Who gets the job referral?

Evidence from a social networks experiment *

Lori Beaman[†] and Jeremy Magruder[‡]

February 2011

Abstract

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

1 Introduction

Social networks influence labor markets worldwide. By now, an extensive empirical literature has utilized natural experiments and other credible identification techniques to persuade us that networks affect labor market outcomes.[1] We also know that a large fraction of jobs are found through networks in many contexts, including 30-60% of U.S. jobs (Bewley, 1999; Ioannides and Loury, 2004). In our sample in Kolkata, India, 40% of employees have helped a friend

*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. (2008); Beaman (2010); Kramarz and Skans (2007); Granovetter (1973); Laschever (2005); Magruder (2010); Munshi (2003); Munshi and Rosenzweig (2006); Topa (2001).

or relative find a job with their current employer. While these analyses have convinced us of the importance of job networks, the empirical literature has had far less to say about why networks are so commonplace. In contrast, theory has suggested several pathways by which firms and job searchers can find social networks beneficial. For example, job seekers can use social network contacts to minimize search costs (Calvo-Armengol, 2004; Mortensen and Vishwanath, 1994; Galeotti and Merlino, 2009); firms can exploit peer monitoring among socially connected employees to address moral hazard (Kugler, 2003); and firms may use referrals as a screening mechanism in order to reduce asymmetric information inherent in the hiring process (Montgomery, 1991; Munshi, 2003).[2] 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.[3] We create short term jobs in a laboratory in the field in urban India and observe how the actual referral process responds to random variation in the incentives to refer a highly-skilled employee. This allows us to determine whether participants have useful information about fellow network members.

We argue that disseminating job information is often not the primary reason that social relationships are formed and maintained. In a developing country setting like the one in this paper, the majority of the literature on networks emphasizes how individuals use network links to improve risk sharing and insure against idiosyncratic shocks (Udry, 1994; Townsend, 1994; Ligon and Schechter, 2010b). Therefore, any empirical investigation of how social networks

[2] Moral hazard is highlighted as a reason for the use of referrals in Bangladeshi garment factories in Heath (2010).

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

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.[4] These spillovers may cause networks to smooth search frictions using network links which do not represent particularly strong job matches. For example, individuals in networks which formed to share risk may not have the right information to identify good job-specific matches, or they may not be inclined to use that information (if they have it) in a way which benefits employers. There may be contingent contracts (or simple altruistic relationships) that encourage an employee to refer a poorly qualified friend rather than the person they believe to be most qualified for the job.[5] 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, 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. There is related evidence of social connections generating inefficiency in the workplace. Bandiera et al. (2007), in the context of a UK fruit farm, show that social connections on the job affect firm productivity, but not by improving performance through peer monitoring. Without incentives, social connections actually drove down productivity due to favoritism.[6] This highlights that we must consider carefully the decision problem faced by an employee who is embedded in a social network, as the network may create incentives counter to the firm's objectives. It is therefore an empirical question whether network-based referrals are useful for the firm, as suggested by Montgomery (1991).

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

[5]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).

[6]There is other work showing positive peer effects in the workplace (Mas and Moretti, 2009), but there are 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 able, the worker's productivity declines.

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 brief tasks which emphasize either 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. The incentives serve as a tool to reveal information held by participants and provide insights into competing incentives outside of the workplace. In order to isolate the effect of the performance pay contract on the selection of referrals, 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.[7]

The controlled setting we create allows us to examine the complete set of on-the-job incentives faced by each of our employees, which would be difficult in a non-experimental setting. We show that there is a tension between the incentives offered by the employer and the social incentives within a network. When individuals in our study receive performance pay, 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

[7]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).

highly skilled, but that these individuals only select highly skilled network members when given a contract in which their own pay is indexed to the referral's performance. Low ability original participants, however, show little capacity to recruit high performing referrals. This result is consistent with the idea that only individuals who performed well on the test can effectively screen network members, and we provide evidence that low ability participants cannot predict the performance of their referrals.[8] We also document that our study participants are aware of these informational advantages: high ability participants are more likely to make a referral if they receive performance pay than low ability participants are, suggesting that the expected return to performance pay is larger for high ability participants.

Finally, we explore which characteristics are sought by the successful high ability, incentivized OPs. High ability, incentivized OPs systematically bring in young, high cognitive ability recruits who are on average 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 describes the context and experimental design, and section 3 provides a theoretical framework to interpret the impact of the exogenous change in the referral bonus scheme. The data is presented in section 4. Section 5 presents the results: OPs' decision to make a referral; 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 whether observable characteristics can explain performance, are analyzed in section 6 and section 7 concludes.

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

2 Context and Experimental Design

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 (Baran and Zingales, 2009; 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 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 identify individuals for casual labor jobs by monitoring the characteristics of the referrals, the relationships between the original participants and their referrals, and the performance of the referrals at the "job." By varying the types of financial incentives provided to our short-term employees, we observe aspects of the decision-making that occurs within networks, and the tradeoffs network members face when making referrals. Providing cash bonuses to existing employees for referrals is an established practice in many firms, including some firms which index these bonuses to referral performance (Lublin, 2010; Castilla, 2005). We can also 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 are therefore a laboratory counterpart to a mechanism firms can use with long-term employees.[9]

[9] 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 the researchers proposed a financial incentive scheme for managers.

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 discussed above. In the experiment, we directly observe a job network allocating jobs, in this case the position of being a laboratory subject. Participants receive Rs. 135 ($3.00) payment in the first round of the study, which is higher than the median daily income for the population in this study (Rs. 110). 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.

