Job Referrals and Strategic Network Formation

Experimental Evidence from Urban Neighbourhoods in Ethiopia

Marc Witte*

November 2018

Latest version

Abstract

In this paper, I study the behavioural motivations underlying job referral decisions. In a field experiment in Addis Ababa, Ethiopia, I show that when choosing who to refer for a job, individuals trade off personal benefits and altruistic considerations, with important implications for the efficiency and equity of the referral process. Using complete data on urban social networks and generating a panel of real work and referral opportunities over multiple sessions, I first show that workers rely on reciprocity. This leads to both significant on-the-job productivity losses and persistence in the exclusion of less connected individuals. This dynamic reciprocity is reduced under incentivised referrals, where workers are paid according to the output of their referral, which makes them screen more productive workers. Second, I find that peripheral workers use job referrals strategically to enlarge their network: they are more likely to establish new and reciprocated links, with connections persisting after 18 months. I show that these findings are consistent with a network-based job referral model where individuals trade off social payoffs and altruistic considerations. My findings suggest that conventional job referrals through social networks can reinforce labour market inequalities and prevent less socially connected individuals from getting access to jobs. However, when given referral opportunities, these individuals can manage to escape exclusion even in the long-run. Policy-makers could exploit this and provide subsidised temporary jobs as linking opportunities, with the goal of alleviating long-term youth unemployment.

Key words: *social networks, job referrals, job search, field experiment.*

*University of Oxford. Corresponding author: marc.witte@economics.ox.ac.uk.

I am grateful to Girum Abebe, Johannes Abeler, Stefano Caria, Marcel Fafchamps, James Fenske, Simon Franklin, Douglas Gollin, David Huffman, Clément Imbert, Pramila Krishnan, Julien Labonne, Christopher Roth, Simon Quinn, Margaret Stevens, Christopher Woodruff and Liam Wren-Lewis for helpful comments and to Alemayehu Woldu and Mersha Beletew for outstanding research assistance. I received helpful feedback from seminar participants at Oxford and UEA and from conference participants at NEUDC, SEEDEC, ABCA, RES, Warwick, CSAE, EEA, TADC-LBS and GLAD. Data collection and experimental implementation were funded by the International Growth Centre. The project would not have been possible without the constant support of Rose Page and the Centre for the Study of African Economies (University of Oxford), nor without the support of the Ethiopian Development Research Institute in Addis Ababa. This field experiment was registered in the American Economic Association Registry for randomized control trials under trial number AEARCTR-0002334.

1. Introduction

Job referrals in social networks are ubiquitous, both in developing and developed countries. Across countries, a substantial fraction of workers obtained their job through social contacts — from 13% in the US to 72% in the Philippines. In developing economies, the stronger presence of spatial and informational labour market frictions often increases the dependence on social networks.[1] Referrals in social networks are such a widespread phenomenon because relying on this practice has a variety of payoffs, for both work-seekers (such as reduced search costs) and employers (such as improved peer monitoring or reduced screening costs).[2] What is missing from the empirical literature is whether and how workers benefit from making job referrals. In particular, we lack evidence on which strategies drive the referral decisions of employees and how they affect who gets invited to the job.

This paper investigates if and how workers benefit from referrals, by experimentally varying the monetary and non-monetary benefits of real job referrals between workers. To do this, I set up a casual data entry firm in Addis Ababa, Ethiopia, where workers can refer other people from their social network for the job. I randomly vary the conditions of this job referral process along two dimensions. The first dimension is a monetary incentive payment: some workers receive a performance incentive according to the work performance of the individual they refer, while others do not get incentivised. The second dimension is the degree to which the invited contact is aware of who invited them, with referrals either completely open (taking place in the field), partly anonymous (where invited workers can reciprocate the referral, but do not know the referee) or completely anonymous (no reciprocity possible). These two treatment variations enable me to cleanly identify how workers trade off monetary and non-monetary payoffs and altruistic considerations while referring others for jobs. Generating a panel of workers' referral decisions over multiple work days, my experiment is the first in the literature that allows the analysis of mechanisms behind repeated referral interactions.

Identifying the relationship between job referrals and the social network position of the invited worker poses several empirical challenges. A first problem lies in the measurement of networks: researchers usually require complete network data (Chandrasekhar and Lewis, 2011), which are difficult to obtain, not just in developing coun-

[1]See appendix figure A2. Studies investigating social networks and labour market outcomes include Hellerstein et al. (2011), Topa (2011), Kramarz and Skans (2014), Schmutte (2015) for developed economies and Wahba and Zenou (2005), Munshi and Rosenzweig (2006), Serneels (2007), Beaman and Magruder (2012), Heath (2018) for developing countries.

[2]First, referrals reduce search costs for job-seekers (Calvó-Armengol, 2004, Galeotti and Merlino, 2014). Second, from the perspective of the employer, relying on worker referrals can improve peer monitoring (Kugler, 2003, Heath, 2018) and the screening of productive workers (Montgomery, 1991, Munshi, 2003, Beaman and Magruder, 2012).

tries. In order to overcome this, and before setting up the experimental firm, I collect complete and detailed social network data from 16 urban neighbourhoods for a total of 739 individuals. This includes comprehensive dyadic information on all social connections within the neighbourhood as well as aggregate information on out-of-neighbourhood contacts. Second, measuring the causal impact of networks on any type of outcome is difficult, as networks represent the endogenous outcome of processes happening over time, and an individual's network position is likely correlated with unmeasured ability and other confounders. Several recent studies investigating the impact of networks on labour market outcomes have provided exogenous shocks to network formation, e.g. by introducing a set of individuals to one another and inducing them to interact for some period.[3] In this paper, I take a different approach, with the expansion of the network chosen by the individual rather than the researchers, under exogenously varying job referral conditions.[4] This enables me to estimate the causal impact of different types of social interactions (through referrals) on individual network position. Third and lastly, various considerations simultaneously drive job referral decisions observed in reality. This limits our ability to identify the strategic motives in the network formation (or reinforcement) process. By randomly varying certain components of the referral payoff, my experiment overcomes this challenge and allows me to cleanly disentangle the mechanisms at play.

The field experiment uses a reduced 2x3 design (see table 1): the standard and 'natural' case is an un-incentivised, open job referral. This is the control condition with which I compare the other treatments. In the first margin of variation, I introduce an incentive treatment, meaning that the worker making a job referral gets a linear financial reward based on the performance of the invited worker. The second margin of variation (applied only to un-incentivised referrals) is the openness of the referrals, in which I introduce partly anonymous or fully anonymous referrals. Partly anonymous referrals only allow the invited worker to reciprocate the job invitation without knowing who invited them. In fully anonymous referrals, the invited worker does not know who invited them and cannot reciprocate. The experiment is conducted over three rounds (or work stints) with two rounds of job referral decisions, yielding a novel individual-level panel of productivity and referral data which enables me to analyse repeated interactions.

This paper contains four key findings, which are in line with a simple theoretical framework that relates a job referral's performance payoff, its social (non-monetary) payoff and the productivity of the referred worker. First, the job referrals are characterised by a substantial degree of reciprocity when conducted openly and without

[3] For example Feigenberg et al. (2013), Fafchamps and Quinn (2016), Cai and Szeidl (2017), Abebe et al. (2017), Vasilaky and Leonard (2018).

[4] I follow Beaman and Magruder (2012) and Beaman et al. (2018) in this aspect.

referee incentives.[5] A casual worker who invites another neighbourhood contact to the day job is almost 30 percentage points more likely to be re-invited by this same person, compared to being invited by a random worker. Importantly, mutual job referring comes at a substantial cost for productivity. Reciprocally invited workers produce between five and ten percent less output than non-reciprocally invited individuals, despite always being paid the same piece rate for each produced unit. This suggests that higher social payoffs prevent individuals from referring the most productive workers to the job. The high level of reciprocity also suggests that individuals with even temporary or one-off referral opportunities have much larger chances of receiving casual work in the future.

Second, this norm of direct reciprocity is substantially reduced under the randomly introduced incentivised referral treatment, where the number of reciprocal referrals decreases by approximately 20% (or 6.5 percentage points). This finding demonstrates the trade-off between performance payoffs and social (non-pecuniary) payoffs, as shown in my model. I also find that when given the performance incentive, workers refer more productive individuals to the job, a finding that is consistent with Beaman and Magruder (2012). If the other worker is invited under the incentivised referral treatment, he or she performs approximately 18% better on the job, compared to workers referred without incentives. This shows that workers are able to screen more productive individuals in local, relatively homogenous neighbourhood labour markets. The performance of workers invited without incentives is statistically indistinguishable from randomly invited workers.

My third finding exploits the exogenous variation in the 'openness' of the referral. I find that the network centrality of the referred workers varies substantially depending on whether the referral is open or anonymous. Under the standard open referral treatment, more central individuals are invited to the job. This is in line with my model, where referrals are not just made with the immediate expectation of direct reciprocity, but also with regard to future network benefits, such as access to information.[6] As a consequence, individuals without many connections are left jobless. This behaviour cannot be explained by central individuals being more able workers: I find no significant correlation between an individual's network centrality and her on-the-job productivity. In contrast, under the referral treatments which are partly anonymous or anonymous, less central but equally productive workers are selected. Under partly anonymous referrals, this is because immediate (and rival) reciprocity is the only network benefit that matters, which is maximised by selecting non-central workers. Under fully anonymous referrals, where any network benefits are eliminated, only other-regarding

[5] In line with Fehr and Gächter (2000), Gächter and Falk (2002).

[6] Similarly Bala and Goyal (2000) predict star-shaped networks (with all connections leading to the central node) in the case of bilaterally flowing non-rival goods such as information.

preferences matter, which explains why the least central workers are chosen.

Lastly, I show that there is persistence in exclusion. Approximately 9% of respondents report not knowing any other respondent in the area by name, and a quarter of respondents only know up to two people.[7] These peripheral individuals who are not (randomly) allocated a day job and a referral opportunity remain excluded from the casual labour market. This could provide an explanation for the existence of long-term joblessness and detachment from local labour markets. On the other hand, isolated or peripheral individuals who are (randomly) given referral opportunities manage to use these opportunities as a device to connect to larger parts of the social network. They are more than twice as likely to invite a previously unknown worker to the job than individuals who are already well-connected at baseline. These new work links are neither unilateral nor short-lived. On the contrary, the other worker, now a new link in the peripheral worker's network, is more likely to reciprocate and provide the peripheral worker with a return job, compared to a connection that already existed at baseline. 18 months after the experiment, almost 40% of individuals who did not know each other at baseline are still in contact, demonstrating that a very light-touch intervention can have long-run effects on connectedness in local labour markets. As opposed to a recent literature characterising the potential costs of relying on social networks for hiring (Calvó-Armengol and Jackson, 2004, Galenianos, 2018, Beaman et al., 2018), I identify one-off job-referrals as a way of escaping exclusion. This finding has implications for how policymakers could think of alleviating long-term youth unemployment, for example through the provision of subsidised temporary jobs as linking opportunities.

My paper contributes to different strands of literature. Closest to my paper is a field experiment by Beaman and Magruder (2012), run in Kolkata, India. The authors exogenously vary the monetary incentives of a one-off job referral process. In line with my results they find that under incentivised referral treatments participants refer high-ability workers to tasks. I extend their study in at least two important ways: first, Beaman and Magruder (2012) only study a one-off interaction, preventing them from describing strategic repeated game interactions which are characteristic of urban and casual labour markets. By observing repeated rounds of referral decisions, I can demonstrate that reciprocity is an important motivation for referrals. Since the dynamic nature of reciprocal referrals increases workers' social pay-offs and leads to lower work performance of the invited worker, I can show that the monetary compensation needed to incentivise workers to choose the optimal referral from the firm's perspective is likely higher than a one-time interaction suggests. Second, as they do not observe who misses out on job referrals, Beaman and Magruder (2012) cannot directly speak to important

[7]For these peripheral nodes, the average number of all social contacts *including* individuals living outside the neighbourhood is four, and hence equally low.

aspects of labour market exclusion. By collecting complete network data prior to the experiment, I can relate the job referral decisions in my experiment to the full, real-life social network of the workers and describe how peripheral workers strategically enlarge their network.

I also contribute to the broader literature on the role of networks in developing countries. Referrals to jobs illustrate the broader phenomenon that social networks shape labour markets and individual outcomes — not just in developing countries, but on a global level. People help each other find employment (Granovetter, 1995), and individuals are more likely to be employed if their friends are (Topa, 2001, Calvó-Armengol and Jackson, 2004). Despite global record levels of urbanisation (Cohen, 2006), social networks in urban areas, particularly in developing countries, are relatively understudied. Most empirical work in development economics either focusses on village communities (e.g. Banerjee et al., 2014, Breza and Chandrasekhar, 2015, Chandrasekhar et al., 2014, Cai et al., 2015) or other social connections, such as family networks (e.g. Magruder, 2010), where geographic proximity or kinship mediate network effects. This is an important gap in the literature, as we know from mostly descriptive studies that job information and referral networks are often situated in urban neighbourhoods (Hipp et al., 2012, Hellerstein et al., 2014, Rothstein, 2018, Caria et al., 2018). Similar to Rigol et al. (2017), I show that screening knowledge can be activated in the context of urban communities.

Lastly, this study relates to a smaller literature on strategic network formation in the context of laboratory or lab-in-the-field games. This literature has shown that individuals link up with more central participants in laboratory games, where the social network structure is artificially imposed either by anonymous computer games without direct human interactions (Conte et al., 2009) or by ascribing participants a random network position (Caria and Hassen, 2013). I provide the first empirical evidence that a similar strategic network formation process takes place within actually existing, real-life social networks, where links are formed 'in the field'.

The remainder of this paper is structured as follows: in section 2, I develop a conceptual framework of job referrals in social networks. Section 3 presents the detailed experimental design and protocol, while section 4 discusses its implementation in the local context and provides summary statistics. Section 5 shows the empirical specifications, the results of which are presented and discussed in sections 6 and 7. Section 8 concludes.

2. Conceptual framework

In this section, I present a stylised model of job referrals in social networks. My model initially builds on the theoretical frameworks presented in Bandiera et al. (2009) and Beaman and Magruder (2012). I then extend the latter version by explicitly modelling the different payoffs individuals can receive their social network.

2.1. Job referrals in social networks

Consider the employed worker i, who has the opportunity to make one job referral to another worker j from her neighbourhood contacts this period. The decision of the worker is: who should I refer to the job? The referral contract is specified by the employer.

Each worker and referral can perform on the job with ability $\theta_i \in \{\theta_i^H, \theta_i^L\}$. For simplicity, ability is modelled as a binary that can take the values *high* (H) or *low* (L), so that there are two types of contacts, those of high ability and those of low ability.[8] The employer offers the worker a referral contract, consisting of an incentive payment (P_i) based on the referred contact j's performance. In particular, P_i is the amount that i will get if j produces output, and P_i is linear in j's output, i.e. $P_i = P_i(\theta_j) = p_i\theta_j$. I assume that worker i can directly observe her contact's ability $\theta_i \in \{\theta_i^H, \theta_i^L\}$.[9]

Every worker i expects monetary pay-offs from the employer for referring contact j. The payoffs are a function of the contract type P_i and the referral's ability θ_j. Workers also expect a social payoff σ_{ij} from the contact j. The social payoff can be interpreted as an actual cash transfer from j to i or the inclusion of j's income in i's utility (following Bandiera et al. (2009) and Beaman and Magruder (2012)), but it can also include immediate and future network benefits, such as reciprocal job referrals or access to information. Given the two stylised types of friends, the two social payoffs are $\sigma_{i1} \in \arg\max_j(\sigma_{ij}|\theta_j = \theta_j^H)$ for the high-ability type and $\sigma_{i2} \in \arg\max_j(\sigma_{ij}|\theta_j = \theta_j^L)$ for the low-ability type.

When selecting the high-ability contact under an incentivised referral scheme, i will receive the payoff $P_i + \sigma_{i1}$, whereas for selecting the low-ability type she will receive σ_{i2}. Assuming that every worker i among her contacts c_i has a choice between the high-ability and the low-ability type, the worker compares the best high-ability type with the

[8]In the appendix section B, I derive the same framework with a continuous ability parameter.

[9]Alternatively, I could change the model to the scenario that every worker observes a signal of her contacts' ability, $\hat{\theta}_i \in \{\theta_i^H, \theta_i^L\}$ with the probability of accuracy β_i, where $P(\theta = \theta^H|\hat{\theta} = \theta^H, i) = P(\theta = \theta^L|\hat{\theta} = \theta^L, i) = \beta_i$, so that $\beta_i \in [0.5, 1]$. This case is presented in appendix section B. The implications of the model stay the same.

best low-ability type. As a consequence, worker i will make a referral to a high-ability type if

$$P_i > \sigma_{i2} - \sigma_{i1}. \tag{1}$$

In other words, worker i will refer the high-ability type if the performance payment is larger than the differences in social benefits. If the worker only has low-ability contacts, she will choose the contact j who provides her with the highest social benefits σ_{ij}. So far, this model follows very closely the version of Beaman and Magruder (2012).

2.2. Decomposing the social pay-off

Consider an employed worker i who has c_i contacts. The social payoff that worker i receives from referring contact j can be modelled explicitly: $\sigma_{ij} = \pi_{ji} + b_{ji} + \omega_{ij} m_j$, where π_{ji} are direct monetary payments from j to i, b_{ji} are network benefits flowing from j to i, (such as reciprocal referrals, network goods or information) and m_j is j's income in i's utility with j-specific weight ω_{ij} (i.e. altruism). The network benefits b_{ji} can be expressed as $b_{ji} = \rho_{ji} + \kappa_{ji}$, where ρ_{ji} are direct and reciprocal referrals from j to i, and κ_{ji} are further (including future) benefits of being connected to j, such as having access to j's information. As a result, the full expression for the social payoff becomes $\sigma_{ij} = \pi_{ji} + \rho_{ji} + \kappa_{ji} + \omega_{ij} m_j$. In the following, I will assume that the direct monetary payments π_{ji} are zero, since I will not observe them during my experiment.

The full expression for the referral payoff with continuous ability is:

$$\begin{aligned} \Pi_{ij} =& P_i(\theta_j) + \sigma_{ij} \\ =& P_i(\theta_j) + \rho_{ji} + \kappa_{ji} + \omega_{ij} m_j \\ =& p_i \theta_j + \rho_{ji} + \kappa_{ji} + \omega_{ij} m_j \end{aligned} \tag{2}$$

where the last steps assumes that P_i is linear in j's output.

As equation 2 shows, the social payoff σ_{ij} and the performance incentive P_i are substitutes. In particular, referral contracts that include a performance incentive should *ceteris paribus* make workers i select referrals j with lower social payoffs σ_{ij}, either through lower network benefits b_{ji}, a lower utility weight ω_{ij}, or both.

2.3. The role of network centrality

Every worker i has c_i connections in her local neighbourhood network. She also observes the number of connections of her contacts, c_j.[10] For simplicity, I interpret a worker i's number of connections c_i as her degree centrality.[11] A contact j's network centrality is going to affect the social payoff σ_{ij} through different channels: $\sigma_{ij} = \rho_{ji}(c_j) + \kappa_{ji}(c_j) + \omega_{ij}m_j(c_j)$.

First, the likelihood of receiving a direct reciprocal referral from contact j *ceteris paribus* decreases with her centrality c_j, so that $\frac{\partial \rho_{ji}}{\partial c_j} < 0$. This is because a contact j with fewer connections c_j is more likely to select i when it comes to direct reciprocity. Second, the likelihood of receiving future benefits from being connected to contact j increase with j's centrality, so that $\frac{\partial \kappa_{ji}}{\partial c_j} > 0$. Intuitively, this is because a central contact j has access to more network goods, which increases the likelihood that they get shared with i.[12] As a consequence, the relationship between all network benefits and a referral's network centrality is unclear, and depends on whether the direct reciprocity or future benefits dominate:

$$\frac{\partial b_{ji}}{\partial c_j} > 0 \text{ if } \frac{\partial \kappa_{ji}}{\partial c_j} > -\frac{\partial \rho_{ji}}{\partial c_j}. \tag{3}$$

When network benefits b_{ji} and direct monetary payments π_{ji} are zero, the social payoff reduces to other-regarding preferences: $\sigma_{ij} = \omega_{ij}m_j$. I assume that the weight ω_{ij} with which j's income m_j is included in i's utility is decreasing in j's network centrality: $\frac{\partial \omega_{ij}}{\partial c_j} < 0$. This means that under only other-regarding preferences ($\sigma_{ij} = \omega_{ij}m_j$), the social pay-off is maximised by referring the least central workers to the job:

$$\frac{\partial \sigma_{ij}}{\partial c_j} < 0. \tag{4}$$

This is in line with several models of social preferences and inequality aversion (Fehr and Schmidt, 1999, Bolton and Ockenfels, 2000, Charness and Rabin, 2002).

