Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes

Author(s): Patrick Bayer, Stephen L. Ross and Giorgio Topa
Source: *Journal of Political Economy*, Vol. 116, No. 6 (December 2008), pp. 1150-1196
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/10.1086/595975
Accessed: 12-05-2016 01:54 UTC

Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes

Patrick Bayer

Duke University and National Bureau of Economic Research

Stephen L. Ross

University of Connecticut

Giorgio Topa

Federal Reserve Bank of New York

We use a novel research design to empirically detect the effect of social interactions on labor market outcomes. Using Census data on residential and employment locations, we examine whether individuals residing in the same city block are more likely to work together than those in nearby blocks. We find evidence of significant social

We are grateful for helpful suggestions and comments from Joe Altonji, Pat Bajari, Ed Glaeser, Kevin Lang, Rob McMillan, David Neumark, Stuart Rosenthal, Wilbert van der Klaauw, Ken Wolpin, the editor, and two anonymous referees and seminar participants at the American Economic Association meeting, Boston College, Brown, Columbia, Cornell, Econometric Society, Miami, Federal Reserve Bank of New York, New York University, Southern Methodist, Stanford, and Yale. Shihe Fu, Anupam Nanda, and Chris Huckfeldt have provided excellent research assistance. We are grateful to the Department of Housing and Urban Development, the Federal Reserve Bank of New York, and the Center for Real Estate and Urban Economic Studies at the University of Connecticut for financial support. The research in this paper was conducted while we were Special Sworn Status researchers of the U.S. Census Bureau at the Boston Census Research Data Center. Research results and conclusions expressed are those of the authors and do not necessarily reflect the views of the Census Bureau. This paper has been screened to ensure that no confidential data are revealed. The views and opinions offered in this paper do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve System, the U.S. Department of Housing and Urban Development, or any other agency of the U.S. government.

[*Journal of Political Economy*, 2008, vol. 116, no. 6]

interactions. The estimated referral effect is stronger when individuals are similar in sociodemographic characteristics. These findings are robust across specifications intended to address sorting and reverse causation. Further, the increased availability of neighborhood referrals has a significant impact on a wide range of labor market outcomes.

I. Introduction

The relevance of social networks and local interactions for economic outcomes has been increasingly recognized by economists in a variety of contexts.[1] An important strand of this literature has focused on the detection and measurement of social interactions that operate at the level of the residential neighborhood.[2] The proper identification of such neighborhood effects is complicated, however, by the nonrandom sorting of households into neighborhoods and the likely presence of unobserved individual and neighborhood attributes.[3] The resulting correlation in unobservables among neighbors can lead to serious bias in the estimation of social effects in the absence of a research design capable of distinguishing social interactions from these alternative explanations.

In this paper, we propose a new empirical strategy for identifying neighborhood effects that is based on isolating block-level variation in the characteristics of neighbors within narrowly defined neighborhood reference groups.[4] In particular, using Census data that detail the city block on which each individual in the Boston metropolitan area resides, we compare outcomes for neighbors who reside on the same versus nearby blocks. The key identifying assumption underlying this design (which is testable on observable attributes) is that there is no block-

[1] Some recent examples include crime (Glaeser, Sacerdote, and Scheinkman 1996; Bayer, Pintoff, and Pozen 2008); welfare program participation (Bertrand, Luttmer, and Mullainathan 2000); the adoption of new technologies (Burke, Fournier, and Prasad 2004; Conley and Udry 2005; Bandiera and Rasul 2006); peer effects in education (Hoxby 2000; Sacerdote 2001; Zax and Rees 2002; Zimmerman 2003); and knowledge spillovers and economies of agglomeration (Glaeser et al. 1992; Jaffe, Trajtenberg, and Henderson 1993; Audretsch and Feldman 1996). For a more extensive review of the literature, both theoretical and empirical, see Brock and Durlauf (2001).

[2] Case and Katz (1991) explore the role of neighborhood effects on several behavioral outcomes using a spatially autoregressive model. Crane (1991) also looks at neighborhood influences on social pathologies, focusing on nonlinearities and threshold effects. See Durlauf (2004) for a recent review of the literature and Jencks and Mayer (1990) for a survey of the older literature on neighborhood effects.

[3] See Manski (1993) and Moffitt (2001) for a general discussion of the identification of social interactions in the presence of correlated unobservables.

[4] As we explain below, we consider two different definitions of a reference group of nearby blocks: the census block group and the 10 nearest blocks. The census block group is a geographic area defined by the Census Bureau that represents the next level of geographic aggregation from an individual city block.

level correlation in unobserved attributes among block residents, after taking into account the broader neighborhood reference group.

We use this approach to study the impact of neighborhood referrals on labor market outcomes. Rather than focusing on more general forms of neighborhood effects, we exploit the fact that our restricted Census data set characterizes the precise location of both an individual's place of residence and place of work to study the propensity of neighbors to work together. Specifically, we examine the propensity of a pair of individuals to work in the same location, comparing such propensities for pairs of individuals who reside on the same versus nearby blocks. We take the propensity to work in the same location as an indication that one member of the pair provided a referral (or more generally information) to the other member about jobs available in her place of work.

Our results indicate the existence of significant social interactions at the block level; on the basis of our most conservative estimates, residing on the same versus nearby blocks increases the probability of working together by over 33 percent. As a consequence, individuals are about 6.9 percentage points more likely to work with at least one person from their block of residence than they would be in the absence of referrals. This result is robust across various specifications intended to address the possibility of sorting into specific blocks within neighborhoods and reverse causation (i.e., the idea that referrals may flow in the opposite direction, from friends and acquaintances in the workplace to residential opportunities).

Our identification strategy relies crucially on the absence of correlation in unobserved traits at the block level within a neighborhood reference group. We conduct a number of exercises that suggest that bias arising from sorting within these reference groups is minimal. First, we examine mobility rates at the block level and find that the housing market is quite thin at low levels of geography. Maybe most convincingly, we examine the degree to which the observed block-level sorting on these attributes would alter the likelihood that individuals on the same versus nearby blocks work together. Remarkably, we find that block-level sorting on observables would actually predict a slight reduction in the propensity to work together, suggesting that sorting at this level does not appear to be directly related to employment outcomes. Finally, we show that our findings are robust to the inclusion of individual fixed effects. This specification controls for a form of sorting on unobserved attributes and thus gives us even greater confidence in our findings.

Our analysis also indicates that there is considerable variation in the likelihood of referrals across different types of worker pairs. We estimate, for example, that a referral is significantly more likely among pairs of high school graduates, pairs of young adults, and pairs in which members have children of a similar age. This analysis of heterogeneous referral effects serves a second purpose in our analysis. In particular, it allows us to develop an individual-specific measure of the availability of referral opportunities on each block in the metropolitan area. The resulting estimate of match quality provides a novel measure of neighborhood quality based on the specific match between an individual's characteristics and those of her neighbors. We include this measure in a series of standard regressions for labor force participation, employment, hours and weeks worked, wages, and earnings. The results of this portion of our analysis reveal that neighborhood referral effects tend to have a (statistically and economically) significant positive impact on several labor market outcomes under consideration; a one-standard-deviation increase in the match quality, for example, raises expected hours worked per week by 1.8 hours and earnings by 3.4 percent for the average male individual in our preferred specification. For females, the earnings effect is weak, but expected labor force participation increases by about 3.4 percentage points.

In addition to providing new evidence on the importance of neighborhood referrals for labor market outcomes, our analysis also demonstrates the potential strengths of the general research design that we introduce in this paper. In a manner that deals directly with the correlation of individual and neighbor characteristics (e.g., due to sorting), this design allows for the identification of neighborhood effects operating (i) through a specific mechanism, (ii) for a broad population and a wide variety of subsets of that population, and (iii) for individuals who have resided in a neighborhood for a variety of tenure lengths. The applicability of this design extends to the study of neighborhood effects in other spatial contexts (e.g., other metro areas, specific types of neighborhoods), on specific populations (e.g., youths), and for alternative outcomes (e.g., education, health, welfare participation, bankruptcy, and home foreclosures), provided that the researcher can demonstrate that the within–reference group correlation in observable neighbor characteristics does not contribute significantly to outcomes, thereby ensuring that the key identifying assumption on unobserved characteristics is at least plausible.

The remainder of the paper is organized as follows. Section II places this paper in the context of the broader literature on neighborhood effects and employment referrals. Section III describes the data set that we have assembled for the Boston metropolitan area. Section IV describes our research design and presents evidence concerning the orthogonality of the block-level variation in individual and neighbor characteristics. We also discuss several extensions of our methodology designed to deal with additional issues related to identification. We report our empirical findings in Sections V and VI and conclusions in Section VII.

II. Previous Literature

This paper is situated in a broader literature that aims to identify neighborhood effects. An important line of research in this literature relies on a random component of neighborhood choice induced by special social experiments. Popkin, Rosenbaum, and Meaden (1993) pioneered this approach using data from the Gautreaux Program conducted in Chicago in the late 1970s, which gave housing vouchers to eligible black families in public housing as part of a court-imposed public housing desegregation effort.[5] Most notably, Katz, Kling, and Liebman (2001) and Ludwig, Duncan, and Hirschfield (2001) have used the randomized housing voucher allocation associated with the Moving to Opportunity (MTO) demonstration to examine the impact of relocation to neighborhoods with much lower poverty rates on a very wide set of individual behavioral outcomes including health, labor market activity, crime, education, and more. Especially in the case of MTO, the advantages of this approach are clear: the randomization inherent in the program design in principle ensures a clean comparison of treatment and proper control groups.

There are, however, important limitations in the extent to which the treatment effects identified through relocation are informative about the nature of general forms of neighborhood effects per se. First, individuals studied must be eligible for a relocation program in the first place; this typically implies that the resulting sample is special (i.e., so as to be a resident in public housing) and may not be as sensitive to neighborhood effects as other individuals. More generally, even if the eligible population is representative of the target population, the results of an experiment based on a small sample may not extend to broader populations because of the strong possibility that general equilibrium effects may arise in that case. Second, the experimental design involves relocation to new neighborhoods that are, by design, very different from baseline neighborhoods; this implies that the identified treatment effect measures the impact of relocating to a neighborhood in which individuals initially have few social contacts and in which the individuals studied may be very different from the average resident of the new neighborhood. In this way, the treatment effects identified with this design are necessarily a composite of several factors related to significant changes in neighborhoods that are not easily disentangled.[6]

A second broad approach seeks to deal with the difficulties induced

[5] Similarly, Oreopoulos (2003) and Jacob (2004) study the impact of relocations arising from administrative assignment to public housing projects in Toronto and from the demolition of the public housing projects in Chicago, respectively.

[6] Moffitt (2001) and Sobel (2006) present detailed discussions of the potential pitfalls of using randomized experiments in the study of neighborhood effects.

by correlation in unobserved attributes at the neighborhood level by aggregating to a higher level of geography. Evans, Oates, and Schwab (1992), Cutler and Glaeser (1997), Ross (1998), Weinberg (2000, 2004), Card and Rothstein (2007), and Ross and Zenou (2008) identify the effect of location on outcomes using cross-metropolitan variation. For example, Cutler and Glaeser analyze the impact of segregation within a metropolitan area on a variety of outcomes including education, labor market activity, and teenage fertility. Evans et al. use metropolitan area poverty rates as an instrument for neighborhood-level poverty. Again, the advantages of this approach are clear: aggregation certainly eliminates the problem of correlation in unobservables among neighbors (although potential correlation in unobservables at the metropolitan level becomes an issue). The effects identified through aggregation, however, include not only the average neighborhood effects operating in a metropolitan area but also any broader consequences of living in a segregated or high-poverty metropolitan area.[7] Thus, the strict interpretation of the estimated effects as neighborhood effects requires the assumption that metropolitan segregation does not directly affect outcomes.[8]

The research design developed in this paper can be viewed as the converse of designs based on cross–metropolitan area variation. That is, instead of aggregating to the metropolitan level, we disaggregate below the level of the neighborhood to isolate block-level variation in neighbor attributes. While the strict identification of neighborhood effects with the cross–metropolitan area design requires the assumptions of no metropolitan effects and no correlation in unobservables at the metropolitan level, strict identification with our design requires the assumptions that social interactions among neighbors are very local in nature—operating at the level of the block—and that there is no correlation in unobservables across blocks within reference groups.[9] Thus, we offer a complementary approach to the existing literature that allows researchers to identify a wide range of causal neighborhood effects using an alternative set of assumptions (testable on the observables) than have been used in previous studies.[10]

[7] More residentially segregated metropolitan areas might be associated, e.g., with increased racial taste-based discrimination in the labor market, in the application of criminal justice, etc. as a result of decreased levels of regular interracial contact in residential neighborhoods.

[8] It is important to point out that Cutler and Glaeser (1997) do not claim that the effects identified in their analysis are strictly neighborhood effects.

[9] We provide some evidence regarding the very local nature of social interactions in Sec. IV.

