

# How Do Workers Choose Referrals?

Lori Beaman and Jeremy Magruder

---

Job referrals are an important way that people find work across the globe. We use an experiment in Kolkata, India to ask whether people can identify highly skilled referrals, and when they are willing to do so (instead of referring a closer, but less skilled friend).

that job referrals are important in labor markets, we have much less evidence on even very basic questions about how they work, such as whether the use of referral systems results in good hires.

Theory has suggested two reasons to suspect that networked hires may be high ability: first, high-ability workers may associate primarily with other high-ability people, so that their referrals are mechanically high quality. Second, if networks are heterogeneous, then employees may be willing and able to screen on their employer's behalf. That is, if they know which of their friends represent a particularly good fit for a job opening, they may be willing to use that information and make a referral.

In a recent paper forthcoming in the *American Economic Review*, we test whether employees have useful information about their network members, and under what conditions they are willing to use that information. In our view, both the capacity and willingness of employees to screen on their employer's behalf can be questioned.

Social networks are formed for many reasons, at least some of which are unrelated to a person's potential as an employee for a specific job. Moreover, in order for employees to be useful in choosing which of their friends would make the best employee, it must be the case that whatever incentive they have to screen on the part of the employer can overcome whatever incentives they have to use that job opportunity in a different way.

For example, in much of the literature on networks in developing countries, the role of networks to insure against risky events would suggest that people may make claims on

the next good opportunity—like a job opening—that their friends learn of, and offer in exchange for a promise of sharing the next good opportunity that they learn of. This possibility would mean that new job referrals go to people who have the strongest claim to the next good opportunity rather than to the person who would be best at the job. Moreover, it is surely the case that altruism plays an important role in social networks, giving incentives to refer the network member one cares about the most rather than the friend who you believe to be the best worker or best match for that job.

*Participants in the study, regardless of ability level, were about 50% more likely to bring in a coworker, and 50% less likely to bring in a relative, when they were given performance pay*


We examine this tradeoff using a laboratory experiment in the field in peri-urban Kolkata, India. The idea behind our experiment is that once a person enters our laboratory, completes a task, and is paid for it, we have essentially given him a casual day-job. We can therefore allow our employee to make referrals into our laboratory, randomize his incentive to find a high quality worker, and observe actual referral choices.

This controlled framework creates several advantages: first, while the measurement of worker ability takes place within a laboratory environment, the actual behavior which we are measuring takes place when the worker goes

Globally, a large fraction of workers are hired through job references. An extensive literature, both in sociology and in economics, has documented the importance of social connections as a means of job search. For example, 30-60% of jobs in the United States are found through referrals. In Kolkata, India we find that 45% of our sample has found a job for a friend. A large and growing empirical literature in economics has more recently utilized a variety of quasi-experimental variations in the size or quality of job networks to confirm what survey data has shown: Job networks are empirically important.

While it is clear that job networks matter, we know much less about how they work and what the consequences of the use of informal social connections in the hiring process are to either firms or unemployed job seekers. In fact, while we have substantial evidence

Figure 1. Sample Cognitive Task



home and determines which of his friends to refer. Thus, we are using the lab to measure an unmonitored behavior. Second, it allows us full control of the incentive contract. We know precisely what incentives our employees face to make a strong referral.

The set up of our experiment is as follows: Original Participants (OPs) were invited to perform two hours of work in our lab, for which they were guaranteed 135 rupees in payment. To ensure a sample of OPs who were representative of labor market participants, we recruited OPs at every third household in neighborhoods surrounding our lab; this resulted in a sample of OPs who were young, active labor-market participants, with an average age of 31 and a 96% employment rate.

Once these OPs arrived in the lab, they completed a survey and then were asked to complete a task which emphasized cognitive skill. A minority of OPs were also assigned to a task which emphasized pure effort rather than cognitive ability, specifically the task of filling plastic baggies with precisely 20 peanuts each. As only a minority of OPs were assigned to this treatment, we restricted our analysis to the cognitive treatment.

More specifically, they were asked to help arrange a set of colored swatches in a square according to several sets of logical rules in order to help design quilts. For example, one

set of rules gave OPs 16 swatches, four each of four different colors, and asked them to arrange the swatches so that no more than one of each color was in each row and each column.

A slightly more difficult puzzle is depicted in figure 1. Once again, OPs were asked to place one of each color in each row and each column; however, now they were also required to keep the diagonal cells fixed as in the top panel of figure 1. A sample solution is in the bottom panel. For each of these puzzles, our supervisor noted whether the OP reached the correct solution, how long it took the OP to solve the puzzle, and how many times the OP thought he had reached a correct solution when he was in fact incorrect.

Upon completion of the task, OPs were asked to identify a friend who “would be good at the task they just completed,” and were offered a finder’s fee for bringing in such a friend. The exact terms of the finder’s fee contract were randomly assigned to OPs. A random selection of OPs were offered fixed finder’s fees, which will pay the same amount regardless of their referral’s performance on the task, while others were offered performance incentives, where the OP’s pay may be indexed to the referral’s performance. In both cases, referrals themselves were paid a fixed 135 rupees, regardless of their performance.

