Who gets the job referral?

Evidence from a social networks experiment *

Lori Beaman† and Jeremy Magruder‡

Abstract

We use recruitment into a laboratory experiment in Kolkata, India to analyze how social networks select individuals for jobs. The experiment allows subjects to refer actual network members for casual jobs as experimental subjects under exogenously varied incentive contracts. We provide evidence that some workers, those who are high ability, have useful information about the abilities of members of their social network. However, the experiment also shows that social networks provide incentives to refer less qualified workers, and firms must counterbalance these incentives in order to effectively use existing employees to help overcome their screening problem.

1 Introduction

Social networks influence labor markets worldwide. By now, an extensive empirical literature has utilized natural experiments and other credible identification techniques to persuade us that networks affect labor market outcomes.[1] We also know that a large fraction of jobs are found through networks in many contexts, including 30-60% of U.S. jobs (Bewley, 1999; Ioannides and Loury, 2004). In our sample in Kolkata, India, 45% of employees have helped a friend or relative find a job with their current employer. While these analyses have convinced us of the importance of job networks, the empirical literature has had far less to say about

*We thank the Russell Sage Foundation for funding part of the study, Prasid Chakraborty and the SRG team for outstanding fieldwork, and Katie Wilson for excellent research assistance. We also thank two anonymous referees, David Figlio, Alain de Janvry, Dean Karlan, Ethan Ligon, Ted Miguel, Rohini Pande, Betty Sadoulet, Laura Schechter and numerous seminar participants for comments. All errors are our own.

†Northwestern University. Email: l-beaman@northwestern.edu.

‡University of California-Berkeley. Email: jmagruder@berkeley.edu

[1] See for example Bayer et al. (2008); Beaman (2010); Kramarz and Skans (2007); Granovetter (1973); Laschever (2009); Magruder (2010); Munshi (2003); Munshi and Rosenzweig (2006); Topa (2001).

why job networks are so commonplace. In contrast, theory has suggested several pathways by which firms and job searchers can find social networks beneficial. For example, job seekers can use social network contacts to minimize search costs (Calvo-Armengol, 2004; Mortensen and Vishwanath, 1994; Galeotti and Merlino, 2009); firms can exploit peer monitoring among socially connected employees to address moral hazard (Kugler, 2003); and firms can use referrals as a screening mechanism to reduce asymmetric information inherent in the hiring process (Montgomery, 1991; Munshi, 2003).[2] Theory has also suggested a potential cost to relying on social networks to address these labor market imperfections: the use of networks in job search can perpetuate inequalities across groups in the long-run (Calvo-Armengol and Jackson, 2004). This paper provides experimental evidence on one of the mechanisms by which networks may generate surplus to counterbalance this cost, by examining whether social networks can and will provide improved screening for firms.[3] We create short term jobs in a laboratory in the field in urban India and observe how the actual referral process responds to random variation in the incentives to refer a highly-skilled employee. This allows us to determine whether participants have useful information about fellow network members.

We argue that disseminating job information is often not the primary reason that social relationships are formed and maintained. In a developing country setting like the one in this paper, the majority of the literature on networks emphasizes how individuals use network links to improve risk sharing and insure against idiosyncratic shocks (Udry, 1994; Townsend, 1994; Ligon and Schechter, 2010). Therefore, any empirical investigation of how social networks can influence labor markets must grapple with the fact that an individual may rely on his or her network in a variety of contexts, and there are likely spillovers from one context to another

[2] Moral hazard is highlighted as a reason for the use of referrals in Bangladeshi garment factories in Heath (2010), and Castilla (2005) highlights that on-the-job social connections provide support to new recruits using data from a call center in the U.S..

[3] We do not rule out reduced search costs and peer monitoring as additional reasons networks influence labor markets.

(Conley and Udry, 1994). These spillovers may cause networks to smooth search frictions using network links which do not represent particularly strong job matches. For example, individuals in networks which formed to share risk may not have the right information to identify good job-specific matches, or they may not be inclined to use that information (if they have it) in a way which benefits employers. There may be contingent contracts or simple altruistic relationships that encourage an employee to refer a poorly qualified friend over the person they believe to be most qualified for the job. Several studies have suggested that particular family relationships may be quite important in job network contexts (Loury, 2006; Magruder, 2010; Wang, 2011), and Fafchamps and Moradi (2009) argues that referrals in the colonial British army in Ghana lowered the quality of recruits due to referee opportunism. In a related context, Bandiera et al. (2007) show that without incentives, social connections decreased productivity due to on-the-job favoritism in a UK fruit farm. We must therefore consider carefully the decision problem faced by an employee who is embedded in a social network, as the network may create incentives counter to the firm's objectives.

This study examines the job referral process in Kolkata, India, using a laboratory experiment which exploits out-of-laboratory behavior. We set up a temporary laboratory in an urban area, and create jobs in an experimental setting by paying individuals to take a survey and complete a cognitively intensive task. Our employees are drawn from a pool of active labor market participants and are offered a financial incentive to refer a friend or relative to the job. While everyone is asked to refer a friend who will be highly skilled at the job, the type of referral contract and amount offered is randomized: some are proposed a fixed payment while others are offered a guaranteed sum plus a contingent bonus based on the referrals' performance (performance pay). The referrals are not themselves given any direct financial incentive to perform well. The incentives serve as a tool to reveal information held by participants and provide insights into competing incentives outside of the workplace. In order to isolate the effect of the performance pay contract on the selection of referrals, all individuals in performance pay treatments are informed that they will receive the full performance bonus before their referrals complete the task.

The controlled setting we create allows us to examine the complete set of on-the-job incentives faced by each of our employees, which would be difficult in a non-experimental setting. We show that there is a tension between the incentives offered by the employer and the social incentives within a network. When individuals in our study receive performance pay, so that their finder's fee depends on their referral's performance, they become 7 percentage points less likely to refer relatives, who are more integrated into our respondents' risk-sharing networks according to the survey data. This is a large change since less than 15% of individuals refer relatives. They are also 8 percentage points more likely to refer coworkers.

Analysis of referrals' actual performance in the cognitive task treatments shows that high performing original participants (OPs) are capable of selecting individuals who are themselves highly skilled, but that these individuals only select highly skilled network members when given a contract in which their own pay is indexed to the referral's performance. Low ability original participants, however, show little capacity to recruit high performing referrals. This result is consistent with the idea that only individuals who performed well on the test can effectively screen network members, and we provide evidence that low ability participants cannot predict the performance of their referrals.[4] We also document that some of our study participants are aware of these informational advantages: high ability participants are more likely to make a referral if they receive performance pay than low ability participants are, suggesting that the expected return to performance pay is larger for high ability participants. Finally, while young,

[4]Low ability participants may also have a lower network quality, an alternative hypothesis we can not rule out as we discuss in section 3.

well educated and high cognitive ability referrals perform best at the task, these observable characteristics cannot explain this productivity premium. This suggests that the information being harnessed by these high ability types is difficult for the econometrician to observe, and may be difficult for prospective employers as well.

The paper is organized as follows. The next section describes the context and experimental design, and section 3 provides a theoretical framework to interpret the impact of the exogenous change in the referral bonus scheme. Section 4 presents the results: OPs' decision to make a referral; the relationship between the OP and the referral; referral performance on the cognitive task, and how OPs anticipated their referrals to perform. Whether observable characteristics can explain performance is analyzed in section 5, and section 6 concludes.

