Who gets the job referral?

Evidence from a social networks experiment *

Lori Beaman[†] and Jeremy Magruder[‡]

October 2010

Abstract

We use recruitment into a laboratory experiment in Kolkata, India to analyze how social networks influence the selection of individuals into jobs. We examine whether individuals in social networks can screen fellow network members for particular abilities for an employer, and how this capacity changes with employee skills. The experiment allows subjects to refer members of their social networks to subsequent rounds of the experiment and varies the incentive schemes offered to these participants. We present evidence that individuals face a tradeoff between choosing the most qualified individual for the job and the individual who is ideal from the perspective of their social network, and that referral choice is affected by recruitment incentives. When faced with performance pay, individuals are more likely to refer co-workers and less likely to refer family members. High ability participants who are offered performance pay recruit referrals who perform significantly better on a cognitive ability task and also prove to be more reliable as evidenced by their choices in the trust game and performance on an effort task.

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

*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 David Figlio, Alain de Janvry, 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. (2005); Beaman (2009); Kramarz and Skans (2007); Granovetter (1973); Laschever (2005); Magruder (2010); Munshi (2003); Munshi and Rosenzweig (2006); Topa (2001).

through networks in many contexts, including 30-60% of U.S. jobs (Bewley, 1999; Ioannides and Loury, 2004). In our sample in Kolkata, India, 40% of employees have helped a friend or relative find a job with their current employer. The empirical literature, however, has had much less to say about why network jobs 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 may use referrals as a screening mechanism in order to reduce asymmetric information inherent in the hiring process (Montgomery, 1991; Munshi, 2003). Theory has also suggested a potential cost to relying on an informal institution like 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.[2] 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.

Bandiera et al. (2007), in their experiment with a fruit farm in the UK, show that social connections on the job do affect firm productivity but not by improving performance through peer monitoring. Without incentives, social connections actually drove down productivity due to favoritism.[3] This highlights that we must consider carefully the decision problem faced

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

[3]There is other work showing positive peer effects in the workplace (Mas and Moretti, 2009), but there is often caveats in that literature as well. Bandiera et al. (2010), for example, only find positive spillovers when an individual works alongside a friend who is more able than himself. When working alongside a friend who is less

by an employee who is embedded in a social network, as the network may create incentives counter to the firm's objectives. The potential incentive problem will likely be even stronger in environments where the network connections which assist in job search are more multi-purposed and stable than the connections of migrant college students (as in Bandiera et al. (2007)). In this paper, we advocate that a broader view of job networks is essential if we are to test the theoretical paths and interpret the common finding that networks matter. In order for firms to be able to use employees' social networks to improve recruitment processes, an employee who is requested to make a referral must act in the firm's interest. Suppose the firm's objective was to enhance productivity through recruiting better quality workers. If ability is highly correlated within networks, then a referral from a good employee is very likely to generate a good match for the firm. However, if there is heterogeneity in ability within networks, workers must be able to identify good matches and be willing to refer them.

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 tasks which alternatively emphasize cognitive skills or effort. Our employees 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 the possibility of a bonus based on the referrals' performance (performance pay). The referrals are not themselves given any direct financial incentive to perform well. This allows us to see which incentives lead to high quality workers, which incentives lead to the systematic referral of relatives, and which individuals respond best to these incentives. In order to isolate

able, the worker's productivity declines.

the effect of the performance pay contract on the selection of referrals, we give individuals in the performance pay treatments, once they have returned to the lab, the maximum payment from the range stipulated in their contract.[4] Finally, all participants are invited to play a series of dictator and trust games at the end of the second round of the experiment.

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 risk and insure against idiosyncratic shocks (Udry, 1994; Townsend, 1994; Ligon and Schechter, 2010b). Therefore, any empirical investigation of how social networks can influence labor markets must grapple with the fact that an individual may rely on their network in a variety of contexts, and there are likely spillovers from one context to another.[5] 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 needy, but poorly qualified friend rather than the person they believe to be most qualified for the job.[6] Several studies (Loury, 2006; Magruder, 2010; Munshi and Rosenzweig, 2010) have suggested that particular family relationships may be quite important in job network contexts. Further, if a firm is looking for an individual with a specific skill set,

[4]If there are side payments in which the original participant indexes the payment to the referral as a function of his performance, we would conflate the selection effect with the impact of the indirect incentives created by such a side contract. The experimental design is similar in spirit to Karlan and Zinman (2009) and Cohen and Dupas (2009).

[5]Conley and Udry (1994) discuss how different economic networks are interconnected in Ghana and the importance of multi-dimensionality of networks.

[6]In the context of risk sharing, other-regarding preferences, including notions of trust, altruism, and reciprocity, have been suggested as reasons why informal arrangements can persist in the absence of formal contracts (Foster and Rosenzweig, 2001; Ligon and Schechter, 2010a).

it may be unlikely that the individual who provides the employee with the best risk sharing is also the one with the best match in terms of skills. It is therefore an empirical question whether network-based referrals are useful for the firm, as suggested by Montgomery (1991).

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. In doing so, we are able to show that there is a tension between the incentives constructed by the employer and the social incentives within a network. When individuals in our study receive performance pay, they become 8 percentage points more likely to refer coworkers and 8 percentage points less likely to refer relatives. This is a large change since less than 15% of individuals refer relatives. Second, 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 can not predict the performance of their referrals.[7]

Finally, we explore which characteristics are sought by the successful high ability, incentivized OPs. A feature of the experimental design is that referrals have no direct financial incentive to perform well, as the experiment will pay them a fixed sum regardless of how well they do. This contract structure means that the original participants - especially those in performance pay treatments - need to select a referral who will perform well in a setting with a moral hazard problem. We provide evidence that OPs in performance pay treatments help

[7] Low ability participants may also have a lower network quality, an alternative hypothesis we can not rule out as we discuss in section 2.2.

solve this problem by recruiting inherently reliable referrals. First, their referrals transfer more money, in the role of player 2, to their partners in the trust game. Such transfers have traditionally been interpreted as a measure of trustworthiness (Karlan, 2005; Schechter, 2007a). We argue that it also reflects reliability and dependability. Second, high ability original participants in the performance pay cognitive treatments refer individuals who perform better on a second task, emphasizing effort, even though there is no incentive for a strong performance for either the OP or the referral. High ability, incentivized OPs also systematically bring in the type of young, high cognitive ability recruits who on average are successful at the task. However, neither these nor other observable characteristics can explain the productivity premium their referrals enjoy. 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 provides a theoretical framework for interpreting the experiment. The experimental design is described in detail in section 3, and the implementation of the experiment and data is presented in section 4. The results are presented in two sections. Section 5 shows how OPs respond to incentives in terms of their relationship with the referral, referral performance on the cognitive task, and how OPs anticipated their referrals to perform. The characteristics of the referrals, including measures of their reliability, are analyzed in section 6. We conclude in section 7.

2 Network Incentives in Job Allocation

2.1 Context

The experiment is designed to test if networks recruit individuals with match-specific skills and if so, under what conditions. Building on other studies which have examined the capacity of laboratory experiments to predict outcomes outside of the laboratory (Karlan, 2005; Ashraf, Karlan, and Yin, 2006; Zhang, 2010), our study utilizes the actual recruitment process into a laboratory experiment in the field to observe behavior which occurs outside of the laboratory.

The general setup of the experiment is that an initial pool of randomly selected subjects are 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 work to 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 employees, we observe aspects of the decision-making that occurs within networks and the tradeoffs network members face when making referrals. While there are some firms who offer cash bonuses for referrals, we think of the financial incentives used in this experiment as analogous to the incentives generated by the long-term relationship between the firm and the employee. If an employee is concerned about his reputation, he has an incentive to refer a good person. Financial incentives is therefore the laboratory counterpart to this mechanism firms use with long-term employees.[8]

Our study takes place in urban Kolkata. 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. Moreover, social networks are an important part of job search in this context as already discussed. In the experiment, we directly observe a job network allocating jobs, in this case the position of

[8]We can not determine in this paper whether most firms in fact solve the employee incentive problem. While reputation concerns may help alleviate the problem, evidence by Bandiera et al. (2009) shows that a similar incentive problem did exist in a UK fruit farm until LSE researchers proposed a financial incentive scheme for managers.

being a laboratory subject. Participants receive payment in the first round of the study, Rs. 135 ($3.00), which is higher than the median daily income for the population in this study (Rs. 110). Participants in the second round can earn even more. While the experiment can not mimic employee referrals for permanent, salaried positions, it does generate real world stakes and offers what could be viewed as one additional temporary or uncertain employment opportunity among many available 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.