As a consequence, I can write the full referral payoff function as:

$$\begin{aligned} \Pi_{ij} =& p_i\theta_j + \sigma_{ij} \\ =& p_i\theta_j + \rho_{ji}(\overset{-}{c_j}) + \kappa_{ji}(\overset{+}{c_j}) + \omega_{ij}(\overset{-}{c_j})m_j \end{aligned} \tag{5}$$

Let us assume for simplicity that the elements of the social payoff function only depend

[10] Again, I could model that i observes a signal $\hat{c_j}$ of c_j, but I omit this for simplicity.

[11] The implications of this model are the same for most other measures of centrality, such as betweenness centrality or Katz centrality.

[12] This is certainly the case for non-rival goods such as information, while the implications for rival goods are more complicated, see Bala and Goyal (2000).

on j through c_j. In that case, I can rewrite equation 5 as:

$$\Pi_{ij} = p_i\theta_j + \rho_i(\overset{-}{c}_j) + \kappa_i(\overset{+}{c}_j) + \omega_i(\overset{-}{c}_j)m_j \tag{6}$$

As experimenter, I vary several parameters of this model. First, and similar to Beaman and Magruder (2012), I randomly introduce the incentive payment P_i that worker i can receive for her referral's performance. Second, I exogenously vary the network benefits b_{ji} and their components, by enabling partly anonymous and anonymous referrals to workers from other neighbourhoods.[13] Third, I vary the signal of j's centrality, γ_i, with the treatment that shrouds the centrality information for the referrals.

2.4. Deductions from the framework

Proposition 1 *An increase in P_i increases the benefits of referring a high-ability type (equation A2).*

What I expect to observe is a larger number of high-ability workers, as expressed by better job performance.

Proposition 2 *Performance payoffs P_i and social payoffs σ_{ji} are substitutes in determining the likelihood of a worker being referred (equation 2).*

What I expect to observe is that an increase in P_i will lower the social payoffs required from referral j for the same overall referral payoff Π_{ij}. This could be expressed through lower direct monetary payments π_{ji}, lower network benefits b_{ji}, a lower utility weight ω_i — or a combination of these.

Proposition 3 *An increase of ρ_{ji} compared to κ_{ji} in the network benefits b_{ji} will lead to the selection of referrals with lower centrality (equation 3).*

Proposition 4 *When network benefits b_{ji} and monetary benefits π_{ji} are switched off, workers will select referrals with the lowest centrality (equation 4).*

[13] While the partly anonymous referrals to other neighbourhoods decrease the likelihood of future benefits, κ_{ji}, and encourage the opportunity for direct referrals, ρ_{ji}, the anonymous referrals set the complete expression b_{ji} to zero.

3. The work experiment

I conduct a field experiment to investigate how workers benefit from making job referrals in their social networks. In particular, I look at the mechanisms driving job referrals in social networks. The design of the experiment allows me to directly speak to the propositions of my conceptual framework from section 2, with important implications for how social networks affect equality and efficiency in local, informal labour markets.

The design of the experiment follows four principal steps. In a first step, I create a casual day labour market: I randomly invite young workers and job seekers from urban neighbourhoods to work on a real and paid data entry job. Second, there are several consecutive work days (or rounds) in the experiment, which is not announced to the participants ex ante. Third, after every work day, workers are given the opportunity to refer individuals from their local social network to the same job the next day. Finally, the conditions of this 'referral treatment' vary randomly along two principal dimensions: 1) incentivised vs un-incentivised referrals and 2) in-person vs gradually anonymised referrals. In addition to that, some individuals are invited randomly[14] as an additional comparison category. Additionally, in a sub-treatment I also vary the signal of the centrality of the potential referrals by shrouding the centrality information available to the inviting workers. The data entry job performed by the workers relates to real and large payoffs (paid as piece rates), is conducted individually, and requires some endurance and cognitive effort.

The participants of the experiments are young urban dwellers aged 18 to 29 who are not in permanent employment or education in order to guarantee that the participants spend sufficient time in their local neighbourhoods. Before the start of the experiment, I take a census of all eligible individuals in a given neighbourhood and survey them with a detailed baseline questionnaire, including dyadic information about interactions with every other participant in the neighbourhood. More details on the sampling frame and participant characteristics can be found in section 4.

3.1. Experiment design

This section describes the design of the experiment. The experiment is conducted over three consecutive work days with two rounds of job referral decisions, which gives me an individual-level panel of productivity and referral data. Its structure is presented in figure 1. I randomly invite two-thirds of all eligible individuals from the neighbourhood to the first work session, while the remaining third does not get an invitation

[14] By a random number generator on a computer.

(for the first work session). Among those workers participating in the first session, half are randomly allocated to the control referral treatment, and the other half of participant workers get a treated referral opportunity. The treatment variations follow a reduced 2x3 design, displayed in table 1. The control referral treatment (C) is an un-incentivised, open job referral, similar to 'stabdard' job referrals mostly observed in reality. Workers in this group can invite someone from their neighbourhood network to the next work session without incentives and openly 'in the field'.[15] Treatment variation 1 (T1) introduces an incentive treatment, meaning that the worker making a job referral gets a financial reward which is linear in the performance of the invited worker. This treatment positively 'shocks' the monetary referral payoffs described in section 2. The second margin of variation shocks the openness of the referrals. Workers in the control referral condition can additionally make partly anonymous or fully anonymous referrals to other workers. Under partly anonymous referrals (T2), the invited worker does not know the identity of the inviting worker, but can reciprocate the referral, while under fully anonymous referrals (T3), reciprocation is now allowed. More details on the treatment variations follow below.

Every participant can only make one referral to someone in her neighbourhood, excluding referring oneself. The second work session is then comprised by individuals referred to by the first round workers. Of all these workers present in the second round, I again randomly provide half of the group with a control referral and the other half with a treated referral. Accordingly, the third work session is again composed of referred workers.[16]

[Figure 1 about here.]

[Table 1 about here.]

Open 'control' referrals (C): As figure 1 shows, the open (or control) referral opportunities are given after session one and two to a random half of the participant workers. Participants can refer to any other eligible and surveyed person from their own neighbourhood of residence, i.e. anyone in their local social network. This includes individuals who are also working during the same day, workers who did not attend,

[15]Hence, the referral consists of both an informational and 'appointment-like' dimension. Even though it is not enforced, everyone in the referral treatment group always invited another worker to the next session.

[16]In some neighbourhoods, the treated referral (T1) is replaced by a 'no referral' treatment (i.e. workers in this group cannot invite a person). In subsequent sessions, the residual 50% of workers are randomly invited from the neighbourhood. I introduce this sub-variation to have an additional comparison category that gets rid of worker selection. I control for this minor variation in experimental design with neighbourhood and session fixed effects.

and jobless individuals who got randomised into the control group in the first round. Participants can only invite one other person. As a consequence, the 'most popular' person in a neighbourhood cannot receive referrals from all participant workers – only the first invitation counts.[17] The actual referral process takes place outside of the work sessions, i.e. through personal interactions of workers "in the field". Before leaving from work at the end of the first or second session, workers who were randomised into the open referral treatment are given the full list of their neighbourhood social network, including the other residents' names and network degree. This (signed) list serves as a ticket for the referral during the next session and ensures that real-life contact between the inviting worker and the referred worker has taken place.[18]

3.1.1. Treatment variation: incentivisation of referral (T1)

I exogenously vary the performance incentive of the referral treatment: a random half of participants i receives an incentivised referral opportunity ($P_i > 0$), meaning that whatever their referral j earns during the next session is matched for them ($P_i = P_j$). Every other aspect of this referral treatment is identical to the open control referral (C), where the participant workers do not receive any payment related to the performance of the worker they invite (P_i=0). This treatment variation allows me to test proposition 1, i.e. analysing whether the introduction of P_i increases the number of high-ability workers, as expressed by better job performance.

3.1.2. Treatment variation: visibility of referral (T2 & T3)

The treatment variation in referral visibility shocks the components of the social payoff σ_{ij} in my conceptual model. The treatments are based on within-subject variation for workers in the open referral group (C) (see table 1). In these two treatments, I vary the extent to which the referred worker can reciprocate and see who invited him/her. In order to avoid treatment contamination, meaning that these gradually anonymised treatments do not remain anonymous but become public knowledge, referrals in T2 and T3 are made to real workers living in other neighbourhoods. These workers are not within the participant's own social neighbourhood network, but live in virtually identical economic circumstances, so that information on who made the anonymous

[17] This potentially attenuates my parameter estimates for degree centrality, if everyone would have wanted to invite people with higher degree centralities that were already invited by someone else. In a sense, this problem is similar to assortative matching on markets, where ideally everyone would like to e.g. marry the most intelligent partner, but only one person gets to do so.

[18] An exemplary own neighbourhood referral list can be found in the appendix at the end of this document. All other referral lists introduced in the following paragraph can also be found in the appendix.

referrals will not diffuse through the network.[19]

1. **Partly anonymous referral to other neighbourhoods (T2)**: From a list of (anonymised) individuals from a different neighbourhood (that is similar to the one the worker lives in), workers choose who to invite to the day job. The referral list contains information on the potential referral's age, sex and number of neighbourhood connections.[20] The salience of network centrality is thus constant between C, T1, T2, and T3. The mobile phone number of the experiment worker who is making the job referral is given to the person she invites, and the invited worker is explicitly primed to reciprocate or 'thank' the inviting worker through the research team by selecting her for a similar job. Only one such reciprocal referral can be made. As a consequence, this treatment encourages immediate reciprocity ρ_{ji} concerns, whereas future network benefits κ_{ji} between the inviting and the invited worker are essentially switched off and cannot be enforced by the inviting worker.

2. **Anonymous referral to other neighbourhoods (T3)**: This referral is similar to the partly anonymous referral to other neighbourhoods, but no information about the inviting worker is passed on and reciprocity is not possible. The referral is therefore made completely anonymously.

During the experiment, the order of T2 and T3 is randomised over sessions to avoid order effects.

3.1.3. Additional treatment variation: centrality shrouding (T4)

Since I investigate the implications of worker referrals on labour market equality, one potential test is whether participants refer workers in their social network who have a higher degree centrality, i.e. more links to other nodes. In reality, an individual's degree centrality is not randomly assigned and likely correlates with other hard-to-observe or unobservable personal characteristics (e.g. sociability, physical attractiveness, intelligence). In order to distinguish between such personal traits and the isolated effects of a person's connectedness, I introduce a separate exogenous sub-treatment based on between-worker variation. In some neighbourhoods, a random half of participants invites a worker based on a referral list where the true number of connections (i.e. the degree centrality) is shrouded. The remaining workers get the usual referral lists, i.e. stating the number of connections explicitly for each person j in the neighbourhood. In practice, the referral with the shrouded degree centrality means that the referral list contains two numbers of friends for each potential referral j, one true number and one

[19] The within-subject design is chosen to maximise statistical power.

[20] The referral list consists of real-life individuals similar to the original workers, and the referral choice of the original experiment workers is implemented in reality, meaning that these 'other neighbourhood' referrals are real.

random false number. As I elicit every person i's belief about j's number of friends at baseline, I control for these baseline beliefs in my regressions in section 5.

4. Context, implementation, and summary statistics

4.1. Job information networks in a large Sub-Saharan African city

Social networks are an important component of everyday life and determinant of socioeconomic outcomes, both in developed and developing countries. Among other things, they are known to influence physical and mental health (Christakis and Fowler, 2007, 2009), educational achievements (Sacerdote, 2001), and labour market outcomes (Ioannides and Datcher Loury, 2004, Beaman and Magruder, 2012, Fafchamps and Moradi, 2015), which is also the focus of this study. In developing countries, they often fulfil additional roles as financial safety nets (Collins et al., 2009) or providers of insurance (Banerjee and Duflo, 2007). Broadly speaking, social networks can usually be characterised as facilitating the spreading of information and/or financial flows.[21] In this study, I focus on job-referral networks (also: job search or job information networks) which can be classified as a combination of both, since the information about a job vacancy often comes with a monetary value of a potential job (Munshi, 2003, Beaman, 2012). Given the importance of social networks for various economic outcomes in developing countries, surprisingly little is known empirically about the formation of such networks. A small number of papers look at network formation explicitly, focussing on why individual linkages happen, with the example of immigrant communities in Italy (Comola and Mendola, 2015), new village settlements in Zimbabwe (Barr et al., 2015), or access to savings in Nepal (Comola and Prina, 2017).

Ethiopia's strong GDP growth rate of 10% on average over the last ten years makes the country one of the fastest growing in Sub-Saharan Africa (The World Bank, 2014). At the same time, while rural areas are traditionally dominated by subsistence agriculture, most of the country's young urban population is out of permanent or formal employment. Unemployment rates are particularly high for young people who graduated from high school or higher education institutions, despite widely reported shortages of qualified employees by Ethiopian firms[22] – suggesting a problem with matching job seekers to vacancies. In addition to that, research has emphasised the importance of referral networks in the Ethiopian urban labour market (Serneels, 2007, Caria and Hassen, 2013), with a large share of (formal and informal sector) positions being filled

[21] Chuang and Schechter (2015) provide an extensive summary of the literature.
[22] Source: data from light manufacturing and service sector firms surveyed in Addis Ababa in 2015, (cf. Abebe et al., 2016).

through job seekers' social networks. Extending the network position of unemployed urban residents could therefore increase their chances of finding a job. In line with a recent strand of literature (Franklin, 2016, Abebe et al., 2016), this paper hypothesises that urban unemployment is at least partly caused by information frictions about existing work opportunities. My experiment tries to bridge this gap, enabling young job-seekers to modify their social connectedness by providing them with real day jobs to pass through their existing network.

In line with other empirical studies, the neighbourhood networks in this paper display large heterogeneity in the numbers of links by individual network member (as shown in section 4.4). Given the importance of job contacts for labour market outcomes, unconnected individuals risk isolation from job information sources. Adding to this is the fact that labour market information is exceptionally centralised and geographically restricted within the city. Most public and white collar jobs are only publicly advertised on centrally located vacancy boards (Franklin, 2016, Abebe et al., 2016). As a consequence, job information is more likely to have to gradually 'trickled' through existing network structures rather than being accessed by all job-seekers simultaneously.[23] My experiment makes use of the disproportional importance that informal networks play in the Ethiopian labour market, by using full network data on all eligible participants from a spatially restricted area of an urban neighbourhood.

4.2. Neighbourhood and participant sampling frames

The experiment was conducted in Addis Ababa between December 2016 and February 2017, while the network census and baseline surveying already started in November 2016 (more information on the timeline can be found in the appendix, section A). This paper focusses on local neighbourhood networks of young casual workers and job-seekers. The choice of sample is motivated by the existing empirical evidence on job search networks in developing countries (Wahba and Zenou, 2005), in particular in the Ethiopian context (Serneels, 2007), where informal labour markets are often dominated by local information exchange and direct referrals to casual jobs. I can observe such a behaviour directly in my baseline sample: 95% of those individuals who ever approached another person for help in getting a job asked directly in their neighbourhood, almost always involving an exchange of relevant information. In 40% of these cases, the exchange of job information led to the provision of actual jobs through a direct referral. Actual work and information about work is often obtained through

[23] Internet penetration in Ethiopia overall was at 11.5% in 2015, but substantially higher in Addis Ababa (The World Bank, 2016). However, the few online or mobile phone services for job-seekers that have been established recently are still very small in size, both in terms of users as well as job vacancies on offer.

these neighbourhood networks among people out of permanent education or employment. Further research on the nature of young job-seekers' social connections across space (Abebe et al., 2016, Caria et al., 2018, Witte, 2017) shows that job search in Addis Ababa predominantly takes place in pairs of not permanently employed individuals, who share information about job vacancies and take trips to the centrally located vacancy boards in rotation.

In a first step, I randomly select 16 urban neighbourhoods in Addis Ababa from which I later collect baseline data on young unemployed individuals and their social networks. I define as a neighbourhood an enumeration area from the Addis Ababa census, in line with both the literature on urban networks (Bayer et al., 2008) and the Ethiopian Central Statistics Agency.[24]

The neighbourhoods included in the randomisation sample represent the 'average' urban neighbourhood in Sub-Saharan Africa, densely populated, with mostly single-storied buildings without any compounds or gardens. More than 65% of the population of Addis Ababa live in such neighbourhoods.[25] I randomly select a total of 16 neighbourhoods from the sample of over 2000 eligible ones. Figure 2 displays my randomly selected neighbourhoods on a map of Addis Ababa. The exclusion of very central neighbourhoods leads to an almost circle-like distribution of selected neighbourhoods around the city's core.

[Figure 2 about here.]

Within the selected neighbourhoods, I take a complete census (via door-to-door solicitation) of all eligible resident individuals,[26] where eligible individuals 1) permanently live in the selected neighbourhood, 2) are between 18 and 29 years of age (inclusive), and 3) are not in permanent employment or education. Permanent employment is defined as full-time employment at one single employer, including full-time self-employment, going at least one month back. Permanent education is defined as a

[24] From the universe of urban neighbourhoods in Addis Ababa, I exclude neighbourhoods from the randomisation sample that are: 1) very large neighbourhoods (more than five hectares); 2) in the bottom quarter or top quarter of population density, to exclude sparsely populated/industrial areas as well as recently built condominium blocks; 3) in the bottom quarter or top quarter of total population (highly correlated with previous point), for similar reasons; 4) very central neighbourhoods (as defined by a distance of less than two kilometres to the National Stadium), as these are mostly commercial blocks; 5) neighbourhoods from the 'Bole' sub-city, an affluent neighbourhood characterised by very high rates of permanent employment and an overall older population.

[25] The remaining 35% are almost equally divided between multi-story condominium complexes, compound houses/ gated communities or informal shantytowns in the outskirts, often characterised by extreme poverty.

[26] In practice, this means that enumerators went to every household in the neighbourhood and asked if an individual satisfying the eligibility criteria was living in the household. If yes, the eligible individual's name as well as basic demographic and contact details were recorded, to populate the baseline questionnaire's network roster.

full-time degree course at a university, school or training institute.

These eligible individuals are surveyed with a baseline questionnaire, including questions on demographics, job search and employment, attitudes and aspirations, as well as a detailed social network section, containing a range of questions about every other eligible individual living in the neighbourhood. Based on these information, I draw social network maps of every neighbourhood (see next section). My complete sample consists of 739 individuals from 16 urban neighbourhoods who were surveyed during baseline in November/December 2016; for summary statistics see section 4.4.

4.3. Data entry job description and work logistics

Shortly after the completion of the baseline survey in their neighbourhood, individuals are invited to the experiment work session, which is announced to them as a paid 'day job'. The sessions take place separately by neighbourhood, and are held in halls or meeting rooms situated close to the respective neighbourhood. Upon arrival at the work site, I select a random two-thirds of present individuals to participate in the first work session.[27]

The data entry job requires some effort from the workers: participants are handed employee data from a company that they have to sort according to employee work team and age (there are 20 work teams of three fictitious workers that need to be sorted). Similar tasks has been tested in very related settings (Abebe et al., 2016).[28] The work performance is individually incentivised with a piece rate of ten ETB per correctly entered work team, for a maximum of 200 ETB. This is a non-negligible payment ($\simeq$ \$12.00) compared to the median weekly spending of 300 ETB per participants in the sample. After the completion of the task, participants are asked to make the different job referrals explained previously.

The day jobs take place on consecutive days or are spread over a whole day, depending on logistical factors. In the empirical analysis, I include session fixed effects to control for any heterogeneity caused by different sessions. The second and third session of the experiment are unannounced to the participants ex ante. It is, however, possible that participants shape expectations for a third session after the second session has been revealed. The order of the within-worker referral treatments (T2+T3)

[27] The other people are reimbursed for transportation expenses, but not told that there are any further sessions.

[28] It was not possible to conduct physical tasks, which would be somewhat closer to the reality of casual and temporary workers in this context (e.g. while working on construction sites or carrying deliveries to shops). However, the work only depends on individual performance, demands real effort, and does not contain any elements of service-style or customer-facing jobs. In that sense, it is not too far away from the kind of work the individuals in my sample are used to.

is randomised over sessions to get rid of order effects and to obtain unbiased within-individual measurements. The referrals in the own neighbourhoods (C+T1) always takes place between work sessions, in the field. Further details about the experimental timeline and protocol can be found in the appendix, section A.