2 Context and Experimental Design

The setup of the experiment is that an initial pool of subjects is asked to refer members of their social networks to participate in the experiment in subsequent rounds. The idea is that paid laboratory participants are fundamentally day labor. If we draw from a random sample of laborers, and allow these laborers to refer others into the study, we can learn about how networks identify individuals for casual labor jobs by monitoring the characteristics of the referrals, the relationships between the original participants and their referrals, and the performance of the referrals at the "job." By varying the types of financial incentives provided to our short-term employees, we observe aspects of the decision-making that occurs within networks, and the tradeoffs network members face when making referrals. The recruitment process into the laboratory therefore allows us to observe behavior which occurs outside of the laboratory.

Our study takes place in urban Kolkata, India. Many of our subjects work in informal and casual labor markets, where employment is often temporary and uncertain; these conditions are closely approximated by the day-labor nature of the task in our laboratory. Several characteristics of our experiment contribute to the external validity of results. First, our applicant pool are labor force participants from several neighborhoods in Kolkata. 91% of our sample are currently employed, 45% of whom have successfully made referrals at their current job. Our sample therefore constitutes individuals who are actively involved in network hires and reflects a diverse pool of workers, with heterogenous educational levels, ages, and labor market experiences including occupation. This kind of heterogeneity would not have been possible if we worked with one firm.

Second, participants receive Rs. 135 ($3.00) payment in the first round of the study, which is higher than the median daily income for the population in this study (Rs. 110). Our jobs therefore feature real world stakes, which provide strong incentives for participants to take the task seriously. The task itself is an assessment of cognitive ability and described in more detail below. The laboratory reproduced key features of a real world workplace: subjects were asked to complete the task and were closely supervised by a research assistant who provided instructions, allowed time for independent work, and evaluated performance in real time. Thus, while the experiment can not mimic employee referrals for permanent, salaried positions, it does generate real world stakes among workers in an employment environment, and offers what could be viewed as one additional temporary employment opportunity among many in a fluid labor market. Moreover, and important for our interpretations, we have full control over the various static and dynamic incentives provided by the employer.

Finally, providing cash bonuses to existing employees for referrals is an established practice in many firms, including some firms which index these bonuses to referral performance (Lublin, 2010; Castilla, 2005). However, in many employment settings, there are non-monetary incentives to induce good referrals: either positive (the ability to make additional referrals) or negative (the employee's reputation is tarnished if he makes a bad referral). Our experiment with a one-time job opportunity does not replicate this feature of the labor market. The advantage of the experimental design is that we can disentangle employees' ability to identify inherently good workers from other on-the-job dynamics, such as monitoring or competition, and we can think of the financial incentives as serving as a proxy for the incentives generated by the long-term relationship between the firm and the employee. We note that while other employers' non-monetary incentives are likely larger than the financial incentives we provide, so are the social incentives to procure a long-term job for a friend. Thus, in a relative sense, we expect our incentive treatments to generate comparable tradeoffs to those employees in many other contexts face. Given the strong evidence from the employer learning literature and elsewhere[5] that the full package of referral incentives that employers provide are insufficient to solve the problem of asymmetric information (Altonji and Pierret, 2001; Simon and Warner, 1992) we expect that the tradeoffs we measure are characteristic of an important problem in many labor markets.

The following describes the two main parts to the experiment: the initial recruitment and the return of the original participants with the referrals.

2.1 Initial Recruitment

We draw a random sample of households through door to door solicitation in a peri-urban residential area of Kolkata, India. Sampled households are offered a fixed wage if they send an adult male household member to the study site, which is located nearby. Sampling and initial invitations were extended continuously from February through June 2009, during which time

[5]For example, Bandiera et al. (2009) show that a similar incentive problem existed in a UK fruit farm until the researchers proposed a financial incentive scheme for managers.

we successfully enrolled 562 OPs in the cognitive treatment. Of those visited during door to door recruitment, 37% of households sent an eligible man to the laboratory.[6] Participants are assigned an appointment time, requested to be available for two hours of work, and are provided with a single coupon to ensure that only one male per household attends. Upon arrival at the study site, individuals complete a survey which includes questions on demographics, labor force participation, social networks, and two measures of cognitive ability: the Digit Span Test and Raven's Matrices. This initial group (original participants or OPs) faces an experimental treatment randomized along several dimensions. OPs are asked to complete one (randomly chosen) task: one task emphasizes cognitive ability while a second task emphasizes pure effort. The majority of our sample (including all high stakes treatment groups) was assigned to the cognitive task, which we focus on in this paper.[7]

In the cognitive task, participants are asked to design a set of four different "quilts" by arranging a group of colored swatches according to a set of logical rules.[8] The puzzles were designed to be progressively more challenging. A supervisor explains the rules to each participant, who is given a maximum time limit to complete each puzzle. When the participant believes he solved a puzzle, he would signal the supervisor who either lets the participant continue to the next puzzle if the solution is correct, or points out the error and tells the participant to try again, allowing up to three incorrect attempts per puzzle. More detail on the

[6]This participation rate compares well to other comparable studies, such as Karlan and Zinman (2009) who had 8.7% of individuals solicited participate in their experiment and Ashraf (2009) who had a 57% take up rate into a laboratory experiment among a sample of previous participants from a field experiment targeted to microfinance clients.

[7]In the effort task, participants are asked to create small bags of peanuts for 30 minutes. Due to limited resources, 1/3 of our sample was assigned to the effort treatment, and they received either the low stakes performance pay or low stakes fixed fee treatments described below. We did not find mean differences in performance for the referrals of OPs who completed the effort task. However, this may be because the sample is much smaller and does not include the high stakes treatments for OPs.

[8]In one puzzle, for example, the participant must fill in a four by four pattern with 16 different color swatches - 4 swatches of 4 colors - and ensure that each row and column has only one of each color. These puzzles are presented in greater detail in the online appendix. The left side represents unmovable squares in each puzzle and the right panel shows one possible solution.

task is given in the online appendix.

The measure of performance we use takes into account three aspects of performance: the time spent on each puzzle, whether the participant ultimately solved the puzzle, and the number of incorrect attempts. Incorrect attempts are important as proxies for how much supervisory time an employee requires in order to successfully complete a task. To utilize variation from all three components of performance, we use the following metric: a perfect score for a given puzzle is assigned for solving the puzzle in under one minute with no incorrect attempts. Incorrect attempts and more time spent lower the score, and a participant receives a zero if the puzzle is not completed within the allotted time. The score of the four puzzles is then averaged and standardized using the mean and standard deviation of the entire OP sample. We note that the main results are robust to sensible alternate measures of performance (for example, the number of puzzles solved correctly).

At the end of the experiment, individuals are paid Rs. 135 for their participation. They are also offered payment to return with a male friend or family member (a referral) between the ages of 18 and 60. All OPs are specifically asked to return with a reference "who would be good at the task you just completed." A second randomization occurs to determine the amount of payment the OP will receive when he returns with a referral. Payment varies along two dimensions: the amount of pay and whether pay may depend on the referral's performance. Participants are ensured that their payment will be at least a minimal threshold, and given the specific terms of the payment arrangement. OPs are informed of the offer payment immediately prior to their exit from the laboratory.