2.2 Theoretical Framework

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

Individual i is choosing between two friends for a job referral. Each friend has an ability level θ, and a transfer t that he will give to person i if he is selected. This transfer can be thought of as the altruistic benefit i receives from referring that friend, the expected value of the favor the friend will return to i in reciprocation, or a direct monetary transfer.[9]

For each of these friends, i observes $\hat{\theta} \in \{\theta_H, \theta_L\}$. Similarly, the true ability of each friend is $\theta \in \{\theta_H, \theta_L\}$. We assume that $P\left(\theta = \theta_H | \hat{\theta} = \theta_H\right) = P\left(\theta = \theta_L | \hat{\theta} = \theta_L\right) = \beta$, and that $\beta \in [0.5, 1]$.[10] Thus, if $\gamma > .5$, the individual has useful information on his friends' capabilities.

[9] 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.

[10] The assumption that β is not less than 0.5 is an innocuous assumption. If $\beta < .5$, an analogous yet less intuitive problem can be set up where friend 1 provides the transfer and the results described here hold.

Suppose friend 1 has $\hat{\theta} = \theta_H$ and friend 2 has $\hat{\theta} = \theta_L$. Without loss of generality, we normalize friend 1's transfer to 0 and denote $\tilde{t}$ the net transfer that friend 2 will give for being chosen as the referral. In principle, the sign of $\tilde{t}$ is indeterminate. Further suppose the firm offers a financial incentive to refer workers which may be indexed to that worker's productivity, $\pi(\theta)$, where $\pi(\theta_H) \geq \pi(\theta_L)$.

If i selects friend 1, then he will receive in expectation

$$\beta\pi(\theta_H) + (1-\beta)\pi(\theta_L) \tag{1}$$

While if i selects friend 2, he will receive in expectation

$$\tilde{t} + \beta\pi(\theta_L) + (1-\beta)\pi(\theta_H) \tag{2}$$

Comparing these two expressions, i will select friend 1 if

$$\pi(\theta_H) - \pi(\theta_L) > \frac{\tilde{t}}{2\beta - 1} \tag{3}$$

The framework helps illustrate what can and can not be identified in the experiment. There are three necessary conditions for performance pay to sponsor a change in the choice of referrals: (i) there must be a tradeoff between ability and transfers, i.e. $\tilde{t} > 0$; (ii) the performance pay regime must be sufficiently high stakes compared to the network-based transfers $\tilde{t}$ and the skill heterogeneity in the network (θ_H versus θ_L); and (iii) individuals must be able to screen network members, i.e. $\gamma > .5$. We discuss each in turn.

If $\tilde{t} < 0$, then the employee's network incentives are aligned with the employer's, and friend 1 will be chosen as the referral regardless of the incentives provided by the firm. In contrast, if $\tilde{t} > 0$, there is a tradeoff between choosing the friend who an individual believes to be high ability and the friend who will provide a transfer with certainty but is believed to be low ability. This type of tradeoff is crucial to the experiment inducing changes in referral behavior based on performance pay incentives. In a more general model with a continuum of referral types, this type of tradeoff will be present whenever the highest (perceived) ability member is not the member who provides the largest transfer.

Second, the level of incentives provided by the employers relative to the return within the network are important. Friend 2 will be referred if the employer does not generate a sufficient premium in expectation for referring friend 1 (who i believes to be higher ability). For example, if the employer offers a fixed fee payment for referrals, then $\pi(\theta_H) = \pi(\theta_L) = \pi$. Friend 2 will then always be chosen, even if π is increased or decreased. This implies that experimental manipulation of a fixed finder's fee should result in no changes in referral behavior. Therefore, i must receive performance pay, where $\pi(\theta_H) > \pi(\theta_L)$ in order to be persuaded to refer friend 1, who is believed to be high ability. The performance pay also has to be sufficiently high stakes relative to the network-based transfers. This requires that $\tilde{t}$ be sufficiently small relative to $\pi(\theta_H) - \pi(\theta_L)$, and that there be sufficient heterogeneity in ability between friend 1 and friend 2. As networks become more homogeneous (so that $\theta_H - \theta_L \to 0$), individuals become less willing to give up transfers and require stronger performance incentives. As $\tilde{t}$ increases, i is less likely to refer the high ability friend.

Finally, if i has no information as to his friend's abilities, then $\beta = 0.5$ and regardless of the structure of the incentive contract, i receives a larger payment in expectation by referring friend 2. Therefore the final necessary condition is that $\beta > .5$. As information increases and γ becomes closer to 1, individual i will be increasingly willing to give up larger transfers and receive lower performance premia in order to refer friend 1 under a performance pay system. In terms of the experiment, this implies that a group with better information about their network members will be more likely to respond to performance pay incentives.

If we observe no differences in referral performance or referral characteristics between fixed and performance pay treatments, we will not be able to determine which of these three conditions have not been met. However, if we observe that referrals recruited under performance pay do perform better, this is evidence that networks both have valuable information about their members' abilities and face a tradeoff between referring the best person for the job and the best person "for the network".

The model also has predictions for interpreting heterogeneous treatment effects. If we observe that a sub-group (notationally, group A) of our OPs respond to our incentives by bringing in highly skilled referrals and that another group (say, group B) does not, then we will have identified that information and tradeoffs do exist at least for group A, i.e. that $\beta_A > 0.5$ and $\tilde{t}_A > 0$. However, the response of referral performance to OP incentives will not help us identify whether group A has greater information than group B, access to a network with a greater dispersion in member ability, or lower tradeoffs in terms of network transfers provided by low ability members. Any investigation into these channels will have to rely on supplementary information (for example, differences between groups in the ability to predict referral performance, which would be indicative of differences in information).

3 Experimental Design

As already described, the experiment consists of multiple rounds. In the first round, a random sample of men between the ages of 18 and 65 is invited to take a survey and perform either an effort-intensive or cognitive-intensive task for a fixed wage. They are offered incentives to refer another individual to the study. In subsequent rounds, the initial participants return with their referrals, who complete a survey and both tasks. Finally, all participants play a series of trust and dictator games and respond to a short concluding survey.

This section describes the three main parts to the experiment: the initial recruitment; the return of the original participants with the referrals, and the economic games.

3.1 Initial Recruitment

We draw a random sample of households 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 until June, 2009. 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 are asked to complete a survey on demographics, labor force participation and social network members. In addition, the survey includes two measures of cognitive ability: the Digit Span Test and Raven's Matrices.[11] The initial group (OPs) faces an experimental treatment randomized along several dimensions. OPs are then asked to complete one of two (randomly chosen) tasks: a task emphasizing cognitive ability or a second task emphasizing pure effort. The effort task asks participants to create small bags of peanuts for 30 minutes and is similar in spirit to the effort task in Jakiela (2009). This task was chosen to mimic the relatively simple and repetitive tasks which are often required in market work.[12] In the cognitive task, participants are asked to design a set of four different "quilts." In each quilt, the participant is asked to arrange group of colored swatches according to a set of logical rules[13]. More detail on this exercise is given

[11] These two measures have been validated by psychologists and are increasingly used in household surveys in developing countries.

[12] At the same time, we did not want to use a task which is commonly done among this population, as it would be too easy to simply refer an individual who has precisely that job. We are instead interested in networks' ability to identify individuals with skills which are difficult for firms to observe.

[13] 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. Participants (both OPs and referrals) are given one of two sets of analogous puzzles at random, allowing us to confirm empirically that in fact referrals do not perform better if they are given the same puzzles as their OPs. Puzzle type is used as a control in all regressions. These puzzles are presented in greater detail in the appendix. The left side represents

in the appendix. We have data on whether the participant gets the puzzle correct, the time it takes to complete the puzzle correctly, and the number of times he indicates to the experimenter that he thinks he has gotten the puzzle correct. Both tasks take place in separate rooms, so OPs assigned to the cognitive task do not observe the details of the effort task, and vice versa.

At the end of the experiment, individuals are paid Rs. 135 for their participation. They are also invited to return with a male friend or family member (a referral) between 18 and 60 and offered to be paid for the reference. 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[14].

There are seven treatment groups as demonstrated in the table below:

	Cognitive Task	Effort Task
Performance Pay	Low: Rs. 60-80 High: Rs. 60-110	Low: Rs. 60-80 -
Fixed Payment	Very Low: Rs. 60 Low: Rs. 80 High: Rs. 110	- Low: Rs. 80 -

For OPs doing the effort task, they are offered either Rs. 60-80 depending on performance pay or Rs. 80 as a fixed finder's fee. In all cases, the exact level of performance for the given task is specified in the offer. In the cognitive treatments, there are multiple levels.

unmovable squares in each puzzle and the right panel shows one possible solution.

