to missing demographic information. The final sample of annual job changers used to estimate the models presented in the text contains 2,198,659 workers.

Research File for Quarterly Job Changers

The sample for quarterly job changers is defined differently, and is a little more complex. Starting from the Master Sample, I take the following steps to define the quarterly job changers and the population at risk for a quarterly job change.

1. Restrict MS to workers employed full-time full-year in 2002–2003.
2. Merge quarterly work history for each worker associated with their 2002 and 2003 dominant job.
3. For workers who change dominant job between 2002 and 2003, identify the exact quarter when they began employment with the 2003 dominant job. This is the quarter of transition.

Using this restrictive definition of job changes results in a sample of 899,147 workers.

To define the reference population, in each quarter, I construct a sample of workers in each block who were employed on their dominant job in that year for the full quarter. This means they were employed on the dominant job in the quarter before, the current

quarter, and the quarter after. I construct block-level characteristics in each quarter from this reference population.

In the empirical work, I restrict attention to workers who live on blocks with at least 10 workers in the reference population from which to compute block-level characteristics. The final used for research includes 816,138 quarterly job changers.

References

- Abowd, J. M., Creecy, R. H. and Kramarz, F. (2002). Computing person and firm effects using linked longitudinal employer-employee data, *Technical Report TP-2002-06*, LEHD, U.S. Census Bureau.
- Abowd, J. M., Kramarz, F. and Margolis, D. N. (1999). High wage workers and high wage firms, *Econometrica* **67**(2): 251–333.
- Abowd, J. M., Lengermann, P. and McKinney, K. L. (2003). The measurement of human capital in the U.S. economy, *Technical Report TP-2002-09*, LEHD U.S. Census Bureau.
- Abowd, J. M., McKinney, K. L. and Schmutte, I. M. (2010). How important is endogenous mobility for measuring employer and employee heterogeneity? Cornell University School of Industrial and Labor Relations.
- Abowd, J. M., Stephens, B. E., Vilhuber, L., Andersson, F., McKinney, K. L., Roemer, M. and Woodcock, S. (2009). The LEHD infrastructure files and the creation of the Quarterly Workforce Indicators, *in* T. Dunne, J. B. Jensen and M. J. Roberts (eds), *Producer Dynamics: New Evidence from Micro Data*, Vol. 68 of *Studies in Income and Wealth*, The University of Chicago Press, chapter 5, pp. 149–234.
- Altonji, J. G., Elder, T. E. and Taber, C. R. (2005). Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools, *Journal of Political Economy* **113**(1): 151–184.
- Andersson, F., Burgess, S. and Lane, J. (2008). Recent immigrants, social networks and labor markets. October 19, 2008.
- Bagnoli, M. and Bergstrom, T. (2005). Log-concave probability and its applications, *Economic Theory* **26**(2): 445–469.

- Bandiera, O., Barankay, I. and Rasul, I. (2009). Social connections and incentives in the workplace: Evidence from personnel data, *Econometrica* **77**(4): 1047–1094.
- Bayer, P., Ross, S. L. and Topa, G. (2008). Place of work and place of residence: Informal hiring networks and labor market outcomes, *Journal of Political Economy* **116**: 1150–1194.
- Beaman, L. and Magruder, J. (2011). Who gets the job referral? Evidence from a social networks experiment, *American Economic Review* .
- Beaman, L. and Magruder, J. (2012). Social networks and the dynamics of labor market outcomes: Evidence from refugees resettled in the U.S., *Review of Economic Studies* .
- Bewley, T. F. (1999). *Why Wages Don't Fall During a Recession*, Harvard University Press.
- Bjelland, M., Fallick, B., Haltiwanger, J. and McInterfer, E. (2008). Employer-to-employer flows in the united states: Estimates using linked employer-employee data, *Technical Report Working Paper 13867*, National Bureau of Economic Research.
- Blume, L. E., Brock, W. A., Durlauf, S. N. and Ioannides, Y. M. (2011). Identification of social interactions, Vol. 1 of *Handbook of Social Economics*, North-Holland, pp. 853 – 964.
- Brock, W. A. and Durlauf, S. N. (2001). Interactions-based models, *Handbook of Econometrics*, Elsevier.
- Bye, B. V. and Judson, D. H. (2004). Census 2000 testing, experimentation, and evaluation program synthesis report, results from the Administrative Records Experiment in 2000, *Technical Report TR-16*, U.S. Census Bureau, Washington, DC 20233.
- Cahuc, P. and Fontaine, F. (2002). On the efficiency of job search with social networks, *Technical Report 583*, IZA.