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

2.1 Initial Recruitment

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

Digit Span Test and Raven's Matrices.[10] The initial group (original participants or OPs) faces an experimental treatment randomized along several dimensions. OPs are asked to complete one of two (randomly chosen) tasks: a task emphasizing cognitive ability or a second task emphasizing pure effort. The majority of our sample (which included all high stakes treatment groups) was assigned to the cognitive task, which we focus on in this paper.[11] In that task, participants are asked to design a set of four different "quilts." In each quilt, the participant is asked to arrange a group of colored swatches according to a set of logical rules.[12] More detail on this exercise is given in the Appendix and Appendix Figure 1. We observe 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 payment for making 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

[10] These two measures have been validated by psychologists (Snow et al., 1984) and are increasingly used in household surveys in developing countries.

[11] In the effort task, participants are asked to create small bags of peanuts for 30 minutes, which is similar in spirit to the effort task in Jakiela (2009) and other *real effort* tasks such as the administrative letters used in Konow (2000) and anagrams in Charness and Villeval (2009). Due to limited resources, 1/3 of our sample was assigned to the effort treatment, and they received either the low stakes performance pay or low stakes fixed fee treatments described below. We did not find mean differences in performance for the referrals of OPs who completed the effort task. However, this may be because the sample is much smaller and does not include the high stakes treatments for OPs.

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

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.[13]

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

Contract	Fixed Component	Performance Component	N of OPs
Low Performance Pay	60	0-20	116
High Performance Pay	60	0-50	136
Very Low Fixed Pay	60	0	71
Low Fixed Pay	80	0	117
High Fixed Pay	100	0	122

There are two performance pay levels: the high stakes treatment varies between between Rs. 60 and 110 total pay while the low performance pay is Rs. 60-80. As fixed finder's fees, OPs are randomly offered either Rs. 60, 80 or 100. In all cases, the exact contract, including the requisite number of correct puzzles needed for a given pay grade, is detailed in the offer. All participants are asked to make an appointment to return with a referral in a designated three day window. In what follows, we denote the initial participation (where we recruit OPs into the laboratory) as round one, and the return of the OPs with referrals as round two.

2.2 Return of OPs with referrals

When the original participants return with their referrals, the referrals fill out the survey and perform both the effort and the cognitive ability tasks. In order to minimize the potential for OPs to cheat by telling their referrals the solutions to the puzzles, we developed two sets of puzzles which are very similar, and we randomized which set was used in each laboratory session.[14] As referrals who experienced the same puzzle set as the OP perform no better on

[13]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.

[14]The type of puzzle used is included as a control in all tables.

their tasks than referrals with the opposite set, we are convinced that this type of "cheating" is minimal. 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 ultimately 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.[15] While referrals perform the tasks and complete the survey, OPs fill out a short interim survey about the process they went through in recruiting referrals. 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. Following, these referrals and OPs complete a set of dictator and trust games, which are described at greater length in Beaman and Magruder (2011).

3 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, which is adapted from Bandiera et al. (2009). 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 and what we can identify in the experiment.

[15]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.

Employee i has the opportunity to make a job referral. In making a referral, i would choose from an ambient network of friends, each of whom have an inherent ability at the job, $\theta_j \in \{\theta^H, \theta^L\}$. In return for making a referral, his employer offers him a contract consisting of a fixed fee (F_i) and a performance incentive (P_i), where he will receive P_i if he correctly selects a high ability friend. He observes a signal of each friend's ability, $\hat{\theta}_j \in \{\theta^H, \theta^L\}$. For simplicity, that signal is accurate with probability β_i, that is, $\mathbb{P}\left(\theta = \theta^H | \hat{\theta} = \theta^H, i\right) = \mathbb{P}\left(\theta = \theta^L | \hat{\theta} = \theta^L, i\right) = \beta_i$. Naturally, β_i may be heterogeneous among employees, and we assume $\beta_i \in [0.5, 1]$ for all i.[16]

Employee i's expected monetary payoffs from referring a particular friend are a function both of his contract type (F_i, P_i), his signal of the selected friend's ability $\left(\hat{\theta}_j\right)$, and the accuracy of that signal. Following Bandiera et al. (2009) and Prendergast and Topel (1996), i also receives a payment σ_{ij} from referring friend j. This payment can be interpreted as an actual cash transfer or as a weighted inclusion of $j's$ income in $i's$ utility.[17] Since there are two ability "types" of friends, it is without loss of generality to focus on the decision between friend 1, for whom $\sigma_{i1} \in argmax(\sigma_{ij} | \hat{\theta}_j = \theta^H)$ and friend 2, for whom $\sigma_{i2} \in argmax(\sigma_{ij} | \hat{\theta}_j = \theta^L_j)$. Finally, i also has the option of declining to make a referral. Suppose the effort of making a referral will cost him c_i.[18]

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

$$F_i + \beta_i P_i + \sigma_{i1} - c_i \tag{1}$$

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

$$F_i + (1 - \beta_i) P_i + \sigma_{i2} - c_i \tag{2}$$

[16] 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.

[17] 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.

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

Comparing these two expressions, i will select friend 1 if

$$P_i > \frac{\sigma_{i2} - \sigma_{i1}}{2\beta_i - 1} \tag{3}$$

He will further choose not to make a referral if

$$c_i > F_i + max\left\{\beta_i P_i + \sigma_{i1}, \left(1 - \beta_i\right) P_i + \sigma_{i2}\right\} \tag{4}$$

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

As experimenters, we exogenously vary F_i and P_i. Equation 3 makes clear that variation in F_i should not affect the optimal referral choice (as F_i is a common payment to all potential referrals). This is a simple empirical implication of the model we will take to the data. F_i does, however, increase the willingness of agents to participate in the referral process.