[10] Only one contemporaneous study, Grinblatt, Keloharju, and Ikaheimo (2004) on automobile consumption, has used variation arising from location in very local neighborhoods as a source of identification. This study assumes that the composition of an individual's 10 closest neighbors is exogenous after conditioning on a neighborhood made

Our paper also contributes to a vast literature on both neighborhood and referral effects in the labor market (see Ioannides and Loury [2004] for an excellent survey of this literature). For instance, Weinberg, Reagan, and Yankow (2004) use longitudinal data from the National Longitudinal Survey of Youth to study the impact of neighborhood quality on employment outcomes. Rees and Shultz (1970), Corcoran, Datcher, and Duncan (1980), Holzer (1988), Blau and Robbins (1990), Blau (1992), Granovetter (1995), Addison and Portugal (2002), and Wahba and Zenou (2005) all document the importance of referrals and other informal hiring channels in the labor market, using both U.S. and non-U.S. data.[11] Additional studies including Datcher (1983), Devine and Kiefer (1991), Marmaros and Sacerdote (2002), and Loury (2006) find evidence that use of informal networks increases the quality of the match as captured by job tenure or earnings.[12]

Moreover, considerable evidence exists to suggest that the use and impact of job information networks varies across demographic groups, which is consistent with our own findings. According to Ioannides and Loury (2004), the evidence on usage differences is mixed in general but suggests that younger, lower-educated, and male workers are more likely to use informal job networks. In terms of relative productivity, Bortnick and Ports (1992) find that these networks are slightly less productive for women than for men. Holzer (1987), Bortnick and Ports (1992), and Korenman and Turner (1996) find that such networks are substantially less productive for African Americans.

III. Data

The data for our analysis are drawn from a restricted version of the 1990 U.S. Census of Population for the Boston metropolitan area. For the full (1-in-7) sample of individuals who filled out the long form of the census, these data contain the complete set of variables that are available in the public-use version of the Census Public Use Microdata

up of the 50 closest neighbors. Also see Ioannides and Zabel (2008) for a model of housing demand and neighborhood choice that uses two levels of geography.

[11] The use of informal channels such as referrals by employers can be rationalized as a means to reduce the uncertainty regarding the quality of a prospective employee. Montgomery (1991) was the first to formally model a labor market in which both formal and informal hiring channels coexist. Focusing more closely on the information exchange among workers, Calvo-Armengol and Jackson (2004) analyze an explicit network model of job search in which agents receive random offers and decide whether to use them themselves or pass them on to their unemployed contacts depending on their own employment status and current wage.

[12] See Elliot (1999) and Loury (2006) for counterexamples in which the use of informal networks led to lower wages. Of course, the lower wages may be associated with increased match quality on desirable job attributes causing the individual to accept a lower wage as a compensating differential.

Sample but, in addition, detail each individual's residential and employment locations down to the census block level. In addition to these geographic variables, the Census Bureau also provides a wide range of sociodemographic information: age, gender and marital status, education, race, family structure, and residential tenure as well as information on labor market outcomes including labor force status, weeks and hours worked, and salary and wage income if employed.

With regard to the geographic structure of the data, census blocks correspond roughly to actual city blocks; they are typically rectangular regions delimited by the four intersections that constitute the corners of the block.[13] Our sample consists of approximately 25,500 census blocks arranged into 2,565 census block groups, that is, an average of 10 blocks per block group. The distribution of blocks per block group is depicted in figure 1; the median number of blocks per block group is eight, and about 95 percent of all block groups have 20 blocks or fewer. Census block group is used as our primary definition of a neighborhood.[14] For robustness we also perform our analysis using an alternative reference group for a given block, defined as the set of 10 closest blocks to that block using physical distance between block centroids.

It is the precise geographical information for each individual in these restricted Census data that provides the backbone of our research design, permitting us to isolate the block-level variation in neighbor exposure by conditioning on reference group fixed effects. The first stage of our analysis considers the propensity of a pair of individuals to work in the same location, comparing this propensity for a pair that lives on the same versus nearby blocks. For this portion of our analysis, we construct a sample that contains individuals (i) who are currently employed, (ii) who are between 25 and 59 years of age, (iii) who do not work at home, and (iv) for whom the Census data on place of work have not

[13] Notice that this definition implies that census blocks are not constituted as the set of buildings that face each other on the same street. To the extent that social interactions are also strong between residents on opposite sides of the same street, a comparison of interactions between individuals who reside on the same census block vs. other blocks in the same block group will tend to understate the increased effect of immediate neighbors since those on the opposite side of the same street will count in the control group. For some blocks, however, one may argue that the opposite holds: streets may effectively act as dividers of local communities, and interactions may be strongest in the alleys and courtyards connecting the rear sides of buildings on the same block. In either case, our research design should detect (although may understate) particularly local interactions provided that the block group contains a reasonable number of blocks.

[14] The Census Bureau Statistical Participant Areas Program provides individuals in the local community a major role in the development of census block group and census tract definitions. The participant guidelines explicitly enable participants to draw on local knowledge and consider features that might "unify a community" in order to "better encompass similar community patterns" (U.S. Census Bureau 1997, 9).

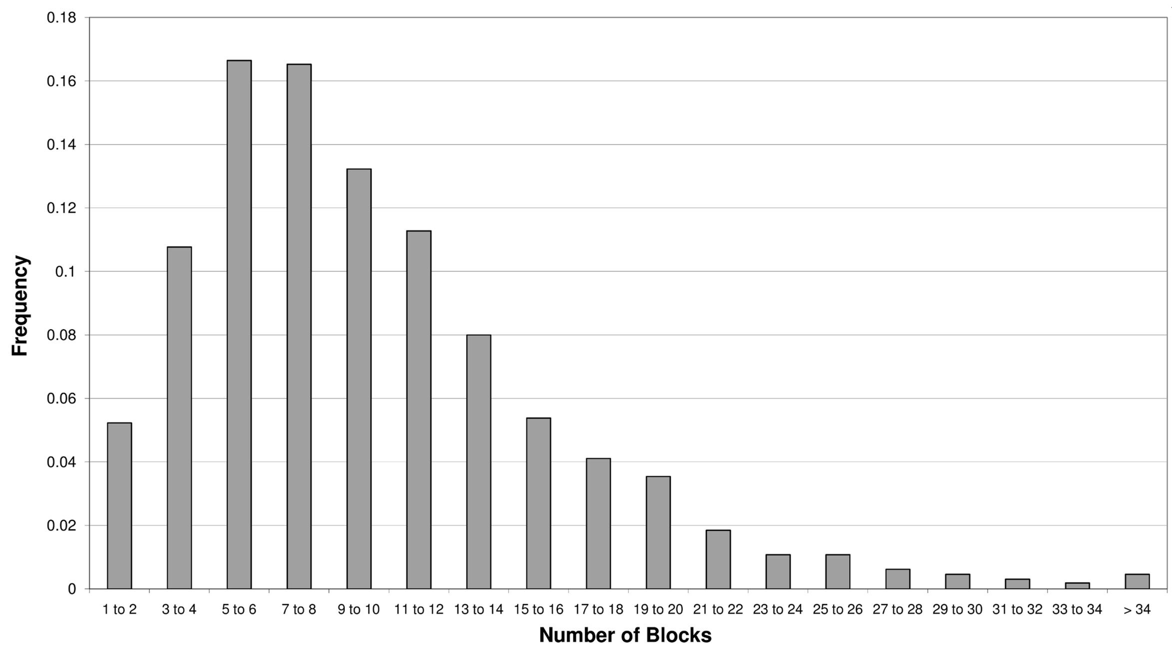

FIG. 1.—Distribution of blocks per block group

been imputed.[15] The total number of workers in the census sample who meet these criteria is 129,175 (5.1 per block, 50 per block group). Figure 2 reports the corresponding histogram of workers meeting these criteria across blocks.[16]

In constructing a sample of pairs for our analysis, we apply two additional criteria, selecting all pairs that (i) reside in the same reference group within the Boston metropolitan area and (ii) do not belong to the same household. Overall, the samples contain 2,037,600 and 2,671,270 pairs that meet all the above criteria for the block group and alternative reference group samples, respectively. Columns 1 and 2 of table 1 characterize these samples of matched pairs, reporting the percentage of pairs that fit the description in the row heading: depending on the definition of reference group, at least one member of roughly 65–72 percent of the pairs has children; about 15–20 percent of pairs match two single individuals.[17]

Examining the characteristics of the samples of pairs shown in table 1 highlights two key dimensions of heterogeneity in which our study will be limited because of the small size of the corresponding sample in the Boston metro area. In particular, (i) less than 1 percent of all pairs reflect a match between two high school dropouts, and (ii) only 0.5–2.6 percent of all pairs reflect a match between two nonwhite workers. Given the nature of the samples, it is not surprising that our analysis tends to be more precise in other dimensions of individual heterogeneity including age, the presence of children, education (aside from high school dropouts), gender, and marital status.

For the second stage of our analysis, which examines the impact of neighborhood characteristics on labor market outcomes including labor force participation and employment, we add to the sample those prime-age (25–59) individuals who are not currently employed; this sample has 163,594 observations.[18] Table 2 reports summary statistics for this sample. Column 1 reports the sample frequencies for each individual characteristic, and the remaining five columns report labor

[15] Currently employed refers to the reference week in the calendar year 1990 used by the census. We focus on prime-age adults in this paper so as to avoid empirical issues related to labor market participation vs. continued schooling of youths and young adults. We drop all individuals for whom place of work is imputed for obvious reasons. Finally, individuals who work at home are deleted because their presence would create a mechanical correlation in our estimates; i.e., a pair of such individuals must by definition work on the same block only if they also live on the same block.

[16] In the analysis below, we consider specifications that limit the analysis to blocks with five or more sample workers. Our results remain stable when we use all blocks.

[17] It should be noted that the sample contains only a small fraction of Asians and Hispanics, and so these two groups are combined. Specifications in which these groups are separated yield very similar results.

[18] We again limit the sample used in each labor market outcome equation to individuals for whom the corresponding dependent variable has not been imputed.

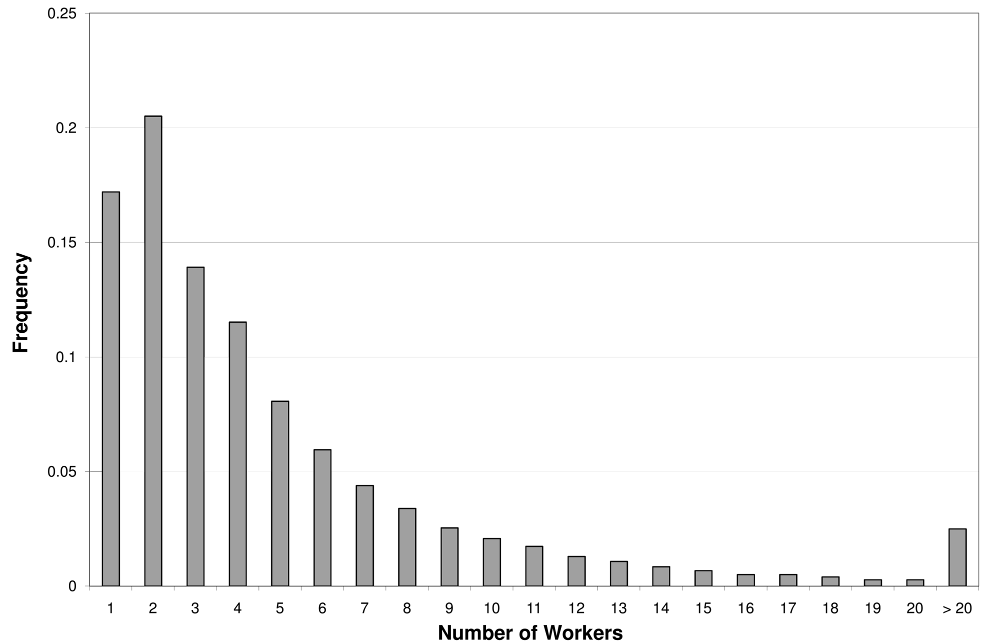

Fig. 2.—Distribution of sampled workers per block

market and commuting information: the fraction of individuals who are currently employed, average weeks worked in the previous year, average hours worked per week in the previous year, average earnings for the sample of individuals who were fully employed in the previous year, and average commute for those who are currently employed.[19] College graduates, married males, and whites display the strongest attachment to the labor force, with respect to employment rates as well as hours and weeks worked. These groups also tend to work the farthest away from home. However, high school dropouts and married females tend to have weak labor force attachment and work close to home when employed.