The experiment was completed when the OPs returned with their referrals. At that time, OPs and referrals were informed that they would in fact be paid the maximum payment possible under their contract, so that when referrals completed the task no one (neither the OP nor the referral) had a direct incentive to exert differential effort depending on their treatment status.

The motivation for this change in contracting terms was that we were particularly concerned that if OPs shared their finder’s fees with referrals, then referrals in performance treatments may have additional incentives to exert effort. This concern, which seemed of first order importance in the experimental design, seemed to have little practical effect as very few of our participants—OPs or referrals—reported any intention of sharing their payments with each other.

There were four necessary conditions for performance incentives to induce OPs to change the person they chose as a referral: first, they had to face tradeoffs between the person who they would most like to refer for social reasons (e.g., altruism or because of claims on the next good opportunity) and the one who they believed to be the highest ability. Second, OPs must have had some information as to their friends’ abilities, so that the performance incentives were more valuable if they referred a friend who they believed to be higher ability. Third, social incentives could not be too large, so that they overwhelmed the expected return (from us) to referring a socially suboptimal but high-ability friend. Finally, networks needed to be heterogeneous: if all of an OP’s friends were the same ability, then the OP would not change his referral choice regardless of incentives.

OPs with different levels of information may have also responded differentially to treatment in their decisions to return with a referral, and we present evidence that high-ability OPs were

more likely to make a referral when faced with performance incentives. This was simultaneously a potentially useful result for employers, behavior consistent with high-ability OPs having better information, and a potential attrition problem for the analysis of the model's other predictions.

We solved this problem by using a Heckman selection correction for the main analysis, which can control for systematic attrition of this type, provided one can identify a random source of variation which affects attrition but not performance on the puzzles. In our case, we used the fact that rain in Kolkata can be quite severe, which occurs at random on a particular day of the year. A day of rainfall made travelling to our laboratory less appealing, yet should not have the ability to impact a referral's ability to solve logic puzzles.

We focused on two types of analyses. First, did OPs change the types of relationships that they brought into the lab? If OPs responded to performance pay by choosing more socially distant referrals, we interpreted that as evidence that they were trading off social incentives for our cash incentives. Second, did OPs succeed in bringing in higher ability referrals when they received performance pay? For all analysis, we looked for differential effects for higher ability OPs: ability has been identified in many peer effects and networks studies as an important dimension of heterogeneity.

Our results indicate that all OPs respond to network incentives by changing the types of relationships that they bring in. OPs, regardless of ability level, were about 50% more likely to bring in a coworker, and 50% less likely to bring in a relative, when they were given performance pay (both results are statistically significant).

We interpret this result as evidence that all OPs respond to performance incentives by bringing in more socially distant network

members. This interpretation is supported by data we collected on gifts and loans: relatives play a dominant role in gift and loan networks, while coworkers are only marginally involved in either gifts or loans.

Interestingly, despite the fact that our OPs were all, on average, responding to performance incentives by bringing in more socially distant referrals, they did not, on average, bring in higher ability referrals when they were given performance incentives. This result, however, hides some important heterogeneity: high-ability OPs respond to performance incentives by bringing in significantly higher ability referrals.

This result tells us that high ability OPs, at least, meet all four of our necessary conditions: they had information about the skills of their network members, they faced social tradeoffs so that they would rather not use that information, those social tradeoffs were not overwhelming in magnitude, and their networks were heterogeneous in ability.

We also document that at least one reason that we did not see these effects for low-ability OPs is that low-ability OPs did not appear to have very much information about their referrals' ability levels. We asked all of our OPs how well they expected their referrals to perform, and while high-ability OPs' predictions were strongly and significantly correlated with their referrals' performance, there was no relationship between the predictions made by low-ability OPs and how their referrals actually performed.

This study offers a first look inside how individuals make the choice of who to refer for new job opportunities, and test of screening models as a motivation for referrals. We were able to both confirm the capacity to screen (at least among high-ability people) and the presence of social tradeoffs which could limit successful screening. In future work, we hope to explore more deeply exactly how a network

“works” in terms of determining how it allocates new opportunities and the consequences of that allocation.

Suggested Citation:

Beaman, L. and J. Magruder. 2012. "How Do Workers Choose Referrals?" *ARE Update* 15(4):5-7. University of California Giannini Foundation of Agricultural Economics.

Lori Beaman is an assistant professor in the Department of Economics at Northwestern University. She can be reached by e-mail at [l-beaman@northwestern.edu](mailto:l-beaman@northwestern.edu). Jeremy Magruder is an assistant professor in the Department of Agricultural and Resource Economics, University of California, Berkeley. He can be contacted by e-mail at [jmagruder@berkeley.edu](mailto:jmagruder@berkeley.edu).

For additional information,  
the authors recommend:

Beaman, L. and J. Magruder “Who Gets the Job Referral? Evidence from a Social Network Experiment.” *The American Economic Review*, forthcoming.

Ioannides, Y.M. and L. Datcher Loury. “Job Information Networks, Neighborhood Effects, and Inequality.” *Journal of Economic Literature* 42 (2004):1056-93.

[www.aeaweb.org/articles.php?doi=10.1257/0022051043004595](http://www.aeaweb.org/articles.php?doi=10.1257/0022051043004595).