A second main empirical implication of the model is that there are four necessary characteristics for performance pay to change the choice of optimal referral: (a) networks must be heterogeneous, so that i observes friends with both types of signals; (b) there must be tradeoffs between network incentives and employer incentives ($\sigma_{i2} - \sigma_{i1} > 0$); (c) the tradeoffs must not be too large relative to P_i; and (d) employee i must have information, so that $\beta_i > 0.5$.

In the experiment, if we observe a change in referral choice in response to performance incentives for some group of respondents, we will be able to conclude that those group members have all four of those characteristics. However, if a group does not change their referral choice in response to performance pay, we will not know which characteristics are missing.

There are several characteristics of heterogeneity in this model. In addition to the contract which is randomized, $\beta_i, \sigma_{i1}, \sigma_{i2}$, and c_i may all vary across individuals. We will discuss certain aspects of the role of heterogeneity in the model in detail in section 5. Here we further note that variation in social payments $(\sigma_{i1}, \sigma_{i2})$ and costs of participation (c_i) affect both the participation decision and the referral choice when participants face either a zero or positive performance pay component. In contrast, information (β_i) only affects these decisions when there is a positive performance pay component. This fact will help us disentangle whether heterogeneous treatment effects are most likely to reflect differences in information or differences in social payments or costs of participation.

Finally, participants face a joint participation and referral choice problem. The contract structure influences both decisions, and this generates a challenge to estimate the referral choice problem, our main interest in this paper. This is discussed in section 5.1.

4 Data and Descriptives

Over the period of the study, we successfully enrolled 562 OPs in the cognitive treatment. 72% of our OPs returned with a referral, so that 407 referrals participated 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.

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, whether the participant ultimately got the puzzle correct, and the number of incorrect attempts. We believe each of these to be an important signal of performance. The relevance of the first two aspects are self-evident; incorrect attempts are important as proxies for how much supervisory time an employee requires in order to successfully complete a task. To utilize variation from all three important performance measures, we use the following metric. A perfect score for a given puzzle is assigned for solving the puzzle in under one minute with no incorrect attempts and 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. This is one metric which parsimoniously combines the key factors of interest; however, the main results are robust to sensible alternate measures of performance (for example, the number of puzzles solved correctly).

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. 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]

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 Members Screen?

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

We will observe whether an OP makes a referral and an objective estimate of that

[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 .18 of a standard deviation is unlikely to significantly alter the results.

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

$$y_{ij} = \beta_0 + \phi_i + X_i\gamma + \epsilon_{ij} \tag{5}$$

where y_{ij} could represent participation in the experiment, the relationship between the OP and referral, or the referrals performance, while ϕ_i represents the OP's treatment categories and X_i are OP characteristics.

The model also suggested that different forms of heterogeneity in the underlying parameters of the decision problem may impact participation and referral choice in different ways. Of course, we cannot directly measure the σ_{ij}, c, or β parameters that our OPs respond to in order to test this model directly. Still, one important dimension where others have found heterogeneity in social effects is worker ability (Bandiera et al., 2010; Fafchamps and Moradi, 2009), which accords with the theoretical assumptions in Montgomery (1991) and Munshi (2003).[20] If high ability workers receive a more accurate signal of their network members' ability, i.e. β is larger, then they will recruit higher ability referrals when given a performance pay incentive, and also be more likely to participate when offered performance pay. Therefore, we also investigate whether OP ability is an important dimension of heterogeneity.

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

$$y_{ij} = \delta_0 + \delta_1\theta_i + \phi_i + X_i\gamma \ + \varepsilon_{ij} \tag{6}$$

and

[20] 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.

$$y_{ij} = \delta_0 + \delta_1 \theta_i + \sum_{k \in low,high} \delta_{2k} perf_{ik} * \theta_i + \phi_i + X_i \gamma + \epsilon_{ij} \quad (7)$$

where θ_i is OP i's ability, $perf_{ik} * \theta_i$ is the interaction of an indicator for whether the OP was in a performance pay treatment with stakes k and the OP's ability as measured by the OP's performance of the task in phase 1 of the experiment. ϕ and X are defined as before. Since ability may be related to any of the underlying parameters, we will need to rely on theoretical restrictions across the referral choice and participation equations to indicate which dimensions of underlying OP heterogeneity create the referral patterns that we observe.

5.1 Returning with a Referral

As was made explicit in the theoretical framework, OPs face extensive and intensive margin choices. On the extensive margin, they may choose whether or not to return 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. None of the treatment indicators are individually significant and, while not shown in the table, they are not jointly significant.

As discussed above, OP ability is a natural and observable dimension to look for heterogeneity highlighted in the model. OP ability may be related to any of the underlying modelling parameters, including differential information, costs of participation or social payments. However, these parameters are associated with different predicted patterns in terms of the decision to make a referral. Most particularly, heterogeneity in c_i (the costs of making a referral) or in σ_{i1} and σ_{i2} (the incentives provided by the network) would be associated with differential participation in response to both the performance payment level and the level of the fixed payments. In contrast, heterogeneity in information levels, β_i, only affects participation through changing the expected return to performance pay. This is straightforward in the model: if there is no performance pay - as in the fixed fee treatments - β is irrelevant in both the participation and referral choice decisions. Thus, if OP ability is a proxy for information, we should see more able OPs participate at different rates in response only to changes in performance incentives and not in changes to fixed fees.