4.4. Baseline summary statistics

4.4.1. The neighbourhood network graphs

Figure 3 depicts the social network graphs for all neighbourhoods, displaying connections between individuals as directed links. A link is established in this picture when an individual replied positively to the question "Do you know [$name\ of\ j$]?". The networks, in particular the number of isolated individuals who do not know anyone else from their network, vary considerably in shape – from eight completely excluded nodes in figure 3, network (b) to zero in a few others. Note that these network graphs present linkages as long as one individual claims knowing the other, which is a more lenient condition than the requirement that both individuals state a connection, in which case the share of isolated links would increase considerably. It is a well-known phenomenon in network data that a significant proportion of links are only reported unilaterally (Comola and Fafchamps, 2014). Moreover, the network graphs give a first 'flavour' of how dense these social networks are. Only rarely do we observe different separated sub-clusters (network components) of groups of linked individuals. Mostly, isolated components consist of single nodes without connections to the main network component.

[Figure 3 about here.]

Appendix figure A5 displays a different type of network for the same neighbourhood, this time only linking individuals that have exchanged job information or referrals with one another. It becomes immediately clear that these job networks are less dense than the networks characterised by individuals knowing each other. In particular, the number of isolated individuals, i.e. people who have never received job information from or given job information to anyone in their network is much larger, with 20% on average and reaching up to 55% in one outlier neighbourhood (figure A5 , network (m)). Generally, the density of the job information and referral exchange seems to closely track the density of the "knowledge" networks – with neighbourhoods where most individuals know each other also displaying a larger number of job information dyads.

4.4.2. Network-level summary statistics

The first panel of table 2 displays network-level summary statistics calculated over all 16 neighbourhood networks in this paper. The number of individuals per neighbourhood network varies between 27 and 61, with an average of 46, and overall little variation. The average number of general links (two individuals knowing each other by name or nickname) between nodes per network is quite large at 355, resulting in an average degree centrality of 9.31. This means that on average across all networks, every individual is *directly* connected to over nine other nodes.[29] The variation of the total number of arcs across networks is very large, pointing to substantial differences in how 'dense' the neighbourhood networks are. Indeed, the network density (calculated as the ratio of the number of arcs to the number of possible arcs in a network) varies strongly from 0.05 to 0.43. As one can assume in these neighbourhood networks, the average network transitivity is very large at 0.61, with a maximum of 0.99, meaning that, on average, the chance that two individuals in a neighbourhood share a common friend is 61%. In terms of reciprocity r (i.e. the ratio of the number of links pointing in both directions to the total number of links), it is clear that the network is far from being purely bidirectional ($r = 1$), with an average of 0.27, and never exceeding 0.43. Taken together, the average statistics for reciprocity and transitivity suggest that my networks are very dense, but this density is not always covered by bi-directional connections.[30] There is little variation in the average eigenvector centrality, as it can only be defined for networks with one single component (only six of the 16 networks in my sample consist of one single component). Lastly, the networks on average consist of 2.88 separate components (i.e. unconnected parts), which is mainly driven by one outlier network with many isolated individuals (and ten separate components). Six of the 16 networks consist of one single component and as a consequence have no completely excluded nodes.

[Table 2 about here.]

The second panel of table 2 displays the same summary statistics for the neighbourhoods, this time only for the job information and referral networks, i.e. those networks where two individuals are linked when they have exchanged jobs or job information in the past. As already visually presented in appendix figure A5, the job exchange networks are less dense than the broadly defined social networks, with accordingly lower

[29] This is a lower bound: qualitative post-experiment interviews revealed that a substantial share of individuals ($> 50\%$) knows each other by face only, without knowing name or nickname of the other individual.

[30] I can overcome any concerns about the directionality of links by defining a link as undirected, i.e. it is established if either i knows j **or** j knows i.

levels of degree centralisation. Interestingly, the average network transitivity is even higher in these job networks. Similarly, the level of reciprocity in the job networks is slightly higher than in the general case. The number of components per network is also much larger than in the general case, reflecting the fact that a larger share of individuals are isolates or only in a dyadic job exchange relationship with a single other person.

Appendix table C3 shows the equivalent network characteristics to those of table 2 separately for each of the 16 networks, giving an idea of the substantial variation in network measures across networks.

4.4.3. Social interactions within neighbourhood networks

Moving from network-level summary statistics to individual-level social interactions, table 3 presents information on the density of social interactions within the neighbourhood networks. The average number of links that each individual has is approximately 9.5,[31] of which she considers three-quarters as relatively central (i.e. as someone to approach if one wanted to spread information about an event in the whole neighbourhood, following Banerjee et al. (2014)). On average, individuals spend a total of 275 hours per month in their local network, which leads to a substantial average of nine hours per day.[32] Individuals travel to the city centre with approximately a quarter of their connections. Among all social connections, an individual shares job information with approximately 30% of them, supporting the notion that talking about jobs is an important part of these neighbourhood networks. In fact, of those individuals who ever asked another person for help getting a job at baseline, 95% asked directly in their neighbourhood, and in 96% of those cases also received meaningful information on jobs. In 40% of these cases, the exchange of job information led to the provision of actual jobs through a direct referral. This demonstrates that actual work, not just information about it, is often obtained through neighbourhood networks. The experiment is thus conducted in an environment that is salient for job referrals, and does not expose individuals to an artificial 'laboratory' situation. The exchange of money is less prevalent among links, and takes place for fewer than 20% of connections.[33]

[31] This number is different from the network-level summary statistics presented above, because those network-level observations are not weighted by the number of individuals per network.

[32] The maximum number of hours spent with individuals from the network exceeds the total number of any given month. This is because for individual i, the time is summed over all other individuals j, which does not take into account that individuals also spend time in groups from the same network, in which case the time for every other individual j is added separately to the total time for individual i.

[33] I observe small differences in the bi-directional variables for visiting someone's home as well as job and money exchanges. Since the networks are complete and symmetric, these variables should in principle be exactly the same. For instance, for every individual i saying she lends money to individual

When comparing these social interactions to the total size of the network, I find that every individual knows about 18% of individuals in the network, with substantial variation across individuals. Incidentally, this average of 0.18 is comparable to the network-level density (0.16) presented in table 2, whereas the extreme values and standard deviations are very different, as measurement in table 3 is on individual-level. Overall, the number of individuals known directly by a respondent, i.e. those people she or he is connected to with degree one, can reach up to 96% of the whole network.

Appendix figure A1 shows standardised coefficients from a multivariate regression of local network centrality on individual characteristics. It shows that being male, part of an ethnic majority in the neighbourhood and months lived in the neighbourhood are the strongest predictors of centrality.

[Table 3 about here.]

4.4.4. Individual-level data and balance

Lastly, I present individual-level summary statistics. In appendix table D4, I describe the whole sample of surveyed individuals, and compare the sample of individuals participating at the work sessions with those who were surveyed, but never participated. It is clear that individuals who came to the sessions consist of a different sample than those who did not, either because they did not show up during the first session, or because they were not invited by anyone during the rounds. Most importantly, the non-attenders are more likely to be in (non-permanent) work, less likely to search for (additional) work, and are less connected to other individuals in the sample. Overall, 25% of individuals have worked in the last seven days, with more than 80% having some kind of paid work experience. The remaining individuals are first-time job seekers. Approximately two-thirds of the individuals still live with their parents and have finished school on average three years ago.

Table 4 shows the summary statistics for those individuals who were randomly allocated to the referral treatment groups during the first sessions of the work experiment. Almost all characteristics are balanced across treatment groups, indicating that randomisation was successful. The test for joint orthogonality of all characteristics is insignificant, with a high p-value of 0.88.

[Table 4 about here.]

j, j should indicate he receives money from i. However, in both cases of job and money exchanges, the values are tilted slightly towards individual i making the visit or giving the money/information, indicating that individuals may be slightly more likely to remember when playing the active role in an interaction.

5. Empirical strategy

I run two principal versions of regressions in order to test the different propositions of my model. The first type are cross-sectional regressions, either separately for each round of the work experiments or pooled together, depending on the question of interest, coupled with the baseline information. Exploiting the random variation in the type of referral treatment assignment – i.e. whether participles are invited to the first work session, whether they receive a referral opportunity, and what kind of referral – I describe overall trends in the referral strategies and test propositions 1, 3 and 4. The second type of regressions are dyadic regressions run on an adjacency matrix of the social neighbourhood network, looking at direct relationships between every pair of individuals in the neighbourhood. In the following, I will briefly present a typical regression specification for both (pooled or separate) cross-sectional or dyadic data analysis, while more specific deviations from these regressions will be introduced and discussed together with the results in section 6.[34]

Measures of network centrality: Different measures of network centrality have different characteristics, but all of them try to identify what the 'most important' node in a social network is. My main measure is degree centrality, i.e. the number of links an individual has. In most of my analysis, a link is defined in two principal ways: for general social networks, it is defined as one person i knowing another person j by name. I mostly focus on directed networks, i.e. a link between two individuals i and j is defined either by i claiming to know j, $i \to j$, or vice versa: $j \to i$. For job information and referral networks, a link is defined as one person i having exchanged job information or referrals with another person j in the past, again defined (from i's perspective) as out-degree ($i \to j$) or in-degree ($j \to i$). In addition to that, I also look at further measures such as between-centrality, Katz centrality or clustering centrality, which I introduce in section 6.

I have three main outcomes of interest:

1. A worker i's on-the-job productivity, as expressed by correctly entered work teams: P_i.
2. A worker i's network degree centrality, as expressed by the number of (in or our) links at baseline.
3. Reciprocity: a indicator taking the value of one if worker j refers worker i to the job in period $t-1$ (i.e. $R_{ji,t-1} = 1$), and i refers j in return in period t (i.e.

[34] I have registered this experiment together with the primary outcomes and analysis of interest in the American Economic Association Registry for randomized control trials under Trial number AEARCTR-0002334.

$R_{ij,t} = 1$).

5.1. Cross-sectional analysis

I run cross-sectional regressions as displayed in specification 7 to test 1) whether the performance incentive makes workers refer more productive individuals to the job (proposition 1) and 2) whether the degree of anonymity of the referral leads to the selection of less central individuals (proposition 3 and 4):

$$y_i = \rho_0 + \rho_1 * T_j^k + \theta X_{ij} + u_i \tag{7}$$

where y_i is the outcome of interest of the referred worker i (work productivity or centrality) and T_j^k is the vector of the k different referral treatments of worker j. X_{ij} is a matrix of possible control variables, including session fixed effects and workers' characteristics.

As an example, to evaluate the impact of the incentive treatment (T1), the cross-sectional regression would take the following form:

$$P_i = \rho_0 + \rho_1 * T_j^{incentivised} + \rho_2 * T_j^{unincentivised} + \theta X_{ij} + u_i \tag{8}$$

In this example, $\rho_1 = 0$ tests whether individuals who were referred to the job by an incentivised[35] referee j perform better on the job than those invited randomly, while the test $\rho_1 = \rho_2$ tests whether workers referred under incentives perform differently than workers referred without incentives for the referee. Both referral indicators are randomly assigned to the referring worker, and hence allow unbiased estimation of the referral contract on work performance. I interpret the performance in the effort-demanding work task as a proxy for overall ability. In Beaman and Magruder (2012), the comparable test to $\rho_1 = \rho_2$ is rejected, with $\rho_1 > \rho_2$.

I further use cross-sectional regressions to look at the impact of the openness or anonymity of referral treatment on the centrality of selected workers. Those regressions are introduced in the respective paragraphs of section 6.4.

[35]The incentive payment means matching the performance payment that i receives for the person j who made the referral.

5.2. Dyadic regressions

I run dyadic least squares regressions to learn about persistent differences in labour market outcomes and productivity as well as exclusion from referral networks. The dyadic shape of my data helps to characterise bilateral relationships between two individuals and trace those over the repeated work days. All dyadic regressions have their standard errors clustered at the dyad level, i.e. at the level ij.[36] In particular, I employ dyadic regression to test the implications of proposition 2, namely a lower prevalence of reciprocal referrals under the incentivised referral treatment condition.

$$\begin{aligned} R_{ij2} =& \beta_0 + \beta_1 R_{ji1} + \beta_2 * Referral_{j1}^{random} + \beta 3 * Referral_{j1}^{incentivised} \\ &+ \beta 4 * Referral_{j1}^{incentivised} * R_{ji1} + \theta X_{ij} + u_{ij}, \end{aligned} \tag{9}$$

where R_{ij2} is an indicator variable for whether individual i invited person j in session two (and R_{ji1} for whether individual j invited person i in session one). $\beta_1 > \beta_2$ tests whether i is more likely to refer back to j than a randomly selected member from the neighborhood (reciprocity). $\beta_4 = 0$ tests whether the prevalence of reciprocity is the same under incentivised referrals (proposition 2 suggests it is lower).

6. Empirical results

In this and the following section, I present the results of the work experiment. First, I start by showing the data on the open control referral (C). Second, I evaluate the impact of the incentivisation treatment (T1), by testing whether incentivised referrals lead to higher work productivity (proposition 1). In the same step, I introduce the dynamic structure of the experiment to test the effect of reciprocal referring on productivity and the overall impact of learning in the workshop (proposition 2). Third, I test propositions 3 and 4 by exploiting the variation in the openness of the referral (open / partly anonymous (T2) / anonymous (T3)). In a final step, I look at treatment effects on the network structure, focussing particularly on how peripheral and isolated workers strategically use the referral opportunity, and what the impacts of shocking the centrality signal are.[37]

[36] As a robustness check, I also run all dyadic regressions as a multinomial logit (following de Rooij and Kroonenberg (2003)). All findings based on multinomial logit are virtually identical to the standard least-squares results, in terms of size, sign and significance of the coefficients, and can be obtained on request.

[37] I look at heterogeneous effects of the incentive treatment by the degree of the invited and the inviting worker in appendix section E.2.

6.1. The control condition: open, unincentivised referrals

Figure 4 displays the cumulative distribution of the workers' performance over the sessions of the day work: the different rounds indicate the performance of all individuals participating in the task for the first, second, or third time, respectively. With every additional work session they participate in, the workers get better at completing the task without errors. The median payment (calculated as performance score $\times 10$) for an individual's first work stint is 120 ETB, and increases to 150 ETB (170 ETB) in stint two (stint three). The cumulative distribution function for participating in the day work for the third (second) time stochastically dominates the cumulative distribution function for participating for the second (first) time. I can strongly reject equality of the distributions with a two-sample Kolmogorov-Smirnov test; the *p*-value for all three comparisons is smaller than 0.01. The improvement in work performance shown in this graph can be explained by three different factors: on-the-job learning, selection into the work, and incentivised selection into the work.[38] The fact that the type of work completed allows at least some kind of within-worker improvement over time through 'learning on the job' needs to be taken into account when looking at the impact of referral type on job performance, as some workers will have already been on the job in a previous session.

[Figure 4 about here.]

6.2. Variation in the monetary pay-off of the referral (T1)

I test proposition 1, namely that an increase in the incentive payment P_i will lead to a selection of better-performing individuals, with regression 8. Due to the piece-rate payments the individuals invited to the work by the previous session's worker all face private incentives to do well in the data entry job. This is true for all individuals, regardless of the 'referral contract', i.e. the treatment type that the inviting worker received. Since the referral treatment status under T1 (incentivised referral vs non-incentivised referral) is randomly allocated to the previous session's workers, all differences in their referrals performance can be attributed to the inviting worker's extra incentive to invite particularly well-performing individuals.[39]

[38] Disentangling these three mechanisms will be left for later versions of the paper.

[39] While Beaman and Magruder (2012) look at two different types of incentive payments called "Low-stakes performance pay" and "High-stakes performance pay", I only have one payment category. Since the performance payments of the original participant in my work sessions is matched to be 100% of the referral's performance payment, my incentive payment could reasonably be characterised as "high". Another difference is that Beaman and Magruder (2012) do not pay their experimental subjects according to work performance, which leads to overall small incentives to perform well during the task. In my design, workers are paid by their performance and arguably is more realistic

Similar to Beaman and Magruder (2012), I find that incentivising the previous session's worker to bring in a more productive individual to the work sessions leads to an improved performance of the referral (table 5). In qualitative terms, referrals who were invited by an incentivised worker score approximately 2.5 points higher on the performance schedule than referrals invited by workers invited randomly. In relative terms, this means an improvement of approximately 18% compared to the control referral performance (C). Importantly, this finding holds across different specifications.[40] In addition, workers invited without referee incentives do not perform any better than workers invited randomly. As a consequence, I can reject the hypothesis that workers invited under incentives perform as well as workers referred without incentives in my preferred specification ($p = 0.03$).

Since some workers have already worked in the first round (and the results are based on performances in the subsequent sessions), the positive effect on incentivised referrals could partly be driven by on-the-job learning, if the incentivised referral disproportionately made workers refer individuals who had already worked in the first round (or if learning is asymmetric between the workers invited under different treatments). Even though there is no statistical difference between the incentivised and the un-incentivised referral on whether the invited worker had already participated previously, I can also conduct a more thorough test of learning. In order to separate out the learning effect from the ability effect, I impute the workers first-round productivity, i.e. the performance on the job when they worked the first time. Column (4) of table 5 shows that the effect of the incentivised referral is even larger, indicating that on-the-job learning if anything works against pure ability: on-the-job learning of workers invited without incentives decreases the productivity gap by a third over time.

These primary findings go beyond Beaman and Magruder (2012)'s and demonstrate that workers can be properly incentivised to screen more productive individuals for a job — both in general (the main result in Beaman and Magruder, 2012), and when doing so in their local network. This is the case even when their referral's performance is independently incentivised. This finding presents a higher threshold than screening more productive individuals for an un-incentivised work task, and minimised the impact of side payments among worker and referral, as referrals have sufficient private incentives to perform well.

[Table 5 about here.]

to real-life day-work markets.

[40] I can rule out that my results are caused by workers who were incentivised in previous rounds telling their referrals about how to do the work best, since a) the work materials are kept in the session rooms, and b) the content of the task varies over the rounds.

6.3. Exploiting the referral panel: repeated interactions

6.3.1. Reciprocity and monetary incentives

Incentivised referrals contracts lead to better worker performance, compared both to workers referred without referee incentives and workers chosen for the job at random. I now exploit the multiple rounds of my experiment to look at dynamic referral contracts, a novel contribution that departs from Beaman and Magruder (2012)'s work. In particular, the dynamic structure of the experiment allows me to look at one specific component of the social payoff σ_{ij}, namely reciprocal referrals ρ_{ji}. Proposition 2 states that an increase in the performance payment P_i should lead to lower social payoffs σ_{ij}. I therefore expect to see a lower prevalence of reciprocal referrals under the incentivised referral treatment condition.

The empirical results of specification 9 (and variants) are presented in table 6. Both specifications regress an indicator variable for whether worker i invited person j after the second work day (i.e. to the third work session) on whether j had previously referred i to the work. In the first panel, I present the results of "simplistic" version of specification 9, which is run on the partially selected sample of workers present in the second session. I find that a prior work referral from j to i is a very strong predictor of whether subsequently i re-invites j — it increases the chances of the latter by a massive 29 percentage points. However, in order to causally interpret these coefficients, I need to compare this coefficient to the chances of being invited by a randomly selected neighbourhood worker i (a subsample of workers is randomly invited by a computer to the second work session). The right panel of table 6 shows that after including the dummies for being referred by a random worker, the results on the reciprocity indicator Referral^1_{ji} remain virtually unchanged. In fact, the coefficient comparison of Referral^1_{ji} to the likelihood of being invited by a random individual is statistically highly different throughout. This suggests that, compared to a random member at the day work sessions, workers are much more likely to refer a person to the job who has previously invited them.

This finding points toward strong reciprocity at the work place in this context, and stands in a marked contrast to the finding by Conte et al. (2009) from computer-based lab-referral games played with European students — only a minority of players opted for a reciprocal referral strategy in that context. In the local labour markets of this paper, where individuals from real social networks provide each other with work opportunities, reciprocity in job referrals accounts for a large share of the strategic network interactions.

[Table 6 about here.]

The third and rightmost panel of regressions of table 6 shows that the norm of reciprocal referring is driven by non-incentivised referrals. For incentivised referrals after round one, the effect on reciprocity is still positive, but smaller than for the un-incentivised treatment. In particular, when interacting the round referral$^1_{ji}$ with whether j was allocated a random incentivised treatment, the effect on reciprocity is about a fifth smaller than for the case of non-incentivised treatments. This is evidence in line with proposition 2: while a round one referral from j to i is still an important predictor of a round two referral from i to j, the prevalence of such direct reciprocity is 20% lower under the incentivised treatment. This supports the idea that an increase in the performance payment P_i leads to a decrease in the reciprocal referrals ρ_{ji} element of the social payoff σ_{ij}.