- Calvo-Armengol, A. and Jackson, M. O. (2007). Networks in labor markets: Wage and employment dynamics and inequality, *Journal of Economic Theory* **132**(1): 27–46.
- Calvo-Armengol, A. and Zenou, Y. (2005). Job-matching, social networks and word-of-mouth communication, *Journal of Urban Economics* **57**(3): 500–522.
- Cameron, C. A., Gelbach, J. B. and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors, *Review of Economics and Statistics* **90**: 414–427.
- Case, A. E. and Katz, L. F. (1991). The company you keep: The effects of family and neighborhood on disadvantaged youths, *Technical report*, National Bureau of Economic Research.
- Cingano, F. and Rosolia, A. (2012). People I know: Job search and social networks, *Journal of Labor Economics* **30**(2): 291 – 332.
- Combes, P.-P., Duranton, G. and Gobillon, L. (2008). Spatial wage disparities: Sorting matters!, *Journal of Urban Economics* **63**: 723–742.
- Conley, T. G. and Topa, G. (2002). Socio-economic distance and spatial patterns in unemployment, *Journal of Applied Econometrics* **17**: 303–327.
- Conley, T. G. and Topa, G. (2007). Estimating dynamic local interactions models, *Journal of Econometrics* **140**: 282–303.
- Conley, T. G. and Udry, C. R. (2010). Learning about a new technology: Pineapple in Ghana, *American Economic Review* **100**: 35–69.
- Damm, A. P. (2009). Ethnic enclaves and immigrant labor market outcomes: Quasi-experimental evidence, *Journal of Labor Economics* **27**(2): pp. 281–314.

- Dolnick, S. (2011). Many immigrant's job search starts in chinatown, *New York Times* p. A19.
- Dustmann, C., Gritz, A. and Schnberg, U. (2011). Referral-based job search networks, *IZA Discussion Papers 5777*, Institute for the Study of Labor (IZA).
- Eeckhout, J. and Kircher, P. (2011). Identifying sorting-in theory, *Review of Economic Studies* **78**(3): 872–906.
- Elliott, J. R. (2001). Referral hiring and ethnically homogeneous jobs: How prevalent is the connection and for whom?, *Social Science Research* **30**(3): 401 – 425.
- Firpo, S., Fortin, N. M. and Lemieux, T. (2009). Unconditional quantile regressions, *Econometrica* **77**(3): 953–973.
- Flinn, C. J. and Heckman, J. J. (1983). Are unemployment and out of the labor force behaviorally distinct labor force states?, *Journal of Labor Economics* **1**(1): 28–42.
- Fontaine, F. (2008). Why are similar workers paid differently? the role of social networks, *Journal of Economic Dynamics and Control* **10**: 1016.
- Goel, D. and Lang, K. (2009). The role of social ties in the job search of recent immigrants, *CLSRN Working Papers clsrn_admin-2009-12*, UBC Department of Economics.
- Gruetter, M. and Lalive, R. (2009). The importance of firms in wage determination, *Labour Economics* **16**: 149–160.
- Hellerstein, J. K., McInerney, M. and Neumark, D. (2011). Neighbors and coworkers: The importance of residential labor market networks, *Journal of Labor Economics* **29**(4): 659 – 695.