IV. Empirical Design—Detecting Referral Effects

Given the structure of the data set just described, it is straightforward to characterize our general research design. Our primary analysis explores the propensity for two individuals to work in the same location, comparing this propensity for a pair that lives in the same block with that of a pair that lives in the same reference group but not on the same block. The implementation of this design is straightforward and can be summarized in the following equation:

$$W_{ij}^{b} = \rho_g + \alpha_0 R_{ij}^{b} + \varepsilon_{ij}, \tag{1}$$

where i and j denote two individuals who reside in the same reference group (census block group or alternative reference group) but not in the same household, W_{ij}^{b} is a dummy variable that is equal to one if i and j work in the same census block, R_{ij}^{b} is a dummy variable that is equal to one if i and j reside in the same census block, and ρ_g denotes the residential reference group fixed effect; this is the baseline probability of working in the same block for individuals residing in the same reference group. The statistical test of the null hypothesis that no local

[19] The census provides information on current employment and labor force participation as well as the location of current workplace at the time of the survey in April 1990. Information on earnings, hours, and weeks are reported for the previous year. Fully employed in 1989 refers to any individual who worked at least 45 weeks and at least 30 hours per week.

TABLE 1
COMPOSITION OF PAIRS RESIDING IN SAME NEIGHBORHOOD REFERENCE GROUP

			Percentage Who Work in Same Location		
	Pairs Residing in Same Census Block Group (1)	Pairs Residing in 10 Closest Blocks Reference Group (2)	Reside in Same Census Block Group but Not Same Block (3)	Reside in 10 Closest Blocks but Not Same Block (4)	Reside on Same Block (5)
Full sample	100.00	100.00	.36	.38	.94
Both high school dropouts	.53	.67	1.03	.79	ND
Both high school graduates	15.50	14.84	.47	.50	1.33
Both college graduates	36.41	37.56	.34	.37	.98
High school dropout–high school graduate	4.75	5.12	.51	.49	.82
High school dropout–college graduate	4.95	5.36	.29	.30	ND
High school graduate–college graduate	37.87	36.46	.30	.32	.82
Both age 25–34	14.70	17.89	.36	.37	1.72
Both age 35–44	11.02	10.27	.33	.39	.77
Both age 45–59	9.55	8.17	.42	.44	.65
Age 25–34 and age 45–59	20.01	19.95	.34	.37	.72
Age 35–44 and age 45–59	19.86	17.64	.38	.40	.63
Age 25–34 and age 35–44	23.27	24.63	.33	.34	.91
Both single males	3.01	4.27	.36	.40	.77
Both single females	4.62	6.06	.40	.36	.54
Single male–single female	7.17	9.71	.33	.36	.43
Both married males	14.69	12.10	.35	.39	1.61
Both married females	8.07	7.00	.52	.50	1.78

Married male–married female	21.58	18.21	.29	.30	1.35
Single male–married female	7.87	8.38	.33	.36	.54
Single male–married male	10.12	10.47	.38	.44	.85
Single female–married female	10.03	10.61	.41	.39	.64
Single female–married male	12.84	13.19	.31	.34	.48
Both have children	26.98	22.55	.36	.39	1.58
Both have children age 0–5	3.37	2.88	.33	.37	3.10
Both have children age 6–12	4.12	3.39	.37	.42	2.50
Both have children age 13–17	2.85	2.29	.42	.44	1.57
Both have children age 18–24	3.01	2.62	.48	.51	ND
No children	27.71	34.57	.37	.40	.72
Both white	86.51	82.64	.35	.37	.77
White–black	3.59	4.67	.30	.35	1.34
White–Asian/Hispanic	8.31	10.07	.39	.43	1.69
Both minority	.51	2.62	.47	.43	2.38
Both 2 years in residence or less	3.18	4.13	.35	.45	1.70
One $\leq$ 2 years, one > 2 years in residence	25.60	28.47	.35	.38	1.24
Both > 2 years in residence	71.22	67.40	.36	.37	.73
Both fully employed	74.21	81.78	.37	.37	.84
One not fully employed, one fully employed	23.81	17.26	.31	.38	1.17
Both not fully employed	1.98	.97	.50	.61	2.07
Both own house	54.93	44.50	.34	.35	.44
Both rent	14.52	21.80	.46	.49	2.38
One rents, one owns	30.54	33.69	.35	.34	.42

Note.—The full sample for census block groups includes 2,037,600 pairs of currently employed, prime-age (25–59) adults who reside in the same block group but not in the same household within the Boston metropolitan area in 1990. The full sample for the alternative neighborhood reference groups includes 2,671,270 pairs of currently employed prime-age (25–59) adults who reside in the 10 closest blocks to a given block but not in the same household within the Boston metropolitan area in 1990. For the type of pair denoted in the row heading, the table describes the fraction of such pairs in the full sample and the propensity of such pairs to work together (on the same block) for individuals in the same block group but not in the same block, those in the 10 closest blocks but not in the same block, and those on the same block, respectively. All figures are expressed as percentages. ND indicates that a value was not disclosed because the number of individuals who work in the same block is less than 75.

TABLE 2
SAMPLE OF PRIME-AGE ADULTS IN BOSTON METROPOLITAN AREA

		FULL SAMPLE			SAMPLE FULLY EMPLOYED IN 1989	SAMPLE CURRENTLY EMPLOYED IN 1990
	PERCENTAGE OF SAMPLE (1)	Percentage Currently Employed (1990) (2)	Average Weeks Worked in 1989 (3)	Average Hours per Week in 1989 (4)	Average Earnings in 1989 in $1,000s (5)	Average Commute Distance (6)
Full sample	100.0	75.1	40.5	34.9	34.3	7.0
High school dropout	10.2	55.4	31.5	27.7	22.0	5.3
High school graduate	42.0	71.9	39.4	33.1	26.9	6.5
College graduate	47.8	82.0	43.4	37.9	42.1	7.6
Age 25–34	38.2	74.8	40.8	35.9	28.7	6.9
Age 35–44	31.6	77.0	41.1	35.0	37.3	7.3
Age 45–59	30.2	73.4	39.6	33.3	38.3	6.8

Single male	17.5	73.0	41.7	38.0	31.2	6.4
Single female	20.0	75.1	40.6	34.0	26.5	5.9
Married male	30.6	86.7	46.7	42.9	46.8	8.6
Married female	31.9	65.0	33.9	25.9	24.0	6.2
Has no children	48.1	77.3	42.2	36.9	32.5	6.7
Has children	51.9	73.0	39.0	33.0	36.2	7.3
Has children age 0–5	19.7	68.3	36.9	31.8	37.6	7.9
Has children age 6–12	20.5	71.6	37.6	31.6	37.6	7.2
Has children age 13–17	15.6	75.8	40.0	33.5	37.1	6.9
Has children age 18–24	17.0	74.7	40.4	33.8	33.3	6.7
White	87.9	76.7	41.2	35.3	35.3	7.2
Black	5.1	63.7	37.2	32.4	25.8	5.4
Asian/Hispanic	7.0	63.1	34.9	31.6	26.7	5.5
In residence $\leq$ 2 years	16.7	73.7	39.9	36.2	31.4	6.8
In residence > 2 years	83.3	75.3	40.7	34.6	34.9	7.0
Employed < 45 weeks in 1989	31.0	40.2	16.2	19.6	ND	5.8
Employed $\geq$ 45 weeks in 1989	69.0	90.8	51.5	41.7	34.3	7.2

Note.—The full sample includes 163,594 prime-age (25–59) adults who reside in the Boston metropolitan area in 1990. For the type of individual denoted in the row heading, the table describes the fraction of such individuals in the full sample, the fraction currently employed in 1990, average weeks worked in 1989, average hours per week in 1989, average earnings for those fully employed in 1989, and average commute distance for those currently employed, respectively. For the purposes of examining earnings throughout the paper, fully employed in 1989 refers to any individual who worked at least 45 weeks and at least 30 hours per week; there are 113,575 such individuals for whom earnings are not imputed.

social interaction effect exists is simply a test of whether the estimated coefficient $\alpha_0 = 0$.[20]

The inclusion of reference group fixed effects in equation (1) is designed to control for any correlation in unobserved attributes among individuals residing in the same neighborhood. Such correlation can arise because of positive sorting into neighborhoods or because of unobserved factors present in those neighborhoods, for example, similar access to the urban transportation network (see Manski [1993] and Moffitt [2001] for a detailed discussion of these issues).

In interpreting α_0 as a social interaction effect, therefore, we are implicitly making two key assumptions to achieve identification. First, while individuals are able to choose their residential neighborhood (reference group), there is no correlation in unobserved factors affecting work location among individuals residing on the same block within a reference group. The plausibility of this assumption is motivated by two considerations. First, the thinness of the housing market at such small geographic scales—for instance, the vast majority of block groups in our sample are less than 0.10 square mile in area—restricts an individual's ability to choose a specific block versus a wider neighborhood. Second, it may be difficult for individuals to identify block-by-block variation in neighbor characteristics at the time of purchase or lease. That is, while an individual may have a reasonable sense of the sociodemographic structure of the neighborhood more generally, that variation across blocks within a neighborhood is less easily observed a priori.

That the housing market is relatively thin at any particular point in time at the block level is supported by an analysis of mobility rates in the census. In our sample, only 11 percent of the blocks have an owner-occupied unit that changed owners in the 2 years prior to the census. Given that the census is a 1-in-7 sample and assuming a uniform probability for a house to be on the market in this 2-year period, this implies that the chances that any owner-occupied unit is available on a given

[20] The sampling scheme, which is based on drawing matched pairs of individuals who reside in the same reference group, makes it very difficult to compute appropriate standard errors for our estimates. In fact, suppose that individuals *a* and *b* work in the same block. Suppose further that individuals *b* and *c* work in the same block. Then, by transitivity, individuals *a* and *c* must also work in the same block. As a consequence, if we compute standard errors via the basic ordinary least squares formula, we may tend to understate their size because we are not taking into account this inherent correlation structure in the data. There is also the reasonable concern of heteroskedasticity across reference groups that may bias standard errors in fixed effects analyses. To address these issues, all standard errors are estimated on the basis of pairwise bootstraps. It should be noted that some concerns have been raised concerning pairwise bootstrap in small samples (Horowitz 2001). While our sample is quite large, we have a very small number of ones in our dependent variable, which may create similar problems. We verified the accuracy of the pairwise bootstraps by also estimating standard errors using a pairwise bootstrap with the HC_3 correction and with a wild bootstrap (Mammen 1993; Flachaire 1999, 2005).

block within a given 3-month period is only about 11 percent.[21] Thus, it may be difficult for households searching in a given time frame to select a house on a particular block. Moreover, the fact that households have heterogeneous tastes for particular housing attributes implies that the availability of a suitable house on a given block at any point in time is likely to be much lower.

The second key assumption underlying our research design is that a significant portion of interactions with neighbors are very local in nature, that is, occur among individuals on the same block. A well-established literature in sociology documents the extent to which individual social networks are local in a geographic sense.[22] Most relevantly to our approach, Lee and Campbell (1999) use data from a 1988 survey of Nashville to look at social ties with immediate neighbors. Their definition of "micro-neighborhoods" is similar to ours: they use "partial face blocks . . . . Each site is made up of 10 adjacent housing units, five on either side of the street" (126). They find that 31 percent of these immediate neighbors are judged close or very close by respondents. Further, they specifically ask respondents to whom they would turn for help in finding a job. About 13 percent of helpers in these networks resided in the respondents' micro-neighborhoods; 73 percent resided elsewhere in Nashville; the residual 14 percent were not Nashville residents. To the extent that individuals do have some interaction with neighbors on surrounding blocks rather than on the same block, our design will provide only a lower bound on the overall strength of local interactions—measuring only the difference between these very local and broader effects. In this way, the design will allow us to detect interactions provided that they are significantly stronger at closer distances but may understate the strength of those interactions.

Specification with individual fixed effects.—The analysis of block-level sorting on observable individual attributes presented below suggests that sorting within neighborhood reference groups (census block group or 10 closest blocks) is minimal. To further assess concerns about sorting on the basis of unobserved individual characteristics, we also consider a generalization of equation (1) that includes individual fixed effects

[21] The comparable figure for renter-occupied units for blocks that contain at least one rental unit in our sample is 45 percent. This suggests that it is generally easier, although far from certain, for renters to find housing on a specific block within a particular search window.

[22] In a study of Toronto residents in 1978, Wellman (1996) finds that 42 percent of yearly contacts in individual networks took place with neighbors who lived less than 1 mile away. Guest and Lee (1983) perform a similar analysis for the city of Seattle and find that for about 35 percent of respondents the majority of their nonkin social contacts resided in the same local community. Otani (1999) uses 1986 General Social Survey data for the United States (in a study comparing Japan and the United States) and finds that roughly one in five contacts listed in individual networks are physical neighbors.

for each member of the pair rather than the block group.[23] In particular, since each worker appears multiple times in our sample of pairs, we can estimate

$$W_{ij}^b = \lambda_i + \lambda_j + \alpha_0 R_{ij}^b + \varepsilon_{ij}, \tag{2}$$

where i and j again denote two individuals who reside in the same neighborhood reference group and λ_i and λ_j represent individual fixed effects.[24]

The inclusion of individual fixed effects in equation (2) allows us to deal to some degree with block-level sorting on the basis of unobserved attributes. In particular, if certain workers were more likely to work with others from their neighborhood for unobserved reasons (e.g., because they are employed in jobs very close to home) and these workers tended to sort themselves onto similar blocks within the neighborhood reference group, our baseline analysis would misattribute their increased propensity to work together to the fact that they live on the same block. Of course this bias could just as easily go in the opposite direction if workers who tend not to work in the same place as others from the same neighborhood (e.g., white-collar workers commuting long distances) were more likely to sort onto the same block (perhaps because of the types of homes that they purchase).