[14]Both the group of OPs who responded to our solicitation to come into the study and the group of OPs who return with referrals are selected samples. The selection of OPs into the study mimics the selection that an employer would face after providing notice of a new casual job; as such this selected sample mimics a "selected sample" of employees and it does not confound inference. The selection of OPs to return with referrals will be explored carefully in what follows.

The highest performance pay offers between Rs. 60 and 110 while the low performance pay is Rs. 60-80 as in the effort task. As fixed finder's fees, OPs are randomly offered either Rs. 60, 80 or 100. All participants are asked to make an appointment to return with a referral in a designated three day window[15]. 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.

3.2 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. 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 randomized which set was used in each laboratory session.[16] A key feature of this study is that both OPs and referrals have no private incentive to perform well on either task. However, there may be unobserved side payments between the OP and the referral. 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 be sure that any unobserved side payments are not indexed to referral performance (creating a private incentive for referrals to try harder), all OPs are paid the maximum amount within the pay range they were told, eliminating any motivation for such a side payment. Participants in the cognitive task performance pay-high category, for example, are all paid 110 Rs.[17] While referrals perform the tasks and complete the

[15] Paying a financial reward to employees to make referrals may not seem representative of how most firms operate. However, we have anecdotal evidence that some Indian firms do pay finder's fees and evidence from the U.S. (Castilla, 2005) that firms pay financial rewards to employees and the rewards can be tied to performance. Finally, the financial reward in this case proxies for other, more diffuse, returns to making a good referral such as the opportunity to refer additional people to the firm in the future or a positive reputation in the eyes of one's supervisor.

[16] The type of puzzle used is included as a control in tables focused on the cognitive task.

[17] The experimental protocol states that both the OP and referral are informed of the good news before the referral performs either task. This eliminates the incentive for OPs to indirectly incentivize their referrals' performance.

survey, OPs fill out a short interim survey about the process they went through in recruiting referrals. Both the OP and the referral are informed when they arrive at the lab that there is an additional opportunity to earn more money by participating in a round of economic games after the referral has completed his tasks.

3.3 Economic Games

Once the referrals have completed the tasks, each OP and referral play two versions of the dictator game and two versions of the trust game. Every OP plays four economic games: two dictator and two trust games, one with his own referral and one randomly selected referral.[18] In this paper, we focus primarily on the version of the games where the OP and referral play together. In both sets of games, respondents are informed of the identity of their partner, and so we are able to observe whether, for exmaple, referrals play the games differently with the OPs who brought them in than they do with other OPs. However, everyone's decisions within the games are private information and are never shared with their partners.

Economic games, particularly dictator and trust games, have been widely used in development economics (Schechter, 2007a; Barr and Genicot, 2008). In the dictator game, player 1 is given Rs. 160 ($3.50) and asked how many Rupees he would like to give to player 2, who is specified by name. The literature has typically found that proposers give a non-zero amount to their partners in both anonymous and non-anonymous versions of the dictator game (**?**).

The protocol for the trust game is as follows: player 1 is given Rs. 80, with the option of dividing it between himself and his partner. His partner receives triple the amount donated, and, in turn, faces the choice of how much to return to the first actor. To keep the first player's action private, player 2 answers how much money he would give back for any possible donated

[18]The order of games, in terms of play with the random or referral partner, is randomized.

amount. While trust games have been played extensively in many contexts, a demonstrated limitation is that the first player's trust decision is confounded by risk preferences and altruism (Cox, 2004; Karlan, 2005; Schechter, 2007b). However, since surplus is not created when the second player returns some of the asset, the second player's decisions can be viewed as an estimate of "trustworthiness", as confirmed empirically in Karlan (2005). Since player two's decisions also likely reflect altruism, we will use each player two's decisions in the dictator game to proxy for altruism. Trustworthiness may be closely related to reliability or dependability, evaluated in this context by the effort a friend exerts in the absence of monitoring by the OP or direct private incentives for the referral. The sociology literature suggests that reliability and trustworthiness are important traits an employee seeks in a referral, especially in an environment with limited trust overall (Smith, 2008).

At the end, each participant gets an independent die roll to determine which game determines his pay. Participants are not told which game is used to determine their payment, and therefore the play of their partners (referral or otherwise) is difficult to infer.

Finally, participants were administered a post-experiment survey which asks whether the participant anticipates sharing any of his earnings. There is no evidence that OPs provide indirect incentives to their referrals, as discussed in section 3.2. Zero OPs report that they will share their finder's fees with their referral. Moreover, there is little evidence of side payments in general: only 14 referrals (out of 598) reported they would share their payment with their OP.

4 Data and Descriptives

Over the period of the study, we enrolled 850 OPs. 71% of our OPs returned with a referral, amounting to 607 referrals participating in our study. Given that respondents had to leave, find a referral, and return with the referral on another day, we believe this is a very low rate of attrition which reflects the value of the jobs we are providing. Nonetheless, it will have implications for our analysis and we discuss attrition at each stage of our study for our analysis in greater detail below.

There are multiple aspects of performance on each of the two tasks. The measure of performance we use for the cognitive ability task takes into account three aspects of performance: the time a participant spent on each puzzle, the number of incorrect attempts and whether the participant ultimately got the puzzle correct. We use a single metric in order to incorporate these three important components using the following functional form. A perfect score for a given puzzle, which would involve getting the correct answer in under one minute with no incorrect attempts, has a value of 20. Incorrect attempts and more time spent to get the correct answer lowers the score. The participate 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. Performance on the effort task is relatively straight-forward. We observe the number of bags of peanuts counted during three consecutive 10 minute intervals, and the number of bags with the correct number of peanuts in bags which were chosen at random. We use the normalized total number of bags created as our measure of peanut performance. Both of these metrics are somewhat arbitrary; reassuringly, the main results here are robust to sensible alternate measures of performance.

Table 1 shows a number of characteristics of original participants from the baseline survey of OPs and round 1 performance as a function of treatment type. Since treatment groups 2 and 3 are very similar and we have a smaller sample size for treatment group 3, in all the tables presented we pool the two treatments together as a cognitive task low fixed payment group. Overall, the randomization created balance on observed characteristics. One exception is that OPs in the high powered incentives treatment group performed worse on the cognitive task compared to OPs in other treatments.[19] Due to attrition, the sample of OPs who participate in round 2 is highly balanced. We note that this complicates the interpretation of the attrition results, especially with respect to heterogeneous responses by OP initial performance.

The average OP in the sample is approximately 30 years old, and 34% of the initial subjects are young, between 18 and 25. Households tended to send an adult son within the age range to participate in the study: only 32% of OPs are heads of households. Almost all of the participants in the study are literate.

5 Can Networks Screen?

5.1 Returning with a Referral

Which original participants chose to bring a referral back to the laboratory? Approximately 70% of OPs returned to the laboratory with a referral. Table 2 shows how the decision to return with a referral is a function of treatment type. The first column suggests that participation does not vary significantly with treatment type in the full sample. None of the treatment indicators are individually significant and, while not shown in the table, they are not jointly significant. Performance pay may induce differential attrition by ability. In particular, low ability individuals who are given high stakes incentives may be less likely to participate in the study. Indeed, column (2) demonstrates that individuals with a high initial test score randomly assigned a high stakes performance pay offer are more likely to recruit a referral. Columns (3) and (4) split the sample into high ability participants, those with normalized test scores greater than 0, and low ability participants, those with normalized scores less than 0. These columns

[19] As randomization was done on a rolling basis through the roll of a die, it was not possible to use stratification or pair-wise matching, as described by Bruhn and McKenzie (2008). Note, however, that the correlation between OP performance and referral performance is only .15. Therefore even a relatively large imbalance such as .25 of a standard deviation is unlikely to significantly alter the results.

reveal divergent responses to performance pay based on an individual's initial performance: high ability OPs are more likely to participate while low ability OPs are less likely. These results indicate that the type of incentive provided by an employer will affect the type of employee who chooses to engage in the recruitment process and make a referral. If only high ability employees bring in good referrals, as we investigate in the next section, performance pay may serve as a way for firms to screen employees to solicit referrals from and induce self-selection.

While attrition is an interesting outcome in its own right, differential attrition between rounds 1 and 2 in the study could bias analysis of other outcomes. Table 2 shows evidence of selective attrition, even conditional on observable characteristics described in a footnote in the table, since low ability OPs are less likely to participate in the study if they are randomly assigned the performance pay treatment. While the lowest possible payment in the performance pay treatment is equal to the lowest fixed payment, Rs. 60, in the cognitive treatment, we can not rule out *ex ante* a negative behavioral response to performance pay.