6.3.2. Reciprocity and the implications for productivity

Is the strong norm of reciprocal referrals to the day jobs is detrimental for productivity? A related literature shows that in-group biases and ethnic differences can prevent an efficient allocation of workers to jobs.[41] Indeed, in my experiment, a large share of the reciprocity is driven by co-ethnic workers referring each other to the job, raising the question of how this affects output. Moreover, proposition 2 demonstrates the trade-off between social payoff σ_{ij} and the selection of referees. An increase in P_i leads to a selection of more productive workers as well as a decrease in the prevalence of reciprocal referral arrangements. In line with proposition 2, can I detect a relationship between reciprocal referrals, i.e. referrals with a high social payoff σ_{ij}, and productivity?

In table 7, I show the results of regressing work performance scores of the third (and last) work session on whether the worker i has been invited reciprocally, randomly (i.e. by the computer), or differently (e.g. by another worker, but not in a reciprocal fashion). I find that compared to "normally" invited workers (i.e. those workers who were invited by someone working in a previous session, but not as a reciprocal act), both randomly and reciprocally invited workers perform worse. Since "normally" invited workers might be selected positively on ability, the true comparison of interest here is whether reciprocally invited workers perform worse than randomly invited ones.[42] Across specifications, I find that reciprocal referring comes at a substantial cost.

[41] Alesina and Ferrara (2005), Anderson (2011), Hjort (2014), Burgess et al. (2015), Fisman et al. (2017); while others argue that referrals allow firms to mitigate moral hazard on the workers' side (Heath, 2018).

[42] Reciprocally invited workers have, by definition, worked in at least one previous session, and are thus familiar with the work task. This could mean that they perform automatically better, and figure 4 demonstrates that learning on the job indeed takes place. Alternatively, it is also thinkable that workers who do a similar job twice may get bored and perform worse, even though this should prevented by the performance-based payment and is consequently not supported by the data shown in figure 4. Regardless I control for both directions of repeated work by imputing the first-session

On average, reciprocally invited workers perform 0.5 performance points worse than randomly invited ones, which is equivalent to a productivity decrease of 5% over a baseline control mean score of 10.4. The difference between reciprocal and random worker performance is statistically significant across specifications. Compared to non-random "normally" referred workers, which is a more natural comparison in a setting of real labour markets (where workers are not invited randomly), the decrease in performance is even more substantial, with a decrease of 1.25 points (11%) over a control mean of 11 performance points.

Importantly, the result that reciprocal links are less productive still holds when controlling for whether a referral was incentivised (last column), and the difference between reciprocally and not-reciprocally invited individuals under incentivised referrals is also negative and significant. In terms of my model, this is the ingredient of the social payoff σ_{ij} that makes individuals select lower-performing individuals compared to the direct monetary incentive payoff: the immediate reciprocation of referrals, ρ_{ji}. Overall, the findings suggests that strong norms of reciprocal job referring leads to important productivity losses in these local labour markets, because the reciprocally referred worker often does not seem to be the most productive one.

[Table 7 about here.]

6.4. Variation in the social pay-off of the referral (T2 & T3)

This subsection describes the impact of changing the visibility of the job referrals induced by treatments 2 and 3. In a first step, I describe the correlation between measures of network centrality and the likelihood of being referred to the job under the open referral (C). Second, I test how varying the importance of κ_{ji} and ρ_{ji} in the network benefits function impacts the centrality c_j of the referral (propositions 3 and 4).

6.4.1. Open referrals: do individuals refer more central individuals from their networks to the job?

In the upper panel of appendix table E5, I regress a binary indicator of whether individual i referred individual j to the job in open referrals (C) on the baseline number of connections of individual j, her degree centrality in her neighbourhood network, and on a variety of further network centrality measures. The regressions are run on the cross-sectional data (i.e. one observation per individual). Importantly, the coefficients

performance of all workers in the sample, which directly captures individual work ability.

in these regressions cannot be interpreted in a causal sense, since network centrality correlates with a host of other characteristics of individuals.

When the different network centrality measures enter the regressions separately, both out-degree and in-degree centrality are significant predictors of whether an individual receives a referral to a job, but the coefficient for in-degree centrality is 20% larger. An increase of j's in-degree centrality by one person (i.e. being known by one more individual in the neighbourhood network) would increase her chances of being referred to the job by 1.2 percentage points, or approximately 4%. When including both in- and out-degree as regressors, the in-degree centrality remains significant and more than twice as large as the out-degree. In the regression that includes all network centrality measures simultaneously (last column), a similar picture emerges: in-degree centrality is the relevant measure predicting the chances of being referred to a job, i.e. the number of individuals claiming to know you (rather than you claim to know). This makes intuitive sense: since the individuals j are referred to the job (by individuals i), it matters how many people the referral j is known by — which is j's in-degree centrality. It is less important how many other people in the neighbourhood j claims to know — j's out-degree centrality — because the inviter i is making the referral. It is noteworthy that the coefficients on eigenvector centrality and clustering are both economically meaningful, but subject to large variation and hence not statistically significant. This is partly the case because eigenvector centrality is only defined for single-component networks, which only holds for six out of my 16 networks.

I repeat the same analysis for job network centrality in the bottom panel of appendix table E5. The findings are very similar to those for general networks and discussed in appendix section E.3.

6.4.2. Why do workers refer more central nodes? Varying the visibility of the referral

A worker's in-degree centrality both in general and job networks positively correlates with her likelihood of being referred to the job. At the same time, I find only limited evidence that central workers perform better on the job, even under incentivised referrals. This rules out the idea that inviting a central person to the job increases the direct monetary pay-off for the referee through the incentivised treatment or side payments.

In this subsection, I present three pieces of evidence showing that more central workers are more likely to get job referrals in their neighbourhood networks due to the importance of potential future network benefits κ_{ji} compared to immediate reciprocity concerns ρ_{ji}. First, I demonstrate that when workers refer individuals under partial anonymity, they invite less central individuals compared to the open referrals. Second, I show that this tendency increases when all potential network benefits are switched off under full anonymity. Lastly, I present short-run follow-up data collected a few days after the work experiment, where workers name reasons for why they made the referrals. I show that self-stated other-regarding preferences increase with the amount of anonymity in the treatment.

In a first piece of evidence, I test proposition 3, stating that an increase of ρ_{ij} compared to κ_{ji} in the network benefits b_{ji} will lead to the selection of referrals with lower centrality. To do this, I use a sub-treatment exploiting within-subject variation. I compare the referrals workers make in their own neighbourhood to those made partly or fully anonymously to other neighbourhoods, where I control which information are made available to the participant about personal characteristics of the unknown referrals.[43] The partly anonymous referral treatment in other neighbourhoods explicitly encourages reciprocity ρ_{ij}, but is unlikely to generate any lasting future exchanges κ_{ji}. In a second step, I also test proposition 4, which posits that when network and monetary benefits are switched off, workers will select referrals with the lowest centrality. The fully anonymous referral treatment in other neighbourhoods completely rules out reciprocity, further benefits or monetary payments ($\rho_{ij} = \kappa_{ji} = \pi_{ji} = 0$), so that the social payoff reduces to the inclusion of j's income in i's social payoff: $\sigma_{ij} = \omega_i m_j$. Both of these tests are conducted with the help of the following regression:

$$\begin{aligned} Centrality_j = &\alpha_0 + \alpha_1 \cdot Referral_i^{Own} + \alpha_2 \cdot Referral_i^{Other,p.anon.} \\ &+ \alpha_3 \cdot Referral_i^{Other,anon} + \theta X_j + u_{ij}, \end{aligned} \tag{10}$$

where $Referral_i^{Own}$ is a dummy for the referral in i's own neighbourhood, $Referral_i^{Other,p.anon.}$ indicates partly anonymous referrals in other neighbourhoods, and $Referral_i^{Other,anon}$ stands for anonymous referrals in other neighbourhoods. Propositions 3 and 4 suggest that $\alpha_1 > \alpha_2 > \alpha_3$.[44]

Table 8 displays the results. When comparing the job referrals that individuals made to other workers in their own social network with the job referrals these same individuals made towards (unknown) individuals from other, similar neighbourhoods, I find contrasting patterns. In the own neighbourhood networks, referred workers are positively selected on centrality, whereas the opposite is the case in other neighbourhoods. The difference between partly anonymous referrals in other neighbourhoods (enabling reciprocity) and anonymous referrals in other neighbourhoods is negative and statisti-

[43]The partly anonymous referral encourages the referred worker to reciprocate. During the anonymous referral, the invited worker does not know anything about who provided the job referral.

[44]The comparison of α_1, α_2, and α_3 is based on experimental variation in referral type, whereas the comparison of each coefficient to the omitted category (non-selected workers) is based on selection via referral.

cally significant.

Why do workers invite significantly less central individuals when referring out of their neighbourhood and more central individuals in their own network? Proposition 3 suggests that for individuals in other neighbourhoods, reciprocity is more likely to be expected from individuals with fewer connections. In the case of fully anonymous referrals, network and monetary benefits can be ruled out, so that only other-regarding preferences (the inclusion of j's income in i's utility) remain. As a consequence and in line with proposition 4, individuals with the fewest connections are referred.

[Table 8 about here.]

To sum up, this subsection presents further evidence on why the workers refer more central individuals to the job, drawing from follow-up data collected a few days after the work sessions. First, I show descriptive evidence of the expectations workers have about central individuals' earnings and job information. I then analyse the self-reported reasons why workers made a specific referral.

The left panel of figure 5 shows that workers overwhelmingly think that individuals in their neighbourhood with more network connections are more likely to have higher monthly earnings. The cumulative distributions for higher number of connections monotonically dominate the lower ones. The right panel of the same figure then shows that the respondents expect that individuals from their neighbourhood with more connections also have access to more information about job opportunities. While the increase is monotonic, it does not seem to be linear, with a large increase in the information ranking happening between five and ten connections.

[Figure 5 about here.]

Finally, figure 6 displays the reasons workers mention in the follow-up interview about why they made a referral to a specific person. The left panel includes more general answer categories. In line with what I find in the previous sections, workers are more likely to claim they referred a more central person to the job in the main referral to individuals from their own neighbourhood, compared to the referrals made to individuals from other neighbourhoods. The reasons given by the respondent for referrals made to workers in other neighbourhoods do not differ significantly by the type of referral (partly or fully anonymous). Both in the partly and fully anonymous referral type, workers claim to have been guided a lot by the other person's gender or age. In the own neighbourhood, many respondents claimed to have been mostly guided by whether another workers was a previous connection.

The right panel of figure 6 zooms in on the answers that specifically mentioned the referrals centrality. Of those people who claimed to have made a referral to a more central person, most state that the main reason for that are the job information central people have, followed by the idea that central people would make easy social connections. Of the workers claiming they referred a less central individual, more than half state pity as the main reason, in particular in the other neighbourhood treatments, followed by the idea that it would be easy to connect to those individuals.

These self-reported follow-up findings are in line with the theoretical reasons presented in section 2: connecting to central people in the own network dominates connecting to peripheral nodes due to the long-run information benefits (κ_{ji}) associated with central workers, whereas in other networks immediate reciprocity concerns (ρ_{ji}) or other-regarding preferences ($\omega_i m_j$) dominate.

[Figure 6 about here.]

7. Discussion

In this section, I discuss the implications of my previous findings on equality in local labour markets. I have shown in the previous section that open, un-incentivised referrals – the ones most likely encountered in the field – lead to the selection of more central individuals to jobs and are characterised by a large degree of reciprocity. Both of these factors lead to the exclusion of peripheral individuals in the local network from jobs. In this section, I focus on these individuals without many network links to understand how they could be included in the labour market.

I provide three pieces of evidence describing the ways in which peripheral workers can or cannot get access to referral networks. First, by looking at the effect of the centrality shrouding treatment (T4), I show that connectedness in these networks is not simply described by individual position (centrality) in the network, but that correlated characteristics of central individuals seem to matter. Second, I demonstrate that peripheral individuals are more likely to form new relationships through the work experiment. Third and lastly, I provide evidence that these new referral links are more likely to be reciprocated by the other worker (compared to connections that already existed at baseline) and that a substantial proportion of new connections are still in place one and a half years after the experiment.

7.1. Degree centrality vs correlated confounders – shrouding treatment

In order to test whether underlying characteristics of central workers determine their job prospects rather than their pure network centrality, I introduce the sub-treatment that exploits the exogenous experimental variation in whether individual j's degree centrality was made salient to the workers or not (i.e. the centrality shrouding sub-treatment T4).[45]

Table 9 shows the results. I find that those workers j who were invited through job referrals that made their network centrality salient (revealed degree centrality on the referral lists) have on average three more connections than the average network member. Workers j who were invited with concealed job referrals (i.e. their degree centrality was not a salient characteristic during the referral process) are also more central than the average network member, with on average two more connections (which is statistically not significant). The comparisons to the average network member are based on non-random variation, as both types of referrals induce selection. On the other hand, the comparison of the coefficients on C and T4 is identified through experimental variation. The difference of one connection on average between the two treatments may bear some economic significance, but I cannot reject that the coefficients for the two different treatments (revealed vs shrouded degree centrality) are equal ($p = 0.695$); in fact, they seem reasonably close to each other. This supports an interpretation where network centrality bears no additional meaning beyond the individual characteristics that are associated with it. The additional value individuals receive from referring more central individuals to the work sessions can be almost fully explained by their innate characteristics, rather than by their pure position in the network alone.

[Table 9 about here.]

7.2. Can peripheral individuals establish new links or is exclusion persistent?

My previous findings indicate a significantly positive correlation between an individual's network centrality and their likelihood of being referred to the job. This effect is driven by an individual's in-degree centrality, i.e. by how many other nodes in the network she is known, not by how many nodes she knows. This suggests that the long-run benefits of being connected to person j, κ_{ji}, dominate the network benefits b_{ij} compared

[45] I pool the experimental data to run the following regression:

$$Centrality_j = \alpha_0 + \alpha_1 \cdot Referral_i^{Open} + \alpha_2 \cdot Referral_i^{Shrouded} + \theta X_{ij} + u_j, \quad (11)$$

where $Referral_j^{Shrouded}$ indicates a dummy for the shrouded degree centrality treatment. I test the hypothesis $\alpha_1 = \alpha_2$ to detect if degree centrality has a separate effect beyond correlated confounders.

to the short-run reciprocity concerns ρ_{ji}. At the same time, I find very strong and large results for reciprocity in the day labour market even for central individuals.

These two findings both point to a disadvantage for individuals who are not highly central in their local social network and who do not have the opportunity to provide jobs or favours themselves. In fact, in all of the previous regressions I controlled for whether two individuals knew each other at baseline (i.e. were linked prior to the experiment) — and in all cases, baseline linkage is a very strong predictor of a referral from one worker to another. In conjunction with the visual display of the baseline networks in figure 3, the evidence overwhelmingly draws a rather negative outlook for individuals without many connections in their local job network — in particular, when in-degree centrality is the driver here. An individual's in-degree centrality is something individuals are less likely to be aware of and in control of than out-degree centrality (which is reported by themselves).

In this next step, I test if peripheral individuals can use the intervention to generate new links when (randomly) given the opportunity to distribute jobs in their network. I run the following regression:

$$R_{ij}^{new} = \zeta_0 + \zeta_1 Referral_i + \zeta_2 Referral_i \cdot PERIPHERAL_i + \theta X_{ij} + u_{ij} \tag{12}$$

where R_{ij}^{new} is an indicator variable for a newly established link (via referral) between individuals i and j (who were thus not linked at baseline). $Referral_i \cdot PERIPHERAL_i$ is an interaction between the exogenous referral treatment and whether the individual i is at the periphery of her local network. This regression can be run on the pooled data from the different experiment rounds. The test $\zeta_1 = \zeta_2$ shows whether peripheral individuals i are more likely to establish new links with other people from their neighbourhood than non-peripheral people.[46] The first panel of table 10 presents the empirical findings.

The results show that treated individuals (i.e. those making a referral in the first work session) from the periphery are significantly more likely to establish a connection to individuals not known at baseline, by approximately 1.5 percentage points. The coefficient is small, but represents a sizeable increase by one and a half times compared to non-peripheral treated workers. Establishing new connections is rare, but it happens more than twice as often for people at the periphery of the general networks. It is important to remember the potentially high costs involved when approaching unknown

[46] I run this regression for different measures of being at the network periphery. I start by defining as peripheral those individuals with less than two connections in the baseline network. My findings are robust to choosing different thresholds, such as less than one, three or four connections, see figure 7.

individuals and presenting them with an invitation to a work opportunity. Given these qualification, the coefficient should arguably display a lower bound and be indicative of the fact that peripheral individuals are more likely to make new connections rather than relying on their (few) existing ones.

[Table 10 about here.]

The second panel of table 10 displays the same regression for the job information networks, i.e. a link between two individuals is only established when they have shared jobs or job information in the past (cf. figure A5). Consequently, the share of peripheral individuals becomes considerably higher. Regardless, we still find a significant effect of 1.2 percentage points for treated individuals, which is roughly a doubling of the coefficient for treated, but non-peripheral individuals. These results indicate that individuals with few or no connections in both the overall social network as well as the job exchange network employ the referral opportunity to connect with new individuals. They use the 'windfall' invitation to a real and well-paid day job for a strategic expansion of their social network.

In figure 7, I show that both previous results are not driven by a mechanical relationship of how I define periphery. Even for higher thresholds of what constitutes being at the periphery of a network, the positive and statistically significant interaction effect between referral treatment and peripheral status remains: these workers are significantly more likely to establish a new link to other workers from their neighbourhood. For the narrower definition of job exchange networks, the effects are even more balanced across similar levels of periphery, albeit on a lower level overall.

[Figure 7 about here.]

Now, with peripheral individuals trying to expand their network by referring individuals previously not in their network, are they successful in consolidating these new links? That is, once they have provided a new connection with a day job opportunity, do those individuals pay back the favour through reciprocal referring, as previously found for the whole sample (table 6)? Table 11 shows the results for general and job exchange networks. I find that for both types of networks, peripheral individuals that invite a new connection to the day job are not less likely to be reciprocated than non-peripheral individuals. In overall social networks (upper panel), peripheral individuals that use the referral to establish a link to a previously entirely unknown person are more likely to be re-invited by that new connection to the job during the next session. This is both in comparison to a randomly invited worker, but also in comparison to peripheral individuals that did not establish a new link in the previous session (but rather invite one of their few baseline connections). So in the case of overall networks, peripheral peers that make new connections not only manage to consolidate that new link by being re-invited, they are actually better off than similar peripheral nodes that do not refer an unknown individual after the first work session. The situation is slightly different for the job networks (in table 11, bottom panel). Again, peripheral individuals that establish a new job connection with their referral are more likely to be re-invited by that new job contact than a randomly invited worker. This is the lower bound of the effect, stating that the peripheral individuals are not worse off by reaching out to new connections. When comparing peripheral individuals who refer their existing job contacts compared to peripherals who connect to new job contacts, I cannot reject the null hypothesis that both coefficients on the likelihood of being re-invited in round two are the same. In other words, once a peripheral individual receives a 'windfall' job referral opportunity and uses it to connect to new parts of the job exchange network, she is not worse off compared to the situation had she given the job to existing contacts. In both cases, she is very likely to be reciprocated for the job referral. However, given the additional value of making a new job connection, reaching out and inviting non-contacts to the job is likely to be a dominant strategy for the peripheral individual.

[Table 11 about here.]

7.3. The long-run effects of the day job intervention

In the previous section, I have established that workers who are peripheral in their social network are more likely to use the treatment to establish new connections in the neighbourhood, and that these new connections also get reciprocated immediately (ρ_{ji}). However, it is unclear whether these new connections also manage to generate long-run network benefits for the workers who are peripheral at baseline (κ_{ji}), or whether the new links break after a certain time.

In order to look at the long-run impacts, I re-interview the respondents 18 months after the work experiment and ask about the links that were newly created in the work experiment. First, I identified a total of 84 individuals who made a referral to a previously unknown person in the experiment. Of those 84, 75 could be reached, which equals a low attrition rate of 10.5% after 18 months. Of those respondents who could not be reached, six migrated out of Ethiopia and the rest did not have any functioning telephone number. One respondent refused participation in the long-run follow-up. The left panel of figure 8 presents the 18 months follow-up data on new connections: of those 74 individuals who made a referral to a previously unknown person and could be reached during the follow-up, 56 (75%) still know the person, 31 (42%) are still in contact with that person, and 18 (24%) are still in a job-exchange relationship with that person, with 18 (24%) of these people giving information and 11 (15%) receiving information about jobs. These numbers indicate that a low-touch and short intervention can generate new and lasting relationships between individuals.