- Holzer, H. (1988). Search method use by unemployed youth, *Journal of Labor Economics* **6**: 1–20.
- Ioannides, Y. M. and Loury, L. D. (2004). Job information networks, neighborhood effects, and inequality, *Journal of Economic Literature* **42**: 1056–1093. Priority: normal.
- Iranzo, S., Schivardi, F. and Tosetti, E. (2008). Skill dispersion and firm productivity: An analysis with matched employer-employee data, *Journal of Labor Economics* **26**: 247–285.
- Kinsler, J. (2012). Assessing Rothstein’s critique of teacher value-added models, *Quantitative Economics* **3**(2): 333–362.
- Krauth, B. (2011). Bounding a linear causal effect using relative correlation restrictions, *Discussion Papers DP11-02*, Department of Economics, Simon Fraser University.
- Kugler, A. D. (2003). Employee referrals and efficiency wages, *Labour Economics* **10**: 531–556. Secondary.
- Laschever, R. (2009). The doughboys network: Social interactions and the employment of World War I veterans. Purdue University Department of Economics.
- Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem, *Review of Economic Studies* **60**: 531–542.
- Moffitt, R. A. (2000). Policy interventions, low-level equilibria and social interactions, *Social Dynamics*, MIT Press, pp. 45–82.
- Mok, D., Wellman, B. and Carrasco, J. A. (2009). Does distance still matter in the age of the Internet?’, *Urban Studies* **46**.
- Mortensen, D. T. (2003). *Wage Dispersion: Why Are Similar Workers Paid Differently*, Zeuthen lecture book series, MIT Press.

- Mortensen, D. and Vishwanath, T. (1994). Personal contacts and earnings: It is who you know!, *Labour Economics* **1**: 187–201.
- Nagypál, E. (2005). On the extent of job-to-job transitions. September 2005.
- Neal, D. (1999). The complexity of job mobility among young men, *Journal of Labor Economics* **17**(2): 237–61.
- Nordström Skans, O. and Kramarz, F. (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, P. and Schaefer, S. (2011). Firm/Employee matching: An industry study of american lawyers. November 2010.
- Postel-Vinay, F. and Robin, J.-M. (2002). Wage dispersion with worker and employer heterogeneity, *Econometrica* **70**: 2295–350.
- Rogerson, R., Shimer, R. and Wright, R. (2005). Search-theoretic models of the labor market: A survey, *Journal of Economic Literature* **43**(4): 959–988.
- Rothstein, J. (2010). Teacher quality in educational production: Tracking, decay and student achievement, *Quarterly Journal of Economics* **125**(1): 175–214.
- Shimer, R. (2005). The assignment of workers to jobs in an economy with coordination frictions, *Journal of Political Economy* **113**: 996.
- Simon, C. J. and Warner, J. T. (1992). Matchmaker, matchmaker: The effect of old boy networks on job match quality, earnings, and tenure, *Journal of Labor Economics* **10**(3): 306–30.

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

Wahba, J. and Zenou, Y. (2005). Density, social networks and job search methods: Theory and application to egypt, *Journal of Development Economics* **78**: 443–473.

Woodcock, S. D. (2008). Wage differentials in the presence of unobserved worker, firm, and match heterogeneity, *Labour Economics* **15**(4): 771–793.

D 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	

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	Quarterly Job Changers	Quarterly Job Changers*	Annual Job Changers*
White	0.6572	0.6280	0.6220	0.6498
Black	0.1151	0.1220	0.1205	0.1128
Hispanic Origin	0.1167	0.1369	0.1400	0.1272
Male	0.5098	0.4985	0.4979	0.5888
Born in U.S.	0.8098	0.8098	0.8026	0.8148
Age in 2002	40.5456	35.05848	34.9561	37.1021
N	25,689,739	899,147	816,138	2,206,421

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. ‘Quarterly Job Changers’ are from a sample where I identify 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. Details in the text. The summaries in Columns 3 and 4 are for job changers who lived on blocks where at least 10 other workers contribute data to compute the block-level average ψ .

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

Variable	Raw	Block Group Controls
White	.2915	.0132
Hispanic	.2859	.0125
Born U.S.	.2245	.0114
Age	.0301	.0067
ε	.0038	.0002

$N = 394,305$

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 IV: Estimates of Direct Referral Effects: With and Without Covariates.