We use multiple strategies for dealing with the bias which may arise from differential attrition. In the tables presented in the paper, we employ the 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 serves as an exclusion restriction, as it affects attendance but should not be correlated with performance in our (indoor) laboratory.[20] Estimates are robust to allowing temperature, which is correlated with rainfall, to have a direct effect on performance; that specification is presented in Appendix Table 2. 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 and maximum temperature on each day.[21] While the exact day that an

[20]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.

[21]The daily rainfall and temperature data were downloaded from Weather Underground, `http://www.wunderground.com`.

OP and his referral would have participated is unknown among the attrited population, we do know each OP's window of 3 days in which he had to return with his referral. We therefore use the number of days, from 0 to 3, in each OP's window that it rained. Section 5.2 discusses the strength of the relationship between rainfall and participation.

The Heckman model achieves parsimonious specification of (unobserved) selection bias term through distributional assumptions. As we have no prior as to the validity of these assumptions, we also have performed three types of robustness. First, we relaxed the normality assumption by including a polynomial of the predicted values from a first stage probit in the second stage instead of the inverse Mills ratio as recommended by (Deaton, 1997). All results in this paper are robust to this alternative specification. Second, we can also view participation as an outcome. For example, the task was to solve puzzles correctly and each attritor successfully solved 0 puzzles. We therefore include zeros as their performance (and then normalize accordingly) and analyze performance using OLS. Finally, we can use a median regression and impute zeros for attritors as in Neal and Johnson (1997). The median regression set up assumes attritors have "worse" unobservables than the median subject. Though this has a less clear interpretation compared to interpreting attrition as an outcome of poor performance, it is a weaker assumption than imputing the zeros in OLS. The results presented in this section are robust to any of these approaches. If anything, the coefficients become larger and remain statistically significant. These results are omitted for brevity, but the tables are available from the coauthors upon request.

5.2 Relationship between Referrals and OPs

The model described in section 2.2 highlighted the potential tradeoffs an individual faces when making a referral. As a proxy variable for the social transfer t that a referral is making to an OP, we use their relationship. We anticipate that relatives in particular differ from friends and coworkers in the types of transfers - financial, insurance-oriented or altruistic - exchanged between two individuals. In this section, we analyze whether the relationship between the OP and referral changes when the offered contract from the employer is altered.

Random assignment of the recruitment contract provides a straightforward strategy to analyze how performance pay affects the type of referral an OP recruits:

$$y_{ik} = \beta_0 + \phi_k + X_k\gamma + \epsilon_{ik} \tag{4}$$

where y_{ik} is an indicator for the type of relationship between referral i and OP k and ϕ_k represents the 7 treatment categories. We focus on three salient relationships: co-workers, relatives and friends. X includes demographic characteristics of OPs from round 1 of the study and are described in the footnotes of Table 2. The Heckman specification, shown in Table 3, also includes the inverse mills ratio.

Table 3 shows the relationship between OPs and their referrals as a function of treatment type using the Heckman specification. Columns (1) and (2) demonstrate that rainfall during the OP's window for recruitment significantly lowers the probability that the OP completes the study. While not shown in the table, the marginal effect at the mean of the covariates of the number of days of rain during the OP's referral cycle in specification (1) implies that one extra day of rainfall makes it 21% less likely an OP will return with his 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 (8) examine the three most salient relationships identified in the survey using the Heckman selection model: coworkers, relatives and friends. The relationship measure is based on self-reports from the referral. Columns (3) and (4) shows that individuals assigned to the cognitive high stakes performance pay treatment were almost 8% more likely to refer a coworker. This is a large effect since only 15% of OPs returned with a coworker as their referral.[22] Columns (5) and (6) show that the high stakes group was also less likely to refer a relative. Both results represent an economically significant change given that a small fraction of OPs refer relatives. This result 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 skills, a coworker. Columns (7) and (8) show that there is no change in the probability of referring a friend.[23] Whether the performance pay resulted in higher performing referrals is investigated in the next section.[24]

5.3 Referral Performance and Response to Incentives

Random variation in treatment type provides exogenous variation to analyze whether the treatment type of the OP affects his referrals' performance:

$$y_{ik} = \delta_0 + \delta_1 \theta_k + \phi_k + X_k \gamma \ + \varepsilon_{ik} \tag{5}$$

where y_{ik} is the performance of referral i who was recruited by OP k; θ_k is the OP k's ability, as measured by the OP's performance of the task in phase 1 of the experiment; and ϕ_k is defined as before. If there is positive assortative matching in networks, we would expect that $\delta_1 > 0$.

[22] Coworkers may be difficult to define for those who are self-employed. Estimating the same specification excluding the self-employed creates qualitatively similar results.

[23] This may be due to the fact that the category friend is too broad to pick up changes and may mask changes in degree of friendship.

[24] In the next section, we will emphasize the importance of heterogeneity by OP ability, and many of our main results are the interaction of OP ability with treatment category. Estimating the same specification as in Table 4 column 3 indicates that the results presented in this section are not heterogeneous with respect to OP ability: the interaction term is not significant and the level effect of high stakes performance pay remains very similar in size and significance. We interpret this as evidence that all OPs are responding to the high stakes performance incentive by changing who they bring in as a referral, though there is substantial heterogeneity in how successful they are at identifying a high quality referral as discussed below.

We test whether financial incentives alter a network members' referral choice - i.e. selecting a referral based more on ability than on other criteria - if OPs in the performance pay treatments recruit referrals who perform better than referrals in fixed pay treatments.

Bandiera et al. (2010) find significant heterogeneity in social effects according to worker ability, which accords with the theoretical assumptions in Montgomery (1991).[25] If high ability workers receive a more accurate signal of their network members' ability, β is larger, then they will recruit higher ability referrals when given a performance pay incentive. Moreoever, work by Munshi (2003) suggests that high ability workers will have higher ability social network members to choose from, if properly incentivized. In this spirit, we also evaluate:

$$y_{ik} = \delta_0 + \delta_1\theta_k + \delta_2 perf_k * \theta_k + \phi_k + X_k\gamma + \epsilon_{ik} \tag{6}$$

where $perf_k * \theta_k$ is the interaction of an indicator for whether the OP was in a performance pay treatment and the OP's ability. If high ability OPs respond more to high powered incentives, then we anticipate $\delta_2 > 0$. If the across-the-board assortative matching assumption in Montgomery (1991) is correct, then $\delta_1 > 0$ in this specification as well.

Table 4 shows how OPs responded to the incentives on the cognitive ability task.[26] The first estimates are from OLS in columns (1) through (3) then the Heckman selection model in columns (4) through (7). Correcting for attrition is important: other than the result that

[25] In context of Bandiera et al. (2010), evaluating spillovers from an individual working in close proximity to his or her friend, they found that the average social effect was zero since high ability workers had the opposite response to peers than low ability workers.

[26] Appendix table 1 shows that there is no effect of the treatments on referral performance at the peanut task. This may be the result of the smaller sample size as the effort treatments constitute only 32% of the entire sample. The experimental design may have also made it difficult to observe a differential performance on the effort task. Suppose there are two components of performance: an inherent ability or work ethic and a portion based on effort. Referrals were chosen when OPs knew there was the possibility of performance pay. Accordingly, OPs may have selected people that would be likely to put in effort under a performance pay contract, if they indirectly incentivized their referrals. By removing the performance pay incentive, referrals had less of an incentive to perform well. If there is no inherent effort quality in these referrals, we would not observe any difference in performance. We do, however, observe a differential performance in the effort task among referrals recruited by incentivized cognitive OPs, suggesting that there are inherent qualities which affect performance on the effort task.

an OP's initial score is positively correlated with a referral's test score, there is no significant relationship between treatment type and performance. Using exogenous variation in rainfall reveals much more. Column (4) shows that there is no significant relationship between treatment type and performance in the full sample. However, as seen in column (5), more able OPs recruited higher performing referrals. As discussed in the theoretical example, this is 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 (6), we see that the differential performance of referrals recruited by high ability OPs is actually driven by OPs who face performance pay incentives. Therefore, high ability individuals refer high ability people only when properly incentivized.[27] We discuss column (7) in section 6.