Are these persisting relationships simply indicative of secular changes in connectivity over time? I have two pieces of evidence that suggest differently: first, the right panel of figure 8 compares those individuals who made one new connection in the experiment with those who made two. In line with the idea that these new relationships got formed during the experiment and do not represent secular network change, the individuals who made two new connections are still in contact with a 1.5 times larger proportion of them than the individuals who established only one new link. This holds across most types of interactions. Second, my baseline data suggests that for every month a person has lived in the neighbourhood, she makes 0.012 new connections. Over 18 months, the secular change would thus be 0.216 new connections. My long-run follow-up data on new connections display a number at least twice as large (42% in contact), suggesting that these are in fact new connections generated through the intervention.

[Figure 8 about here.]

Figure 9 displays the number of newly established links that were still in some form of relation 18 months after the intervention, by whether the worker was peripheral at baseline for different periphery threshold. The top figures defines periphery based on the worker's in-degree centrality, the bottom ones based on out-degree. The left figures look at general networks, the right at job networks. Two things become apparent in the graphs. First, peripheral workers have lower social interactions across all categories than non-peripheral workers, regardless of the threshold chosen to define peripheral individuals. Second, while this finding holds across all panels it is stronger for in-degree than for out-degree, and for general networks than for job networks. In particular, peripheral workers in job networks often manage to receive job information to the same extent as non-peripheral workers. They also receive information to the same extent than they give, which is not true for non-peripheral workers.

[Figure 9 about here.]

8. Conclusion

This experimental study investigates the relationship between network centrality, social exclusion and labour market outcomes in the context of urban neighbourhoods in Addis Ababa. I experimentally create day jobs for young individuals without permanent employment and living in densely populated urban neighbourhoods. Upon work completion, some of the young workers get the opportunity to invite their social contacts to the same job. Importantly, this real-life job referral opportunity takes place within the existing social networks of the workers' residential neighbourhoods. The design of this 'job referral treatment' is varied randomly along two key dimensions: 1) incentivised vs un-incentivised referrals and 2) in-person vs gradually anonymised referrals. Repeated work days and referrals enable me to investigate dynamic interaction mechanisms.

This paper makes several contributions to the field. First, I find that the real-life job referrals are characterised by a substantial degree of reciprocity, which is significantly reduced under the incentivised referral treatment. This finding, together with the fact that reciprocally invited workers are less productive, demonstrates the trade-off between performance payoffs and social payoffs derived in a conceptual framework. In line with the model, I find that when given the performance incentive, workers are able to screen more productive individuals in the local neighbourhood networks, using salient information about other network members. The prevalence of reciprocity under the open, non-incentivised referral treatment also implies that individuals with even temporary or one-off referral opportunities have much larger chances of receiving casual work in the future.

In addition, I find that the centrality of the referred workers varies substantially with the type of referral treatment. Under the standard open referral treatment, more central individuals are invited to the job. This again means that individuals without many connections are left jobless. Importantly, this behaviour cannot be not explained by central individuals being more able workers: I find no significant correlation between an individual's network centrality and their on-the-job productivity. In contrast, under the referral treatments which are partly or fully anonymous, less central workers are selected – due to the lower importance of network benefits and dominance of other-regarding concerns.

Finally, I show that there is persistence in social network exclusion. Connectedness in neighbourhood networks predicts the likelihood of obtaining jobs in informal local labour markets, and peripheral individuals who are not (randomly) given a day job and a referral opportunity largely remain excluded. This finding could provide an alternative explanation for long-term unemployment and detachment from the labour market, in particular in the informal and local labour markets of many cities in developing countries. On a more positive note, peripheral individuals who are (randomly) given referral opportunities manage to use these as devices to connect to larger parts of the social and job referral networks, by establishing links with previously unknown individuals. These links are usually confirmed by the other node through reciprocal provision of jobs, and hence remain in place after more than one work session. One and a half years after the experiment, more than 40% of newly generated social links are still in contact, demonstrating the potential of a very light-touch intervention to affect connectedness in social networks in the long-run.

This paper suggests that is possible to design markets that enable socially excluded individuals to integrate into networks. Individuals can permanently overcome exclusion once given an opportunity — this finding has implications for how policy makers could think about alleviating youth unemployment, e.g. through the provision of subsidised temporary job opportunities. A very light-touch intervention has long-run effects on connectedness in the local labour market. This conclusion goes beyond the context of cities in the developing world and could similarly apply to rich countries, where economic development is often said to have brought about an erosion of social interactions (Putnam, 2000).

References

Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin, and S. Quinn (2016). Curse of Anonymity or Tyranny of Distance? The Impacts of Job-Search Support in Urban Ethiopia. Working paper 22409, National Bureau of Economic Research.

Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin, S. R. Quinn, and F. J. Shilpi (2017). Job Fairs: Matching Firms and Workers in a Field Experiment in Ethiopia. Policy Research Working Paper Series 8092, The World Bank.

Alesina, A. and E. L. Ferrara (2005). Ethnic Diversity and Economic Performance. *Journal of Economic Literature 43*(3), 762–800.

Anderson, S. (2011). Caste as an Impediment to Trade. *American Economic Journal: Applied Economics 3*(1), 239–263.

Bala, V. and S. Goyal (2000). A Noncooperative Model of Network Formation. *Econometrica 68*(5), 1181–1229.

Bandiera, O., I. Barankay, and I. Rasul (2009). Social Connections and Incentives in the Workplace: Evidence From Personnel Data. *Econometrica 77*(4), 1047–1094.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2014). Gossip: Identifying Central Individuals in a Social Network. Working paper 20422, National Bureau of Economic Research.

Banerjee, A. V. and E. Duflo (2007). The Economic Lives of the Poor. *Journal of Economic Perspectives 21*(1), 141–168.

Barr, A., M. Dekker, and M. Fafchamps (2015). The Formation of Community-Based Organizations: An Analysis of a Quasi-Experiment in Zimbabwe. *World Development 66*, 131–153.

Bayer, P., S. L. Ross, and G. Topa (2008). Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes. *Journal of Political Economy 116*(6), 1150–1196.

Beaman, L., N. Keleher, and J. Magruder (2018). Do Job Networks Disadvantage Women? Evidence from a Recruitment Experiment in Malawi. *Journal of Labor Economics 36*(1), 121–157.

Beaman, L. and J. Magruder (2012). Who Gets the Job Referral? Evidence from a Social Networks Experiment. *The American Economic Review 102*(7), 3574–3593.

Beaman, L. A. (2012). Social Networks and the Dynamics of Labour Market Outcomes: Evidence from Refugees Resettled in the U.S. *The Review of Economic Studies 79*(1), 128–161.

Bolton, G. E. and A. Ockenfels (2000, March). ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review 90*(1), 166–193.

Breza, E. and A. G. Chandrasekhar (2015). Social Networks, Reputation and Commitment: Evidence from a Savings Monitors Experiment. Working paper 21169, National Bureau of Economic Research.

Burgess, R., R. Jedwab, E. Miguel, A. Morjaria, and G. Padró i Miquel (2015). The Value of Democracy: Evidence from Road Building in Kenya. *The American Economic Review 105*(6), 1817–1851.

Cai, J., A. De Janvry, and E. Sadoulet (2015). Social Networks and the Decision to Insure. *American Economic Journal: Applied Economics 7*(2), 81–108.

Cai, J. and A. Szeidl (2017). Interfirm Relationships and Business Performance. *The Quarterly Journal of Economics*, qjx049.

Calvó-Armengol, A. (2004). Job Contact Networks. *Journal of Economic Theory 115*(1), 191–206.

Calvó-Armengol, A. and M. O. Jackson (2004). The Effects of Social Networks on Employment and Inequality. *American Economic Review 94*(3), 426–454.

Caria, A. S. and I. W. Hassen (2013). The Formation of Job Referral Networks: Experimental Evidence from Urban Ethiopia. Discussion paper 1282, International Food Policy Research Institute.

Caria, S., S. Franklin, and M. Witte (2018). Searching with Friends. Working paper, University of Oxford.

Chandrasekhar, A. and R. Lewis (2011). Econometrics of Sampled Networks. *Unpublished manuscript, MIT.[422]*.

Chandrasekhar, A. G., C. Kinnan, and H. Larreguy (2014). Social Networks as Contract Enforcement: Evidence from a Lab Experiment in the Field. Working paper 20259, National Bureau of Economic Research.

Charness, G. and M. Rabin (2002). Understanding Social Preferences with Simple Tests. *The Quarterly Journal of Economics 117*(3), 817–869.

Christakis, N. A. and J. H. Fowler (2007). The Spread of Obesity in a Large Social Network over 32 Years. *New England Journal of Medicine 357*(4), 370–379. PMID: 17652652.

Christakis, N. A. and J. H. Fowler (2009). *Connected: The Surprising Power of Our Social Networks and How They Shape Our Lives*. Little, Brown.

Chuang, Y. and L. Schechter (2015). Social Networks in Developing Countries. *Annual Review of Resource Economics 7*(1), 451–472.

Cohen, B. (2006). Urbanization in Developing Countries: Current Trends, Future Projections, and Key Challenges for Sustainability. *Technology in Society 28*(1), 63–80.

Collins, D., J. Morduch, S. Rutherford, and O. Ruthven (2009). *Portfolios of the Poor: How the World's Poor Live on $2 a Day*. Princeton University Press.

Comola, M. and M. Fafchamps (2014). Testing Unilateral and Bilateral Link Formation. *The Economic Journal 124*(579), 954–976.

Comola, M. and M. Mendola (2015). Formation of Migrant Networks. *The Scandinavian Journal of Economics 117*(2), 592–618.

Comola, M. and S. Prina (2017). Treatment Effects Accounting for Network Changes. Working paper, Case Western Reserve University.

Conte, A., D. D. Cagno, and E. Sciubba (2009). Strategies in Social Network Formation. Birkbeck Working Papers in Economics and Finance 0905, Birkbeck, University of London.

de Rooij, M. and P. M. Kroonenberg (2003). Multivariate multinomial logit models for dyadic sequential interaction data. *Multivariate Behavioral Research 38*(4), 463–504. PMID: 26777443.

Fafchamps, M. and A. Moradi (2015). Referral and Job Performance: Evidence from the Ghana Colonial Army. *Economic Development and Cultural Change 63*(4), 715–751.

Fafchamps, M. and S. Quinn (2016). Networks and Manufacturing Firms in Africa: Results from a Randomized Field Experiment. *The World Bank Economic Review*, lhw057.

Fehr, E. and S. Gächter (2000). Fairness and Retaliation: The Economics of Reciprocity. *Journal of Economic Perspectives 14*(3), 159–181.

Fehr, E. and K. M. Schmidt (1999). A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics 114*(3), 817–868.

Feigenberg, B., E. Field, and R. Pande (2013). The Economic Returns to Social Interaction: Experimental Evidence from Microfinance. *The Review of Economic Studies 80*(4), 1459–1483.

Fisman, R., D. Paravisini, and V. Vig (2017). Cultural Proximity and Loan Outcomes. *The American Economic Review 107*(2), 457–492.

Franklin, S. (2016). Location, Search Costs and Youth Unemployment: Experimental Evidence from Transport Subsidies. *The Economic Journal*.

Gächter, S. and A. Falk (2002). Reputation and Reciprocity: Consequences for the Labour Relation. *Scandinavian Journal of Economics 104*(1), 1–26.

Galenianos, M. (2018). Referral Networks and Inequality. Working paper, Royal Holloway, University of London.

Galeotti, A. and L. P. Merlino (2014). Endogenous Job Contact Networks. *International Economic Review 55*(4), 1201–1226.

Granovetter, M. (1995). *Getting a Job: A Study of Contacts and Careers*. University of Chicago Press.

Heath, R. (2018). Why Do Firms Hire Using Referrals? Evidence from Bangladeshi Garment Factories. *Journal of Political Economy 126*(4), 1691–1746.

Hellerstein, J. K., M. J. Kutzbach, and D. Neumark (2014). Do Labor Market Networks Have an Important Spatial Dimension? *Journal of Urban Economics 79*, 39–58.

Hellerstein, J. K., M. McInerney, and D. Neumark (2011). Neighbors and Coworkers: The Importance of Residential Labor Market Networks. *Journal of Labor Economics 29*(4), 659–695.

Hipp, J. R., R. W. Faris, and A. Boessen (2012). Measuring 'neighborhood': Constructing Network Neighborhoods. *Social Networks 34*(1), 128–140.

Hjort, J. (2014). Ethnic Divisions and Production in Firms. *The Quarterly Journal of Economics 129*(4), 1899–1946.

Imbens, G. W. (2015). Matching Methods in Practice: Three Examples. *Journal of Human Resources 50*(2), 373 – 419.

Ioannides, Y. M. and L. Datcher Loury (2004). Job Information Networks, Neighborhood Effects, and Inequality. *Journal of Economic Literature 42*(4), 1056–1093.

Kramarz, F. and O. N. Skans (2014). When Strong Ties Are Strong: Networks and Youth Labour Market Entry. *The Review of Economic Studies 81*(3), 1164–1200.

Kugler, A. D. (2003). Employee Referrals and Efficiency Wages. *Labour Economics 10*(5), 531–556.

Magruder, J. R. (2010). Intergenerational Networks, Unemployment, and Persistent Inequality in South Africa. *American Economic Journal: Applied Economics 2*(1), 62–85.

Montgomery, J. D. (1991). Social Networks and Labor-Market Outcomes: Toward an Economic Analysis. *The American Economic Review 81*(5), 1408–1418.

Munshi, K. (2003). Networks in the Modern Economy: Mexican Migrants in the U. S. Labor Market. *The Quarterly Journal of Economics 118*(2), 549–599.

Munshi, K. and M. Rosenzweig (2006). Traditional Institutions Meet the Modern World: Caste, Gender, and Schooling Choice in a Globalizing Economy. *The American Economic Review 96*(4), 1225–1252.

Putnam, R. D. (2000). Bowling Alone: Americas Declining Social Capital. In *Culture and politics*, pp. 223–234. Springer.

Rigol, N., R. Hussam, and B. Roth (2017). Targeting High Ability Entrepreneurs using Community Information: Mechanism Design in the Field. Working paper, Harvard University.

Rothstein, J. (2018, April). Inequality of Educational Opportunity? Schools as Mediators of the Intergenerational Transmission of Income. Working Paper 24537, National Bureau of Economic Research.

Sacerdote, B. (2001). Peer Effects with Random Assignment: Results for Dartmouth Roommates. *The Quarterly Journal of Economics 116*(2), 681–704.

Schmutte, I. M. (2015). Job Referral Networks and the Determination of Earnings in Local Labor Markets. *Journal of Labor Economics 33*(1), 1–32.

Serneels, P. (2007). The Nature of Unemployment among Young Men in Urban Ethiopia. *Review of Development Economics 11*(1), 170–186.

Topa, G. (2001). Social Interactions, Local Spillovers and Unemployment. *The Review of Economic Studies 68*(2), 261–295.

Topa, G. (2011). Labor Markets and Referrals. Volume 1 of *Handbook of Social Economics*, pp. 1193–1221. Elsevier North Holland Amsterdam.

Vasilaky, K. N. and K. Leonard (2018). As Good as the Networks They Keep?: Improving Outcomes through Weak Ties in Rural Uganda. *Economic Development and Cultural Change*.

Wahba, J. and Y. Zenou (2005). Density, Social Networks and Job Search Methods: Theory and Application to Egypt. *Journal of Development Economics 78*(2), 443–473.

Watts, D. J. and S. H. Strogatz (1998). Collective Dynamics of Small-World Networks. *Nature 393*(6684), 440–442.

Witte, M. (2017). Urban Job Networks along a Spatial Dimension: Local Neighbourhoods Vs City-Wide Contacts. Working paper, University of Oxford.

The World Bank (2014). *World Development Indicators 2014*. Washington, DC: World Bank Publications.

The World Bank (2016). *World Development Indicators 2016*. Washington, DC: World Bank Publications.