Among the OPs randomized into the cognitive task, there are 5 treatment groups:

Contract	Fixed Component	Performance Component	N of OPs
Low Stakes Performance Pay	60	0-20	116
High Stakes Performance Pay	60	0-50	136
Very Low Fixed Pay	60	0	71
Low Fixed Pay	80	0	117
High Fixed Pay	110	0	122

There are two performance pay levels: the high stakes treatment varies between Rs. 60 and 110 total pay while the low performance pay is Rs. 60-80. As fixed finder's fees, OPs are randomly offered either Rs. 60, 80 or 110.In all cases, the exact contract, including the requisite number of correct puzzles needed for a given pay grade, is detailed in the offer. All participants are asked to make an appointment to return with a referral in a designated three day window. In what follows, we denote the initial participation (where we recruit OPs into the laboratory) as round one, and the return of the OPs with referrals as round two.

Table 1 shows that the randomization created balance on observed characteristics of original participants from the baseline survey and round 1 performance. One exception is that OPs in the high powered incentives treatment group performed worse on the cognitive task compared to OPs in other treatments.[9] The average OP in the sample is approximately 30 years old, 34% of the initial subjects are between 18 and 25. 78% of OPs are the primary income earner in their household, while 32% are household heads. Almost all of the participants in the study are literate.

2.2 Return of OPs with referrals

When the original participants return with their referrals, the referrals fill out the survey and perform both the effort and the cognitive ability tasks.[10] A key feature of this study is that

[9]As randomization was done on a rolling basis, it was not possible to use stratification. Note, however, that the correlation between OP performance and referral performance is only .15. Therefore even a relatively large imbalance such as .18 of a standard deviation is unlikely to significantly alter the results.

[10]In order to minimize the potential for OPs to cheat by telling their referrals the solutions to the puzzles, we developed two sets of puzzles which are very similar, and we randomized which set was used in each laboratory session. The type of puzzle the OP was given is included as a control in all specifications.

both OPs and referrals have no private incentive to perform well on either task. However, there may be unobserved side payments indexed to referral performance (and creating a private incentive for referrals). The OP, for example, may give part of his finder's fee to the referral to entice a highly qualified network member to participate. To eliminate the incentive for such a side payment, both the OP and referral are informed that the OP will be paid the maximum performance bonus regardless of the referral's performance before the referral performs either task.[11] While referrals perform the tasks and complete the survey, OPs fill out a short interim survey about the process they went through in recruiting referrals.

3 Theoretical Framework

We present a stylized model, similar in spirit to Bandiera et al. (2009), to illustrate the potential tradeoffs an individual faces when asked by his employer to make a referral. By incorporating financial incentives provided by the firm and heterogeneity in imperfect information on the part of the network member, it also highlights how incentives can affect the choice of the referral and what we can identify in the experiment.

Employee i has the opportunity to make a job referral. In making a referral, i would choose from an ambient network of friends, each of whom has an inherent ability at the job, $\theta_j \in \left\{\theta^H, \theta^L\right\}$. In return for making a referral, his employer offers him a contract consisting of a fixed fee (F_i) and a performance incentive (P_i), where he will receive P_i if he correctly selects a high ability friend. He observes a signal of each friend's ability, $\hat{\theta}_j \in \left\{\theta^H, \theta^L\right\}$. For simplicity, that signal is accurate with probability β_i, that is, $\mathbb{P}\left(\theta = \theta^H | \hat{\theta} = \theta^H, i\right) = \mathbb{P}\left(\theta = \theta^L | \hat{\theta} = \theta^L, i\right) = \beta_i$. Naturally, $\beta_i \in [0.5, 1]$, and it may be heterogeneous among employees.

Employee i's expected monetary payoffs from referring a particular friend are a function

[11] This experimental design is similar in spirit to Karlan and Zinman (2009) and Cohen and Dupas (2009).

of his contract type (F_i, P_i), his signal of the selected friend's ability $\left(\hat{\theta}_j\right)$, and the accuracy of that signal. Following Bandiera et al. (2009) and Prendergast and Topel (1996), i also receives a payment σ_{ij} from referring friend j. This payment can be interpreted as an actual cash transfer or as a weighted inclusion of $j's$ income in $i's$ utility.[12] Since there are two ability "types" of friends, it is without loss of generality to focus on the decision between friend 1, for whom $\sigma_{i1} \in argmax(\sigma_{ij}|\hat{\theta}_j = \theta^H)$ and friend 2, for whom $\sigma_{i2} \in argmax(\sigma_{ij}|\hat{\theta}_j = \theta^L_j)$. Finally, i also has the option of declining to make a referral. Suppose the effort of making a referral will cost him c_i.[13]

If i selects friend 1, then he will receive in expectation $F_i + \beta_i P_i + \sigma_{i1} - c_i$. While if i selects friend 2, he will receive in expectation $F_i + (1 - \beta_i) P_i + \sigma_{i2} - c_i$.

Comparing these two expressions, i will select friend 1 if

$$P_i > \frac{\sigma_{i2} - \sigma_{i1}}{2\beta_i - 1} \tag{1}$$

He will further choose not to make a referral if

$$c_i > F_i + max\left\{\beta_i P_i + \sigma_{i1}, (1 - \beta_i) P_i + \sigma_{i2}\right\} \tag{2}$$

We observe three pieces of data which can speak to this model. First, we observe whether the OP chooses to make a referral; second, the relationship between the referral and OP, which we consider a proxy for $\sigma_{i2} - \sigma_{i1}$; third, we observe the referral's ability $\theta_{j.}$

As experimenters, we exogenously vary F_i and P_i. Equation 1 makes clear that variation in F_i should not affect the optimal referral choice (as F_i is a common payment to all potential referrals). This is a simple empirical implication of the model that we will take to the data. F_i does, however, increase the willingness of agents to participate in the referral process. We

[12]Symmetrically we could think of this as a reduction in future transfers i would otherwise have to make to this friend due to other risk sharing or network-based agreements.

[13]It is possible that different referrals require different exertions of effort; for example, it may require more effort to recruit a high ability referral who has better alternate options. Such additional effort is included in the payment term σ_{ij}.

discuss the implications of the joint participation and referral choice problem in section 4.1.

A second main empirical implication of the model is that there are four necessary characteristics for performance pay to change the choice of optimal referral: (a) networks must be heterogeneous, so that i observes friends with both types of signals; (b) there must be tradeoffs between network incentives and employer incentives ($\sigma_{i2} - \sigma_{i1} > 0$); (c) the tradeoffs must not be too large relative to P_i; and (d) employee i must have information, so that $\beta_i > 0.5$. In the experiment, if we observe a change in referral performance in response to performance incentives for some group of respondents, we will be able to conclude that those group members have all four of those characteristics. However, if a group does not change their referral choice in response to performance pay, we will not know which characteristics are missing.

There are several dimensions of heterogeneity in this model. We note that variation in social payments (σ_{i1}, σ_{i2}) and costs of participation (c_i) affect both the participation decision and the referral choice when participants face either a zero or positive performance pay component. In contrast, information (β_i) only affects these decisions when there is a positive performance pay component. This fact will help us disentangle whether heterogeneous treatment effects most likely reflect differences in information or differences in social payments or costs of participation.