As we discuss in more detail below, the inclusion of individual fixed effects has little impact on the overall social interaction effect, α_0, but does in fact change a couple of the estimates from the heterogeneous model that we now present.

Heterogeneous specification.—The initial specifications shown in equations (1) and (2) can easily be extended to include a set of covariates

[23] We would like to thank an anonymous referee for giving us this suggestion.

[24] The individual fixed effects model is estimated using a differencing approach in which 10 individuals k within pair (i, j)'s reference group are matched with each individual in the pair and these matches are mean-differenced over the new i and j pairs, respectively. Specifically,

$$\begin{aligned} D_i^b &= W_{ik}^b - \left(\frac{1}{N_k}\sum_{k'\neq i,j} W_{ik'}^b\right) \\ &= \left(\frac{N_k - 1}{N_k}\right)\lambda_k - \frac{1}{N_k}\sum_{k'\neq i,j,k}\lambda_{k'} + \alpha_0\left[R_{ik}^b - \left(\frac{1}{N_k}\sum_{k'\neq i,j} R_{ik'}^b\right)\right] + \left[\varepsilon_{ik} - \left(\frac{1}{N_k}\sum_{k'\neq i,j}\varepsilon_{ik'}^b\right)\right], \end{aligned}$$

where N_k is the number of individuals k. The differencing of D_j^b from D_i^b eliminates the fixed effects associated with each individual k, as well as the fixed effects associated with the k' individuals. See Arcidiacono et al. (2007) for methods to estimate paired fixed effects of this type in nonlinear models, which are tractable for sample sizes that are smaller than the samples used in this paper.

X_{ij} that describe the pair of individuals (e.g., those summarized in table 1) both in levels and interacted with R^b_{ij}:

$$W^b_{ij} = \rho_g + \beta' X_{ij} + (\alpha_0 + \alpha'_1 X_{ij}) R^b_{ij} + \varepsilon_{ij} \tag{3}$$

and

$$W^b_{ij} = \lambda_i + \lambda_j + \beta' X_{ij} + (\alpha_0 + \alpha'_1 X_{ij}) R^b_{ij} + \varepsilon_{ij}. \tag{4}$$

In this case, the estimated coefficients on the cross terms, α_1, allow us to investigate whether the social interaction effect is weaker or stronger for specific sociodemographic characteristics of the matched pair. There are two aspects to this: first, certain pairs are more likely to interact because of the assortative matching present in social networks: for instance, two individuals of similar age, education, or race or with children of similar age.[25] Second, certain individuals may be more strongly attached to the labor market and may thus provide better referrals or information on jobs, for example, college graduates, married males, or individuals with children. In this case, matches between pairs in which one individual is strongly attached to the labor market and the other is generally more likely to need a referral should also lead to an increased number of referrals.

Examining block-level sorting.—To study whether our first key assumption—that there is no correlation in unobserved factors affecting work location among individuals residing on the same block within a reference group—is reasonable, we analyze the correlation between observable individual and neighbor characteristics at the block level. While this kind of analysis does not prove anything with respect to the importance of potential correlation in unobserved factors, it provides an indication of whether this assumption is at all reasonable.[26]

For each block in the sample, a single prime-age adult is selected randomly, and the characteristics of other individuals who reside in the same block but not the same household are used to construct a measure

[25] See Marsden (1987, 1988) for a discussion of the evidence from the General Social Survey on assortative matching in networks.

[26] This is in the same vein of Altonji, Elder, and Taber (2005): their approach to correcting for selection bias suggests that selectivity in terms of unobserved heterogeneity is in some sense proportional to selectivity on observables.

TABLE 3
Correlation between Individual and Average Characteristics of Neighbors Residing on Same Block

	Sample: Blocks with Five+ Workers in Sample		
	Unconditional (1)	Conditional on Census Block Group (2)	Conditional on 10 Closest Blocks Reference Group (3)
High school graduate	.182	.040	.021
College graduate	.294	.060	.030
Age 45–59	.051	.008	−.020
Age 35–44	.017	−.004	−.031
Age 25–34	.098	.027	−.005
Single female	.110	.033	.014
Single male	.094	.027	.004
Married female	.080	.005	−.015
Married male	.088	.026	.011
Children	.142	.046	.006
Children 0–5	.046	.019	−.007
Children 6–12	.058	.017	−.017
Children 13–17	.048	.015	−.025
Children 18–24	.064	.022	−.014
Black	.593	.054	.017
Asian/Hispanic	.275	.084	.049

Note.—The table reports unbiased estimates of correlation between a series of individual characteristics and the corresponding average characteristics of other individuals residing on the same block but not in the same household. Blocks with fewer than five workers have been dropped from this sample. Column 1 reports unconditional correlation, col. 2 conditions on block group fixed effects, and col. 3 conditions on fixed effects for neighborhood reference groups based on the 10 closest blocks to each block.

of average neighbor characteristics.[27] Table 3 reports the average correlations for our baseline sample of blocks with at least five workers:[28] column 1 reports unconditional correlations, column 2 conditions on block group fixed effects, and column 3 conditions on the alternative

[27] By sampling only one individual per block, we avoid inducing a mechanical negative correlation that would come about if all individuals were used in estimating the correlation between individual and average neighbor characteristics. This negative correlation arises because each individual is counted as a neighbor for all the others in the same block, but not for herself. For estimates of the correlation that do not condition on reference group fixed effects, this bias is inconsequential because an individual's own characteristics contribute very little to the average neighborhood characteristics of others in the full sample. For estimates that condition on reference group fixed effects, however, this negative bias is quite large in magnitude because an individual's own characteristics contribute a significant amount to the average neighborhood characteristics of others within the same reference group. By sampling only one individual per block, we report an unbiased estimate of the correlation between individual and neighborhood characteristics at the block level.

[28] We drop blocks with fewer than five workers for two reasons. First, blocks with a small number of residents are largely nonresidential, and consequently, interactions among neighbors may be limited on such blocks. Second, as we discuss in greater detail below, a measurement error arises related to the use of the 1-in-7 sample of individuals observed in the census to estimate neighborhood effects. In this case, blocks with only a small number of workers may be particularly prone to measurement error.

reference group definition based on the 10 closest blocks. In each case, both the individual and block measures are first regressed on the corresponding variables (e.g., block group fixed effects), and the correlation between the residuals is reported.

The results indicate a significant amount of sorting on the basis of education, race, age, and the presence of children across the neighborhoods of the metropolitan area as a whole. The correlation between whether an individual is a college graduate and the fraction of neighbors who are college graduates is 0.29, whereas that between whether an individual is black and the fraction of black neighbors is 0.59. Columns 2 and 3 provide an explicit test of our identification strategy, providing a measure of sorting on observables within reference groups. As these successive columns clearly demonstrate, the correlation between observable individual and neighbor characteristics falls to near zero as only within–reference group variation is isolated. The inclusion of block group fixed effects reduces the estimated correlations by 70–90 percent for most categories, with a remaining maximum correlation of 0.05 across all characteristics, except for Asian and Hispanic. When the alternative definition of reference groups is used, residual correlations drop even further. This evidence is broadly consistent with that reported by Ioannides (2004), who finds a similar correlation among neighbors' incomes, with neighborhood defined as the 10 closest neighbors or as a census tract.

The magnitude of the remaining correlation between individual and neighbor attributes within reference groups provides clear support for the notion that the amount of sorting on observables within reference groups is less extensive than across the neighborhoods of the metropolitan area as a whole. This evidence is particularly compelling for our identification strategy because a number of these attributes, such as residents' race or the presence of families with children, would be the characteristics of one's immediate neighbors that might be most observable at the time of moving into a new residence. Thus, when these observables are controlled for, it may be the case that within–reference group sorting on other characteristics is even less extensive.

A direct test of the importance of block-level sorting on observables.—While the correlation estimates reported in table 3 are small, they are not identically zero. An obvious question, then, is whether the remaining block-by-block sorting on the basis of observables within reference groups, small though it may be, is enough to significantly increase the propensity of pairs drawn from the same block within a reference group to work together.

To answer this question, we turn to the heterogeneous version of the model presented in equations (3) and (4). In these equations, $\beta' X$ measures how the propensity to work together of two individuals who reside in the same reference group but not on the same block varies with the observable characteristics of the pair. Given an estimate of $\hat{\beta}$, this heterogeneous specification provides a way to test whether the remaining within–reference group correlation between observable neighbor attributes would lead to a significantly higher predicted propensity for pairs on the same block to work together. Specifically, we compare the average $\hat{\beta}'X$ for those pairs that reside on the same block with those that reside on nearby (but not the same) blocks within the reference group. In other words, we apply the $\hat{\beta}$ estimated for pairs not in the same block to pairs in the same block to see whether block-level correlation in the X's alone is sufficient to induce a higher propensity to work together for pairs on the same block.

Given the $\hat{\beta}$ that we estimate for the block group fixed effects model, the predicted propensity for pairs that reside on the same block is 0.343 percent; this is 0.01 percentage points lower than the observed (and predicted) propensity for pairs that reside in the same block group but not on the same block (0.355). Thus, the remaining block-level sorting on observables does not predict any increased propensity for individuals on the same block to work together. This evidence strongly favors the notion that our research design is credible in the face of the small amount of within-block sorting that exists in the data.

Additional specifications and robustness.—As described above, our empirical design relies critically on the assumption that social interactions are especially strong at the block level, whereas households are able to choose a broader neighborhood (block group or other set of nearby blocks) only at the time of the location decision, perhaps because of the thinness of the housing market. While the analysis of correlation between observable neighbor characteristics described above provides assurance that this assumption is reasonable, we also consider the robustness of our results to alternative samples designed to isolate those reference groups that are most homogeneous along a number of dimensions including race, education, and the presence of children in the household. In particular, in each case, we select the 50 percent of reference groups that display the least amount of within–reference group correlation between the corresponding individual and neighbor characteristics and reestimate the baseline model for the restricted sample in order to see if our results are robust across samples.[29]

A separate confounding issue is the possibility that the estimated social interaction effect may be due to reverse causation: workers could receive tips and referrals about residential locations from their coworkers at a

[29] While the resulting analysis obviously changes the nature of the sample, the results described below do provide some reassurance that our baseline results are not sensitive to sorting.

given firm. We address this issue in several ways. First, the empirical focus on the difference between reference group–level and block-level propensities again mitigates this problem because residential referrals are unlikely to result in people residing on exactly the same block, because of the thinness of the housing market at the block level.

We also tackle the potential for reverse causation directly by estimating equations (1) and (2) on a subsample of the data in which both respondents in a given matched pair have lived in that neighborhood for at least 2 years, but one of them was not employed for the full year in the previous year, defined as having worked less than 45 weeks in 1989 (for robustness, we also use a more restrictive definition of not employed for the full year, using 20 weeks in 1989 as the cutoff). In this case, we can be fairly certain that if we see the same individuals working together in the current year, then the referral was among residential neighbors rather than work colleagues. Unfortunately, the census does not contain any direct information on job search activity. Therefore, we use the not employed for the full year in 1989 category as a proxy for the set of individuals who are most likely to have been actively searching for a job last year.[30] The goal of this analysis is to examine whether evidence of referrals is present in this subsample. Importantly, because this subsample is (by construction) very different from the main sample, we do not expect the resulting extent of social interactions to be identical to our baseline results. Finally, we also perform a counterfactual experiment to look at a situation in which a residential referral may be most likely: namely, a sample in which one member of each pair lived in a given block at least 5 years, the other is a recent arrival (less than 2 years in residence), but both workers were employed for the full year in 1989.

V. Results—Detecting Referral Effects

We now present the results of our primary analysis, beginning with an examination of the propensity for two individuals to work together. Table 1 contains summary statistics for our matched pairs sample. As described above, columns 1 and 2 report the fraction of pairs residing in each reference group definition that fit the description in the row heading. Columns 3 and 4 report—for each category—the empirical frequency that two individuals who reside in the same reference group but not on the same block work together. Column 5 reports the probability that two individuals who reside on the same block work together. In this way,

[30] Note that in estimating earnings and wage equations in table 8 below, we condition on a set of individuals who were fully employed in the previous year, defined as having worked at least 45 weeks and at least 30 hours per week. This definition is different from that for not employed for the full year in 1989 used here, which is not based on hours at all.

the first row indicates that the baseline probability of working together for two individuals who reside in the same reference group but not on the same block is 0.36 percent for block groups and 0.38 percent for the 10 closest blocks; this figure rises to 0.94 percent for two individuals who reside on the same block. As we will see below, much of this increased propensity for individuals residing on the same block to work together results from the fact that the sample of individuals who reside on the same block is disproportionately weighted to larger blocks, that is, dense block groups. The inclusion of reference group fixed effects in our main empirical specification ensures that our social referral effects are estimated purely on the basis of comparisons within the same reference group.