A key component of the experimental design is paying the OPs the maximum amount of the performance pay range to disentangle selection of the referral from any indirect performance pay incentive the OP could have given the referral in an out-of-laboratory contract. Essential is that both the OP and the referral were informed of the change so that any informal contract can be renegotiated and the referral not be indirectly incentivized. In order to investigate first whether this part of the protocol was implemented rigorously, especially when the laboratory was busy, and second whether side payment contracts (to the extent they exist) were in fact renegotiated, we ran an additional set of experiments. There are three treatments: the first informed the OP of the good news about his payment but the referral was told nothing; the second was the full information treatment as described in the experimental protocol; and the third paid the OP according to the performance pay contract. Appendix Table 3 shows the

[27] While rainfall may affect the probability an OP-referral pair participate in the second round of the study, rainfall is also correlated with temperature. Rainfall means lower temperature, which may increase performance in a hot climate such as Kolkata. Appendix Table 2 shows the results of the same specification as in Table 4 with an average maximum temperature on referral days an additional covariate. The results are robust to this specification.

results. If there were side payments indirectly incentivizing referrals, we would anticipate that referrals in treatments 1 and 3 would have better performance than those in treatment 2. This is not the case: there are no significant differences across any of the treatments. The standard errors are large, which may be the result of a relatively small sample size even though the number of observations per treatment is approximately 60% of the size of the treatments in the main study. Moreover, the coefficients on the interaction terms in column (2) are negative, the opposite of what would occur if there were side performance pay contracts, and so is the coefficient on the level effect of the no information to referral treatment. This is consistent with the data already discussed on anticipated transfers between OPs and referrals, showing no unobserved out-of-laboratory contracts where the OP pays the referral.

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

We observed in Table 3 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 are changing 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 theoretical example in section 2.3 emphasized, a variety of possible differences between high and low ability OPs could explain why low ability OPs do not bring in higher ability referrals in response to performance incentives. In particular, in order for our experiment to find a performance premium for referrals, the OP needed to have several characteristics: the OP had to know some high ability referrals ($\theta_H > \theta_L$); he had to have information as to the ability of his network members ($\beta > .5$), and he had to face a tradeoff between network incentives and the performance incentives generated by the experiment $\left(\tilde{t} > 0\right)$. If low ability OPs lack any of these characteristics, then they would not have reacted to performance pay by recruiting higher quality referrals.

Unfortunately, we do not observe social transfers or the full ambient network of potential referrals, and so we cannot provide direct evidence on the heterogeneity in quality of available referrals or on the presence and magnitude of social transfers and tradeoffs. However, we can provide evidence on differences in information by OP ability. In particular, 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. OPs were quite optimistic; on average OPs thought their referrals would get 3.5 puzzles correct. Table 5 shows the results of estimating both OLS and a Heckman selection model of referrals' test score performance on anticipated performance. Columns (1) and (2) show that high ability OPs, those with a normalized test score above zero, 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 columns (3) and (4).[28] Thus, low ability OPs do not appear to have useful information on referral's capabilities. 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 5 suggests that a lack of information is at least part of the reason low ability OPs do not respond to performance pay. Moreover, it is consistent with the fact that all OPs adjust their behavior on the margin of relationships between the OP and the referral: if low ability OPs are trying to bring in higher ability referrals, but simply do not have a good understanding of

[28] A caveat applies however since the rainfall instruments are not powerful in the Heckman selection model in the low ability OP sample, as shown in column (4).

which referrals will perform better, then we might expect to see patterns very similar to those presented here.

6 Reliability and Referral Characteristics

Reliability

An important characteristic of a good employee, and therefore a good referral as argued by some sociologists, is reliability and trustworthiness. That is, an OP must be confident that the referral will perform well on the job and realize his potential as a good match, often in the presence of moral hazard. The structure of the contract in the experiment emphasizes this aspect of referral choice. The OP could earn a bonus based on the referral's performance, but the referral himself has no performance incentive. Therefore the OP must select a referral who will exert effort without any monitoring by the OP. As already discussed, the OP can minimize this principle-agent problem by either choosing inherently reliable referrals or by introducing a side contract where the referral is paid by the OP based on his performance.[29] We start this section by investigating whether OPs selected more reliable referrals as a response to incentives.

Reliability is measured two ways in this study: first, by the referral's transfers made as player 2 in the trust game, where referrals similarly take a hidden, unincentivized, and unmonitored action which can directly contribute to the OPs payoff. Second, we also observe cognitive ability treatment referrals perform the effort task. Since no one in the cognitive treatment (neither OPs or referrals) ever believe they are incentivized on their performance on the effort task (and OPs in the cognitive have never even seen this task performed), we also interpret this performance as evidence of effort exerted in a fully unincentivized environment. Table 6 shows how referrals play the trust game with their OPs. Recall that referrals play the

[29] In the theoretical framework presented in this paper, the optimal choice between the two will depend on how reliability is distributed with ability and private transfers.

role of player 2 in the trust game, deciding how much to return to player 1 for each of the eight possible transfers the OP could have made to the referral (the OP was allowed to transfer up to 80 Rupees in 10 Rupee increments). The dependent variable used in this specification is the average transfer across all eight decisions, normalized separately at each Rs 10 interval.

The first two columns show that referrals' transfers in the trust game do not vary significantly across treatment type on average, nor is there a correlation between the OP's puzzle ability and the amount the referral transfers in the trust game. We again focus on heterogenous response to performance pay by OP ability in columns (3)-(5). Column (3) shows that high ability OPs in the high stakes performance pay treatment refer individuals who transfer significantly more to their OPs. The literature usually interprets player 2's decisions as a measure of trustworthiness. However, the measure may conflate trustworthiness with other preferences such as altruism. Therefore, column (4) presents the results of a specification which also includes both the amount the OP and the referral decided to transfer to one another in the dictator game. These measures should capture the altruistic relationship between the OP and referral. The effect of performance pay among high ability OPs continues to be positive and statistically significant. We interpret this as evidence consistent with the trust game measuring the referral's trustworthiness. Recall that each referral played the trust game with his OP and then another randomly selected OP. Column (5) shows that referrals of incentivized, high ability OPs are more trustworthy vis-a-vis random partners, providing further evidence of the interpretation of referral trustworthiness as reliability and not altruism or reciprocity towards OPs who just did them a favor.

As an additional approach to measuring referrals' reliability, we return to Table 4 where in column (7) we show how referrals recruited by OPs in the cognitive ability treatments perform on the peanut task. If incentivized OPs recruited individuals who would perform well and exert effort consistently, as we have argued, they would likely also perform well on the peanut task. This is in spite of the fact that the OPs did not know their referral would be asked to perform the effort task, and there was accordingly no incentive attached to the task. The analysis confirms that referrals recruited by incentivized OPs out-perform their counterparts on the effort task since the interaction between OP ability and high stakes performance pay is positive and significant.

Identifying Good Referrals

As discussed in the introduction, theoretical models suggest that social networks may smooth information asymmetries as employees can identify referrals who are productive in a way which is hard to observe for a prospective employer. Here, we have provided evidence that high ability OPs do identify high ability referrals when properly incentivized. However, we have said little about whether the information they use is something which should be easy or difficult for an employer to observe. In this section we will first describe which observable characteristics are predictors of performance and then look to see if OPs select individuals who perform better than what their observables alone would predict.

Appendix Table 4 looks at other characteristics, including other cognitive tests, age, education and income. Panel A shows how the referral's performance on the puzzle test correlates with these characteristics. Columns (1) and (2) show that the two cognitive ability tests we included in the background survey, the Raven Test[30] and the Digit Span Test, are positively and significantly correlated with puzzle performance. Panel B shows the same specification as used in Table 4 with alternative dependent variables. Given the similarity between the Raven test and the cognitive ability task, it is intuitive that we see evidence of a similar relationship between OP ability and treatment type for the Raven measure as we did for task performance

[30] Since the Raven test asks participants to identify patterns, it is the closest conceptually to the puzzle test.

in Table 4. Education is also positively correlated with referral puzzle performance, as shown in column (3). Similar to the digit span test, however, we do not observe any significant differences in these characteristics by treatment type. Columns (4) and (5) highlight an interesting aspect of the cognitive ability task used in this study. Referrals' age and income are both negatively correlated with puzzle performance.[31] 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.