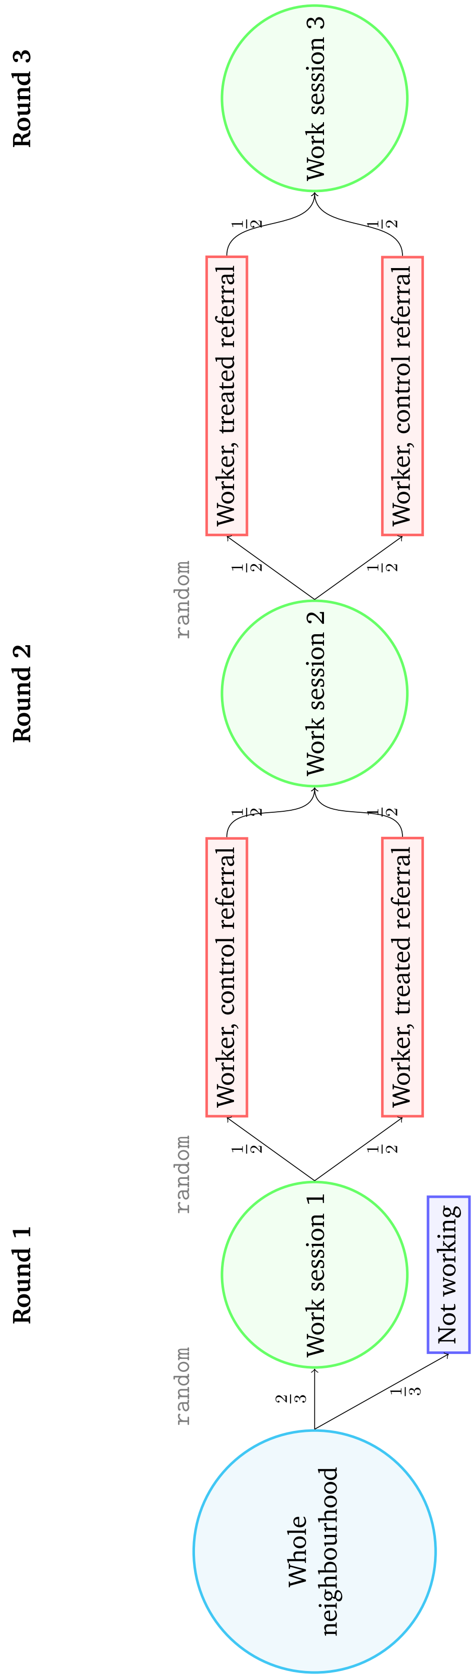

Figure 1: Main experiment structure

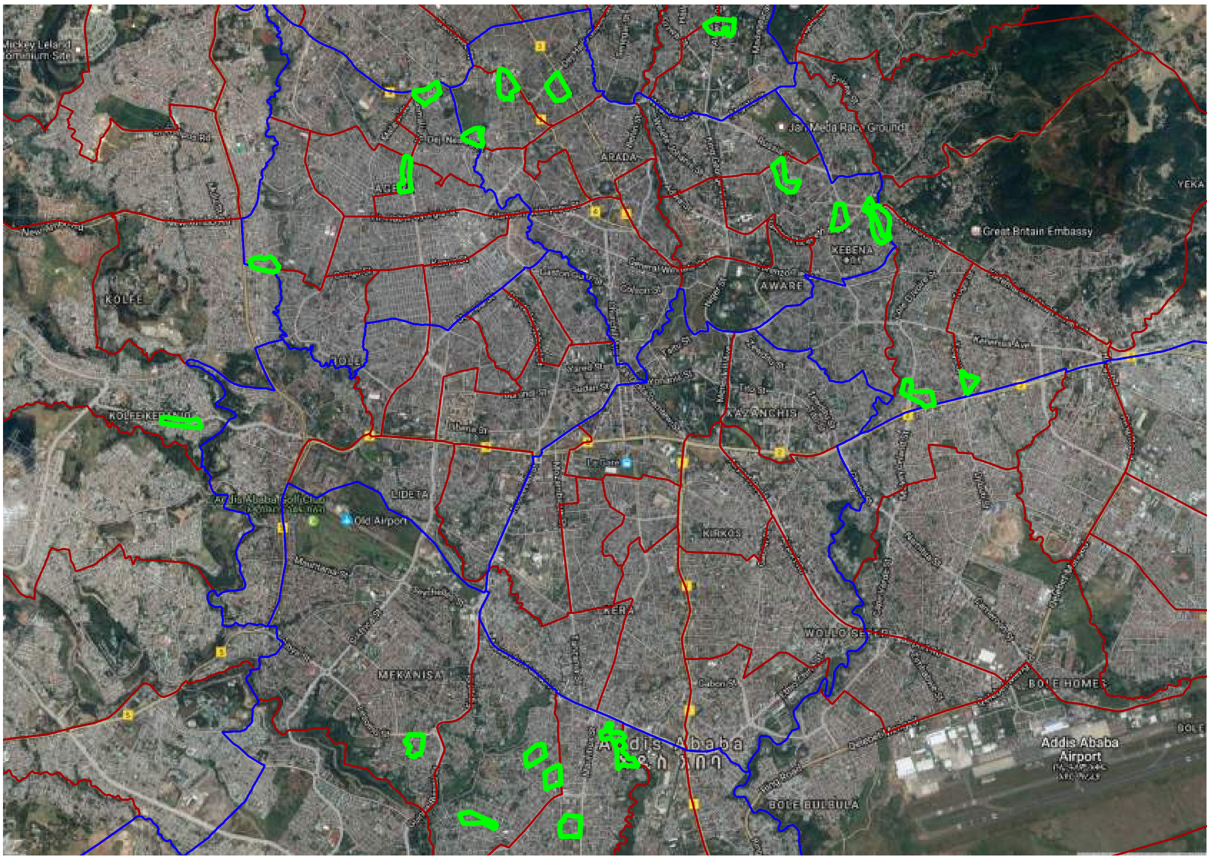

Figure 2: Distribution of randomly selected neighbourhoods in Addis Ababa

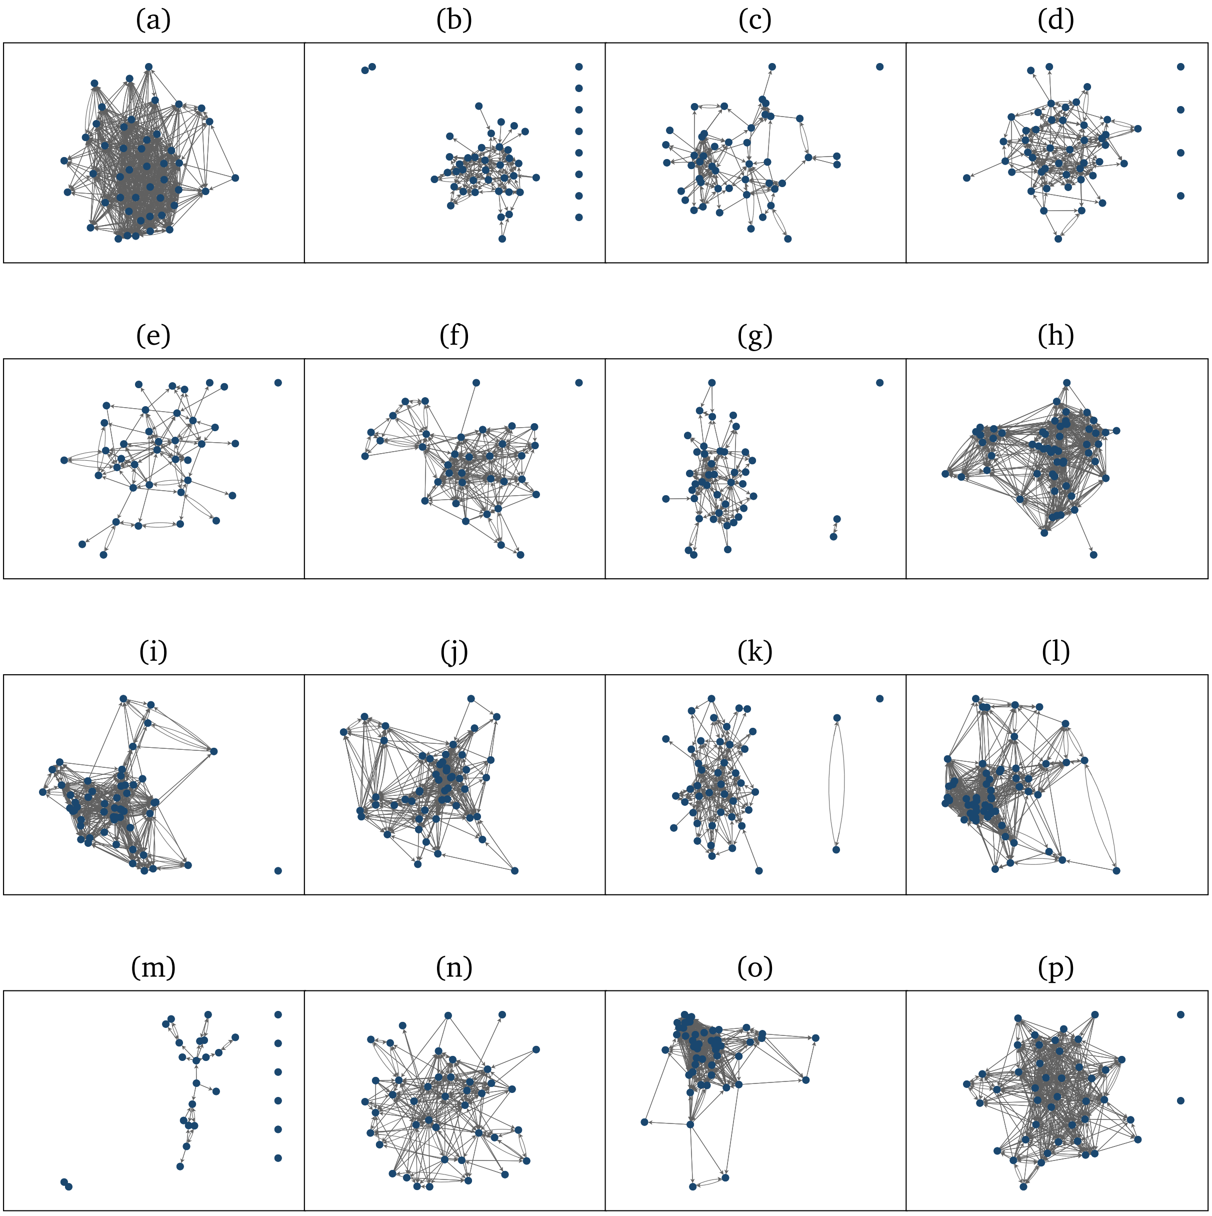

Figure 3: Overview of directed social network graphs at baseline

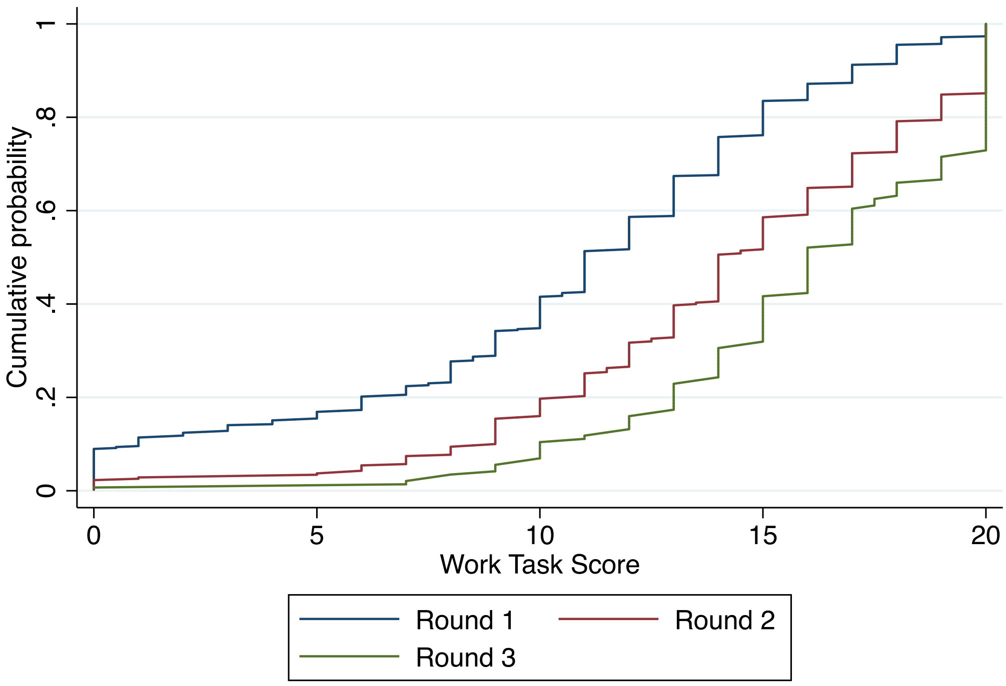

Note: a two-sample Kolmogorov-Smirnov test strongly rejects the equality of the distributions.

Figure 4: Cumulative density function of work task performance over time

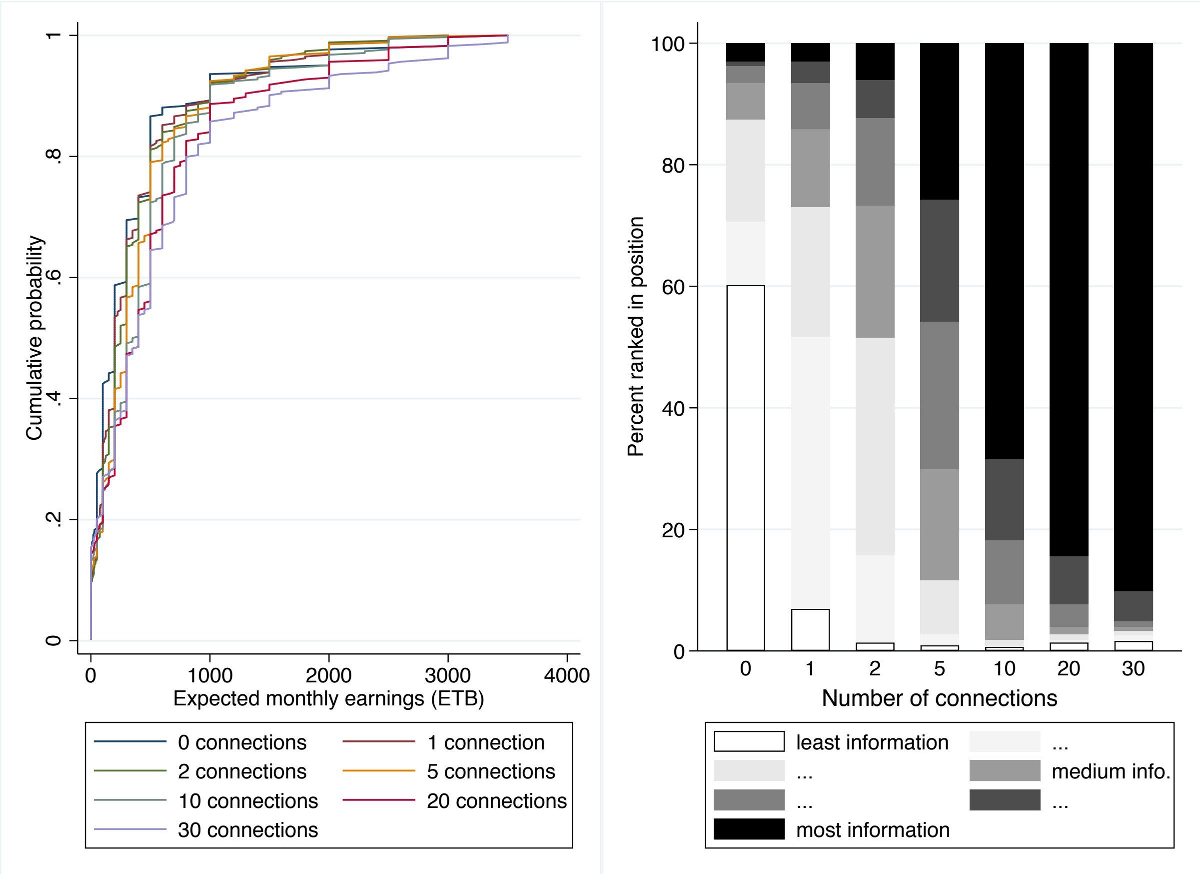

Figure 5: Workers' expectations about correlates of centrality

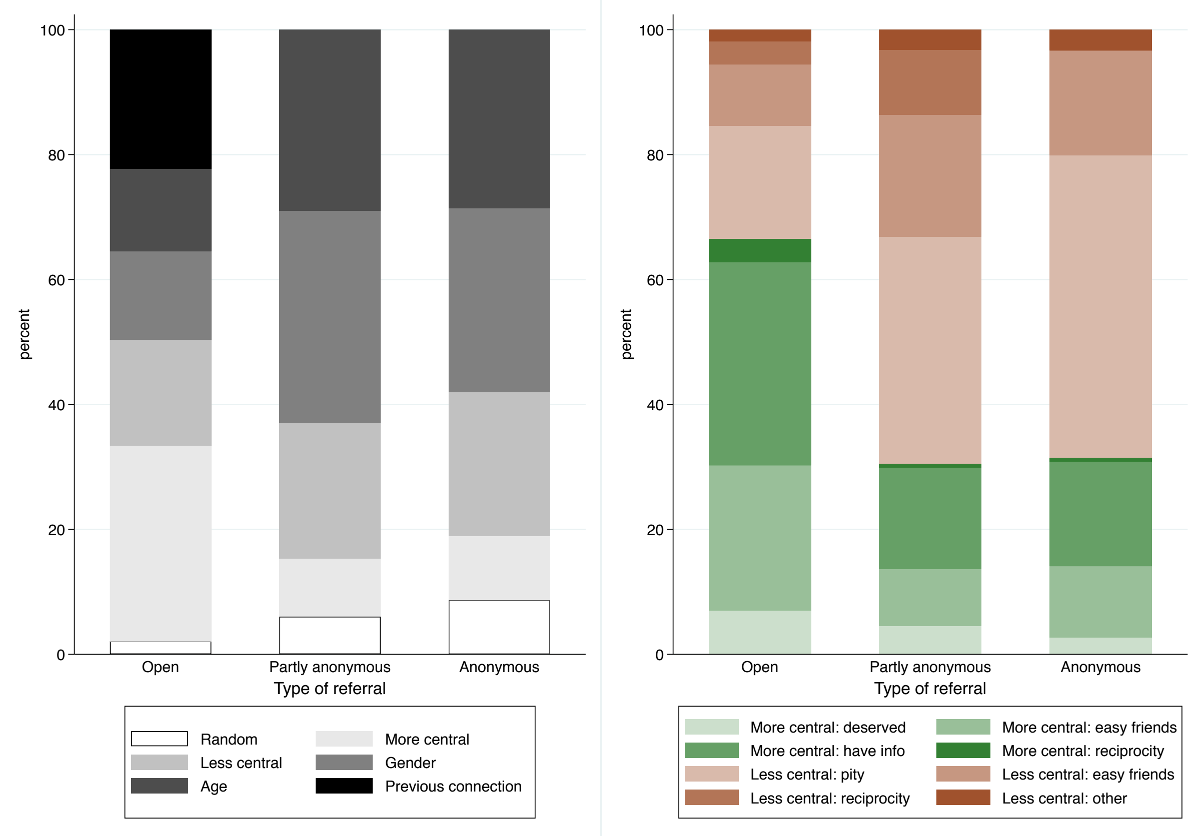

This graph displays data from the short-run follow-up collected by phone. NH: Neighbourhood.

Figure 6: Workers' self-reported reasons for why they made a particular referral

Figure 7: Dyadic regressions: Different periphery thresholds

(a) General networks

(b) Job networks

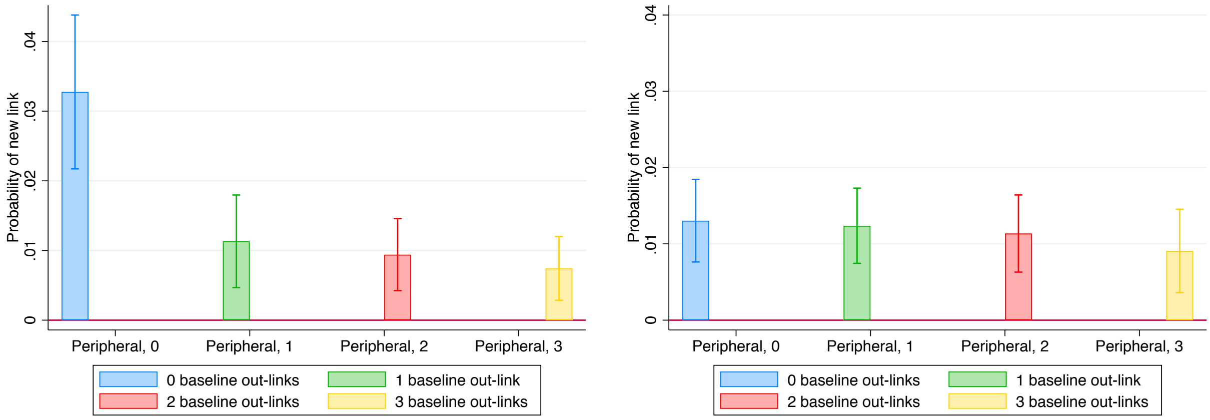

Figure 8: Persistence of newly established social connections after 18 months

(a) General networks (b) General networks - by number of new links

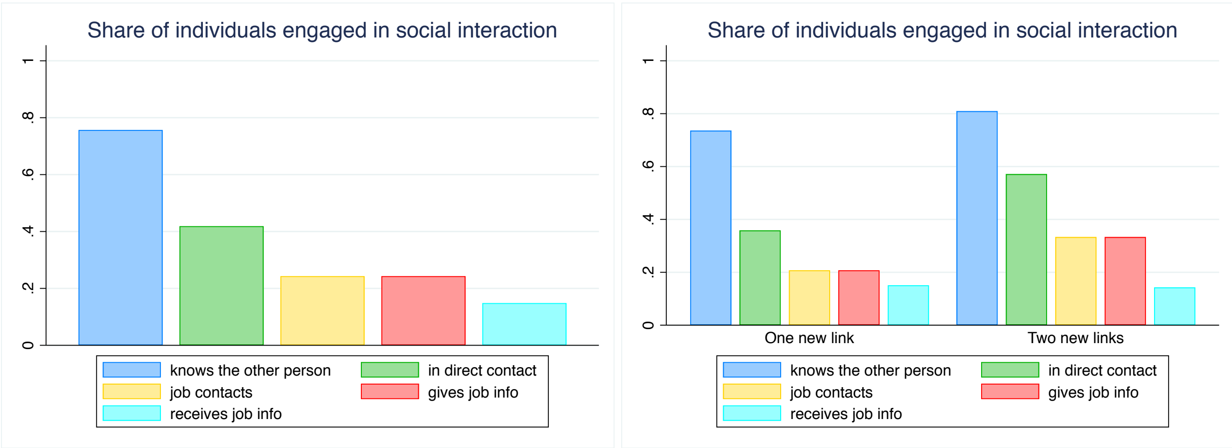

Figure 9: Long-run persistence of new links by baseline peripheral status

(a) General networks – out-degree

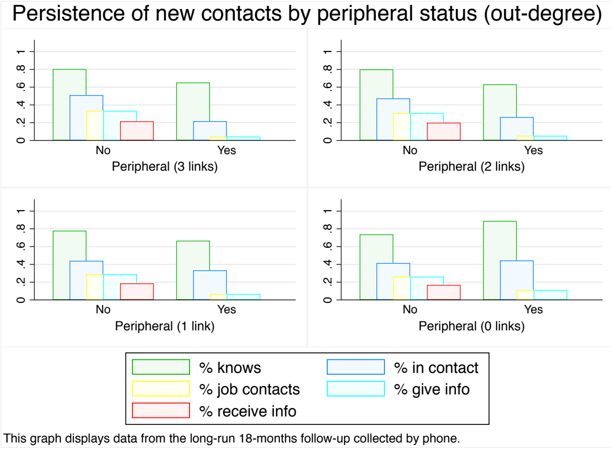

(b) Job networks – out-degree

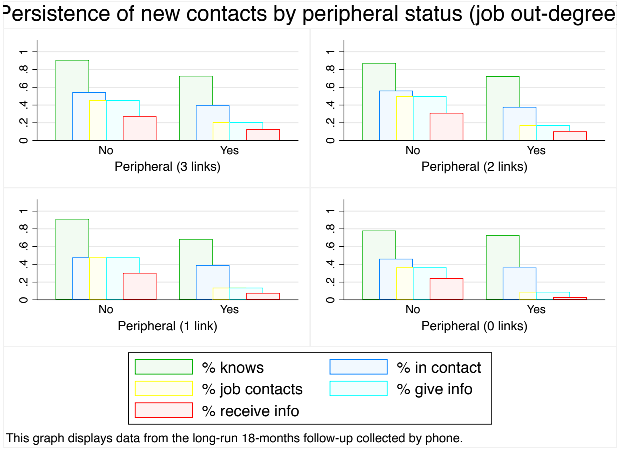

(c) General networks – in-degree

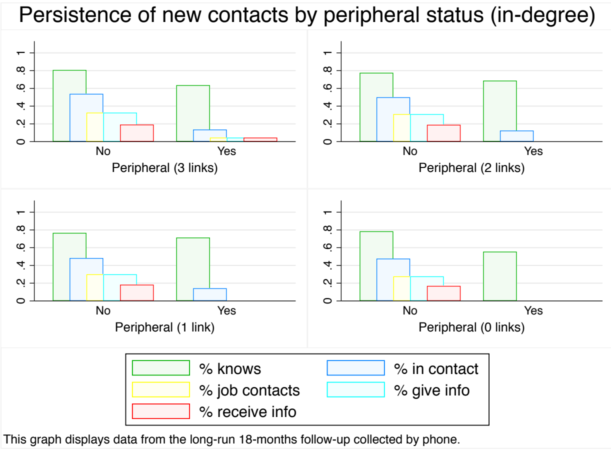

(d) Job networks – in-degree

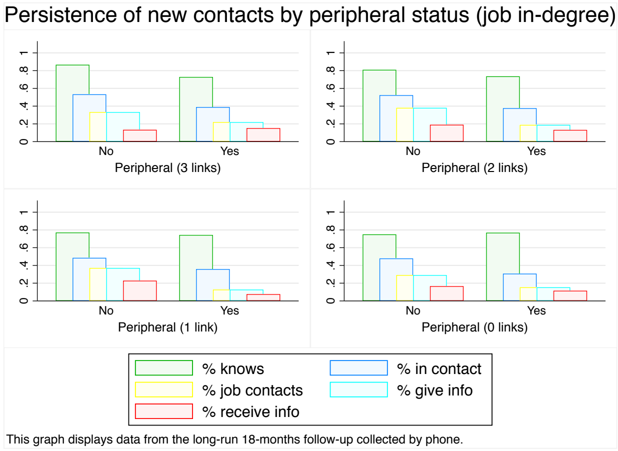

Table 1: Treatment matrix

		Visibility			Shrouding
		Open	Partly anon.	Anonymous	Open+ Cent. shrouded
Incentivisation	No	C	T2	T3	(T4)
	Yes	T1	-	-	-