While the decision to make a referral is not on average influenced by treatment categories or with OP ability, column (2) shows that the high stakes performance pay sharply increases the participation rate among high ability OPs. The result in column (2) is consistent with high ability OPs differing from low abilities in their level of information.[21] However, in a more general model with multiple ability types, OP ability may also be correlated with network quality: that is, the probability of having a high ability individual in his network. This would also generate a higher expected return to performance pay and be consistent with the result in column (2). We will provide more direct evidence on the role of information in section 5.5.

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

Therefore, we use two main strategies to estimate the impact of contract change on referral choice. In our preferred specification, we employ a Heckman two step selection model with a first stage probit and second stage estimation including the inverse mills ratio from the

[21] One can imagine alternate explanations related to psychology and subjective expectations; these will be ruled out below when we demonstrate that not only do high ability OPs expect a higher performance from their referral but that the referral also performs better in practice.

first stage (Heckman, 1976). Rainfall makes a natural exclusion restriction, as it is random and it affects the desirability of travelling to our laboratory but it should not be correlated with performance in our (indoor) laboratory.[22] 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 1. 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.[23] While the exact day that an 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 5.3 discusses the strength of the relationship between rainfall and participation.

A second approach is to combine the participation and referral choice decisions into one outcome of interest. For example, the task was to solve puzzles correctly, and OPs who did not return with a referral successfully solved zero puzzles in the second round. This can be viewed as similar in spirit to an Intent to Treat (ITT) effect. We therefore include zeros as their performance (and then normalize accordingly) and analyze performance using OLS on the full sample. For the results on the relationship between the OP and the referral, OPs who did not make a referral are coded as not bringing in a coworker or a relative. The advantage of this strategy is that we can fully utilize the exogenous, random variation. We present both the Heckman and these full sample OLS results in the main tables.

We have also conducted two additional robustness checks. The Heckman model achieves

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

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

a parsimonious specification of the (unobserved) selection bias term through distributional assumptions. As we have no prior as to the validity of these assumptions, 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). 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 non-participation as an outcome of poor performance, it is a weaker assumption than imputing the zeros in OLS. Results presented in this section are robust to these two alternative approaches. Results from the polynomial specification and median regressions are omitted for brevity, but the tables are available from the coauthors upon request. The tables we present are conservative: in general, the specifications used in the robustness analysis reveal the largest and most significant coefficients.

5.2 Responsiveness to Fixed Fees

Turning to estimation of the referral choice problem, the model predicted that variation in the level of fixed fees should not affect the choice of referral, at least once differential participation rates are properly accounted for. We have several characteristics that could be used to estimate the choice of referral, and those can be broadly categorized as characteristics based on relationships (a proxy for σ_{ij}), or characteristics related to productivity (a proxy for θ_j). Table 3 restricts data to the subsample who received one of the fixed fee treatments and asks whether any of these characteristics are related to the level of the fixed fee payment, which would suggest that the referral decision process is somewhat more subtle and may raise concerns about the validity of the experiment. Once again, we look both at whether the level of the fixed fee payment is related to any observable characteristics of the referral, and also whether more able OPs respond differentially to the level of the payment.

Each column in Table 3 represents a different dependent variable. All estimates are consistent with the theoretical prediction, where each column represents a different dependent variable. First, columns (1) and (2) show that rainfall during the OP's window for recruitment significantly lowers the probability that the OP completes the study, and the joint test of both rainfall variables is above 8. The main results are in columns (3) through (10). Across these eight specifications, only one interaction term finds marginal significance, and joint p-values of the overall effects of fixed fee and interaction terms are never close to significant. For brevity, we omit the specifications which only include the levels of fixed fee treatments on the productive characteristics and only present the results from the Heckman specification.[24] Since the data are consistent with theoretical prediction that fixed fees variation does not alter the referral choice problem, we combine all fixed fee treatments into a single control group in subsequent specifications and test the performance pay treatments against the fixed fee treatments jointly.

5.3 Relationship between Referrals and OPs

The referral choice equation suggested that one important dimension that should change with performance pay is the selection of referrals in terms of the network payoff σ_{ij}. In particular, if OPs respond to performance pay by changing their choice of referral, they should be shifting away from referrals who grant them larger social transfers in favor of those who generate a smaller transfer. Of course, we cannot directly estimate σ_{ij}; here, we focus on two salient relationships: co-workers and relatives.

Table 4 shows the relationship between OPs and their referrals as a function of treatment

[24] Analyses using the alternative specification, in which we interpret OPs who choose not to participate as OPs whose referrals solved zero puzzles and did not bring in a co-worker or relative, never reveal a significant effect and are available from the authors.

type using alternatively the Heckman specification and OLS with the full sample. Columns (1) and (2) demonstrates that rainfall during the OP's window for recruitment significantly affects the participation rate within the full cognitive sample. While not shown in the table, the marginal effect of the number of days of rain during the OP's referral cycle implies that one extra day of rainfall makes it 21% less likely an OP will return with his referral to the laboratory (measured at the mean of the covariates in specification (1)). 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 (10) examine two easy to interpret, relationships identified in the survey: coworkers and relatives. The relationship measure is based on self-reports from the referral, and we anticipate that for both insurance and altruistic reasons, relatives are likely to donate larger social transfers than coworkers. Columns (3) and (4) show the results of the Heckman specification. 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. There is, however, little evidence of heterogeneity in the response to performance pay as shown in column (4). Columns (5) and (6) show the results of OLS on the full sample and show similar results. The estimate of the likelihood of referring a coworker among performance pay OPs declines to around 6% but remains precisely estimated. There is limited evidence again of heterogeneity. The coefficient on the interaction between OP ability and the high stakes performance pay in column (6) is close to significance at the 10% level but this is not a robust result.