	Variable	No Covariates		Pair Covariates	
		(1)		(2)	
		Coeff.	t-Stat.	Coeff.	t-Stat.
Reside on same block	R	0.02	27.81	0.013	36.10
Wage premium of non-changer	$R \times \psi_{ref}$			0.014	63.78
Wage effect of non-changer	$R \times \theta_{ref}$			0.023	46.73
Wage effect of changer	$R \times \theta$			0.018	33.47
Block Group Effects		YES		YES	
Sample Size		1, 558, 436, 893			

The table reports results for two regressions in which an observation is a pair of adults employed between 2002–2003 who live in the same Census block group. The pairs are drawn from 30 large CBSAs described in the text. The first worker in the pair (changer) is observed to change jobs between 2002–2003. The second worker in the pair (non-changer) did not. The sample includes all such pairs within each block group, restricted to block groups with at least ten non-changers. The coefficients have been multiplied by 100 to reflect percentage point changes. Block group heterogeneity is estimated by fixed-effects in both models. In addition to the variables listed, specification (2) also includes controls for race, ethnicity, age, AKM person effect (θ) and AKM firm effect (ψ) for both the changer and non-changer members of the pair.

Table V: 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 VI: Main Offer Function Estimates

Premium on next job, ψ	Quarterly Job Changers			Yearly Job Changers		
	Baseline (1)	(2)	(3)	(4)	(5)	(6)
Initial ψ : ψ_0 (β)	0.46 (.008)	0.45 (.008)	0.45 (.008)	0.45 (.032)	0.45 (.032)	0.37 (.024)
Avg. ψ in block: $\bar{\psi}_{block}$ (γ)	0.33 (.016)	0.10 (.011)	0.10 (.012)	0.11 (.007)	0.09 (.006)	0.08 (.005)
Avg. ψ in block group: $\bar{\psi}_{bg}$ (ϕ)		0.34 (.024)				
white	0.001 (.001)	0.001 (.002)	-0.002 (.001)	-.01 (.005)	-.01 (.004)	-.01 (.005)
Hispanic Origin	-0.02 (0.005)	-0.01 (0.005)	-0.02 (.004)	-.03 (.006)	-.02 (.005)	-.02 (.005)
male	0.03 (0.004)	0.03 (0.004)	0.04 (0.004)	.03 (.003)	.03 (.003)	.04 (.003)
age in 2002	0.01 (.001)	0.01 (.001)	0.01 (0.001)	.01 (.001)	0.01 (.001)	.01 (.001)
Square of age in 2002	-0.00 (.000)	-0.00 (.000)	-0.00 (.000)	-0.00 (.000)	-.000 (.000)	-.000 (.000)
Born in U.S.	0.00 (.003)	0.01 (.002)	0.01 (.003)	.01 (.005)	.01 (.005)	.01 (.004)
θ from wage eqn.	-0.00 (.009)	-0.01 (.010)	-0.01 (.010)	-.00 (.007)	-.01 (.007)	-.01 (.008)
block group controls	no	no	yes	yes	yes	yes
block mean characteristics	no	no	no	no	yes	yes
Industry of origin job	no	no	no	no	no	yes
N	815, 899			2, 198, 659		
R^2	.3149	.3175	.2711	.3050	.3053	.3225

Estimates of the log wage premium, ψ , for quarterly and annual job changers. Standard errors are clustered on 30 MSAs.

Table VII: 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 VI. ρ is the estimated correlation between the errors in the selection equation and the offer function.

Table VIII: Offer Function 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 (.024)	0.37 (.032)	.37 (.024)	0.37 (.032)
Avg. ψ in block: $\bar{\psi}_{block}$ (γ)	0.08 .005	0.16 (.006)	.07 (.006)	0.09 (.006)
Born in U.S. $\times \bar{\psi}_{block}$		-.09 (.020)		
Younger Worker $\times \bar{\psi}_{block}$.04 (.010)	
Older Worker $\times \bar{\psi}_{block}$				-.04 (0.02)
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	.3225	.3226	.3228	.3226


Estimates of the log wage premium, ψ , for workers in the sample of annual job changers. 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, as well as the major NAICS sector of the origin job. ‘Younger workers’ are those between 25–35 years of age in 2002. ‘Older workers’ are those between 45–55 years of age in 2002. For reference, the baseline model presented in column (1) is the same estimate presented in column (6) of Table VI, and is included here for ease of presentation. Standard errors are clustered on 30 MSAs.