The remaining rows of table 1 reveal how these patterns vary with the characteristics of the pair of individuals. First, notice that individuals residing on the same versus nearby blocks show an increased propensity to work together across all the types of pairs characterized in the table. This increased propensity to work together for individuals on the same block versus block group is especially strong for pairs of individuals in which (i) both have children and especially similar-aged young children, (ii) both are married, (iii) both are young, and (iv) both are high school graduates.

Table 1 also makes clear that the propensity that two individuals residing on the same block work together is not a simple monotonic function of the baseline propensity for individuals residing in the same reference group but not on the same block. While pairs of all age combinations residing in the same reference group but not on the same block are about equally likely to work together, pairs of young workers residing on the same block are especially likely to work together. Similarly, while pairs of workers with children in nearby blocks are about as equally likely to work together as pairs without children, the corresponding propensity of pairs with children to work together is more than twice that of those without at the block level.

Baseline specifications.—While table 1 provides suggestive evidence as to the presence and nature of a social interaction effect operating at the very local (block) level, our regression specifications help clarify this evidence since they include reference group fixed effects. This ensures that the estimation of our social interaction effects is based exclusively on comparisons of block- versus neighborhood-level propensities to work together within the same reference group. Table 4 reports the results of three specifications. Columns 1 and 2 report the parameter estimate of the average social interaction effect, α_0, in our baseline model (1) for each definition of neighborhood reference groups. Column 3 reports the results of the estimation of equation (2), which includes individual worker fixed effects, using block group as the ref-

TABLE 4
ESTIMATES OF EMPLOYMENT LOCATION MATCH REGRESSIONS: SPECIFICATION WITHOUT COVARIATES—AVERAGE EFFECTS ONLY

	Sample: Blocks with Five+ Workers					
	Census Block Group (1)		10 Closest Blocks (2)		Census Block Group (3)	
	Coefficient	*t*-Statistic	Coefficient	*t*-Statistic	Coefficient	*t*-Statistic
Reside on same block:						
bmatch	.1200	6.80	.3345	33.02	.1688	33.09
Sample size	1,234,494		2,198,183		1,234,494	
Individual fixed effects	No		No		Yes	
Includes reference group fixed effects	Yes		Yes		Yes	

NOTE.—The table reports results for three specifications of a regression in which an observation is a pair of currently employed, prime-age (25–59) adults who reside in the same neighborhood reference group but not in the same household within the Boston metropolitan area in 1990. In each specification, the dependent variable equals one if both individuals work in the same location (census block) and zero otherwise. All specifications are for a sample that drops blocks with fewer than five workers, which includes 1,234,494 pairs. Column 1 reports results using census block groups as reference groups. Column 2 reports results using the 10 closest blocks as the neighborhood reference group. Column 3 adds individual fixed effects to col. 1. Neighborhood reference group fixed effects are included in all specifications (although these are redundant in the specification that includes individual fixed effects). The coefficients have been multiplied by 100 to reflect percentage changes. Standard errors in all cases are estimated by pairwise bootstraps, and *t*-statistics are reported.

erence group. In all cases, our baseline sample is based on dropping blocks with fewer than five workers.

The estimated social interaction effect is positive and highly statistically significant in each case, indicating a strong additional propensity for two workers living in the same block to also work in the same block (distinct from the residential one), over and above the estimated propensity for matches in their reference group. The magnitude is 0.12 percentage points for the specification that uses census block groups. This effect is sizable: it is roughly 33 percent the size of the baseline propensity to work together for two individuals who reside in the same block group but not on the same block (0.36 percent).[31]

An increased propensity to work with a given neighbor implies a much larger propensity to work with at least one neighbor. For our baseline sample, which restricts the sample to blocks with at least five sampled workers, given the average of 80 individuals per block,[32] an estimated referral effect of 0.12 percentage points translates to approximately a 6.9-percentage-point increase in the propensity that an individual works

[31] As noted above, that this effect is less than the mean difference reported in table 1 suggests that a portion of the difference in means between those residing on the same block vs. those in the same block group but not on the same block was driven by variation across block groups related to population density. See Sec. IV for a discussion of this issue.

[32] While the average number of workers meeting our sample criteria for the match model is only 5.1 workers, the fact that the census is a 1-in-7 sample and that many workers are excluded from our analysis because of the presence of imputed data accounts for the larger average number of actual prime-age workers per block.

with at least one individual on the same block.[33] Thus, the referral effect estimated here is certainly economically meaningful.

The estimated average referral effects are almost three times as large as the baseline effect for the specification based on the alternative definition of neighborhood reference group (10 closest blocks). The increase in magnitude is driven by the fact that the social interaction effect is increasing in population density (we present this finding below). Because census block groups are defined in such a way as to keep the total population of block groups relatively stable across the sample, very few reference blocks are included for dense blocks in our baseline sample. Thus, changing the definition of the reference group to include the 10 nearest blocks has the effect of including many more reference blocks for the densest blocks in the sample, that is, weighting the sample toward dense blocks, thereby increasing the estimated effect size.

The estimated interaction effect is also slightly larger for the specification that includes individual fixed effects (compared to our baseline specification). We take this as a strong sign that our research design is fundamentally solid and controls effectively for block-level sorting on the basis of both observed and unobserved attributes.

Robustness—sorting within block groups and reverse causation.—While the correlation analysis presented in Section IV and the results of the specifications reported in table 4 provide a great deal of reassurance regarding the robustness of our analysis to concerns about the sorting of households across blocks within reference groups, we seek to provide additional evidence that such sorting is not fundamentally driving the results. To this end, as described in Section IV, table 5 reports the results of estimates using subsamples based on the 50 percent of reference groups that exhibit the least amount of block-by-block sorting in three dimensions: education, the presence of children in the household, and race. It is important to note, of course, that these restrictions on the sample change the nature of the set of households for which social interaction effects are identified so that there is no reason to expect the results to be identical to the full specification. In our minds, then, this exercise serves mainly as a broad check regarding block-level sorting.

Table 5 is organized as follows: the columns report estimation results for each homogeneous subsample, whereas the row panels refer to our three main specifications: census block groups, alternative reference groups, and block groups with individual fixed effects. The estimated referral effects remain fairly stable across subsamples, with the exception

[33] For computational ease, this calculation treats the likelihood of working with each neighbor as an independent event. The reported $0.069 = (1 - 0.00355)^{80} - [1 - (0.00355 + 0.0012)]^{80}$, where 80 is the average number of adults on the same block, 0.00355 is the baseline propensity for individuals to work with someone in the same block group, and 0.0012 is our estimated social interaction effect.

TABLE 5
Employment Location Match Regressions for Homogeneous Subsamples

	Reference Groups Most Homogeneous with Respect to:					
	Education (1)		Presence of Children (2)		Race (3)	
	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
	A. Specification without Covariates: Census Block Group Fixed Effects					
Reside on same block:						
bmatch	.0974	6.03	.0883	5.64	.0989	5.95
Sample size	980,548		1,032,037		921,277	
Includes block group fixed effects	Yes		Yes		Yes	
	B. Specification without Covariates: 10 Closest Blocks Reference Group Fixed Effects					
Reside on same block:						
bmatch	.3632	26.42	.3923	29.85	.6677	29.81
Sample size	1,732,535		1,918,414		800,909	
Includes reference group fixed effects	Yes		Yes		Yes	
	C. Specification without Covariates: Census Block Group Geography with Individual Fixed Effects					
Reside on same block:						
bmatch	.1226	19.75	.1048	19.19	.1305	20.71
Sample size	980,548		1,032,037		921,277	
Includes individual fixed effects	Yes		Yes		Yes	

Note.—The table reports results for nine specifications of a regression in which an observation is a pair of currently employed, prime-age (25–59) adults who reside in the same reference group but not in the same household within the Boston metropolitan area in 1990. The dependent variable equals one if both individuals work in the same location (census block) and zero otherwise. Each specification is based on the sample of pairs in blocks with at least five workers. The columns report results for samples of the most homogeneous reference groups in terms of education, the presence of children, and race, respectively. Reference group fixed effects are included in all specifications (although these are redundant in the specifications that includes individual fixed effects). Panel A reports results for specifications that include only census block group fixed effects and an indicator for whether the individuals reside on the same block. Panel B uses the same specifications as panel A but with the alternative neighborhood reference groups based on the 10 closest blocks geography. Panel C reports results for specifications that include block group fixed effects, individual fixed effects, and an indicator for whether the individuals reside on the same block. The coefficients have been multiplied by 100 to reflect percentage changes. Standard errors in all cases are estimated by pairwise bootstraps, and t-statistics are reported.

of the race subsample for the specification using the 10 closest blocks as reference group, where the effect almost doubles. In the specifications using census block groups, the estimated average effects are slightly attenuated relative to table 4, but they remain statistically and economically significant.

In sum, our estimated social interaction effects persist, even in areas that do not experience a significant degree of sorting below the reference group level with respect to characteristics most likely to be observed at the time a household moves into a block. We believe that this set of results further validates our attempt to isolate referral effects from sorting via the general research design proposed in this paper.

Table 6 collects the estimation results of the specifications that address the reverse causation issue. Here again the row panels refer to our three

TABLE 6

Employment Location Match Regressions: Tenure-Based Subsamples

	Both in Residence at Least 2 Years (1)		Both in Residence at Least 2 Years; One Not Employed for Full Year 1989 (2)		Both in Residence at Least 2 Years; One Not Employed for Full Year 1989 (Alternative Definition) (3)		One in Residence Less than 2 Years; Other in Residence at Least 5 Years; Both Fully Employed in 1989 (4)	
	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
	A. Specification without Covariates: Census Block Group Fixed Effects							
Reside on same block: *bmatch*	.1016	4.07	.0918	2.28	.1917	1.92	.0531	1.44
Sample size	846,061		196,167		41,979		130,229	
Includes block group fixed effects	Yes		Yes		Yes		Yes	
	B. Specification without Covariates: 10 Closest Blocks Reference Group Fixed Effects							
Reside on same block: *bmatch*	.2609	15.12	.2801	6.28	.6474	9.27	.1584	5.99
Sample size	1,417,125		230,424		74,227		265,269	
Includes reference group fixed effects	Yes		Yes		Yes		Yes	

	C. Specification without Covariates: Census Block Group Geography with Individual Fixed Effects							
Reside on same block: *bmatch*	.1577	24.73	.1582	8.64	.1949	14.00	.1392	8.61
Sample size	846,061		196,167		41,979		130,229	
Includes individual fixed effects	Yes		Yes		Yes		Yes	

NOTE.—The table reports results for 12 specifications of a regression in which an observation is a pair of currently employed, prime-age (25–59) adults who reside in the same reference group but not in the same household within the Boston metropolitan area in 1990. The dependent variable equals one if both individuals work in the same location (census block) and zero otherwise. Each specification is based on the sample of pairs in blocks with at least five workers. Column 1 is based on a sample that includes only individuals who have lived in their current residence for at least 2 years. Column 2 uses a sample that includes only those individuals who have lived in their current residence for at least 2 years but in which one member of the pair was not employed for the full year in 1989 (defined as employed for 45 weeks or less). Column 3 uses a sample in which not employed for the full year in 1989 is defined as employed for 20 weeks or less. Column 4 reports results for a sample in which one member has lived in his or her current residence less than 2 years and the other member has lived in his or her current residence at least 5 years; both were employed for the full year in 1989. Reference group fixed effects are included in all specifications (although these are redundant in the specifications that include individual fixed effects). In panel A of the table, results are reported for a specification that includes only block group fixed effects and an indicator for whether the individuals reside on the same block. Panel B reports results for a specification that includes only reference group fixed effects and an indicator for whether the individuals reside on the same block using the new 10 closest blocks geography. Panel C adds individual fixed effects to panel A. The coefficients have been multiplied by 100 to reflect percentage changes. Standard errors are estimated by pairwise bootstraps, and *t*-statistics are reported.

main specifications (block groups, 10 closest blocks, and block groups with individual fixed effects), whereas columns 1–3 refer to the different subsamples that aim at isolating instances in which it is more likely that the residential location decision preceded the current job location. Column 4 reports the results of our counterfactual experiment, in which it is relatively more likely that the current job preceded the residential location decision of one member of a given worker pair.