Columns (7) and (8) 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 anticipated skills, a coworker. There is again no evidence of a heterogeneous response by OP ability. The coefficients on high performance pay using OLS with the full sample is less precisely estimated but similar in magnitude as the coefficient in column (5). Therefore, it appears that performance pay induced OPs to refer relatives less often and instead refer coworkers. Finally, in results not reported for brevity, there is no change in the probability of referring a friend.[25] Whether the performance pay resulted in higher performing referrals is investigated in the next section.

5.4 Referral Performance and Response to Incentives

Table 5 shows how OPs responded to the incentives on the cognitive ability task. The first estimates are from the OLS on the sample of participants in columns (1) through (3), then the Heckman selection model in columns (4) through (6) and finally OLS on the full sample in columns (7) through (9). The estimates from OLS with only the sample of those who participate differ significantly from estimates which address differential participation. In columns (1) through (3), only 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 in the Heckman specification 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. This would be consistent with a positive correlation between an OP's ability and the overall ability of the OP's network, or it may represent differential ability to

[25]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.

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 driven by OPs who face performance pay incentives. Therefore, high ability individuals refer high ability people only when properly incentivized, suggesting that the networks of high ability OPs are heterogeneous and that high ability OPs do have the capacity to screen. Columns (7) through (9) show that these results are similar when using OLS on the full sample: performance pay offers results in high ability OPs generating more round two puzzles solved.

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 2 shows the 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 of the level effect of no information and the interactions of OP ability with no information and with actual performance pay are negative, the opposite of what would occur if side performance pay contracts were responsible for our main results. This is consistent with reports in the data on anticipated transfers between OPs and referrals, where zero OPs report an out-of-laboratory contract where the OP pays the referral.

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

We observed in Table 4 that all OPs in the high stakes performance pay treatments respond to incentives by recruiting coworkers more often and recruiting relatives less often. Only high ability OPs, however, recruited referrals who actually performed better on the cognitive task. Thus, while all OPs 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 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 ($\beta_i > .5$), and he had to face a tradeoff between network incentives and the performance incentives generated by the experiment ($\sigma_{i2} - \sigma_{i1} > 0$) and $P_i > 0$ where P_i is sufficiently large. If low ability OPs lack any of these characteristics, then they would not have reacted to performance pay by recruiting higher quality referrals.

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

Table 6 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).[26] 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, consistent with our theoretical restrictions. Moreover, it is consistent with the fact that all OPs adjust their behavior on the margin of relationships between the OP and the referral: low ability OPs are trying to bring in higher ability referrals, but simply do not have a good understanding of

[26]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 network members will perform better.

6 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. We have provided evidence that high ability OPs identify high ability referrals when properly incentivized. However, we have said little about whether the information they use would be easy or difficult for an employer to observe. In this section, we check which characteristics are correlated with strong performance and test whether these observable characteristics can explain the referral performance premium.

Table 7 looks at how performance on the cognitive task covaries with age, education, other cognitive tests and income. The table reveals that referrals who perform well on the cognitive task tend to be young, well educated, and low in income. 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. High scoring referrals also tend to perform well on the cognitive ability tests we included in the background survey, the Raven Test[27] and the Digit Span Test. Given that the Raven and Digit Span tests have been used extensively in the psychology literature on measuring cognitive ability (Snow et al., 1984), this correlation provides reassuring evidence on the validity of our cognitive task.

Given that some easy to observe characteristics like age and education are strong predictors of performance, a natural question is whether an employer could easily acquire this screening without the use of referral networks. For example, age and education could be easily observed on a resume, while the Raven and Digit span test results could be easily determined

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

through a quick prescreening test. 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

$$y_{ij} = \beta_0 + \beta_1\theta_i + \sum_{k\in\{low,high\}} \beta_{2k} perf_{ik} * \theta_i + \phi_i + X_i\gamma + W_j\delta + \epsilon_{ik}$$

where θ_i, ϕ_i, and X_i are respectively OP $i's$ ability, treatment group, and OP characteristics as before, but now we control also for a vector of the referral's characteristics, W_j, 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 8 presents the results of this estimation. In column (1), we reproduce the analysis from Table 5 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 allows for a flexible relationship between these characteristics and productivity: specifically, we add in indicators for the referral's 5-year age group, each education level, and occupational category. While young and well-educated individuals do perform well on the cognitive task, this effect is not related to the performance premium that referrals of high ability, high stakes performance pay OPs enjoy. The remaining columns in Table 8 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. Column (4) includes the referral's income as well. 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.[28] These results suggest that networks are identifying skillsets which may 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. This paper begins to look inside the black box of social networks by directly observing job networks spread temporary jobs in a hybrid laboratory-field experiment environment under a variety of incentive schemes. We find that 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 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

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

of monetary or social payments) and the friend who they think will perform best 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 suggests that some workers could not screen effectively. Moreover, the workers who could screen were only willing to do so when they were directly incentivized, as they faced competing incentives generated by the network itself. 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 study 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 (2010). Social incentives in the workplace. *Review of Economics Studies 77*, 417–59.