Main treatments:

1. Standard referral treatment (C) – comparison condition
2. Performance-incentivised referral treatment (T1)
3. Partly anonymous referral treatment (T2)
4. Anonymous referral treatment (T3)

Additional treatments:

5. Shrouded centrality referral treatment (T4)

Table 2: Neighbourhood network summary statistics

			(1)		
General networks	mean	min	max	sd	count
Individuals	46.19	27.00	61.00	8.28	16
Arcs	355.06	33.00	884.00	247.87	16
Density	0.16	0.05	0.43	0.10	16
In-degree centralisation	0.28	0.11	0.50	0.12	16
Out-degree centralisation	0.31	0.12	0.56	0.11	16
Transitivity	0.61	0.36	0.99	0.16	16
Reciprocity	0.27	0.17	0.43	0.08	16
Average between centrality	42.24	1.30	73.41	17.98	16
Average degree centrality	9.31	1.33	19.24	4.73	16
Average eigenvector centrality	0.12	0.11	0.14	0.01	6
Average Katz centrality	21.60	14.88	35.11	5.91	16
Average Clustering coefficient	0.37	0.14	0.58	0.13	16
Number of components	2.88	1.00	10.00	2.66	16
Job networks	mean	min	max	sd	count
Individuals	46.19	27.00	61.00	8.28	16
Arcs	102.88	11.00	301.00	82.09	16
Density	0.05	0.02	0.15	0.03	16
In-degree centralisation	0.13	0.06	0.21	0.05	16
Out-degree centralisation	0.17	0.03	0.49	0.12	16
Transitivity	0.75	0.21	1.27	0.31	15
Reciprocity	0.31	0.10	0.52	0.10	16
Average between centrality	13.45	0.04	47.26	15.34	16
Average degree centrality	2.13	0.41	6.57	1.62	16
Average Katz centrality	39.64	26.77	51.33	6.75	16
Average Clustering coefficient	0.19	0.00	0.40	0.10	16
Number of components	14.44	3.00	24.00	6.92	16

Notes: The variables labelled are the means of individuals within a network, of which subsequently is taken an average across networks. Centralisation scores are divided by N - 1, where N = number of nodes in a network. This standardization makes sure that centrality scores always range from 0 to 1.

Table 3: Connections within the neighbourhood networks

	(1)				
Number of links within network:	mean	min	max	sd	count
All links	9.56	0.00	51.00	8.55	513
Considered information spreaders	7.20	0.00	366.00	20.27	513
Total hours spent with links in network	274.56	0.00	3600.00	421.06	513
Travel connections	2.64	0.00	21.00	3.78	513
Links visited	4.01	0.00	31.00	5.29	513
Visitors	3.95	0.00	30.00	5.21	513
Job info given to	2.86	0.00	28.00	4.29	513
Job info received from	2.73	0.00	31.00	4.22	513
Money lent to	1.79	0.00	18.00	2.77	513
Money borrowed from	1.70	0.00	18.00	2.61	513
Share of total network:	mean	min	max	sd	count
All links	0.18	0.00	0.96	0.16	513
Considered information spreaders	0.14	0.00	8.32	0.45	513
Travel connections	0.05	0.00	0.43	0.07	513
Links visited	0.08	0.00	0.55	0.10	513
Visitors	0.07	0.00	0.57	0.10	513
Job info given to	0.05	0.00	0.53	0.08	513
Job info received from	0.05	0.00	0.58	0.08	513
Money lent to	0.03	0.00	0.34	0.05	513
Money borrowed from	0.03	0.00	0.34	0.05	513

Notes: The absolute numbers sum all connections of a certain type of individual i in her network. The shares divide the number of connections of a certian type of individual i by the overall network size.

Table 4: Balance of covariates for baseline sample, by round one treatment status

	(1) Referral	(2) Mean of control (SD)	(3) Max pairwise difference	(4) Obs.
Age	0.05 (0.31) [0.96]	23.78 (3.43)	0.01	507
Female	0.06 (0.04) [0.82]	0.43 (0.50)	0.12	513
Worked (7d)	0.02 (0.04) [0.96]	0.20 (0.40)	0.04	513
# of friends	2.14 (1.31) [0.82]	8.52 (7.64)	0.25	513
Wage empl (6m)	0.10* (0.05) [0.82]	0.31 (0.46)	0.20	513
Self empl (6m)	-0.03 (0.03) [0.82]	0.20 (0.40)	0.07	513
Ever worked for pay	0.03 (0.03) [0.82]	0.78 (0.42)	0.07	513
Searched job (7d)	-0.01 (0.04) [0.96]	0.40 (0.49)	0.01	513
Searched job (6m)	0.02 (0.05) [0.96]	0.66 (0.47)	0.04	513
Ever searched wage job	0.03 (0.03) [0.82]	0.49 (0.50)	0.05	513
Searched boards (7d)	0.03 (0.03) [0.82]	0.19 (0.39)	0.07	513
Lives with parents	0.06 (0.06) [0.82]	0.60 (0.49)	0.13	513
Years since school	7.67 (7.20) [0.82]	2.10 (3.62)	0.10	503
Permanently empl	0.00 (0.01) [0.96]	0.00 (0.06)	0.00	513
Temporarily empl	0.01 (0.02) [0.96]	0.08 (0.27)	0.03	513
Casual worker	-0.00 (0.02) [0.96]	0.03 (0.18)	0.01	513
Contract worker	0.03 (0.02) [0.82]	0.03 (0.16)	0.13	513
Self-employed	-0.01 (0.02) [0.96]	0.04 (0.20)	0.03	513
Weekly spending (in ETB)	73.62 (52.50) [0.82]	432.94 (538.63)	0.11	491
Amhara	-0.01 (0.07) [0.96]	0.39 (0.49)	0.01	513
Oromo	0.01 (0.04) [0.96]	0.22 (0.42)	0.03	513
Other ethicity	-0.01 (0.04) [0.96]	0.38 (0.49)	0.01	513
Joint p-value	0.88			

Notes: OLS estimates of individual baseline differences of treatment groups. Outcome variables are listed on the left. Standard errors are in parentheses and are clustered at village level for individual-level outcomes. Stars on the coefficient estimates reflect unadjusted p-values. Minimum q-values are in brackets. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level. All monetary values are displayed as converted from Ethiopian birr (ETB) to 2015 USD, with an exchange rate of 21.5 ETB per 1 USD, the average rate from December 2015 to March 2016. In column 5, we calculate the pairwise difference between the two group means and divide this by the standard deviation of the variable, following Imbens (2015). The last row shows joint significance of the coefficients in the corresponding column from SUR estimation.

Table 5: The impact of incentivised referrals on work performance

	Work performance	Work performance	Work performance	Work performance net of learning
Incentivised referral (T1)	2.354***	2.386***	2.385***	3.842***
	(0.855)	(0.839)	(0.840)	(1.313)
Invited by computer	1.033	-0.140	0.00142	0.032
	(0.718)	(0.741)	(0.794)	(0.722)
Constant	15.26***	14.29***	15.44***	15.44***
	(0.672)	(0.686)	(0.934)	(1.032)
NH fixed effects	Yes	Yes	No	No
Day fixed effects	No	Yes	No	No
Session fixed effects	No	No	Yes	Yes
Number of observations	664	664	664	332
p-value for H0: Incentivised referral$_{ji}$=Invited by computer$_i$	0.237	0.025	0.040	0.004

Notes: This model pools work performances from sessions two and three. Session fixed effects implicate neighbourhood and day fixed effects. Standard errors in parentheses. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table 6: Dyadic regressions: Individuals refer strongly reciprocally

	(1)	(2)			(3)	
	Referral$_{ij2}$	Referral$_{ij2}$	Referral$_{ij2}$	Referral$_{ij2}$	Referral$_{ij2}$	Referral$_{ij2}$
Referral$_{ji1}(C)$	0.286***	0.287***	0.280***	0.279***	0.279***	0.288***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Random$_{j1}$		0.00469**	0.00507**	0.0104***	0.0104***	0.0105***
		(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Incentivised$_{j1}$ (T1)					0.00141	0.00325
					(0.00)	(0.00)
Incentivised$_{ji1}$ (T1)						-0.0644***
						(0.02)
Constant	0.00986***	0.00910***	0.0201***	0.0279***	0.0279***	0.0278***
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)	(0.01)
Baseline controls	No	No	Yes	Yes	Yes	Yes
NH fixed effects	No	No	No	Yes	Yes	Yes
Number of observations	17280	17280	17280	17280	17280	17280
p-value for H0: Referral$_{ji1}$=Random$_{i1}$		0.000	0.000	0.000	0.000	0.000
p-value for H0: Referral$_{ji1}$=Incentivised Referral$_{ji1}$						0.000

Notes: This table presents the results from specification (9). Both specifications regress an indicator variable for whether worker i invited person j after the second work day on whether j had previously referred i to the work. Standard errors in parentheses. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table 7: Dyadic regressions: reciprocal referring and costs for productivity

	Work performance			
	Score_i	Score_i	Score_i	Score_i
Reciprocally invited$_i$	-1.265***	-1.209***	-1.154***	-1.114***
	(0.13)	(0.13)	(0.13)	(0.13)
Randomly invited$_i$	-0.795***	-0.655***	-0.767***	-0.758***
	(0.11)	(0.11)	(0.13)	(0.13)
Incentivised and reciprocally invited$_i$				-1.763**
				(0.77)
Incentivised invited$_i$				4.085***
				(0.57)
Constant	12.62***	14.94***	17.58***	17.49***
	(0.06)	(0.37)	(0.45)	(0.45)
Baseline controls	No	Yes	Yes	Yes
NH fixed effects	No	No	Yes	Yes
Number of observations	9990	9990	9990	9990
p-value for H0: $\text{Score}_{i,reciproc.}=\text{Score}_{i,random}$	0.001	0.000	0.021	0.034
$\text{Score}_{i,reciproc.,incentiv.}=\text{Score}_{i,reciproc.}$				0.000

Notes: The work scores display the first-session performance of the workers, which directly captures individual work ability rather than learning on the job or other repeated work considerations. Standard errors in parentheses. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table 8: Pooled cross-sections: Centrality of referred workers compared to compared to whole network

	Number of connections of invited worker
Referral in own neighbourhood (C)	2.644***
	(0.88)
Partly anonymous referral in other neighbourhood (T2)	-3.742***
	(0.68)
Anonymous referral in other neighbourhood (T3)	-4.444***
	(0.62)
Constant	12.30***
	(0.08)
Number of observations	991
p-value for H0:	
Open referral in own neighbourhood= Open referral in other neighbourhood	0.000
Open referral in other neighbourhood= Anonymous referral in other neighbourhood	0.039

Notes: OLS regression of j's degree centrality on the (exogenous) type of referral treatment. Pooled data from all work sessions. Standard errors in parentheses. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table 9: Pooled cross-sections: Open vs shrouded degree centrality referral

	Degr. Centr.$_j$
i invited j with standard referral (C)	2.923**
	(1.17)
i invited j with concealed referral (T4)	2.019
	(1.41)
Constant	9.838***
	(0.63)
Number of observations	480
p-value for H0: Degr. Centr.$_j$ (open referral)=Degr. Centr.$_n$ (concealed referral)	0.695

Notes: OLS regression of j's degree centrality on the (exogenous) type of referral treatment. Standard errors in parentheses. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table 10: Dyadic regressions: Persistent exclusion from general and job networks

	New Link$^{1}_{ij}$	New Link$^{1}_{ij}$	New Link$^{1}_{ij}$
General networks:			
Treatment$_i$	0.0101***	0.0108***	0.0100***
	(0.00)	(0.00)	(0.00)
Peripheral$_i$*Treatment$_i$	0.0174***	0.0165***	0.0142***
	(0.00)	(0.00)	(0.00)
Constant	-1.41e-16	0.00552	0.00754
	(0.00)	(0.00)	(0.01)
Baseline controls	No	Yes	Yes
NH fixed effects	No	No	Yes
N	17280	17280	17280
Job networks:			
Treatment$_i$	0.0128***	0.0136***	0.0133***
	(0.00)	(0.00)	(0.00)
Peripheral$_i$*Treatment$_i$	0.0127***	0.0121***	0.0112***
	(0.00)	(0.00)	(0.00)
Constant	-1.91e-15	0.00718	0.00771
	(0.00)	(0.01)	(0.01)
Baseline controls	No	Yes	Yes
NH fixed effects	No	No	Yes
N	17280	17280	17280

Notes: In this table, an individual i is defined as peripheral when having fewer than two out-degree connections. Standard errors in parentheses. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table 11: Dyadic regressions: Peripheral individuals in general and job networks benefit from reciprocity

Specification	(1)	(2)		
General networks	Referral^2_{ij}	Referral^2_{ij}	Referral^2_{ij}	Referral^2_{ij}
Peripheral_j*$\text{New Referral}^1_{ji}$	0.235***	0.235***	0.236***	0.234***
	(0.04)	(0.04)	(0.04)	(0.04)
Peripheral_j*$\text{Any Referral}^1_{ji}$	-0.0559	-0.0559	-0.0607*	-0.0617*
	(0.04)	(0.04)	(0.04)	(0.04)
$\text{New referral}^1_{ji}$	-0.0988***	-0.0988***	-0.0807***	-0.0821***
	(0.02)	(0.02)	(0.02)	(0.02)
$\text{Any Referral}^1_{ji}$	0.320***	0.321***	0.307***	0.307***
	(0.01)	(0.01)	(0.01)	(0.01)
Random_i		0.00470**	0.00512**	0.0103***
		(0.00)	(0.00)	(0.00)
Constant	0.0100***	0.00928***	0.0197***	0.0285***
	(0.00)	(0.00)	(0.01)	(0.01)
Baseline controls	No	No	Yes	Yes
NH fixed effects	No	No	No	Yes
Number of observations	17280	17280	17280	17280
p-value for H0:				
$\text{Any Referral}^1_{ji}$+$\text{New referral}^1_{ji}$+ Peripheral_j*$\text{New Referral}^1_{ji}$=0 (=$\text{Total effect}^{per}_{new}$)		0.000	0.000	0.000
$\text{Total effect}^{per}_{new}$=$\text{Random}_i$		0.000	0.000	0.000
$\text{Total effect}^{per}_{new}$=$\text{Total effect}^{per}_{any}$		0.008	0.003	0.003
Job networks				
Peripheral_j*$\text{New Referral}^1_{ji}$	-0.0247	-0.0247	-0.0205	-0.0178
	(0.03)	(0.03)	(0.03)	(0.03)
Peripheral_j*$\text{Any Referral}^1_{ji}$	0.0469	0.0469	0.0415	0.0383
	(0.03)	(0.03)	(0.03)	(0.03)
$\text{New referral}^1_{ji}$	0.0149	0.0149	0.0263	0.0243
	(0.02)	(0.02)	(0.02)	(0.02)
$\text{Any Referral}^1_{ji}$	0.261***	0.262***	0.248***	0.249***
	(0.02)	(0.02)	(0.02)	(0.02)
Random_i		0.00468**	0.00522**	0.0103***
		(0.00)	(0.00)	(0.00)
Constant	0.00960***	0.00887***	0.0175**	0.0262***
	(0.00)	(0.00)	(0.01)	(0.01)
Baseline controls	No	No	Yes	Yes
NH fixed effects	No	No	No	Yes
Number of observations	17280	17280	17280	17280
p-value for H0:				
$\text{Any Referral}^1_{ji}$+$\text{New referral}^1_{ji}$+ Peripheral_j*$\text{New Referral}^1_{ji}$=0 (=$\text{Total effect}^{per}_{new}$)		0.000	0.000	0.000
$\text{Total effect}^{per}_{new}$=$\text{Random}_i$		0.000	0.000	0.000
$\text{Total effect}^{per}_{new}$=$\text{Total effect}^{per}_{any}$		0.272	0.488	0.536

Notes: In this table, an individual i is defined as peripheral when having fewer than two out-degree connections. Standard errors in parentheses. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Appendices

A. Information about experiment

A.1. Stylised experimental timeline

Table A1: Stylised timeline by neighbourhood

Week	*Days*	*Time of day*	*Activity*
Week 0	Day 1-2		Neighbourhood census
	Day 3-7		Baseline survey
Week 1	Day 1	Morning	Work session 1
	Day 1	Afternoon	Field referrals
	Day 1	Evening	Computer referrals
	Day 2	Morning	Work session 2
	Day 2	Afternoon	Field referrals
	Day 2	Evening	Computer referrals
	Day 3	Morning	Work session 3
	Day 3	Afternoon	Field referrals
	Day 3	Evening	Computer referrals
Week 2	Day 1-3		Short-run follow-up survey
Week 80	Day 1-7		Long-run follow-up survey

A.2. Experimental protocol

Table A2: Protocol for exemplary work session

Time	*Activity*
At work session	
8:30 am	- invited participants arrive at location and report to assistants
9:30 am	- randomisation into treatment: in blocks of three (same gender, same education, same age)
10:00 am	- beginning of the work session - participants enter room, introduction and game - distribute two kinds of test sheets in the room (one sheet with data, one sheet for data entry) - participants sign attendance sheet upon entering the room, and sit down separately - brief introduction in English - RA reads Amharic instructions - participants can ask any questions they have
11:00 am	- participants have one and a half hours (90 minutes) time to complete task.
12:30 am	- participants hand in their data sheets - assistants count the number of correctly completed tasks
12:30 am	- three different referral processes: - referral to other neighbourhood - anonymous referral to other neighbourhood - main referral in own neighbourhood, by treatment group
after referral processes	- participants will be given their remuneration: a standard 50 ETB for transportation and lunch expenses, and 0-200 ETB for the task completion.
13:00 am	- participants leave
During afternoon	Referral takes place in the field
7:00 pm	- assistants call participants invited by random treatment - if respondent states that she is already invited, she is replaced by next one on random list

General rules for the work session referrals:

- Participants can only invite one person for the next session, and only someone living in the same neighbourhood (i.e. only someone from the list of people handed out.)
- Every invited person has to state precisely who invited him/her.
- If a person is invited by two separate people, we will count the participant who made the first invitation as the main invitation giver.
- Don't tell people that there will be 3 sessions.
- Don't tell people that they can invite someone in the morning.