Some easy to observe characteristics like age and education are strong predictors of performance. Are OPs then merely selecting individuals with characteristics which predict performance and are easy to observe? 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 productivity characteristics which our high ability OPs are identifying can be explained by the other characteristics in our data. To test this, we regress

$$y_{ik} = \beta_0 + \beta_1\theta_k + \beta_2 perf_k * \theta_k + \phi_k + X_k\gamma + W_i\delta + \epsilon_{ik}$$

where θ_k, ϕ_k, and X_k is $\text{OP}'_k s$ ability, treatment group, and characteristics as before, but now we control also for a vector of the referral's characteristics, W_i, observable in our data. Thus, we will test whether high ability, incentivized OPs bring in referrals who are highly productive in a way which is hard to observe in the data. Table 7 presents the results of this estimation. In column (1), we reproduce the analysis from Table 4 with the sample restricted to observations where we observe all the referral characteristics used in the table. Column (2) adds in characteristics which should be easily observable in a resume and allow for a flexible relationship between these characteristics and productivity: specifically, we add in

[31] The negative correlation between income and puzzle performance is largely driven by the age effect.

indicators for the referral's 5-year age group, each education level, and occupational category. These resume controls do not impact the size or precision of the coefficient estimate, suggesting that what OPs are identifying is something that would be hard to observe in a resume. The remaining columns in table 7 include additional covariates which may be less frequent on a resume but which we can observe and may be correlated with other characteristics observable to prospective employers. Column (3) adds the referral's performance on the Ravens and Digit Span tests, while column (4) includes the referral's transfer during the trust game. Column (5) includes the referral's income as well. As before, referral performance on the cognitive tests is associated with referral performance on the puzzle task. However, in all specifications, β_2 remains statistically significant and the point estimate does not change dramatically. 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 correlated with puzzle task performance. These results suggest that networks are identifying skillsets which may indeed be hard for prospective employers to observe.

7 Conclusion

Job networks are a ubiquitous phenomenon in labor markets, in both developed and developing countries. Individuals serve as explicit references for other individuals and also as conduits of information about new job postings. While a large literature in economics and sociology have described the presence of these networks, we know little about how these networks select referrals. Simply understanding that social connections facilitate job allocation says little about the welfare consequences of this system, both from the perspective of potential workers and firms. This paper begins to look inside the black box of social networks and concretely identifies the efficiency gains from using job networks to spread jobs under a variety of incentive schemes.

First, financial incentives do lead to a change in the type of referral who is chosen: coworkers are more likely to be referred at the expense of relatives. This points to a tradeoff individuals may face between a potentially more productive referral and a referral who has other network-oriented benefits.

The analysis also indicates that while performance pay induces all employees to change who the types of relationships that they share the referral with, only employees with an initially high ability change their referral in a way which boosts productivity. This suggests directly that high ability workers have information on the capabilities of network members and that they face a trade-off between the friend who will reward them most for the referral (either in terms of monetary or social payments) and the friend who they think will perform well on the task. Low ability coworkers, in contrast, do not respond to performance incentives by referring a high quality referral, which could in principle result either because they do not have the capacity to screen their network members effectively or because they do not have enough high ability coworkers in their network to take advantage of the incentive. We further present evidence that low ability workers are unable to predict the performance of their referrals (in contrast to high ability workers, who can do so successfully). This suggests that a lack of information may hamper the effectiveness of low ability individuals.

Taken together, the evidence suggests that at least some individuals have the ability to screen others in the 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 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. Job networks are composed of individual members with heterogeneous capabilities and diverse underlying incentives. This research highlights the importance of a disaggregate understanding of networks which considers individual abilities and the full incentive environment if we want to predict what happens when referrals are allowed to filter through networks. Moreover, while this example focused on the capacity of a network to identify a high quality referral, this conclusion may be relevant for other uses of networks as well, such as social learning. Future research will aim at understanding whether individual incentives or capabilities are important constraints on the capacity of networks to distribute a variety of information, goods, and opportunities.

References

Ashraf, N., D. Karlan, and W. Yin (2006). Tying odysseus to the mast: Evidence from a commitment savings product in the philippines. *Quarterly Journal of Economics 121*, 635–672.

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 (Forthcoming, 2010). Social incentives in the workplace. *Review of Economics Studies*.

Barr, A. and G. Genicot (2008). Risk pooling, commitment and information: An experimental test. *Journal of the European Economic Association 6*, 11511185.

Bayer, P., S. Ross, and G. Topa (2005). Place of work and place of residence: Informal hiring networks and labor market outcomes. *NBER Working Paper 11019*.

Beaman, L. (2009). Social networks and the dynamics of labor market outcomes: Evidence from refugees resettled in the u.s. *Mimeo, Northwestern University*.

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

Bruhn, M. and D. McKenzie (2008). In pursuit of balance: Randomization in practice in development field experiments. *World Bank Policy Research Working Paper No. 4752*.

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 Forthcoming*.

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

Cox, J. (2004). How to identify trust and reciprocity. *Games and Economic Behavior 46*, 260–281.

Deaton, A. (1997). *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. Washington, DC: World Bank and Johns Hopkins University Press.

Foster, A. and M. Rosenzweig (2001). Imperfect commitment, altruism, and the family: Evidence from transfer behavior in low-income rural areas. *Review of Economics and Statistics 83*, 389–407.

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.

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.

Jakiela, P. (2009). How fair shares compare: Experimental evidence from two cultures. *Mimeo, Washington University*.

Karlan, D. (2005). Using experimental economics to measure social capital and predict financial decisions. *American Economic Review 95*, 1688–1699.

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

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. (2005). The doughboys network: Social interactions and labor market outcomes of World War I veterans. *Mimeo, University of Illinois, Urbana*.

Ligon, E. and L. Schechter (2010a). Motives for sharing in social networks. *Mimeo, University of Wisconsin-Madison*.

Ligon, E. and L. Schechter (2010b). 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.

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.

Munshi, K. and M. Rosenzweig (2010). Does political accountability improve leadership competence? *Mimeo, Brown University*.

Schechter, L. (2007a). Theft, gift-giving, and trustworthiness: Honesty is its own reward in rural paraguay. *American Economic Review 97*, 1560–1582.

Schechter, L. (2007b). Traditional trust measurement and the risk confound: an experiment in rural paraguay. *Journal of Economic Behavior and Organization 62*, 272–292.

Smith, S. (2008). *A Question of Access or Mobilization? Understanding Inefficacious Job Referral Networks among the Black Poor*, pp. 157–184. Social Capital: An International Research Program. Oxford: Oxford University Press.

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.

Zhang, J. (2010). Gender difference in competitive inclination and educational outcomes: Evidence from three ethnic groups. *Mimeo, University of California, Berkeley*.

Table 1: Randomization Check

	OP Treat 1	OP Treat 4	OP Treat 5	OP Treat 6	OP Treat 7	Constant	N	P value of joint test
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Age of Resp	-0.460	-0.062	0.625	-0.046	-0.133	29.95	823	0.97
	(1.087)	(1.053)	(1.104)	(1.072)	(1.060)	(0.682)		
Resp is 18-25 Years Old	0.078	-0.002	0.039	0.019	0.005	0.340	823	0.75
	(0.056)	(0.054)	(0.057)	(0.055)	(0.055)	(0.035)		
Resp is literate	0.003	0.004	0.008	0.023	0.017	0.915	823	0.98
	(0.031)	(0.030)	(0.031)	(0.031)	(0.030)	(0.019)		
Resp had 5 or less years of schooling	0.024	0.019	0.025	-0.016	-0.030	0.165	823	0.80
	(0.044)	(0.042)	(0.044)	(0.043)	(0.042)	(0.027)		
Resp had 5-10 yrs schooling	-0.018	-0.071	-0.087	0.020	0.007	0.527	823	0.42
	(0.058)	(0.056)	(0.059)	(0.057)	(0.057)	(0.036)		
Resp was married	-0.025	0.045	-0.036	0.032	0.080	0.484	823	0.41
	(0.058)	(0.056)	(0.059)	(0.057)	(0.057)	(0.036)		
Resp was employed	-0.040	-0.036	-0.038	0.043	-0.031	0.926	823	0.13
	(0.033)	(0.032)	(0.034)	(0.033)	(0.032)	(0.021)		
Ln of Income earned by respondent	-0.329	-0.072	-0.175	0.498	-0.170	7.05	823	0.12
	(0.281)	(0.273)	(0.286)	(0.277)	(0.274)	(0.177)		
Resp is HH Head	-0.029	-0.045	-0.057	0.027	0.006	0.324	823	0.67
	(0.054)	(0.052)	(0.055)	(0.053)	(0.052)	(0.034)		
Number of Ravens Correct	0.058	-0.051	-0.124	-0.019	0.074	1.93	823	0.61
	(0.111)	(0.107)	(0.112)	(0.109)	(0.108)	(0.069)		
Number of Digits Correct	0.603	-0.243	0.021	-0.213	0.360	11.98	823	0.33
	(0.405)	(0.392)	(0.411)	(0.399)	(0.395)	(0.254)		
Ln Income Among Non-Attriting OPs	-0.154	-0.164	-0.092	0.615	-0.261	6.89	594	0.21
	(0.346)	(0.341)	(0.360)	(0.344)	(0.345)	(0.219)		
Puzzle Type*	0.001	0.035	0.005			0.245	562	0.897
	(0.051)	(0.049)	(0.052)			(0.032)		
Normalized Test Score on All Puzzles*	0.066	-0.254	-0.060			0.063	562	0.050
	(0.115)	(0.112)	(0.117)			(0.072)		
Puzzle Test Scores of Non-Attriting Ops*	0.069	-0.078	-0.066			0.058	407	0.726
	(0.132)	(0.130)	(0.138)			(0.084)		
Normalized Score for Peanuts				-0.033		0.024	257	0.792
				(0.125)		(0.088)		