Baran, N., P. S. and L. Zingales (2009). Can we infer social preferences from the lab? Evidence from the trust game. *Mimeo, Northwestern University*.

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

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

Beaman, L. and J. Magruder (2011). Reliability in referral networks: Experimental evidence from India. *Mimeo, University of California, Berkeley*.

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.

Charness, G. and M. Villeval (2009). Cooperation and competition in intergenerational experiments in the field and laboratory. *American Economic Review 99*, 956–978.

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*.

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

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

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.

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

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

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

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*.

Konow, J. (2000). Fair shares: Accountability and cognitive dissonance in allocation decisions. *American Economic Review 90*, 1072–1092.

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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

8 Appendix

The appendix contains two sets of robustness analysis, both of which are discussed at greater length in the main text. Appendix Table 1 reproduces the main results of Table 5 from the main paper, with daily temperature controlled for (as a possible confounder to our exclusion restriction). Appendix Table 2 shows the results of a robustness experiment where participants were either actually paid performance pay or where the OP was not informed that he would receive the maximum of his performance pay range until after the referral had already entered the lab.

Finally, Appendix Figure 1 gives some additional detail on the cognitive task which is the central measurement in this paper. Each subject in our treatment completed one of two similar sets of four puzzles; one of those sets is presented in Appendix Figure 1, with initial conditions in the first column and suggested solutions in the second. Puzzle A gave the subject four swatches of each of four colors and asked the respondent to arrange them so that exactly one of each color was in each row and each column. Puzzle B repeated this request, but began with the diagonal set as presented in the figure. That is, the subject had to again make sure that one of each color was in each row and each column, but had to do so in a way which did not disturb the diagonal. Puzzle C changed the rules slightly: respondents were again given four swatches each of four different colors, but this time were asked to make sure that each row and column contained either two or zero of each color in each row and each column. Moreover, the puzzle began with four corners set as the same color, and the respondents were told they must keep those four corners untouched in their solution. Finally, Puzzle D had the most complex rules. Subjects were given 9 swatches: four each of two colors and one of the third. The rules in this puzzle were different for each color: For the first color (with four swatches), the rule was that swatches of that color could not border any other swatches of the same color, and it was explained that bordering could mean touching on a horizontal edge, a vertical edge, or across a corner. For the second color (also with four swatches), the rule was that the swatches of that color must border exactly two swatches of the same color. For the final color, with only a single swatch, they were told it was free and could be placed anywhere. The solution to this puzzle is unique, and presented here.

Table 1: Randomization Check

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

Notes

1 Each row is the regression results of the characteristics in the title column on the treatments. The regressions include the cognitive treatment sample and the omitted group is the very low fixed treatment in all rows. Column 9 shows the p value for the joint test of significance of all the treatment dummies.

Table 2: Participation in Second Round of survey

	(1)	(2)	
OP Cog Test Score * High Fixed Pay		0.040	
		(0.073)	
OP Cog Test Score * Low Fixed Pay		0.064	
		(0.071)	
OP Cognitive Test Score * High Perf Pay		0.191	***
		(0.071)	
OP Cognitive Test Score * Low Perf Pay		0.026	
		(0.071)	
OP Cognitive Test Score	0.032	-0.040	
	(0.022)	(0.057)	
OP Treatment: High Fixed Pay	0.018	0.024	
	(0.072)	(0.071)	
OP Treatment: Low Fixed Pay	-0.041	-0.040	
	(0.075)	(0.075)	
OP Treatment: High Perf Pay	-0.024	0.016	
	(0.072)	(0.071)	
OP Treatment: Low Perf Pay	-0.056	-0.054	
	(0.077)	(0.076)	
N	545	545	

Notes

1 The dependent variable in all columns is 1 if the respondant returned to the laboratory with a referral. The coefficients are marginal effect estimates from a probit.

2 All columns restrict the sample to OPs in the cognitive ability treatments.

3 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: Fixed Fee Treatments - Referral Choice

			Relationship to OP				Referral Characteristics			
	First Stage		Co-worker		Relative		Test Score	Age	Educ	Raven's Test
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Number of Days with Rainfall during OP's Referral Cycle	-0.522 **	-0.536 **								
	(0.267)	(0.273)								
Rainfall on OP Arrival Day	0.805 **	0.820 **								
	(0.353)	(0.357)								
OP Test Score * High Fixed Pay		0.214		-0.049		-0.021	-0.304	2.181 *	0.704	-0.026
		(0.254)		(0.064)		(0.064)	(0.200)	(1.290)	(0.596)	(0.180)
OP Test Score * Low Fixed Pay		0.220		-0.079		-0.085	-0.139	2.011	0.541	-0.137
		(0.239)		(0.066)		(0.065)	(0.202)	(1.336)	(0.601)	(0.184)
OP Cognitive Test Score		-0.150		0.022		0.039	0.196	-0.554	-0.254	0.084
		(0.203)		(0.055)		(0.054)	(0.168)	(1.106)	(0.500)	(0.154)
OP Treatment: High Fixed Pay	-0.009	-0.027	0.010	0.013	-0.024	-0.031	0.072	1.215	-0.397	0.027
	(0.250)	(0.255)	(0.057)	(0.057)	(0.056)	(0.056)	(0.179)	(1.139)	(0.509)	(0.158)
OP Treatment: Low Fixed Pay	-0.143	-0.153	0.055	0.061	0.009	0.013	0.192	-0.357	0.222	0.214
	(0.242)	(0.244)	(0.059)	(0.059)	(0.058)	(0.057)	(0.183)	(1.186)	(0.528)	(0.163)
N	310	310	310	310	310	310	310	310	310	310
p value from joint test of treatment and treatment interactions			0.801	0.880	0.912	0.932	0.865	0.505	0.686	0.588
Chi^2 statistic: joint test of rainfall variables	8.118	8.289	8.118	8.289	8.118	8.289	8.289	9.596	6.503	8.289
Mills: Coefficient			-0.199	-0.189	0.115	0.098	0.864	1.425	0.072	0.173
Mills: SE			0.166	0.165	0.164	0.164	0.507	3.118	1.665	0.466
N Censored Obs			81	81	81	81	81	83	82	81

Notes

1 The excluded treatment category is the very low fixed treatment. All columns include additional covariates as described in Table 2.