To begin, we focus on the first row panel. Our estimated referral effects range from 0.09 to 0.19 for columns 1–3. Again, the sampling schemes in columns 2 and 3 reduce the possibility of reverse causation since we are considering workers who are more likely to have made a transition to full employment during the past year and whose residential tenure is longer than 2 years. At the same time, by looking at pairs in which one was employed for the full year but the other was not, we are focusing on instances in which it is most likely that a referral or information exchange actually took place. The largest estimated effect, 0.19, occurs for the sample in which one individual had worked less than 20 weeks last year, an individual who quite likely needed a labor market referral during the last year. Instead, the estimated effect for our counterfactual (col. 4) is 0.05, or roughly half the size of the effect for most samples, and is no longer significant at the 5 percent level. This pattern is consistent with our interpretation of the estimated effect as the result of a job referral mechanism. The other panels are broadly consistent with this pattern, again with the estimated effect in column 4 being smaller and less statistically significant than in the previous columns, which are less likely to be consistent with reverse causation.[34]

Heterogeneous referral effects.—Table 7 reports our estimation results for the heterogeneous specifications described in equations (3) and (4). The structure of the table is identical to that of table 4, with each column reporting results for our three main specifications (alternative definitions of reference groups and individual fixed effects). All results pertain to the baseline sample of blocks with at least five workers. To enhance the readability of the table, only the coefficients for the interaction terms (i.e., those interacted with whether the two workers live on the same block, *bmatch*) are reported in table 7.[35]

The vast majority of the estimated interaction effects are robust across the three main specifications with a couple of key exceptions that are highlighted by our choice of the excluded category for each set of

[34] Weinberg et al. (2004) also find little evidence of reverse causation. They look at hours worked before and after a move to try to assess the possibility of an exogenous change in employment status taking place before a move; they find that hours are flat in the years preceding a move but increase after a move to neighborhoods with higher employment and better access to jobs.

[35] The results for the level coefficients are available from the authors on request.

characteristics.[36] We focus first on the robust findings. The results for education, age, the presence of children of similar ages, gender and marital status (except pairs of married females), and race (except Asian/Hispanic pairs) are very stable across all three specifications. These robust results imply that stronger interactions occur (i) for pairs in which both individuals are high school versus college graduates; (ii) for pairs in which both have children, and especially those with elementary or secondary school-aged children of the same age; (iii) between the youngest adults in the sample; and (iv) for married males relative to other gender–marital status combinations.

In general, these findings are broadly consistent with two common empirical findings in the existing literature on social networks and on informal hiring channels: that there is strong assortative matching within social networks and that referrals can occur only when at least one member of the pair is well attached to the labor market (see, e.g., Corcoran et al. 1980). That referral effects are stronger for high school than for college graduates is consistent with two other common results in the referrals and the social networks literatures. One is that informal hiring channels are used more intensively by individuals with less education (see Corcoran et al. 1980; Topa 2001); the other is that the more educated tend to have more spatially dispersed social networks.[37] This result also suggests that reverse causation (referrals from jobs to residences) is not a primary driver of our findings. Since college-educated workers are more likely to be spatially footloose and thus more likely to take a job in a location that subsequently requires a residential move, we might have expected to see large effects for college-educated workers if reverse causation were a major issue.

Our finding that referral effects are stronger for younger pairs of workers is also consistent with the literature on job networks. Corcoran et al. (1980) report that use of informal hiring channels declines with age, and Granovetter (1995) finds in his study of Boston that workers are more likely to have found their first jobs through informal rather than formal means (relative to subsequent jobs). Ioannides and Loury (2004) also discuss similar findings in their survey.

The results presented in table 7 differ across specifications in ways that highlight the impact of including individual fixed effects and the

[36] A negative intercept for the specification with covariates means that the effect is negative (but barely statistically significant) for the left-out category: this is for matches between Asians/Hispanics and blacks in which one person is a high school graduate and the other is a college graduate, one person is 25 years old and the other is 35, and both live on very small (rural) blocks. Such a category is a very tiny portion of all pairs in the sample. The estimated social interaction effect is estimated to be positive for over 99 percent of pairs observed in the data for each specification shown in table 7.

[37] See Fischer (1982) for evidence from northern California and Kadushin and Jones (1992) for evidence from a 1988 survey of New York City residents.

TABLE 7
ESTIMATES OF EMPLOYMENT LOCATION MATCH REGRESSIONS: SPECIFICATION WITH COVARIATES

		Sample: Blocks with Five+ Workers					
		Census Block Group (1)		10 Closest Blocks (2)		Census Block Group (3)	
	Variable	Coefficient	t-Statistic	Coefficient	t-Statistic	Coefficient	t-Statistic
Reside on same block	*bmatch*	.0427	.09	−.9431	−5.60	.0969	1.93
Both high school dropouts	*bmatch* × *hsd_hsd*	−.0284	−.15	−.1764	−.73	−.0991	−1.20
Both high school graduates	*bmatch* × *hsg_hsg*	.1851	3.80	.2622	6.64	.1025	5.02
Both college graduates	*bmatch* × *clg_clg*	.0124	.17	−.0981	−3.59	.0214	2.80
High school dropout–high school graduate	*bmatch* × *hsd_hsg*	−.0606	−1.47	.0174	.38	.0373	2.96
High school dropout–college graduate	*bmatch* × *hsd_clg*	−.1103	−2.28	−.1363	−3.05	−.0055	−.35
Both have children	*bmatch* × *child_m*	.0779	.71	.2999	4.90	−.0301	−2.14
Both have children age 0–5	*bmatch* × *c05_05*	.3742	1.77	.5852	5.32	.0687	3.61
Both have children age 6–12	*bmatch* × *c612_612*	.3070	2.18	.4956	5.40	.0659	1.88
Both have children age 13–17	*bmatch* × *c1317_1317*	.3190	2.57	.3773	3.30	.2345	8.38
Both have children age 18–24	*bmatch* × *c1824_1824*	−.1522	−.86	−.4504	−3.97	−.0615	−2.16
Both age 25–34	*bmatch* × *a25_25*	.2094	2.62	.3110	10.39	.2049	13.44
Both age 35–44	*bmatch* × *a35_35*	−.0063	−.09	−.1440	−4.12	.1585	13.13
Both age 45–59	*bmatch* × *a45_45*	.0166	.24	−.0729	−1.50	.1102	6.40
Age 25–34 and age 45–59	*bmatch* × *a25_45*	.0167	.37	.0116	.34	.0583	3.71
Age 35–44 and age 45–59	*bmatch* × *a35_45*	−.0126	−.16	−.1181	−2.40	.0649	5.22

Both single males	*bmatch* × *sm_sm*	.8642	2.39	.1906	1.63	−.1010	−4.08
Both single females	*bmatch* × *sf_sf*	.6649	2.08	.0916	1.00	−.2654	−11.05
Single male–single female	*bmatch* × *sm_sf*	.6135	1.73	−.0561	−.61	−.2611	−13.62
Both married males	*bmatch* × *mm_mm*	1.2782	4.46	.8214	9.45	.0934	4.15
Married male–married female	*bmatch* × *mm_mf*	.6091	3.53	.3404	3.63	−.2870	−11.54
Single male-married female	*bmatch* × *sm_mf*	.5544	1.90	.0127	.12	−.2773	−10.31
Single male-married male	*bmatch* × *sm_mm*	.8861	2.51	.2532	3.45	−.1321	−5.61
Single female-married female	*bmatch* × *sf_mf*	.7149	2.52	.1895	2.45	−.1566	−4.76
Single female-married male	*bmatch* × *sf_mm*	.6892	1.92	.0858	1.05	−.2012	−9.10
Both white	*bmatch* × *wht_wht*	−.7491	−2.78	−.3965	−3.32	.0277	.57
Both black	*bmatch* × *bl_bl*	−.7038	−2.80	−.4563	−3.67	.1182	1.48
White–black	*bmatch* × *bl_wht*	−.7858	−2.63	−.3367	−2.04	.1252	2.42
White–Asian/Hispanic	*bmatch* × *ashi_wht*	−.4746	−1.76	−.1571	−1.21	.0422	.67
Combined time in residence	*bmatch* × *lngth*	−.0017	−.87	.0048	2.16	−.0019	−1.63
Minimum time in residence	*bmatch* × *lngth_min*	.0048	1.44	−.0026	−.71	.0077	3.92
Moved within 5 years of each other	*bmatch* × *lngth_win_5*	.0387	1.38	.1562	5.14	.0351	2.04
Block size (population)	*bmatch* × *blocksize*	−.0010	−.09	.0227	39.07	.0001	.49
Sample size		1,234,494		2,198,183		1,234,494	
Individual fixed effects		No		No		Yes	
Includes reference group fixed effects		Yes		Yes		Yes	

Note.—The table reports results for three specifications of a regression in which an observation is a pair of currently employed, prime-age (25–59) adults who reside in the same reference group but not in the same household within the Boston metropolitan area in 1990. In each specification, the dependent variable equals one if both individuals work in the same location (census block) and zero otherwise. All specifications are for a sample that drops blocks with fewer than five workers, which includes 1,234,494 pairs. Column 1 reports results using census block groups as reference groups. Column 2 reports results using the 10 closest blocks geography. Column 3 adds individual fixed effects to col. 1. Reference group fixed effects are included in all specifications (although these are redundant in the specification that includes individual fixed effects). The coefficients have been multiplied by 100 to reflect percentage changes. Standard errors in all cases are estimated by pairwise bootstraps, and *t*-statistics are reported.

role of geography. With regard to the results that include individual fixed effects, the coefficient estimates for gender–marital status and race relative to the excluded groups (both married females and both Asian or Hispanic, respectively) are largely much smaller in magnitude and often change sign. This suggests that the large (negative and positive, respectively) effects found for these two excluded groups in the specification without individual fixed effects are largely driven by the way in which individuals in these groups sort across blocks. Notice, however, that in both cases the relative ordering and magnitude of the estimated effects remain consistent across the included categories. Thus, the inclusion of individual fixed effects seems to be helpful in controlling for unobserved block-level sorting related to employment for married females (without kids) and Asians and Hispanics.

Comparing the results using the alternative definitions for a neighborhood reference group also reveals another key difference. In particular, the estimated coefficient on block size (population) is essentially zero when the census block group is used as the reference group and substantially positive when the 10 nearest blocks are used. The key difference between these specifications is the way in which very dense blocks are treated in the samples. In this way, the large block size interaction effect presented in table 7 is consistent with the larger overall interaction effect presented in table 4 when the reference group includes the 10 nearest blocks; both results suggest that the magnitude of our referral effect is sharply increasing in density, especially for the most dense blocks in the sample.

VI. Labor Market Outcomes

Having analyzed the impact of local interactions on job referrals, we conduct a second portion of our analysis that is designed to study whether such referrals have an impact on labor market outcomes more generally. In particular, given the characterization of how the strength of social interactions related to job referrals (i.e., the propensity to work together) varies with the attributes of a pair of individuals identified in the first portion of our analysis, we explore whether an individual's labor market outcomes are related to the idiosyncratic quality of the strength of the potential networks available on her block. Specifically, we estimate a series of labor market outcome regressions that include a measure of match quality defined at the individual level along with controls for individual and average neighbor characteristics (measured at the block level) as well as reference group fixed effects.

This portion of our analysis has two goals. First, since we detect informal hiring effects indirectly, it serves as a check on the plausibility of the first portion of our analysis. Second, by focusing on outcomes, we hope to provide a better sense of the magnitude of our estimated network effects. It is certainly possible that referrals may be more likely among neighbors but may have little effect on labor market outcomes, that is, that without the referral the individual would find a comparable job through another search method. Thus, it is important to be able to say something about the impact of an individual network's potential on outcomes.

For this analysis, the unit of observation is an individual rather than a pair. For the employment and labor force participation outcomes, the econometric model is a linear probability model.[38] For all other outcomes, such as weeks worked, hours per week worked, wages, and earnings (in logs), we use a simple linear regression.

We then add—for each model specification—a network quality proxy variable for each individual, which is constructed by examining that individual's matches with other adults in her block, using the coefficient estimates α_1 from the estimation of equations (3) and (4). Specifically, the average match quality for individual i, Q_{ib}, is constructed using a sample of all possible pairings of individual i with other individuals who reside in the same block b and do not belong to the same household. For each pair, a linear combination M_{ij} of the pair's covariates is created using the estimated parameters from the interaction of these variables with R^b_{ij} in equations (3) and (4): $M_{ij} = \hat{\alpha}_1' X_{ij}$. Then, Q_{ib} is computed as the mean value of M_{ij} over all matches for individual i:

$$Q_{ib} = \frac{1}{|N_{ib}|} \sum_{j \in N_{ib}} M_{ij}, \tag{5}$$

where N_{ib} is defined as the set of other individuals who reside on the same block b but not in the same household as individual i.

We would generally expect individuals with good potential matches in their block—high value of Q_{ib}—to have better labor force outcomes on average, after controlling for the direct effect of their individual characteristics and block-level fixed effects. The resulting specification is given by

$$E_{ib} = \theta_b + \delta_3' Q_{ib} + \delta_4' X_{ib} + u_{ib}, \tag{6}$$

where θ_b denotes a block-level fixed effect and X_{ib} is a vector of individual attributes (the same as those used in the workplace clustering specification).