4 Can Network Members Screen?

The model described in section 3 highlighted the potential tradeoffs an individual faces when making a referral. This framework suggested that contract type should influence referral behavior in terms of the choice of referral and also whether the OP will find it worthwhile to make a referral at all.

We will observe whether an OP makes a referral and an objective estimate of that referral's ability. We also will observe the relationship between the OP and his referral, which we interpret as a proxy for the social transfer. Since contract type is randomly assigned, we can use a straightforward strategy to analyze how performance pay affects the type of referral an OP recruits:

$$y_{ij} = \beta_0 + \phi_i + X_i\gamma + \epsilon_{ij} \tag{3}$$

where y_{ij} could represent participation in the experiment, the relationship between the OP and referral, or the referrals performance, while ϕ_i represents the OP's treatment categories and X_i include OP characteristics detailed in Table 2 and week fixed effects to eliminate any secular trends.

The model also suggested that different forms of heterogeneity in the underlying parameters of the decision problem may impact participation and referral choice in different ways. Of course, we cannot directly measure the σ_{ij}, c, or β parameters that our OPs respond to in order to test this model directly. Still, one important dimension where others have found heterogeneity in social effects is worker ability[14], which accords with theoretical assumptions in Montgomery (1991). If high ability workers receive a more accurate signal of their network members' ability, i.e. β is larger, then they will recruit higher ability referrals when given a performance pay incentive, and also be more likely to participate when offered performance pay. Therefore, we also investigate whether OP ability is an important dimension of heterogeneity.

In this spirit, and derived from the theory above, we also estimate:

$$y_{ij} = \delta_0 + \delta_1\theta_i + \sum_{k\in\ low,high} \delta_{2k} perf_{ik} * \theta_i + \phi_i + X_i\gamma + \epsilon_{ij} \tag{4}$$

where θ_i is OP i's ability, as captured by the OP's normalized test score (described in section 2.1)

[14]See for example Bandiera et al. (2010); Fafchamps and Moradi (2009); Yakubovich and Lup (2006) and Mas and Moretti (2009).

on the cognitive task. $perf_{ik} * \theta_i$ is the interaction of an indicator for whether the OP was in a performance pay treatment with stakes k and the OP's test score. ϕ and X are defined as before. Since ability may be related to any of the underlying parameters, we rely on supplemental data and theoretical restrictions across the referral choice and participation equations to indicate which dimensions of underlying OP heterogeneity create the referral patterns that we observe.

4.1 Returning with a Referral

As was made explicit in the theoretical framework, OPs face extensive and intensive margin choices. On the extensive margin, they choose whether or not to return with a referral. 72% of our OPs returned with a referral, so that 407 referrals participated in round 2. This high participation rate we believe reflects the value of the jobs we provided.

The model shows that an increase in the fixed component of the finder's fee should induce more OPs to return with a referral. Increases in the performance pay component will affect the participation decision depending on the information signal the OP has about their potential referrals. In table 2 we look at the impact of the fixed component using 2 different strategies. Column (1) shows the simplest specification, equation 3. We do not observe any differences in the high fixed or low fixed treatment categories compared to the excluded group, the very low fixed treatment. However, as shown in section 2.1, there are very few observations in the very low fixed pay group.[15] In order to increase power to test for whether OPs who expected to receive 110 Rs returned to the laboratory more frequently than OPs that expected to receive only 60-80 Rs, column (2) expands the control group and presents an alternative specification which looks at differential behavior only among individuals who seem likely to have expected 110 Rupees: those in the high fixed wage treatment, and those in the high performance pay

[15]The very low fixed group was by design smaller than the other groups due to budget constraints.

treatment who did well on the task themselves. The performance pay offer detailed that only the OPs who returned with a referral who got 3 or more puzzles correct would be guaranteed at least 100 Rs, so that if OPs measured expectations by their own performance, those who solved 2 or fewer puzzles correctly may have anticipated a low return.[16] Column (2) shows that in this specification, the high fixed treatment group is about 8 percentage points more likely to participate in round 2, and this effect is statistically indistinguishable from the return rate among the high performing high stakes group, who may have had similar expectations.

In the model, heterogeneity in information levels, β_i, only affects participation through changing the expected return to performance pay. Thus, if OP ability is a proxy for information, we should see more able OPs participate at different rates in response only to changes in performance incentives, but not to changes in fixed payments. Column (2) showed that high ability OPs in the high stakes performance pay treatment had a high participation rate in round 2. However, OP ability may be correlated with other underlying modelling parameters as well, such as the incentives provided by the network. If OP ability were correlated with heterogeneity in c_i (the costs of making a referral) or in σ_{i1} and σ_{i2} (the incentives provided by the network), it would be associated with differential participation in response to both the performance payment level and the level of the fixed payments. We therefore estimate equation 4 to test whether the heterogenous response by ability also occurs in the fixed treatments.

Column (3) shows that the high stakes performance pay sharply increases the participation rate among high ability OPs, but there is no heterogeneous effect among the other treatment groups. The result in column (3) is consistent with high ability OPs differing from low ability OPs in their level of information but not in their costs of participation or the network incentives. However, in a more general model with multiple ability types, OP ability may also

[16]The offer stated 4 puzzles would earn the OP 110 Rs, 3 puzzles 100 Rs, 2 puzzles 85 Rs and 2 or fewer puzzles would generate 60 Rs. Therefore we are assuming that OPs own performance is correlated with the signal they receive about their network members or the quality of their network.

be correlated with network quality: that is, the probability of having a high ability individual in his network. This would also generate a higher expected return to performance pay and be consistent with the result in column (3).[17] We will provide more direct evidence on the role of information in section 4.5.

While the participation decision yielded our first test for the presence of network information, differential participation rates between rounds 1 and 2 in the study could also bias the estimation of the referral choice equation. In fact, both theory and our empirical work suggested that participation in round 2 is related to key parameters of interest and treatment type. Simulations of the model (not presented here) suggest that even in the simplest case, where social incentives, information and participation costs are all independently distributed, the direction of the bias in estimating the interaction of β with performance pay on the sub-sample of round 2 participants cannot be signed.

Therefore, we use two main strategies to estimate the impact of contract change on referral choice. In our preferred specification, we employ a Heckman two step selection model with a first stage probit and second stage estimation including the inverse mills ratio from the first stage (Heckman, 1976). Rainfall makes a natural exclusion restriction, as it is random and it affects the desirability of travelling to our laboratory,[18] while not being correlated with performance in our (indoor) laboratory.[19] The weather data we have available includes an indicator for whether there was non-zero rainfall on each day of the study as well as the mean

[17]The data is suggestive, however, that many low ability individuals are likely to know high ability workers. In the fixed treatments, in which there is the least incentive to recruit high ability workers, we see that OPs in the bottom quartile of the performance distribution are as likely to bring in a referral who performs in the top quartile as the second quartile. While imperfect, this is suggestive that network quality alone may not be the binding constraint for low ability OPs.

[18]As there may be selectivity into the first round of the study, we also include an indicator for whether there was rainfall on the day the OP participates in round 1. We find that OPs who join the study on rainy days are less likely to attrit in the subsequent round, consistent with the hypothesis that OPs who attend despite the presence of rain are more committed to returning with a referral.