2 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.

3 Columns (1) and (2) show probit coefficients, not marginal effects.

4 Relative and co-worker are dummy variables indicating the relationship between the Original Participant and the referral. Columnns (3)-(10) are heckman two step estimates with the rainfall variables from columns (1) and (2) used as exclusion restrictions. The first stage is shown in columns (1) and (2) with the F test of joint significance of the two rainfall variables.

Table 4: Relationship between OP and Referral

	First Stage		Co-worker				Relative			
			Selection Model		OLS Full Sample		Selection Model		OLS: Full Sample	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Number of Days with Rainfall during OP's Referral Cycle	-0.644 ***	-0.651 ***								
	(0.204)	(0.207)								
Rainfall on OP Arrival Day	0.455 *	0.490 **								
	(0.242)	(0.247)								
OP Cog Test Score * High Perf Pay		0.460 ***		0.008		0.049		0.023		0.019
		(0.166)		(0.048)		(0.030)		(0.049)		(0.031)
OP Cog Test Score * Low Perf Pay		-0.057		0.059		0.042		-0.001		-0.009
		(0.162)		(0.042)		(0.031)		(0.042)		(0.032)
OP Cognitive Test Score		0.027		-0.021		-0.012		-0.003		0.000
		(0.092)		(0.024)		(0.018)		(0.024)		(0.018)
OP Treatment: High Perf Pay	-0.069	0.087	0.079 **	0.076 *	0.056 *	0.062 **	-0.070 *	-0.072 *	-0.048	-0.044
	(0.151)	(0.165)	(0.039)	(0.039)	(0.029)	(0.030)	(0.040)	(0.040)	(0.030)	(0.031)
OP Treatment: Low Perf Pay	-0.138	-0.142	0.007	0.010	0.004	0.003	0.068	0.065	0.037	0.037
	(0.161)	(0.162)	(0.043)	(0.043)	(0.032)	(0.032)	(0.044)	(0.044)	(0.033)	(0.033)
N	562	562	562	562	562	562	562	562	562	562
Mean			0.145				0.132			
SD			0.352				0.339			
Chi^2 statistic: joint test of rainfall variables	12.743	12.743	12.743	13.056			12.743	13.056		
Mills: Coefficient			-0.082	-0.155			-0.071	-0.008		
Mills: SE			0.144	0.134			0.150	0.137		
N Censored Obs			155	155			155	155		

Notes

1 The excluded category is the fixed fee performance treatments.

2 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.

3 Columns (1) and (2) show probit coefficients, not marginal effects.

4 Relative and co-worker are dummy variables indicating the relationship between the Original Participant and the referral. Columnns (3), (4), (7) and (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. Columns (5), (6), (7) and (8) use the the full cognitive treatment sample and OPs who did not bring in a referral are recorded as not having brought in a Co-worker or a Relative.

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

Table 5: Task Performance and Treatment Type

	Referral Cognitive Ability Task Performance								
	OLS			Selection Model			OLS: Full Sample		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
OP Cognitive Test Score * High Perf Pay			0.163			0.370 **			0.346 ***
			(0.125)			(0.159)			(0.128)
OP Cognitive Test Score * Low Perf Pay			0.102			0.065			0.037
			(0.124)			(0.138)			(0.133)
OP Cognitive Test Score		0.092 *	0.028		0.152 **	0.036		0.123 **	0.027
		(0.053)	(0.071)		(0.071)	(0.079)		(0.057)	(0.075)
OP Treatment: High Perf Pay	-0.090	-0.079	-0.080	-0.135	-0.107	-0.084	-0.072	-0.045	-0.004
	(0.118)	(0.118)	(0.118)	(0.157)	(0.151)	(0.131)	(0.126)	(0.126)	(0.127)
OP Treatment: Low Perf Pay	0.144	0.147	0.147	0.068	0.077	0.078	0.014	0.019	0.013
	(0.129)	(0.128)	(0.128)	(0.172)	(0.164)	(0.144)	(0.136)	(0.136)	(0.135)
N	407	407	407	562	562	562	562	562	562
Mean	0.061								
SD	0.999								
Chi^2 statistic: joint test of rainfall variables				12.743	13.449	13.056			
Mills: Coefficient				1.356	1.301	1.123			
Mills: SE				0.561	0.514	0.432			
N Censored Obs				155	155	155			

Notes

1 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.

2 The dependent variable in columns (1)-(9) is the referrals' normalized performance on the cognitive task.

Table 6: 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
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 using the sample of OPs who returned with a referral, and columns (2) and (4) are estimates from a heckman two step selection model.

3 Columns (1) and (2) restrict the sample to high ability OPs: those with a normalized test score greater than 0. Columns (3) 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 The 56 OPs 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 5.