In the context of traditional linear-in-means social interaction models (of the type described in Manski [1993] and Moffitt [2001]), the block

[38] We have also performed our analysis using a multinomial logit specification, with three discrete outcomes: out of the labor force, unemployed, and employed. The results are qualitatively very similar.

fixed effect θ_b in equation (6) can be thought of as replacing block averages of individual outcomes and attributes ($\bar{E}_b$ and $\bar{X}_b$, respectively); one could in fact rewrite (6) as

$$\begin{aligned} E_{ib} &= \delta_0 + \delta_1'\bar{E}_b + \delta_2'\bar{X}_b + \delta_3'Q_{ib} + \delta_4'X_{ib} + u_{ib} \\ &\equiv \theta_b + \delta_3'Q_{ib} + \delta_4'X_{ib} + u_{ib}. \end{aligned} \tag{7}$$

In such models, endogenous and exogenous social effects (δ_1' and δ_2') are typically not separately identified without imposing some additional structure. In our case, the block fixed effects capture the influence of mean neighborhood outcomes and observables without attempting to decompose these two effects.[39] Our variable of interest Q_{ib} then identifies the additional influence of social interactions that are heterogeneous across individuals on the basis of our specific mechanism for detecting labor market referrals, that is, the propensity to work together.

Finally, it is important to point out a limitation of this exercise. In particular, what are actually identified by the first-stage analysis are the types of pairs that are more likely to work together as a result of the strength of the referral effect between the pair. As discussed above, we expect this effect to be large in two cases: (i) when a pair is more likely to interact within their residential neighborhood and (ii) when one person is well attached to the labor market and the other is likely to need a referral. In this way, for a person who is not well attached to the labor market, the measure of match quality described here should do a good job of characterizing the quality of matches in a neighborhood. For a person better attached to the labor market, however, our match quality variable may actually measure neighborhoods in which such a person provides rather than receives referrals. In this way, to the extent that our estimated social interaction effects in the first stage of our analysis are driven by the asymmetry in labor market attachment rather than by the strength of neighborhood interactions, our analysis of the effect of match quality on labor market outcomes is likely to understate the benefits of improved matches.

Labor market outcome regressions.—We now present a series of labor market outcome regressions based on each of the specifications of the matched pairs equations reported in table 7. Each regression includes a set of individual and average neighbor characteristics for each sociodemographic characteristic included in the work match specification as well as a set of reference group fixed effects. The three broad columns of table 8 report the effect of a one-standard-deviation increase in match quality on labor market outcomes for specifications corresponding to

[39] Using block fixed effects is also preferable since mean outcomes and attributes may be measured with error at the block level, and block attributes may influence outcomes in a nonlinear fashion.

the three columns of table 7. Separate results are reported for males and females. In this table, we report only the coefficient estimates associated with match quality for the sake of expositional clarity.[40] Note also that the number of observations varies across specification because of the number of observations with imputed dependent variables in each case; we drop such observations from the analysis.

For the specifications based on the baseline sample with census block groups, match quality has a positive and (statistically and economically) significant impact on all dependent variables under consideration except for wages.[41] For this specification, a one-standard-deviation increase in match quality raises labor force participation by about 3.3 percentage points, average hours worked per week by about 1.8 hours, and earnings by 3.4 percentage points for males. The results for females are similar to those for males in terms of labor force participation and employment conditional on participation, whereas they are slightly larger for weeks worked, hours worked, and earnings. In this way, our estimated referral effects are indeed associated with an improvement in labor market outcomes especially as it concerns participation in the labor market and the intensity of that participation.[42] The latter is especially true for women, for whom the increase in hours and weeks worked more than compensates for the drop in hourly wages, so that there is still a positive effect on earnings.

A subset of these results is broadly confirmed in the specification with the alternative reference group definition and with individual fixed effects (cols. 2 and 3), although the size of the estimated effects is generally smaller. In interpreting the magnitudes of the results in table 8, one should keep in mind that, as shown in the first row of the table, the estimated standard deviation of match quality across blocks is significantly smaller in the individual fixed effects specification. This means that even though the coefficients on match quality in the set of labor market outcome regressions are similar in magnitude to those obtained for the other specifications, the reported effect of a one-standard-deviation increase in match quality is much smaller, as seen in the table.

For males, an increase in match quality has a statistically significant

[40] The estimation results for the full sets of individual and block-level covariates are quite standard and are available from the authors on request. The first two dependent variables refer to labor market outcomes for the week preceding the census survey. The last four variables represent labor market outcomes for the preceding year. Earnings and wage regressions are run for the sample of individuals who were fully employed in the previous year, defined as having worked at least 45 weeks and at least 30 hours per week.

[41] Standard errors are corrected for clustering at the block level in all labor market outcome regressions reported in the paper.

[42] Recall from our discussion above that this analysis will tend to understate the benefits of improved match quality at the block level since the quality of local matches will typically be overstated for individuals who generally provide referrals.

TABLE 8
Effect of Match Quality on Labor Market Outcomes: Effect of a One-Standard-Deviation Increase in Block-Level Match Quality

		Sample: Blocks with Five+ Workers					
		Census Block Group (1)		10 Closest Blocks (2)		Census Block Group (3)	
	Observations	Coefficient	*t*-Statistic	Coefficient	*t*-Statistic	Coefficient	*t*-Statistic
Standard deviation of match quality (%)		.294		.401		.079	
				A. Males			
Labor force participation	128,916	.033	6.60	.008	1.94	.000	.15
Employment	128,916	.032	5.52	.013	2.81	.004	1.23
Weeks worked last year	118,679	1.526	5.24	.995	4.14	.238	1.60
Hours worked per week	118,679	1.801	6.85	.721	3.33	.309	2.27
Log(earnings)	84,773	.034	2.67	.037	3.59	.020	3.23
Log(wage)	70,215	−.014	−1.51	.006	.77	.020	4.07

	B. Females						
Labor force participation	128,916	.034	5.57	.011	1.96	.009	2.87
Employment	128,916	.036	5.72	.012	1.87	.008	2.49
Weeks worked last year	118,679	2.162	6.88	.398	1.28	.020	.12
Hours worked per week	118,679	2.417	8.99	−.042	−.16	.223	1.62
Log(earnings)	84,773	.081	5.30	.031	1.97	−.006	−.85
Log(wage)	70,215	−.045	−5.11	−.036	−4.38	.004	.88
Individual fixed effects in employment match model		No		No		Yes	

NOTE.—The table reports results for three specifications of six labor market outcome regressions. The labor market outcomes are labor force participation status in 1990, current employment in 1990, weeks worked in 1989, average hours worked per week in 1989, the log of 1989 earnings, and the log of 1989 hourly wage. For the first four of these outcome measures, respectively, the sample consists of all prime-age (25–59) adults who reside in the Boston metropolitan area in 1990. For the last two outcomes, the sample consists of all such individuals who were fully employed in 1989. In these earnings and wage regressions, fully employed refers to individuals who worked at least 45 weeks and at least 30 hours per week. All specifications use a sample that drops blocks with fewer than five workers, which includes 128,916 individuals. In all cases any observations for which the census imputed the dependent variable were dropped. Census block fixed effects are included in all labor market outcome regressions along with controls for the full set of characteristics reported in table 2 associated with race, education, age, sex, marital status, time in residence, and presence of children. For each dependent variable a single regression was estimated allowing the match quality coefficient to vary by gender; the reported number of observations for each specification refers to the total number of both men and women. The coefficients reported in the table characterize the effect of a one-standard-deviation increase in match quality on the corresponding labor market outcome. For the three specifications reported, match quality was constructed using the estimated coefficients from the corresponding employment match model presented in table 4. The standard deviation of match quality across blocks for each specification is shown in the first row. Standard errors are corrected for clustering at the block level, and *t*-statistics are reported.

positive effect on hours worked and earnings across all three specifications: the earnings effect ranges from 2 to almost 4 percentage points. For female workers, labor force participation and employment conditional on participation are positively affected by match quality across all specifications; here the effects range from 1 to about 3 percentage points.

It is not surprising that the availability of potential referrals as measured by our match quality variable has a differential impact on labor market outcomes for male and female workers. Several studies in the literature on informal hiring channels find that both usage and productivity of referral networks vary across gender. Bradshaw (1973), Ports (1993), and Rosenbaum et al. (1999) all find that unemployed women are less likely to use informal job networks. Loury (2006) finds that female contacts have a lower impact on outcomes than male impacts; since personal networks are assortative along gender lines, this implies that referrals tend to be less productive for female than for male workers. Bortnick and Ports (1992) similarly find that referral networks are less productive for females.

The magnitudes of the labor force participation and employment effects estimated in table 8 are generally consistent with the increased propensity to work with at least one neighbor in the same block estimated in the corresponding employment match models presented in table 4. Consider, for example, the results presented for our primary specification (census block group, no individual fixed effects, sample includes blocks with five or more workers). In this case, a one-standard-deviation increase in match quality in the employment match model corresponds to a 10.1 percent increase in the probability that an individual works with at least one neighbor at the mean.[43] Given that one person in a match is providing the referral, this in turns implies an increase in the propensity to find a job through a neighborhood referral of 5.0 percent. This number corresponds to the increased propensity to work with someone on exactly the same block and therefore provides a lower bound on the number of neighborhood referrals more generally.

When compared to the employment effect estimated for the corresponding sample (3.2 percent for men, 3.6 percent for women), this then suggests that at most 65–75 percent of referrals result in the em-

[43] As discussed above, match quality is measured with error because of the 1-in-7 nature of the census sample. As a result, the measured standard deviation reported in table 8 overstates the true variation in match quality. For the specification described here, we estimated (using Monte Carlo simulations) that the true standard deviation of match quality is about 0.18 percentage points (as compared to the measured standard deviation of 0.29 reported in table 8). Following the same procedure as in the example worked out in n. 33, an 0.18-percentage-point increase in working with each neighbor leads to a $0.101 = (1 - 0.00355)^{80} - [1 - (0.00355 + 0.0018)]^{80}$ increase in the likelihood of working with at least one neighbor. Note 33 presents more details regarding this calculation.

ployment of an individual who would not be employed in the absence of the referral; the other 20–30 percent of neighborhood referrals go to individuals who would find employment through another search method. Again, because the denominator in this calculation is expected to be understated but the numerator is not, the actual fraction of referrals that result in a noninframarginal employment may be much less.[44]

VII. Conclusion

This paper aims to detect and measure the importance of neighborhood referrals on labor market outcomes by using a novel data set and identification strategy. Using Census data that detail the exact block of residence and workplace for a large sample of prime-age workers in the Boston metro area, we identify social interactions by comparing the propensity of individuals on the same versus nearby blocks to work together. We find significant evidence of social interactions: residing on the same block increases the probability of working together by over 33 percent. This finding is robust across various specifications intended to address biases caused by sorting below the reference group level and housing market referrals exchanged between people who work together as well as to the introduction of detailed controls for sociodemographic characteristics and individual fixed effects. Furthermore, the relationship between sociodemographic characteristics and the strength of social interactions makes sense. Social interactions tend to be stronger when the match involves individuals who are likely to interact because they are similar in terms of education, age, and the presence of children, which is consistent with the notion of assortative matching in social networks. Interactions also appear to be stronger when they involve at least one type of individual who is strongly attached to the labor market, leading to stronger interactions when both members of the pair are married males.

In the second half of our analysis we use our heterogeneous referral effect estimates to construct an individual-specific measure of the availability of referral opportunities on her block of residence. Even after we control for individual attributes, observable block attributes, and reference group or individual fixed effects, this measure is a statistically significant determinant of several labor market outcomes across all our specifications. In terms of economic magnitude, a one-standard-devia-

[44] We expect the labor market outcome regressions to provide an estimate of the ultimate impact of all actual referrals from the neighborhood including individuals in both the same and nearby blocks. In particular, with limited sorting within block groups, expected match quality for an individual with others in the same block group is the same as their actual block match quality. Consequently, the block-level index for match quality is likely to capture the effect of referrals both within the block and from neighboring blocks.

tion increase in referral opportunities raises expected earnings by 2.0–3.7 percentage points for men and labor force participation by 0.9–3.4 percentage points for women.

More generally, this paper provides a new approach for examining the effect of social interactions based on variation in geographic scale. In presenting the results related to neighborhood referrals and labor market outcomes, we also provide direct evidence on the reasonableness of this new design by testing whether its key assumptions hold on observable characteristics. In particular, we demonstrate that on the basis of their observable characteristics, pairs of individuals residing on the same block would actually be slightly less likely to work together than pairs in the same reference group but not on the same block. This provides strong evidence that our research design is likely to be robust to within–reference group sorting. Further, the inclusion of individual fixed effects allows us to control for a form of sorting on unobservables.

This evidence also suggests that the research design proposed in the paper may be useful in a variety of contexts. For example, in the case of welfare participation, the block of residence is unlikely to greatly influence access to public service providers after the reference group is controlled for. More generally, this design might be extended to the study of neighborhood effects in specific contexts (e.g., specific types of neighborhoods), on specific populations (e.g., youths), and for alternative outcomes (e.g., education, teenage fertility, health, welfare participation, bankruptcy, and mortgage delinquency), provided that the researcher can demonstrate that the within–reference group correlation in observable neighbor characteristics is zero, thereby ensuring that the key identifying assumption on unobserved characteristics is at least reasonable. In future work, we also intend to extend this analysis to young adults for whom neighborhood contacts might be an especially important source of job referrals.

References

Addison, John T., and Pedro Portugal. 2002. "Job Search Methods and Outcomes." *Oxford Econ. Papers* 54 (3): 505–33.

Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber. 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *J.P.E.* 113 (1): 151–84.

Arcidiacono, Peter, Gigi Foster, Natalie Goodpaster, and Josh Kinsler. 2007. "Estimating Spillovers using Panel Data, with an Application to the Classroom." Manuscript, Duke Univ.

Audretsch, David B., and Maryann P. Feldman. 1996. "R&D Spillovers and the Geography of Innovation and Production." *A.E.R.* 86 (3): 630–40.

Bandiera, Oriana, and Imran Rasul. 2006. "Social Networks, and Technology Adoption in Northern Mozambique." *Econ. J.* 116 (514): 862–902.

Bayer, Patrick, Randi Pintoff, and David Pozen. 2008. "Building Criminal Capital behind Bars: Peer Effects in Juvenile Corrections." Manuscript, Duke Univ.

Bertrand, Marianne, Erzo Luttmer, and Sendhil Mullainathan. 2000. "Network Effects and Welfare Cultures." *Q.J.E.* 115 (3): 1019–55.

Blau, David. 1992. "An Empirical Analysis of Employed and Unemployed Job Search Behavior." *Indus. and Labor Relations Rev.* 45 (4): 738–52.

Blau, David, and Phillip Robbins. 1990. "Job Search Outcomes for the Employed and Unemployed." *J.P.E.* 98 (3): 637–55.

Bortnick, Steven, and Michelle Ports. 1992. "Job Search Methods and Results: Tracking the Unemployed." *Monthly Labor Rev.* 115 (12): 29–35.

Bradshaw, Thomas F. 1973. "Jobseeking Methods Used by Unemployed Workers." *Monthly Labor Rev.* 96 (2): 35–40.

Brock, William A., and Steven N. Durlauf. 2001. "Interactions-Based Models." In *Handbook of Econometrics*, vol. 5, edited by James J. Heckman and Edward Leamer. Amsterdam: Elsevier Sci.

Burke, Mary A., Gary Fournier, and Kislaya Prasad. 2004. "The Diffusion of a Medical Innovation: Is Success in the Stars?" *Southern Econ. J.* 73 (3): 588–603.

Calvo-Armengol, Antoni, and Matthew O. Jackson. 2004. "The Effects of Social Networks on Employment and Inequality." *A.E.R.* 94 (3): 426–54.

Card, David, and Jesse Rothstein. 2007. "Racial Segregation and the Black-White Test Score Gap." *J. Public Econ.* 91 (11–12): 2158–84.

Case, Anne C., and Lawrence F. Katz. 1991. "The Company You Keep: The Effects of Family and Neighborhood on Disadvantaged Youths." Working Paper no. 3705, NBER, Cambridge, MA.

Conley, Timothy G., and Christopher R. Udry. 2005. "Learning about a New Technology: Pineapple in Ghana." Manuscript, Grad. School Bus., Univ. Chicago.

Corcoran, Mary, Linda Datcher, and Greg Duncan. 1980. "Information and Influence Networks in Labor Markets." In *Five Thousand American Families: Patterns of Economic Progress*, vol. 7, edited by Greg Duncan and James Morgan. Ann Arbor, MI: Inst. Soc. Res.

Crane, Jonathan. 1991. "The Epidemic Theory of Ghettos and Neighborhood Effects on Dropping Out and Teenage Childbearing." *American J. Sociology* 96 (5): 1226–59.

Cutler, David M., and Edward L. Glaeser. 1997. "Are Ghettos Good or Bad?" *Q.J.E.* 112 (3): 827–72.

Datcher, Linda. 1983. "The Impact of Informal Networks on Quit Behavior." *Rev. Econ. and Statis.* 65 (3): 491–95.

Devine, Theresa, and Nicholas Kiefer. 1991. *Empirical Labor Economics: The Search Approach.* New York: Oxford Univ. Press.

Durlauf, Steven. 2004. "Neighborhood Effects." In *The Handbook of Regional and Urban Economics: Cities and Geography*, vol. 4, edited by V. Henderson and J. F. Thisse. Amsterdam: Elsevier Sci.

Elliot, James. 1999. "Social Isolation and Labor Market Isolation: Network and Neighborhood Effects on Less-Educated Workers." *Sociological Q.* 40 (2): 199–216.

Evans, William, Wallace Oates, and Robert Schwab. 1992. "Measuring Peer Group Effects: A Study of Teenage Behavior." *J.P.E.* 100 (5): 966–91.

Fischer, Claude S. 1982. *To Dwell among Friends: Personal Networks in Town and City.* Chicago: Univ. Chicago Press.

Flachaire, E. 1999. "A Better Way to Bootstrap Pairs." *Econ. Letters* 64 (3): 257–62.

———. 2005. "Bootstrapping Heteroskedastic Regression Models: Wild Bootstrap versus Pairs Bootstrap." *Computational Statis. and Data Analysis* 49 (2): 361–76.

Glaeser, Edward L., Hedi D. Kallal, José A. Scheinkman, and Andrei Shleifer. 1992. "Growth in Cities." *J.P.E.* 100 (6): 1126–52.

Glaeser, Edward L., Bruce Sacerdote, and José A. Scheinkman. 1996. "Crime and Social Interactions." *Q.J.E.* 111 (2): 507–48.

Granovetter, Mark S. 1995. *Getting a Job: A Study of Contacts and Careers*. Cambridge, MA: Harvard Univ. Press.

Grinblatt, Mark, Matti Keloharju, and Seppo Ikaheimo. 2004. "Interpersonal Effects in Consumption: Evidence from the Automobile Purchases on Neighbors." Working Paper no. 10226, NBER, Cambridge, MA.

Guest, Avery M., and Barrett A. Lee. 1983. "The Social Organization of Local Areas." *Urban Affairs Q.* 19 (2): 217–40.

Holzer, Harry J. 1987. "Informal Job Search and Black Youth Unemployment." *A.E.R.* 77 (3): 446–52.

———. 1988. "Search Method Use by Unemployed Youth." *J. Labor Econ.* 6 (1): 1–20.

Horowitz, J. L. 2001. "The Bootstrap." In *Handbook of Econometrics*, vol. 5, edited by James J. Heckman and Edward Leamer. Amsterdam: Elsevier Sci.

Hoxby, Caroline. 2000. "Peer Effects in the Classroom: Learning from Gender and Race Variation." Working Paper no. w7867, NBER, Cambridge, MA.

Ioannides, Yannis M. 2004. "Neighborhood Income Distributions." *J. Urban Econ.* 56 (3): 435–57.

Ioannides, Yannis M., and Linda Datcher Loury. 2004. "Job Information Networks, Neighborhood Effects, and Inequality." *J. Econ Literature* 42 (4): 1056–93.

Ioannides, Yannis M., and Jeffrey E. Zabel. 2008. "Interactions, Neighborhood Selection, and Housing Demand." *J. Urban Econ.* 63 (1): 228–52.

Jacob, Brian. 2004. "Public Housing, Housing Vouchers and Student Achievement: Evidence from Public Housing Demolitions in Chicago." *A.E.R.* 94 (1): 233–58.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Q.J.E.* 108 (3): 577–98.

Jencks, Christopher, and Susan E. Mayer. 1990. "The Social Consequences of Growing Up in a Poor Neighborhood." In *Inner-City Poverty in the United States*, edited by L. Lynn and M. McGeary. Washington, DC: Nat. Acad. Press.

Kadushin, Charles, and Delmos J. Jones. 1992. "Social Networks and Urban Neighborhoods in New York City." *City and Society* 6 (1): 58–75.

Katz, Lawrence F., Jeffrey Kling, and Jeffrey Liebman. 2001. "Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment." *Q.J.E.* 116 (2): 607–54.

Korenman, Sanders, and Susan Turner. 1996. "Employment Contacts and Minority-White Wage Differences." *Indus. Relations* 35 (1): 106–22.

Lee, Barrett A., and Karen E. Campbell. 1999. "Neighbor Networks of Black and White Americans." In *Networks in the Global Village: Life in Contemporary Communities*, edited by Barry Wellman. Boulder, CO: Westview.

Loury, Linda Datcher. 2006. "Some Contacts Are More Equal than Others: Informal Networks, Job Tenure, and Wages." *J. Labor Econ.* 24 (2): 299–318.

Ludwig, Jens, Greg J. Duncan, and Paul Hirschfield. 2001. "Urban Poverty and

Juvenile Crime: Evidence from a Randomized Housing-Mobility Experiment." *Q.J.E.* 116 (2): 655–79.

Mammen, E. 1993. "Bootstrap and Wild Bootstrap for High Dimensional Linear Models." *Ann. Statis.* 21 (1): 255–85.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Rev. Econ. Studies* 60 (3): 531–42.

Marmaros, David, and Bruce Sacerdote. 2002. "Peer and Social Networks in Job Search." *European Econ. Rev.* 46 (4–5): 870–79.

Marsden, Peter V. 1987. "Core Discussion Networks of Americans." *American Sociological Rev.* 52 (1): 122–31.

———. 1988. "Homogeneity in Confiding Relations." *Soc. Networks* 10 (1): 57–76.

Moffitt, Robert A. 2001. "Policy Interventions, Low-Level Equilibria, and Social Interactions." In *Social Dynamics*, edited by Steven N. Durlauf and H. Peyton Young. Washington, DC: Brookings Inst. Press.

Montgomery, James D. 1991. "Social Networks and Labor-Market Outcomes: Toward an Economic Analysis." *A.E.R.* 81 (5): 1408–18.

Oreopoulos, Philip. 2003. "The Long-Run Consequences of Living in a Poor Neighborhood." *Q.J.E.* 118 (4): 1533–75.

Otani, Shinsuke. 1999. "Personal Community Networks in Contemporary Japan." In *Networks in the Global Village: Life in Contemporary Communities*, edited by Barry Wellman. Boulder, CO: Westview.

Popkin, Susan J., James E. Rosenbaum, and Patricia M. Meaden. 1993. "Labor Market Experiences of Low-Income Black Women in Middle-Class Suburbs: Evidence from a Survey of Gautreaux Program Participants." *J. Policy Analysis and Management* 12 (3): 556–73.

Ports, Michelle Harrison. 1993. "Trends in Job Search Methods, 1970–92." *Monthly Labor Rev.* 116 (10): 63–67.

Rees, Albert, and George P. Shultz. 1970. *Workers and Wages in an Urban Labor Market.* Chicago: Univ. Chicago Press.

Rosenbaum, James E., Stefanie DeLuca, Shazia R. Miller, and Kevin Roy. 1999. "Pathways into Work: Short- and Long-Term Effects of Personal and Institutional Ties." *Sociology of Educ.* 72 (3): 179–96.

Ross, Stephen L. 1998. "Racial Differences in Residential and Job Mobility: Evidence concerning the Spatial Mismatch Hypothesis." *J. Urban Econ.* 43 (1): 112–35.

Ross, Stephen L., and Yves Zenou. 2008. "Effort, Location, and Urban Unemployment." *Regional Sci. and Urban Econ* 38 (5): 498–517.

Sacerdote, Bruce. 2001. "Peer Effects with Random Assignment: Results for Dartmouth Roommates." *Q.J.E.* 116 (2): 681–704.

Sobel, Michael E. 2006. "Spatial Concentration and Social Stratification: Does the Clustering of Disadvantage 'Beget' Bad Outcomes?" In *Poverty Traps*, edited by Samuel Bowles, Steven N. Durlauf, and Karla Hoff. New York: Sage Found.

Topa, Giorgio. 2001. "Social Interactions, Local Spillovers, and Unemployment." *Rev. Econ. Studies* 68 (2): 261–95.

U.S. Census Bureau. 1997. "United States Census 2000: Participant Statistical Areas Program Guidelines." Washington, DC: U.S. Dept. Commerce.

Wahba, Jackline, and Yves Zenou. 2005. "Density, Social Networks and Job Search Methods: Theory and Application to Egypt." *J. Development Econ.* 78 (2): 443–73.

Weinberg, Bruce. 2000. "Black Residential Centralization and the Spatial Mismatch Hypothesis." *J. Urban Econ.* 48 (1): 110–34.

———. 2004. "Testing the Spatial Mismatch Hypothesis Using Inter-city Variations in Industrial Composition." *Regional Sci. and Urban Econ.* 34 (5): 505–32.

Weinberg, Bruce, Patricia Reagan, and Jeffrey Yankow. 2004. "Do Neighborhoods Affect Hours Worked? Evidence from Longitudinal Data." *J. Labor Econ.* 22 (4): 891–924.

Wellman, Barry. 1996. "Are Personal Communities Local? A Dumptarian Reconsideration." *Soc. Networks* 18 (4): 347–54.

Zax, Jeffrey S., and Daniel Rees. 2002. "IQ, Academic Performance, Environment, and Earnings." *Rev. Econ. and Statis.* 84 (4): 600–616.

Zimmerman, David J. 2003. "Peer Effects in Academic Outcomes: Evidence from a Natural Experiment." *Rev. Econ. and Statis.* 85 (1): 9–23.