[19]Estimates are robust to allowing temperature, which is correlated with rainfall, to have a direct effect on performance, as shown in Online Appendix Table 2. The daily rainfall and temperature data were downloaded from Weather Underground, `http://www.wunderground.com`.

and maximum temperature on each day. As the exact day that an OP and his referral would have participated is unknown among the attrited population, we use the number of days that it rained in each OP's allotted 3-day window to return with a referral. Section 4.3 discusses the strength of the relationship between rainfall and participation.

A second approach is to combine the participation and referral choice decisions into one outcome of interest. For example, the task was to solve puzzles correctly, and OPs who did not return with a referral successfully solved zero puzzles in the second round. We therefore include zeros for their performance (and then normalize accordingly) and analyze performance using OLS on the full sample. The advantage of this strategy is that we can fully utilize the exogenous, random variation.

4.2 Responsiveness to Fixed Fees

The model predicted that variation in the level of fixed fees should not affect the choice of referral, at least once differential participation rates are properly accounted for. We have several characteristics that could be used to estimate the choice of referral, and those can be broadly categorized as characteristics based on relationships (a proxy for σ_{ij}), or characteristics related to productivity (a proxy for θ_j). Table 3 asks whether any of these characteristics are related to the level of payment among the fixed fee subsample.

In Table 3, the dependent variable is indicated in the column heading. All estimates are consistent with the theoretical prediction. First, columns (1) and (2) show that rainfall during the OP's window for recruitment significantly lowers the probability that the OP completes the study, and the joint test of both rainfall variables is above 8. The main results are in columns (3) through (7). Odd columns show estimates of the level effects of the different fixed fee payments, while even columns also include the interaction terms with OP performance. Across all specifications, the joint p-values of the overall effects of fixed fee and interaction terms are never close to significant. While not shown for brevity, all results are similar using OLS with the full sample. Since the data are consistent with the theoretical prediction that variation in fixed fees does not alter the referral choice problem, we combine all fixed fee treatments into a single control group in subsequent specifications and test the performance pay treatments against the fixed fee treatments jointly.

4.3 Relationship between Referrals and OPs

The referral choice equation suggested that one important dimension that should change with performance pay is the selection of referrals in terms of the network payoff σ_{ij}. In particular, if OPs respond to performance pay by changing their choice of referral, they should be shifting away from referrals who grant them larger social transfers in favor of those who generate a smaller transfer. Of course, we cannot directly estimate σ_{ij}; here, we focus on two salient relationships: co-workers and relatives. We anticipate that for both insurance and altruistic reasons, relatives are likely to donate larger social transfers than coworkers. The idea that relatives engage in more altruistic or risk-sharing arrangements than co-workers is supported by our survey data: over 35% of reported gifts occurred between relatives, while only 2% were between coworkers. High value (at least 500 Rs) gifts and loans demonstrate a similar pattern.

Table 4 shows the relationship between OPs and their referrals as a function of treatment type. Columns (1) and (2) demonstrate that rainfall during the OP's window for recruitment significantly affects the participation rate within the full cognitive sample. One extra day of rainfall within the 3 day referral cycle makes an OP 21 percentage points less likely to return with a referral to the laboratory. Moreover, the instruments jointly have power: the chi squared statistic is over 12 in both specifications. In subsequent tables, only the chi squared statistic from the joint test of significance of the two rainfall variables is shown.

Columns (3) through (6) examine coworkers and relatives, and report estimates from the Heckman specification. Individuals assigned to the cognitive high stakes performance pay treatment were almost 8 percentage points more likely to refer a coworker. This is a large effect since only 12% of OPs in the control group returned with a coworker as their referral. There is limited evidence again of heterogeneity: column (4) shows little evidence of heterogeneity in the response to performance pay.

Columns (5) and (6) show that the high stakes group was also less likely to refer a relative than the fixed fee groups. The result represents an economically significant change given that a small fraction of OPs refer relatives. There is again no evidence of a heterogeneous response by OP ability. Overall, table 4 is consistent with the model's prediction that performance pay may lead to a shift from a preferred reference, in this case a relative, to one with better anticipated skills, a co-worker. Finally, the results, shown in Online Appendix Table 1, are similar using OLS on the full sample. Whether the performance pay actually resulted in higher performing referrals is investigated in the next section.

4.4 Referral Performance and Response to Incentives

Table 5 shows how OPs responded to the incentives using referrals' performance on the cognitive ability task. Columns (1) though (3) show the Heckman selection model and columns (4) through (6) show OLS estimates from the full sample. Column (1) shows that there is no significant relationship on average between treatment type and performance in the Heckman specification. However, as seen in column (2), more able OPs recruited higher performing referrals. This would be consistent with a positive correlation between an OP's ability and the overall ability of the OP's network, or it may represent differential ability to screen. By interacting initial OP ability with performance pay in column (3), we see that the differential performance of referrals recruited by high ability OPs is driven by OPs who face performance pay incentives. Therefore, high ability individuals refer high ability people only when properly incentivized, suggesting that the networks of high ability OPs are heterogeneous and that high ability OPs do have the capacity to screen.[20] Columns (4) through (6) show that these results are similar when using OLS on the full sample: performance pay offers result in high ability OPs generating more round 2 puzzles solved. More detail on the relationship between OP and referral test scores is presented in the online appendix, which presents test score densities by treatment and also demonstrates the relationship between OP and referral test scores by OP-referral relationship and by treatment type.

4.5 Why are high ability OPs different from low ability OPs?

We observed in Table 4 that all OPs in the high stakes performance pay treatments respond to incentives by recruiting coworkers more often and recruiting relatives less often. Only high ability OPs, however, recruited referrals who actually performed better on the cognitive task. Thus, while all OPs change their referral choices in response to changing contractual conditions, only high ability OPs do so in a way which results in higher ability referrals. As the model emphasized, a variety of possible differences between high and low ability OPs could explain why performance incentives did not induce low ability OPs to recruit higher ability referrals: they may not know high ability referrals; they may lack information on the ability of their network members; or the tradeoff between their network incentives and the performance incentives may be too large.[21]

[20]OLS regressions using only the sample of round 2 participants show no significant relationship between treatment type nor heterogeneous effects by OP ability.

[21]Another possibility is that low ability OPs sought out a referral similar to themselves, mistakenly thinking they had performed well themselves. Given that OPs received real-time feedback on their performance, as described in section 2.1, and were told the exact number of puzzles their referral needed to get correct in order

We provide two pieces of evidence that differential information is at least one reason high ability OPs are successful in recruiting high quality referrals while low ability OPs are not. First, Table 2 showed that high ability OPs were more likely to make a referral when they were given performance pay but not when the level of the fixed component varied, which the theoretical model suggested would be due to additional information. However, variation in network quality - which is outside our model - is also consistent with that result. In this section, we supplement this argument with a direct investigation of OP knowledge. During the interim survey, OPs were asked how they expected their referrals to perform. The question was simply “How many puzzles do you think he [your referral] will solve correctly without making any mistakes?” The answer is between 0 and 4 puzzles. On average OPs thought their referrals would answer 3.5 puzzles correctly.