Notes

1 Each row is the regression results of the characteristics in the title column on the treatments. The regressions include the full sample and the omitted group is the Cognitive Task Low Fixed Payment group (which pools together Treatment Groups 2 and 3) in all rows except: Puzzle Type; Normalized Test Score on All Puzzles; Puzzle Test Scores of Non-Attriting OPs; and Normalized Score for Peanuts. For the three rows with the * subscript, only cognitive treatment OPs are included and the excluded category is treatment group 3. For the row Normalized score for Peanuts, the omitted group is treatment group 7 and the sample includes only effort treatment OPs. Column 9 shows the p value for the joint test of significance of all the treatment dummies.

2 The following describes the treatment groups. The first five, OP treatments 1 - 5, are cognitive ability: OP treatment 1 is high fixed payment; OP treatment 2 low fixed payment; OP treatment 3 very low fixed payment; OP treatment 4 high performance pay; OP treatment 5 low performance pay. The final two, OP treatments 6 and 7, are effort task: OP treatment 6 is low fixed payment and OP treatment 7 is low performance pay.

Table 2: Participation in Second Round of survey

	(1)		(2)		(3)		(4)	
OP Test Score * OP Treat: Cog High Perf			0.121	***				
			(0.046)					
OP Test Score			-0.021					
			(0.024)					
OP Treatment: Cog High Fixed Pay	0.044		0.081					
	(0.051)		(0.053)					
OP Treatment: Cog High Perf Pay	0.000		0.079		0.118	**	-0.116	*
	(0.049)		(0.051)		(0.060)		(0.066)	
OP Treatment: Cog Low Perf Pay	-0.021		0.018					
	(0.052)		(0.054)					
OP Treatment: Effort Low Fixed Pay	0.001							
	(0.053)							
OP Treatment: Effort Low Perf Pay	-0.012							
	(0.052)							
Constant	1.009	***			1.114	***	1.056	***
	(0.123)				(0.194)		(0.260)	
N	823		562		312		250	
Sample	ALL		COG		COG HIGH		COG LOW	

Notes

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

2 The cognitive ability fixed performance low and very payment treatments are pooled since we have a small sample size for the very low fixed payment group, and then used as the excluded treatment category.

3 Columns (2)-(4) restrict the sample to OPs in the cognitive ability treatments. Column (3) restricts the sample to high ability OPs (defined as OPs with a normalized test score above 0) while column (4) includes only OPs with a normalized test score less than 0.

4 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 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: Relationship between OP and Referral

	First Stage		Co-worker		Relative		Friend	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Number of Days with Rainfall during OP's Referral Cycle	-0.649 ***	-0.570 ***						
	(0.204)	(0.161)						
Rainfall on OP Arrival Day	0.441 *	0.435 **						
	(0.241)	(0.205)						
OP Treatment: Cog High Fixed Pay		0.088		-0.014		-0.049		-0.019
		(0.174)		(0.044)		(0.047)		(0.065)
OP Treatment: Cog High Perf Pay	-0.030	-0.014	0.077 **	0.077 *	-0.087 **	-0.083 *	0.025	-0.004
	(0.144)	(0.163)	(0.037)	(0.043)	(0.039)	(0.045)	(0.055)	(0.062)
OP Treatment: Cog Low Perf Pay		-0.083		0.002		0.056		-0.078
		(0.171)		(0.045)		(0.048)		(0.066)
OP Treatment: Effort Low Fixed Pay		-0.051		0.030		-0.029		-0.057
		(0.175)		(0.045)		(0.047)		(0.066)
OP Treatment: Effort Low Perf Pay		-0.104		0.007		0.005		-0.097
		(0.174)		(0.045)		(0.047)		(0.066)
N	562	823	562	823	562	823	562	823
Sample	COG	ALL	COG	ALL	COG	ALL	COG	ALL
Mean			0.145		0.132		0.618	
SD			0.352		0.339		0.487	
Chi^2 statistic: joint test of rainfall variables	12.803	15.777						
Mills: Coefficient			-0.074	-0.174	-0.065	-0.038	0.096	0.033
Mills: SE			0.14348	0.12898	0.150	0.137	0.213	0.191
N Censored Obs			155	229	155	229	155	229

Notes

1 The excluded treatment category is cognitive ability, fixed performance low as described in Table 2.

2 Two samples are presented in this table: columns (1), (3), (5) and (7) include only OPs who were randomly assigned to the cognitive treatment. This is indicated in the tables as the 'COG' sample. The remaining columns include the full sample, including cognitive and effort task OPs. Columns using the full sample are indicated by the 'ALL' sample.

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 the probit coefficients, not marginal effects.

5 Relative, co-worker, and friend are dummy variables indicating the relationship between the Original Participant and the referral. Columnns (3) through (8) 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.

6 All columns include additional covariates as described in Table 2.

Table 4: Task Performance and Treatment Type

	Referral Cognitive Ability Task Performance						Referral Peanut Performance
	OLS			Selection Model			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
OP Cognitive Test Score * Cog High Perf Pay			0.131			0.346 **	0.311 *
			(0.120)			(0.156)	(0.161)
OP Cognitive Test Score		0.093 *	0.058		0.151 **	0.055	-0.021
		(0.053)	(0.061)		(0.070)	(0.067)	(0.069)
OP Treatment: Cog High Fixed Pay	-0.136	-0.138		-0.029	-0.037		
	(0.130)	(0.130)		(0.176)	(0.168)		
OP Treatment: Cog High Perf Pay	-0.144	-0.134	-0.118	-0.146	-0.121	-0.104	-0.194
	(0.129)	(0.129)	(0.113)	(0.166)	(0.160)	(0.126)	(0.129)
OP Treatment: Cog Low Perf Pay	0.091	0.093		0.058	0.064		
	(0.138)	(0.138)		(0.178)	(0.171)		
N	407	407	407	562	562	562	562
Sample	COG	COG	COG	COG	COG	COG	COG
Mean	0.061						-0.199
SD	0.999						1.078
Chi^2 statistic: joint test of rainfall variables				12.341	13.069	13.093	13.093
Mills: Coefficient				1.326	1.271	1.115	0.942
Mills: SE				0.560	0.512	0.432	0.448
N Censored Obs				155	155	155	155

Notes

1 All columns in the table are restricted to the sample of cognitive task treatment pairs.

2 All columns also include the individual characteristics of the Original Participant, as defined in Table 2, plus an indicator for version for the puzzles administered to the OP.

3 The dependent variable in columns (1)-(6) is the referrals' normalized performance on the cognitive task. Column (7) uses the normalized performance measure on the peanut task as the depend variable using the identical sample as in columns (1)-(6).

Table 5: OP Ability to Predict Performance

	High Ability OPs				Low Ability OPs	
	(1)		(2)		(3)	(4)
OP's Anticipated Performance: Puzzle	0.214	**	0.190	**	0.031	0.027
	(0.098)		(0.090)		(0.092)	(0.082)
N	202		280		149	226
Sample	COG		COG		COG	COG
Model	OLS		Selection		OLS	Selection
Chi^2 statistic: joint test of rainfall variables			13.908			4.193
Mills: Coefficient			0.989			0.286
Mills: SE			0.412			0.570
N Censored Obs			78			77

Notes

1 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 performance used in Table 4.

2 Columns (1) and (3) are OLS and columns (2) and (4) are estimates from a heckman two step selection model.

3 Columns (1) and (3) restrict the sample to high ability OPs: those with a normalized test score greater than 0. Columns (2) and (4) are restricted to OPs with a normalized test score less than 0.

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

5 There were 56 Ops in the cognitive treatment who responded with 'I don't know' as the response to the question on anticipated performance are dropped from the sample. There are therefore fewer observations in this Table than in Table 4.