Table 7: Other Referral Characteristics

	Raven Test		Digit Span Test		Age		Education		Ln Income	
	(1)		(2)		(3)		(4)		(5)	
Referral Puzzle Performance	0.249	***	1.277	***	-1.261	***	0.763	***	-0.372	**
	(0.048)		(0.185)		(0.396)		(0.164)		(0.150)	
N	402		402		402		402		402	
Mean	2.07		12.56		27.92		8.56		6.65	
SD	(0.92)		(3.70)		(8.53)		(3.28)		(2.87)	

Notes

1 The dependent variable is the variable described in the column heading, and the independent variable is referral puzzle performance as previously described. Coefficients and standard errors are from OLS.

2 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. The mean and standard deviation of each dependent variable among the referral sample is also included.

Table 8: Puzzle Performance with Referral Characteristics

	(1)		(2)		(3)		(4)	
OP Cognitive Test Score * High Perf Pay	0.375	**	0.435	***	0.383	***	0.383	***
	(0.159)		(0.145)		(0.141)		(0.142)	
OP Cognitive Test Score * Low Perf Pay	0.067		0.161		0.212	*	0.208	*
	(0.139)		(0.128)		(0.124)		(0.125)	
OP Cognitive Test Score	0.030		-0.036		-0.053		-0.057	
	(0.079)		(0.073)		(0.071)		(0.072)	
OP Treatment: High Perf Pay	-0.094		-0.135		-0.123		-0.125	
	(0.131)		(0.118)		(0.115)		(0.115)	
OP Treatment: Low Perf Pay	0.078		0.079		0.059		0.069	
	(0.146)		(0.135)		(0.131)		(0.132)	
Referral's Ravens Test Score					0.146	***	0.146	***
					(0.052)		(0.052)	
Referral's Digit Span Score					0.061	***	0.061	***
					(0.013)		(0.013)	
Ln Referral Income							-0.037	
							(0.038)	
N	555		555		555		555	
Referral Controls	NO		YES		YES		YES	
Chi^2 statistic: joint test of rainfall variables	13.021		13.021		13.021		13.021	
Mills: Coefficient	1.110		0.894		0.893		0.920	
Mills: SE	0.429		0.390		0.377		0.379	
N Censored Obs	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: Cognitive Ability Task Performance Robustness

	Selection Model		
	(1)	(2)	(3)
OP Cognitive Test Score * Cog High Perf Pay			0.380 **
			(0.171)
OP Cognitive Test Score * Cog Low Perf Pay			0.073
			(0.148)
OP Cognitive Test Score		0.160 **	0.037
		(0.079)	(0.084)
OP Treatment: Cog High Perf Pay	-0.144	-0.115	-0.089
	(0.171)	(0.166)	(0.141)
OP Treatment: Cog Low Perf Pay	0.059	0.067	0.070
	(0.187)	(0.181)	(0.154)
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)-(3), in addition to OP characteristics as defined in Table 2.

Appendix Table 2: 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 Figure 1: Puzzles

Initial Setup	Proposed Solution

Puzzle A

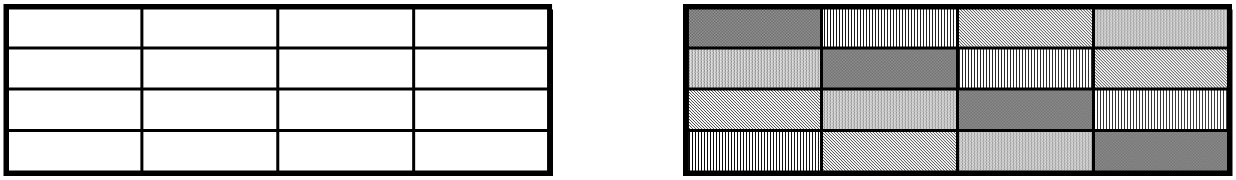

Puzzle B

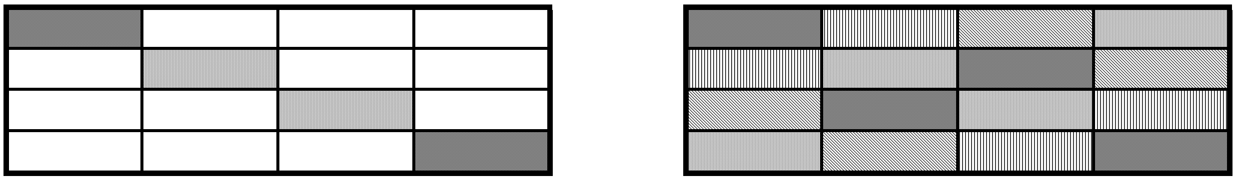

Puzzle C

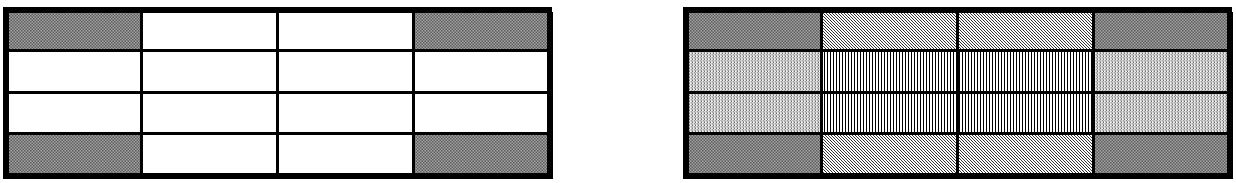

Puzzle D