Table 6 shows the results of estimating a Heckman selection model of referrals’ test score performance on anticipated performance. To ease exposition, OPs are divided into discrete ability groups, where high ability OPs are those with a normalized test score above zero. Column (1) shows that high ability OPs are able to predict their referrals’ ability. The coefficient on anticipated performance implies that if an OP anticipated a perfect score, the referral did on average .8 of a standard deviation better than if the OP expected 0 correct puzzles. Low ability OPs, on the other hand, are not systematically able to predict their referrals’ performance, as shown in column (2).[22] Thus, while it may also be the case that low ability OPs have access to fewer high ability potential referrals or that network-based transfers are larger for these participants, Table 6 suggests that a lack of information on referrals’ capabilities is at least part of the reason low ability OPs do not respond to performance pay. This is consistent with the fact that all participating OPs adjust their behavior on the margin of relationships between

to earn the bonus, we think this is an unlikely explanation.

[22]A caveat applies however since the rainfall instruments are not powerful in the Heckman selection model in the low ability OP sample.

the OP and the referral: low ability OPs are trying to bring in higher ability referrals, but simply do not have a good understanding of which network members will perform better.

5 Identifying Good Referrals

High performing referrals tend to be young and low income, yet well educated and high-scoring on the Ravens and Digit Span tests, as shown in Online Appendix Table 3.[23] OPs therefore had to find referrals who would do well on the task specifically, not just the most successful individual in the network, as income would proxy for.

Can an employer just use these observable characteristics to screen recruits without the use of the network, or are social networks identifying productive, but hard to identify, employees? While we cannot mimic the full range of information that any prospective employer could observe through resumes, interviews, and other recruitment methods, we can at least discuss whether the productive characteristics which our high ability OPs are identifying can be explained by the other characteristics in our data. To test this, we add a variety of other characteristics to the main specification from Table 5, and present those results in Online Appendix Table 4. When we add in controls which should be easily observable in a resume (indicators for the referral's 5-year age group, each education level, and occupational category) and others which could be easily gauged (Ravens and Digit Span tests, income levels), β_2 remains statistically significant, and the point estimate is not substantially affected (changing from 0.370 in the main specification to 0.383 with the full vector of controls). That is, highly skilled, incentivized OPs are bringing in referrals who are highly skilled in ways which are hard to predict by the covariates in our data, even though some of those covariates are highly

[23]Given that the Raven and Digit Span tests have been used extensively in the psychology literature on measuring cognitive ability (Snow et al., 1984), this correlation provides reassuring evidence on the validity of our cognitive task.

correlated with puzzle task performance.[24]

6 Conclusion

Using a hybrid laboratory-field experiment in which temporary jobs were filled through social networks, our results indicate that at least some individuals have the ability to screen others in their networks to enhance firm productivity, and will do so if properly incentivized. This result validates the plausibility of the assumption that employees can help screen for their employer, at least in some contexts. However, we also find evidence that suggests that some workers could not screen effectively. Moreover, the workers who could screen were only willing to do so when they were directly incentivized, as they faced competing incentives generated by the network itself.

References

Altonji, J. and C. Pierret (2001). Employer learning and statistical discrimination. *Quarterly Journal of Economics 116*, 313–350.

Ashraf, N. (2009). Spousal control and intra-household decision making: An experimental study in the philippines. *American Economic Review 99*, 1245–1277.

Bandiera, O., I. Barankay, and I. Rasul (2007). Incentives for managers and inequality among workers: Evidence from a firm level experiment. *Quarterly Journal of Economics 122*, 729–773.

Bandiera, O., I. Barankay, and I. Rasul (2009). Social connections and incentives in the workplace: Evidence from personnel data. *Econometrica 77*, 1047–1094.

Bandiera, O., I. Barankay, and I. Rasul (2010). Social incentives in the workplace. *Review of Economics Studies 77*, 417–59.

Bayer, P., S. Ross, and G. Topa (2008). Place of work and place of residence: Informal hiring networks and labor market outcomes. *Journal of Political Economy 116*, 1150–1196.

Beaman, L. (2010). Social networks and the dynamics of labor market outcomes: Evidence from refugees resettled in the u.s. *Forthcoming, Review of Economic Studies*.

[24] Additionally, the full vector of controls renders the interaction of low-stakes performance pay with OP ability marginally significant, suggesting that high ability OPs in low stakes performance pay groups may also be identifying referrals who are unobservably productive.

Bewley, T. F. (1999). *Why Wages Don't Fall During a Recession*. Cambridge: Harvard University Press.

Calvo-Armengol, A. (2004). Job contat networks. *Journal of Economic Theory 115*, 191–206.

Calvo-Armengol, A. and M. O. Jackson (2004). The effects of social networks on employment and inequality. *American Economic Review 94*, 426–454.

Castilla, E. (2005). Social networks and employee performance in a call center. *American Journal of Sociology 110*, 1243–83.

Cohen, J. and P. Dupas (2009). Free distribution or cost-sharing? evidence from a randomized malaria prevention experiment. *Quarterly Journal of Economics 125*, 1–45.

Conley, T. and C. Udry (1994). Social networks in ghana. *Mimeo, Yale University*.

Fafchamps, M. and A. Moradi (2009). Referral and job performance: Evidence from the ghana colonial army. *BREAD Working Paper 238*.

Galeotti, A. and L. Merlino (2009). Endogenous job contact networks. *Mimeo, University of Essex*.

Granovetter, M. (1973). The strength of weak ties. *American Journal of Sociology 78*, 1360–1380.

Heath (2010). Why do firms hire using referrals? evidence from bangladeshi garment factories. *Mimeo, Yale University*.

Heckman, J. (1976). The common structure of statistical models of truncation, sample selection, and limited dependent variables and a simple estimator for such models. *Annals of Economic and Social Measurement 5*, 475–492.

Ioannides, Y. M. and L. D. Loury (2004). Job information networks, neighborhood effects, and inequality. *Journal of Economic Literature 42*, 1056–1093.

Karlan, D. and J. Zinman (2009). Observing unobservables: Identifying information asymmetries with a consumer credit field experiment. *Econometrica 77*, 1993–2008.

Kramarz, F. and O. Skans (2007). With a little help from my parents? Family networks and youth labor market entry. *CREST Working Paper*.

Kugler, A. (2003). Employee referrals and efficiency wages. *Labour Economics 10*, 531–556.

Laschever, R. (2009). The doughboys network: Social interactions and labor market outcomes of World War I veterans. *Mimeo, University of Illinois, Urbana*.

Ligon, E. and L. Schechter (2010). Structural experimentation to distinguish bewteen models of risk sharing with frictions. *Mimeo, University of Wisconsin-Madison*.

Loury, L. (2006). Some contacts are more equal than others: Informal networks, job tenure, and wages. *Journal of Labor Economics 24*, 299–318.

Lublin, J. (2010). Greasing the inside track to a job. *Wall Street Journal*, D4.

Magruder, J. (2010). Intergenerational networks, unemployment, and persistent inequality in south africa. *AEJ: Applied Economics 2*, 62–85.

Mas, A. and E. Moretti (2009). Peers at work. *American Economic Review 99*, 112–145.

Montgomery, J. D. (1991). Social networks and labor market outcomes: Toward an economic analysis. *American Economic Review 81*, 1408–1418.

Mortensen, D. and T. Vishwanath (1994). Personal contacts and earnings. It is who you know! *Labour Economics 1*, 187–201.

Munshi, K. (2003). Networks in the modern economy: Mexican migrants in the US labor market. *Quarterly Journal of Economics 118*, 549–599.