Table IX: Offer Function Estimates: The Influence of Direct Referrals

Premium on next job, ψ	Baseline			
	(1)	(2)	(3)	(4)
Initial ψ : ψ_0 (β)	0.37 (.024)	0.37 (.024)	.37 (.024)	0.37 (.024)
Avg. ψ in block: $\bar{\psi}_{block}$ (γ)	0.08 .005		.08 (.005)	0.08 (.005)
Move to same job (as a block-neighbor)		-.03 (.010)	-.03 (.010)	-.07 (.016)
Move to same job $\times \bar{\psi}_{block}$.18 (.051)
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	.3225	.3228	.3230	.3232

Estimates of the log wage premium, ψ , for workers in the sample of annual job changers. 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, as well as the major NAICS sector of the origin job. For reference, the baseline model presented in column (1) is the same estimate presented in column (6) of Table VI, and is included here for ease of presentation. Standard errors are clustered on 30 MSAs.

(a)
Spatial Autocorrelation Function: tract-level means



(b)
Spatial Autocorrelation Function: block-level means





Figure 1: Spatial Autocorrelation Function Estimates

(a)
Frequency Distribution



(b)
Cumulative Frequency Distribution




Figure 2: Empirical transition rates between deciles of the employer wage premium (ψ) distribution

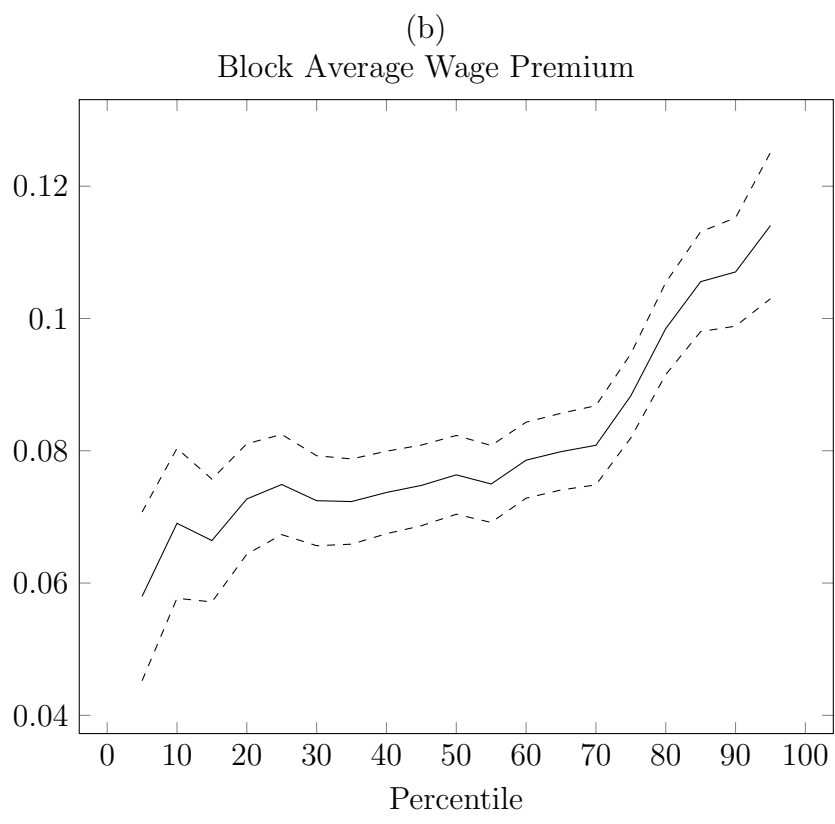



Figure 3: Unconditional Quantile Regression Estimates