Table 6. Trust Game Play and Treatment Type

	Referral Transfers to OP				Referral to Random Partner
	(1)	(2)	(3)	(4)	(5)
OP Test Score * OP Treat: Cog High Perf			0.254 *	0.225 **	0.276 **
			(0.142)	(0.113)	(0.138)
OP Test Score		-0.015	-0.069	-0.091	-0.020
		(0.057)	(0.060)	(0.059)	(0.059)
OP Treatment: Cog High Fixed Pay	0.040	0.048			
	(0.133)	(0.132)			
OP Treatment: Cog High Perf Pay	0.036	0.035	0.038	0.015	0.010
	(0.127)	(0.127)	(0.112)	(0.104)	(0.109)
OP Treatment: Cog Low Perf Pay	-0.075	-0.079			
	(0.137)	(0.137)			
Referral's Transfer during Dictator Game				0.0102 ***	
				(0.0013)	
OP's Transfer During Dictator Game				0.0008	
				(0.0013)	
N	550	550	550	550	550
Sample	COG	COG	COG	COG	COG
Mean of Pre-Normalization Transfer Amount at Rs. 40	56.338				35.253
SD	31.654				28.670
Chi^2 statistic: joint test of rainfall variables	11.238	12.033	12.100	12.199	12.100
Mills: Coefficient	-0.101	-0.001	0.143	-0.410	0.503
Mills: SE	0.457	0.431	0.400	0.387	0.386
N Censored Obs	155	155	155	155	155

Notes

1 All specifications use the heckman selection model. Also included are individual characteristics of the Original Participant, as defined in Table 2.

2 Dependent variable in columns (1)-(4) is an index constructed from each referral's decisions in the trust game for all possible OP transfers. At each 10 Rs interval, the transfer amount was normalized and the index takes the average of the 8 decisions.

3 The dependent variable in Column (5) is an analagous index as in columns (1)-(4) but uses each referral's decisions in the trust game with his randomly paired partner.

4 There are fewer observations in the table due to 12 OP-referral pairs who chose not to participate in the economic games portion of the experiment.

Table 7: Puzzle Performance with Referral Characteristics

	(1)		(2)		(3)		(4)		(5)	
OP Cognitive Test Score * Cog High Perf Pay	0.349	**	0.405	***	0.348	**	0.353	**	0.357	**
	(0.157)		(0.145)		(0.142)		(0.142)		(0.143)	
OP Cognitive Test Score	0.047		0.002		0.000		-0.001		-0.006	
	(0.068)		(0.063)		(0.061)		(0.061)		(0.062)	
OP Treatment: Cog High Perf Pay	-0.101		-0.144		-0.128		-0.122		-0.129	
	(0.125)		(0.115)		(0.112)		(0.112)		(0.113)	
Referral's Ravens Test Score					0.164	***	0.173	***	0.173	***
					(0.054)		(0.054)		(0.054)	
Referral's Digit Span Score					0.061	***	0.061	***	0.061	***
					(0.014)		(0.013)		(0.013)	
Referral's Transfer during Dictator Game							-0.054		-0.053	
							(0.046)		(0.046)	
Ln Referral Income									-0.039	
									(0.039)	
Sample	COG		COG		COG		COG		COG	
N	543		543		543		543		543	
Referral Controls	NO		YES		YES		YES		YES	
Chi^2 statistic: joint test of rainfall variables	12.010		12.010		12.010		12.010		12.010	
Mills: Coefficient	1.002		0.876		0.896		0.902		0.934	
Mills: SE	0.432		0.399		0.387		0.387		0.390	
N Censored Obs	155		155		155		155		155	

Notes

1 All specifications use the heckman selection model. Also included are individual characteristics of the Original Participant, as defined in Table 2.

2 Resume controls include the following characteristics of the referral: (i) indicators for 5 year age groups; (ii) indicators for each educational level and (iii) occupation code. Ln Referral Income is the ln of (referral income+1).

Appendix Table 1: Peanut Task Performance and Treatment Type

	Selection Model		
	(1)	(2)	(3)
OP Peanut Test Score * Peanut Perf Pay			-0.140
			(0.164)
OP Peanut Test Score		0.064	0.141
		(0.112)	(0.107)
OP Treatment: Peanut Performance Pay	-0.242	-0.224	-0.231
	(0.156)	(0.145)	(0.143)
N	257	257	257
Sample	EFFORT	EFFORT	EFFORT
Mean	-0.254		
SD	(0.962)		
Chi^2 statistic: joint test of rainfall variables	3.213	2.567	3.052
Mills: Coefficient	0.972	0.697	0.623
Mills: SE	0.808	0.790	0.754
N Censored Obs	73	73	73

Notes

1 Also included are individual characteristics of the Original Participant, as defined in Table 2.

2 The sample includes OP-referral pairs randomly assigned to the effort task treatments. The excluded treatment group is effort low fixed payment (80 Rs).

Appendix Table 2: Cognitive Ability Task Performance Robustness

	Selection Model		
	(1)	(2)	(3)
OP Cognitive Test Score * Cog High Perf Pay			0.353 **
			(0.165)
OP Cognitive Test Score		0.158 **	0.058
		(0.077)	(0.071)
OP Treatment: Cog High Fixed Pay	-0.016	-0.023	
	(0.192)	(0.186)	
OP Treatment: Cog High Perf Pay	-0.149	-0.124	-0.106
	(0.181)	(0.176)	(0.133)
OP Treatment: Cog Low Perf Pay	0.054	0.060	
	(0.194)	(0.188)	
Sample	COG	COG	COG
N	562	562	562
Chi2 statistic: joint test of rainfall variables	12.341	13.069	13.093
Mills: Coefficient	1.444	1.401	1.195
Mills: SE	0.633	0.587	0.468
N Censored Obs	155	155	155

Notes

1 Temperature on day the referral performed the cognitive ability task is also included in specifications (1)-(4).

2 Also included are individual characteristics of the Original Participant, as defined in Table 2.

Appendix Table 3: Experiment with Full Info, No Info and Perf Pay

	Selection Model	
	(1)	(2)
OP Cognitive Test Score * No Info		-0.106
		(0.251)
OP Cognitive Test Score * Perf Pay		-0.145
		(0.280)
OP Cognitive Test Score		0.236
		(0.188)
Treatment: No Information to Referral (No Info)	-0.103	-0.102
	(0.370)	(0.296)
Treatment: Performance Pay to OP (Perf Pay)	0.128	0.145
	(0.380)	(0.309)
N	193	193
Chi2 statistic: joint test of rainfall variables	8.549	9.024
N Censored Obs	68	68

Notes

1 All specifications include additional covariates as described in Table 2.

Appendix Table 4: Other Referral Characteristics

	Raven Test		Digit Span Test		Education		Age		Ln Income		Dictator Play
	(1)		(2)		(3)		(4)		(5)		(6)
Panel A											
Referral Puzzle Performance	0.249	***	1.277	***	-1.261	***	0.763	***	-0.372	**	2.388
	(0.048)		(0.185)		(0.396)		(0.164)		(0.150)		(1.982)
N	402		402		402		402		402		390
Panel B											
OP Cognitive Test Score * Cog High Perf Pay	0.240	*	0.240		-0.011		-0.106		-0.266		3.708
	(0.129)		(0.525)		(1.059)		(0.448)		(0.402)		(5.773)
OP Cognitive Test Score	0.014		0.074		0.280		0.156		0.050		0.413
	(0.055)		(0.226)		(0.455)		(0.193)		(0.173)		(2.494)
OP Treatment: Cog High Perf Pay	-0.066		0.134		-1.087		0.117		-0.475		3.913
	(0.103)		(0.418)		(0.841)		(0.357)		(0.320)		(4.597)
Sample	COG		COG		COG		COG		COG		COG
N	557		557		557		557		557		545
Mean	2.071		12.545		8.526		28.0		6.664		63.8
SD	(0.927)		(3.714)		(3.301)		(8.6)		(2.858)		(37.8)
Chi^2 statistic: joint test of rainfall variables	13.491		13.491		13.491		13.491		13.491		12.438
Mills: Coefficient	0.056		0.988		0.146		0.549		-0.836		33.232
Mills: SE	0.361		1.464		2.958		1.252		1.121		15.832
N Censored Obs	155		155		155		155		155		155

Notes

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

2 Panel A show OLS estimates while Panel B show estimates from the Heckman two step.

3 The Raven Test measure is on a scale of 1 to 3, capturing the number of patterns identified correctly. The Digit Span Test measure is the number of series repeated correctly. Each respondent did two trials for the Digits Forward Game and two trials of the Digits Backward Game. The maximum correct score is 32.

Appendix Figure 1: Puzzles