Munshi, K. and M. Rosenzweig (2006). Traditional institutions meet the modern world: Caste, gender, and schooling choice in a globalizing economy. *American Economic Review 96*, 1225–1252.

Prendergast, C. and R. Topel (1996). Favoritism in organizations. *Journal of Political Economy 104*, 958–978.

Simon, C. and J. Warner (1992). Matchmaker, matchmaker: The effect of old boys network on job match quality, earnings and tenure. *Journal of Labor Economics 10*, 306–330.

Snow, R., P. Kyllonen, and B. Marshalek (1984). *The topography of ability and learning correlations*, Volume 2 of *Advances in the Psychology of Human Intelligence*, pp. 47–103. Hillsdale, NJ: Erlbaum.

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

Townsend, R. (1994). Risk and insurance in village india. *Econometrica 62*, 539–5915.

Udry, C. (1994). Risk and insurance in a rural credit market: An empirical investigation in northern nigeria. *Review of Economics Studies 61*, 495–526.

Wang, S. (2011). Marriage networks, nepotism and labor market outcomes in china. *Mimeo, New York University*.

Yakubovich, V. and D. Lup (2006). Stages of the recruitment process and the referrers performance effect. *Organization Science 17*, 710–723.

Table 1: Randomization Check - Original Participant Characteristics

	High Fixed	Low Fixed	High Perf	Low Perf	Constant	N	P value of joint test
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Age of OP	-1.508	-1.684	-1.110	-0.422	31.000	562	0.70
	(1.414)	(1.425)	(1.387)	(1.428)	(1.125)		
OP is literate	0.031	0.044	0.032	0.035	0.887	562	0.88
	(0.041)	(0.041)	(0.040)	(0.041)	(0.033)		
OP had 5 or less years of schooling	0.034	0.016	0.029	0.035	0.155	562	0.97
	(0.058)	(0.058)	(0.056)	(0.058)	(0.046)		
OP had 5-10 yrs schooling	0.001	0.031	-0.051	-0.067	0.507	562	0.54
	(0.075)	(0.075)	(0.073)	(0.075)	(0.059)		
OP was married	-0.076	-0.082	-0.006	-0.087	0.535	562	0.53
	(0.075)	(0.075)	(0.073)	(0.075)	(0.059)		
OP was employed	-0.073	-0.052	-0.068	-0.070	0.958	562	0.51
	(0.045)	(0.045)	(0.044)	(0.045)	(0.036)		
Ln of Income earned by OP	-0.644	-0.507	-0.388	-0.491	7.365	562	0.52
	(0.372)	(0.375)	(0.365)	(0.376)	(0.296)		
OP is HH Head	-0.043	-0.022	-0.059	-0.071	0.338	562	0.83
	(0.068)	(0.069)	(0.067)	(0.069)	(0.054)		
OP is Primary Income Earner in HH	-0.084	-0.062	-0.046	-0.090	0.789	562	0.68
	(0.067)	(0.067)	(0.065)	(0.067)	(0.053)		
OP is 18-25 Years Old	0.066	-0.019	-0.014	0.027	0.352	562	0.64
	(0.072)	(0.073)	(0.071)	(0.073)	(0.057)		
Number of Ravens Correct	-0.045	-0.165	-0.153	-0.226	2.028	562	0.45
	(0.142)	(0.144)	(0.140)	(0.144)	(0.113)		
Number of Digits Correct	0.751	0.237	-0.096	0.169	11.831	562	0.37
	(0.518)	(0.522)	(0.508)	(0.523)	(0.412)		
Puzzle Type	-0.022	-0.037	0.012	-0.018	0.268	562	0.92
	(0.065)	(0.066)	(0.064)	(0.066)	(0.052)		
Normalized Test Score on All Puzzles	0.141	0.119	-0.180	0.014	-0.011	562	0.08
	(0.148)	(0.149)	(0.145)	(0.150)	(0.118)		
Puzzle Test Scores of Non-Attriting OPs	0.168	0.163	0.021	0.033	-0.041	407	0.70
	(0.169)	(0.172)	(0.167)	(0.173)	(0.134)		

Notes

1 OPs, or Original Participants, are the respondents who were recruited door-to-door. This table presents mean characteristics for OPs only and excludes (endogenously selected) referrals.

2 Each row is the regression results of the characteristics in the title column on the treatments. The regressions include the cognitive treatment sample and the omitted group is the very low fixed treatment in all rows. Column 7 shows the p value for the joint test of significance of all the treatment dummies.

Table 2: Was a Referral Brought In?

	(1)	(2)	(3)
OP Test Score * High Fixed Pay			0.043
			(0.067)
OPTest Score * Low Fixed Pay			0.064
			(0.068)
OP Test Score * High Perf Pay			0.162 **
			(0.066)
OP Test Score * Low Perf Pay			0.027
			(0.067)
OP Solved 3 or 4 Puzzles in High Perf Pay		0.152 ***	
		(0.055)	
OP Test Score			-0.038
			(0.054)
OP Treatment: High Fixed Pay	0.018	0.077 *	0.021
	(0.066)	(0.046)	(0.066)
OP Treatment: Low Fixed Pay	-0.034		-0.035
	(0.067)		(0.067)
OP Treatment: High Perf Pay	-0.026		-0.003
	(0.064)		(0.065)
OP Treatment: Low Perf Pay	-0.052		-0.051
	(0.067)		(0.067)
N	562	562	562
Mean of Dep Var for Excluded Group	0.761	0.694	0.761
SD	0.430	0.461	0.430

Notes

1 OPs, or Original Participants, are the respondents who were recruited door-to-door.

2 The dependent variable in all columns is 1 if the OP returned to the laboratory with a referral. The coefficients are from a linear probability model.

3 All columns restrict the sample to OPs in the cognitive ability treatments. Very Low Fixed Pay is the excluded group in columns (1) and (3). Column (2) uses very low fixed, low fixed, low performance and high stakes perf pay OPs who correctly solved 2 or fewer puzzles correct as the excluded group. These individuals had the lowest likelihood of expecting to win the bonus since they themselves performed badly.

4 OP Test Score is the metric of cognitive test performance discussed in section 2.1: a perfect score of 20 is awarded for a given puzzle when it is solved in under one minute with no incorrect attempts; incorrect attempts and more time spent lower the score. If a participant does not completed a puzzle within the allotted time, the score is zero. The score of the four puzzles is then averaged and standardized using the mean and standard deviation of the entire OP sample.

5 All columns include additional covariates: indicators for the OP's age group (18-19, 20-24, 25-29, 30-34, 35-39, 40-44, 45-49, 50-54 and 55 and above); highest grade level attained by the OP, the OP's ln of (income +1) in previous month; the type of puzzle the OP was given; the OP's performance on the Raven's Test and Digit Span Test; indicator dummies for week the OP participated in round 1 of the study and an indicator for partication during a weekend.

Table 3: Fixed Fee Treatments - Referral Choice

	First Stage		Relationship to OP: Co-worker		Relationship to OP: Relative		Referral Test Score
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Number of Days with Rainfall during OP's Referral Cycle	-0.166 **	-0.170 **					
	(0.085)	(0.086)					
Rainfall on OP Arrival Day	0.200 ***	0.202 ***					
	(0.064)	(0.063)					
OP Test Score * High Fixed Pay		0.068		-0.049		-0.021	-0.304
		(0.081)		(0.064)		(0.064)	(0.200)
OP Test Score * Low Fixed Pay		0.070		-0.079		-0.085	-0.139
		(0.076)		(0.066)		(0.065)	(0.202)
OP Test Score		-0.048		0.022		0.039	0.196
		(0.064)		(0.055)		(0.054)	(0.168)
OP Treatment: High Fixed Pay	-0.003	-0.008	0.010	0.013	-0.024	-0.031	0.072
	(0.080)	(0.081)	(0.057)	(0.057)	(0.056)	(0.056)	(0.179)
OP Treatment: Low Fixed Pay	-0.046	-0.049	0.055	0.061	0.009	0.013	0.192
	(0.079)	(0.079)	(0.059)	(0.059)	(0.058)	(0.057)	(0.183)
N	310	310	310	310	310	310	310
p value from joint test of treatment and treatment interactions			0.801	0.880	0.912	0.932	0.865
Mean of Dep Var for Excluded Group	0.761		0.130		0.148		-0.068
SD	0.430		0.339		0.359		1.166
Chi^2 statistic: joint test of rainfall variables	8.118	8.289	8.118	8.289	8.118	8.289	8.289
Mills: Coefficient			-0.199	-0.189	0.115	0.098	0.864
Mills: SE			0.166	0.165	0.164	0.164	0.507
N Censored Obs			81	81	81	81	81

Notes

1 OPs, or Original Participants, are the respondents who were recruited door-to-door.

2 The excluded treatment category is the very low fixed treatment. All columns include additional covariates as described in Table 2, and OP Test Score is as defined in Table 2.

3 An OP's "Referral Cycle" is the three days the OP had to choose from to bring in his referral. The exclusion restriction uses the number of days, from 0 to 3, where there was non-zero rainfall among the potential referral days for each OP.

4 Columns (1) and (2) show probit marginal effects.

5 Relative and co-worker are dummy variables indicating the relationship between the Original Participant and the referral. Columnns (3)-(7) are Heckman two step estimates with the rainfall variables from columns (1) and (2) used as exclusion restrictions. The first stage is shown in columns (1) and (2) with the F test of joint significance of the two rainfall variables.

Table 4: Relationship between OP and Referral

	First Stage				Co-worker				Relative			
	(1)		(2)		(3)		(4)		(5)		(6)	
Number of Days with Rainfall during OP's Referral Cycle	-0.207	***	-0.207	***								
	(0.065)		(0.065)									
Rainfall on OP Arrival Day	0.129	**	0.136	**								
	(0.059)		(0.058)									
OP Test Score * High Perf Pay			0.146	***			0.008				0.023	
			(0.053)				(0.048)				(0.049)	
OP Test Score * Low Perf Pay			-0.018				0.059				-0.001	
			(0.051)				(0.042)				(0.042)	
OP Test Score			0.009				-0.021				-0.003	
			(0.029)				(0.024)				(0.024)	
OP Treatment: High Perf Pay	-0.022		0.027		0.079	**	0.076	*	-0.070	*	-0.072	*
	(0.050)		(0.051)		(0.039)		(0.039)		(0.040)		(0.040)	
OP Treatment: Low Perf Pay	-0.046		-0.046		0.007		0.010		0.068		0.065	
	(0.055)		(0.054)		(0.043)		(0.043)		(0.044)		(0.044)	
N	562		562		562		562		562		562	
Mean of Dep Var for Excluded Group	0.761				0.130				0.148			
SD	0.430				0.339				0.359			
Chi^2 statistic: joint test of rainfall variables	12.743		13.056		12.743		13.056		12.743		13.056	
Mills: Coefficient					-0.082		-0.155		-0.071		-0.008	
Mills: SE					0.144		0.134		0.150		0.137	
N Censored Obs					155		155		155		155	

Notes

1 OPs, or Original Participants, are the respondents who were recruited door-to-door.

2 The excluded category is the fixed fee treatments.

3 An OP's "Referral Cycle" is the three days the OP had to choose from to bring in his referral. The exclusion restriction uses the number of days, from 0 to 3, where there was non-zero rainfall among the potential referral days for each OP.

4 Columns (1) and (2) show probit marginal effects.

5 Co-worker (columns (3)-(4)) and Relative (columns (5)-(6)) are dummy variables indicating the relationship between the Original Participant and the referral. These columns show Heckman two step estimates with the rainfall variables from columns (1) and (2) used as exclusion restrictions. The first stage is shown in columns (1) and (2) with the F test of joint significance of the two rainfall variables.

6 All columns include additional covariates as described in Table 2, and OP Test Score is as defined in Table 2.

Table 5: Task Performance and Treatment Type

	Referral Cognitive Ability Task Performance					
	Selection Model			OLS: Full Sample		
	(1)	(2)	(3)	(4)	(5)	(6)
OP Test Score * High Perf Pay			0.370 **			0.346 ***
			(0.159)			(0.128)
OP Test Score * Low Perf Pay			0.065			0.037
			(0.138)			(0.133)
OP Test Score		0.152 **	0.036		0.123 **	0.027
		(0.071)	(0.079)		(0.057)	(0.075)
OP Treatment: High Perf Pay	-0.135	-0.107	-0.084	-0.072	-0.045	-0.004
	(0.157)	(0.151)	(0.131)	(0.126)	(0.126)	(0.127)
OP Treatment: Low Perf Pay	0.068	0.077	0.078	0.014	0.019	0.013
	(0.172)	(0.164)	(0.144)	(0.136)	(0.136)	(0.135)
N	562	562	562	562	562	562
Mean of Dep Var for Excluded Group	-0.068			-0.539		
SD	1.166			1.320		
Chi^2 statistic: joint test of rainfall variables	12.743	13.449	13.056			
Mills: Coefficient	1.356	1.301	1.123			
Mills: SE	0.561	0.514	0.432			
N Censored Obs	155	155	155			

Notes

1 OPs, or Original Participants, are the respondents who were recruited door-to-door.

2 All columns also include the individual characteristics of the Original Participant, as defined in Table 2.

3 The dependent variable in all columns is the referrals' normalized performance on the cognitive task. It is constructed analogously as OP Test Score, which is described in the notes to Table 2.

Table 6: OP Ability to Predict Performance

	High Ability OPs (1)		Low Ability OPs (2)
OP's Anticipated Performance: Puzzle	0.190	**	0.027
	(0.090)		(0.082)
N	280		226
Chi^2 statistic: joint test of rainfall variables	13.908		4.193
N Censored Obs	78		77

Notes

1 OPs, or Original Participants, are the respondents who were recruited door-to-door.

2 The independent variable is the number of puzzles, from 0 to 4, that the OP expects the referral to solve correctly in the allotted time. The dependent variable is the measure of actual referral performance used in Table 5.

3 All estimates are from a heckman two step selection model.

4 Column (1) restricts the sample to high ability OPs: those with a normalized test score greater than 0 while column 2 uses the sample of OPs with a normalized test score less than 0.

5 All columns also include additional covariates of the OP as described in Table 2.

6 There are fewer observations than in Table 5 since there were 56 OPs who responded with 'I don't know' as the response to the question on anticipated performance and are dropped from the sample.