A.3. Work task instructions English and Amharic

English

In front of you, you will find a copy of these instructions, a booklet including 200 forms and a test sheet. Please add your name, surname, and mobile phone number on the first page of the test sheet. Please do not open any of the materials, until you are told to do so. Each page in the booklet contains some basic information about a fictional worker of a factory: name, surname, age, work team. Each work team has five workers assigned to it. Your task will be to copy the information from the test forms into your data entry table. The table has a separate section for each team. So you first have to ensure that you are recording the information of every individual in the right team. Within each team section, you should enter the information of individuals in age order: younger individuals should be entered before older individuals. Your task is completed when you have entered the information of all 200 individuals into your data entry booklet. You have 1 hour 30 minutes time to complete the test. After 1 hour 30 minutes, you have to stop writing and we will collect your data entry booklet. Your performance will be scored according to how many teams you entered correctly.

Amharic

Script for Data Entry Task

ሁለተኛው የዚህ ከሰአት ስራ መረጃ ማጠናቀርን ወይንም አመዘጋገብን ይመለከታል።ለዚህም መረጃ የያዘ ቅጽና መመዝገቢያ ወረቀት ተሠጥቷል።በመረጃ መመዝገቢው የመጀመሪያ ገጽ ላይ ስም ፤ የአባት ስምና የሞባይል ቁጥር በተሰጠው ቦታ ጻፍ/ጻፊ።እስከሚፈቀድ ድረስ የተሠጡት መረጃ የያዘ ቅጽና መመዝገቢያ ባሉበት ይቆዩ።

የመጀመሪያዉ ቅጽ የአንድ ፋብሪካ 120 ሰራተኞች ሙሉ ስም የስራ ቡዱን እና እድሜ የያዘ ነው።እያንዳንዱ የሥራ ቡዱን 3 አባላት ያሉት ነው።ስለዚህም አርባ የስራ ቡዱኖች ቡዱን 1፣ቡዱን 2……ወዘተ በሚል ይኖራሉ።

የዚህ ፈተና ምዘና ስራ በመረጃ ቅጽ ላይ ያለውን መረጃ በተሠጠው መረጃ መመዝገቢያ ቅጽ ሳጥን ውስጥ እንዳየምድባቸውና እድሜአቸው መመዝገብ ይሆናል።መረጃው እያንዳንዱን ሰው በተሠጠው ትክክለኛ የሥራ ቡዱን ዉስጥ በእድሜ ከትንሽ ወደ ትልቁ መመዝብ ይኖርበታል፡ለዚህም የቡድን አንድ አባላትን ከተሰጠዉ የመረጃ ቅጽ ላይ በመሰብሰብ ለቡድን አንድ በተተወዉ ሳጥን ዉስጥ በእድሜ ትንሽ ከሆነዉ በመጀመር እስከ ትልቁ የቡድን አባል በቅደም ተከተል የሚመዘገብ ይሆናል።በተመሳሳይ መልኩ የቡድን 2፣ የቡድን 3፣ የቡድን4፣ እያለ እስከ ቡድን 40 ድረስ መሞላት ይኖርታል።

መመዝገቢያው ለእያንዳንዱ የሠራተኛ ቡዱን የተለየ ሶስት መስመር ያለዉ ሳጥን አለዉ።

ይህ ፈተና የሚያልቀው መረጃውን በ 40 የሥራ ቡድን ዉስጥ መዝግቦ በመጨረስ ይሆናል።ለዚህም 1 ሠዓት ከ00 ደቂቃ ተሠጥተል።የተሠጠው ሠዓት እንዲልቅ መጻፍ የተከለከለ ሲሆን የፈተና ወረቀቱ የሚሠበሠበው ይሆናል።ይህ ፈተና የሚታረመው በትክክለኛ ቅድመ ተከተል እና በትክክለኛ የሥራ መደብ መረጃው ሰለመመዝገብ በማረጋገጥ ይሆናል።የአንድ የስራ ቡድን አባላትን በሙሉ በትክክልለኛ ቡድናቸዉና በትክክለኛዉ ቅደም ተከተል መዝግቦ መጨረስ አንድ ነጥብ የሚያስገኝ ይሆናል።

በተጨማሪም ፈተና ጨርሶ የሚወጣ ተፈታኝ ለቀን ውጪ የተመደበውን ክፍያ መውሠድና ገንዘብ መዉሰድህን የሚያረጋግጥ ፊርማ ለመፈረም ፍቃደኛ መሆን ይኖርብሀል።

አሁን ጥያቄ ካለ መጠየቅ ይቻላል።

A.4. Other neighbourhood referral instructions English and Amharic

1. "Below is a list of (irrecognisable) unemployed young people from a neighborhood that is very similar to the one you live in. You can see the gender, age and how many friends that person has in her/his neighbourhood. Please pick one person who you would like to invite to the same job you did today. We will implement this referral in reality (with a certain probability). We will give the other person your phone number, so that he or she can get back to you to thank you."

ከዚህ በታች የተጠቀሱት ሰዎች ስማቸው ባይገለፅም እንዳንተው/ቺው አይነት አካባቢ የሚኖሩ ስራ ፈላጊዎች ናቸው። የእያንዳንዱ ፆታ፣ ዕድሜ እና በመኖሪያቸው አካባቢ ያላቸውን የጓደኛ ብዛት ተቀምጧል። ከተገለፁት ሰዎች መካከል አንድ ሰው በመምረጥ ዛሬ ለተሳተፍከበት አይነት ስራ እንድትጋብዝ/ዢ እንፈልጋለን። የተጋበዙት ሰዎች በትክክልም ተሳታፊ ይሆናሉ። ለጋበዝከው/ሽው ሰው ስልክህን/ሽን በመስጠት ለግብዣው ምስጋና እንዲያቀርብ/እንድታቀርብ እናደርጋለን።

2. "Below is a list of (irrecognisable) unemployed young people from a neighborhood that is very similar to the one you live in. You can see the gender, age and how many friends that person has in her/his neighbourhood. Please pick one person who you would like to invite to the same job you did today. We will implement this referral in reality (with a certain probability). However, we will <u>not</u> give the other person your phone number."

ከዚህ በታች የተጠቀሱት ሰዎች ስማቸው ባይገለፅም እንዳንተው/ቺው አይነት አካባቢ የሚኖሩ ስራ ፈላጊዎች ናቸው። የእያንዳንዱ ፆታ፣ ዕድሜ እና አካባቢ ያላቸውን የጓደኛ ብዛት ተቀምጧል። ከተገለፁት ሰዎች መካከል አንድ ሰው በመምረጥ ዛሬ ለተሳተፍከበት አይነት ስራ እንድትጋብዝ/ዢ እንፈልጋለን። የተጋበዙት ሰዎች በትክክልም ተሳታፊ ይሆናሉ። ለጋበዝከው/ሽው ሰው ስልክህን/ሽን የምንሰጥ አይሆንም/አንሰጥም።

B. Extensions of the model

B.1. Continuous ability distribution

This extension of the model assumes a continuous distribution of the ability parameter θ_j. Each worker and referral can perform on the job with continuous ability $\theta_i \in [0,1]$. As before, every worker i can directly observe her contact's ability θ_j.

Similar to the binary ability version of the model, every worker i expects monetary pay-offs from the employer for referring contact j as a function of the contract type P_i and the referral's ability θ_j. Workers also expect a social payoff σ_{ij} from the contact j. When selecting a contact under an incentivised referral scheme, i will receive the payoff $\Pi_{ij} = P_i(\theta_j) + \sigma_{ij}$. The referral payoff with P_i being linear in j's output is thus: $\Pi_{ij} = p_i\theta_j + \sigma_{ij}$.

Assuming that every worker i among her c_i contacts has variation in the contact j's ability, the referral decision of the worker is $\arg\max_j(\Pi_{ij} = p_i\theta_j + \sigma_{ij})$. Under referral performance incentives from the employer (P_i), workers i can maximise their payoff by selecting the referral with the highest ability θ_j.

$$\frac{\partial \Pi_{ij}}{\partial P_i} = \theta_j. \tag{A1}$$

As a consequence, proposition 1 changes to:

Proposition A1 *An increase in P_i increases the benefits of referring a contact with higher ability (equation A1).*

What I expect to observe in the experiment is a higher average ability of workers invited under performance incentives, as expressed by better job performance.

The referral payoff $\Pi_{ij} = P_i\theta_j + \sigma_{ij}$ demonstrates that the social payoff σ_{ij} and the performance incentive P_i are substitutes. In particular, referral contracts that include a performance incentive should *ceteris paribus* make workers i select referrals j with lower social payoffs σ_{ij}. Graphically, this could be expressed as shown below for a given referral payoff Π_{ij}.

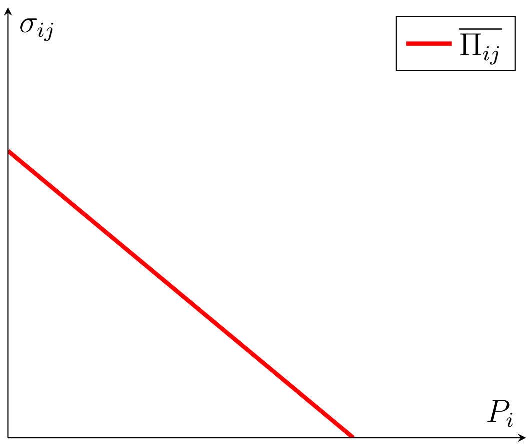

B.2. Workers observe signals of ability

In this version of the model, I assume that the worker i cannot observe her contact's ability directly, but only receives a signal of her contact's ability $\hat{\theta}_i \in \{\theta_i^H, \theta_i^L\}$. The signal is accurate with probability β_i, where $P(\theta = \theta^H | \hat{\theta} = \theta^H, j) = P(\theta = \theta^L | \hat{\theta} = \theta^L, j) = \beta_i$, so that $\beta_i \in [0.5, 1]$. As the subscript i indicates, β_i can differ among workers.[47]

Now, worker i's monetary payoffs for referring contact j are a function of the contract type P_i, the referral's ability θ_j and the accuracy β_i of the signal $\hat{\theta}$. As before, the worker also expects a social payoff σ_{ij}. The decision between a low-ability and high-ability referral is $\sigma_{i1} \in \arg\max_j(\sigma_{ij} | \hat{\theta}_j = \theta_j^H)$ for the high-ability type and $\sigma_{i2} \in \arg\max_j(\sigma_{ij} | \hat{\theta}_j = \theta_j^L)$ for the low-ability type.

When selecting the high-signal contact under an incentivised referral scheme, i will receive the payoff $\beta_i P_i + \sigma_{i1}$, whereas for selecting the low-signal type she will receive $(1-\beta_i)P_i + \sigma_{i2}$. Assuming that every worker i among her contacts c_i has a choice between the high-signal and the low-signal type, the worker compares the best high-signal type with the best low-signal type. As a consequence, worker i will make a referral to a high-signal type if:

$$P_i > \frac{\sigma_{i2} - \sigma_{i1}}{2\beta_i - 1}. \tag{A2}$$

This means that the performance incentive P_i has to exceed the difference in social benefits, adjusted for the probability that the signal is accurate. The remainder of the model is as presented in section 2.

[47]The signal could actually vary for every worker-referral pair: β_{ij}.

C. Network-level summary statistics

Table C3: General and job network-level summary and centrality statistics

General networks:													
#	betweenness	degree	eigenvector	katz	clustering	# of components	nodes	arcs	density	in-degree centralisation	out-degree centralisation	transitivity	reciprocity
1	25.72	19.24	0.140	18.27	0.576	1	46	884	0.427	0.358	0.562	0.995	0.433
2	21.96	3.362	.	35.11	0.215	10	47	158	0.0731	0.170	0.236	0.619	0.206
3	37.51	3.404	.	30.74	0.305	2	47	159	0.0735	0.235	0.124	0.463	0.339
4	50.82	4.245	.	26.53	0.282	5	49	208	0.0884	0.144	0.293	0.608	0.216
5	44.11	2.711	.	20.53	0.140	2	38	103	0.0733	0.147	0.202	0.394	0.212
6	34.44	6.917	.	14.94	0.508	2	36	249	0.198	0.267	0.267	0.362	0.415
7	29.88	3.605	.	27.71	0.225	3	43	154	0.0853	0.205	0.278	0.560	0.225
8	51.12	11.84	0.125	19.16	0.474	1	56	662	0.215	0.336	0.410	0.601	0.304
9	57.94	8.815	.	19.19	0.418	2	54	475	0.166	0.503	0.273	0.544	0.219
10	64.49	7.667	0.120	15.45	0.440	1	51	391	0.153	0.456	0.313	0.687	0.192
11	52.82	4.644	.	17.66	0.348	3	45	209	0.106	0.264	0.334	0.502	0.215
12	73.41	9.902	0.108	19.67	0.436	1	61	604	0.165	0.290	0.493	0.657	0.277
13	1.296	1.333	.	25.63	0.175	8	27	33	0.0470	0.107	0.186	0.844	0.309
14	55.76	5.643	0.131	14.88	0.337	1	42	237	0.138	0.209	0.434	0.564	0.167
15	41.79	12.70	0.120	21.59	0.495	1	53	672	0.244	0.339	0.280	0.527	0.369
16	32.84	10.98	.	18.58	0.483	3	44	483	0.255	0.429	0.310	0.786	0.291

Job networks:													
#	betweenness	degree	eigenvector	katz	clustering	# of components	nodes	arcs	density	in-degree centralisation	out-degree centralisation	transitivity	reciprocity
1	27.37	6.565	.	28.82	0.231	3	46	301	0.145	0.214	0.487	1.063	0.277
2	1.851	1.170	.	45.31	0.124	18	47	55	0.0254	0.0851	0.129	0.260	0.279
3	3.468	1.447	.	44.54	0.241	22	47	68	0.0315	0.101	0.168	0.964	0.388
4	0.388	1.102	.	48.18	0.156	24	49	54	0.0230	0.0829	0.0829	0.844	0.421
5	2.289	0.789	.	36.74	0.105	18	38	30	0.0213	0.0614	0.0336	0.214	0.364
6	10.97	2.611	.	30.88	0.404	5	36	94	0.0746	0.158	0.129	0.545	0.516
7	1.256	1.023	.	41.81	0.116	23	43	43	0.0238	0.146	0.121	0.945	0.397
8	41.50	2.589	.	39.70	0.247	7	56	145	0.0471	0.0817	0.211	0.914	0.318
9	47.26	2.907	.	36.90	0.265	6	54	156	0.0545	0.175	0.213	0.524	0.231
10	16.71	1.961	.	42.91	0.0887	13	51	100	0.0392	0.184	0.246	0.443	0.163
11	1	0.756	.	44.04	0.0815	22	45	34	0.0172	0.0754	0.0522	1.125	0.259
12	23.23	2.574	.	51.33	0.286	16	61	157	0.0429	0.160	0.194	1.271	0.319
13	0.0370	0.407	.	26.77	0	18	27	11	0.0157	0.0636	0.0636	.	0.100
14	5.476	1.595	.	38.66	0.205	16	42	67	0.0389	0.0851	0.160	0.857	0.340
15	26.26	4.887	.	37.23	0.297	8	53	258	0.0936	0.198	0.355	0.640	0.332
16	6.091	1.659	.	40.41	0.157	12	44	73	0.0386	0.151	0.127	0.694	0.259

D. Additional summary statistics, take-up and covariate balance

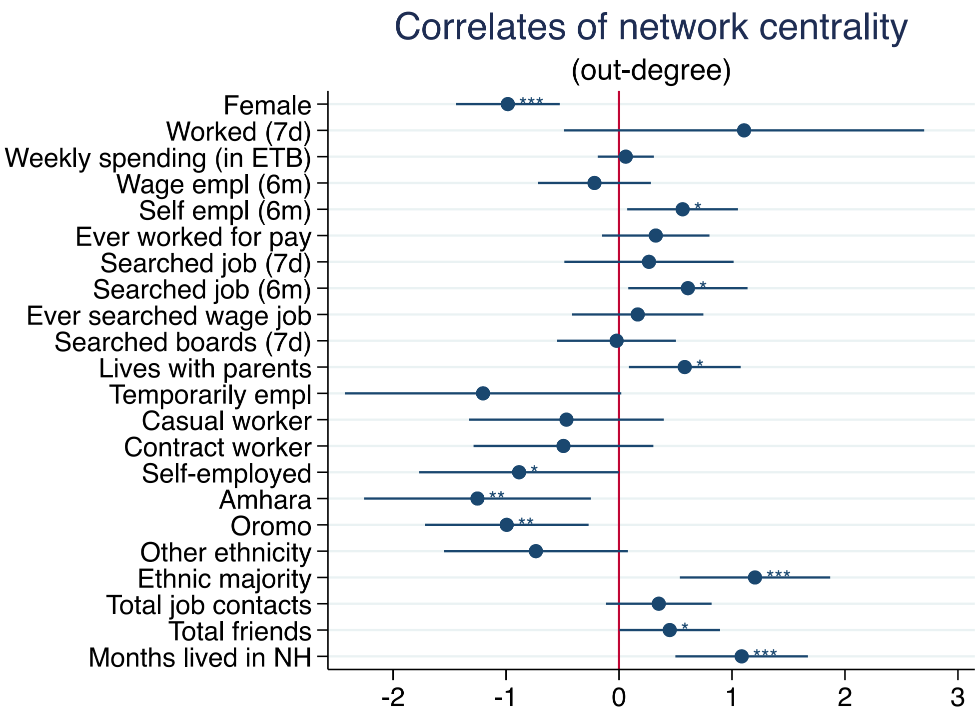

Figure A1: Regression of out-degree centrality on individual covariates
Notes: Standardised regression coefficients. NH=neighbourhood. Imputed missing values for weekly spending (ETB) and number of months lived in NH.

Table D4: Balance of covariates for baseline sample, by attendance

	(1) Attended	(2) Mean of non-attenders (SD)	(3) Max pairwise difference	(4) Obs.
Age	-0.38	24.18	0.11	729
	(0.32)	(3.45)		
	[0.49]			
Female	-0.03	0.49	0.06	739
	(0.06)	(0.50)		
	[0.74]			
Worked (7d)	-0.08*	0.29	0.20	739
	(0.05)	(0.45)		
	[0.27]			
# of friends	2.68***	6.88	0.32	739
	(0.73)	(7.19)		
	[0.05]**			
Wage empl (6m)	-0.06	0.42	0.13	739
	(0.05)	(0.49)		
	[0.49]			
Self empl (6m)	0.02	0.17	0.06	739
	(0.03)	(0.37)		
	[0.74]			
Ever worked for pay	-0.02	0.81	0.05	739
	(0.03)	(0.39)		
	[0.74]			
Searched job (7d)	0.09***	0.30	0.19	739
	(0.03)	(0.46)		
	[0.05]*			
Searched job (6m)	0.08*	0.58	0.18	739
	(0.05)	(0.49)		
	[0.25]			
Ever searched wage job	-0.12***	0.63	0.24	739
	(0.04)	(0.48)		
	[0.05]*			
Searched boards (7d)	0.06**	0.15	0.15	739
	(0.02)	(0.35)		
	[0.08]*			
Lives with parents	0.11**	0.52	0.23	739
	(0.05)	(0.50)		
	[0.18]			
Years since school	3.41	2.39	0.05	724
	(3.63)	(3.89)		
	[0.57]			
Permanently empl	-0.00	0.00	0.01	739
	(0.01)	(0.07)		
	[0.93]			
Temporarily empl	-0.07*	0.15	0.22	739
	(0.03)	(0.36)		
	[0.23]			
Casual worker	0.00	0.03	0.01	739
	(0.02)	(0.17)		
	[0.93]			
Contract worker	-0.01	0.05	0.05	739
	(0.02)	(0.22)		
	[0.74]			
Self-employed	-0.01	0.04	0.03	739
	(0.01)	(0.21)		
	[0.81]			
Weekly spending (in ETB)	-181.80	650.88	0.14	702
	(169.33)	(2170.03)		
	[0.55]			
Amhara	-0.04	0.43	0.08	739
	(0.04)	(0.50)		
	[0.57]			
Oromo	0.03	0.20	0.07	739
	(0.02)	(0.40)		
	[0.49]			
Other ethicity	0.01	0.37	0.02	739
	(0.04)	(0.48)		
	[0.89]			
Joint p-value	0.17			

Notes: OLS estimates of individual baseline differences across attendance. Outcome variables are listed on the left. Standard errors are in parentheses and are clustered at village level for individual-level outcomes. Stars on the coefficient estimates reflect unadjusted p-values. Minimum q-values are in brackets. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level. All monetary values are displayed as converted from Ethiopian birr (ETB) to 2015 USD, with an exchange rate of 21.5 ETB per 1 USD, the average rate from December 2015 to March 2016. In column 5, we calculate the pairwise difference between the two group means and divide this by the standard deviation of the variable, following Imbens (2015). The last row shows joint significance of the coefficients in the corresponding column from SUR estimation.

E. Additional results

E.1. Cross-country evidence on obtaining jobs through social networks

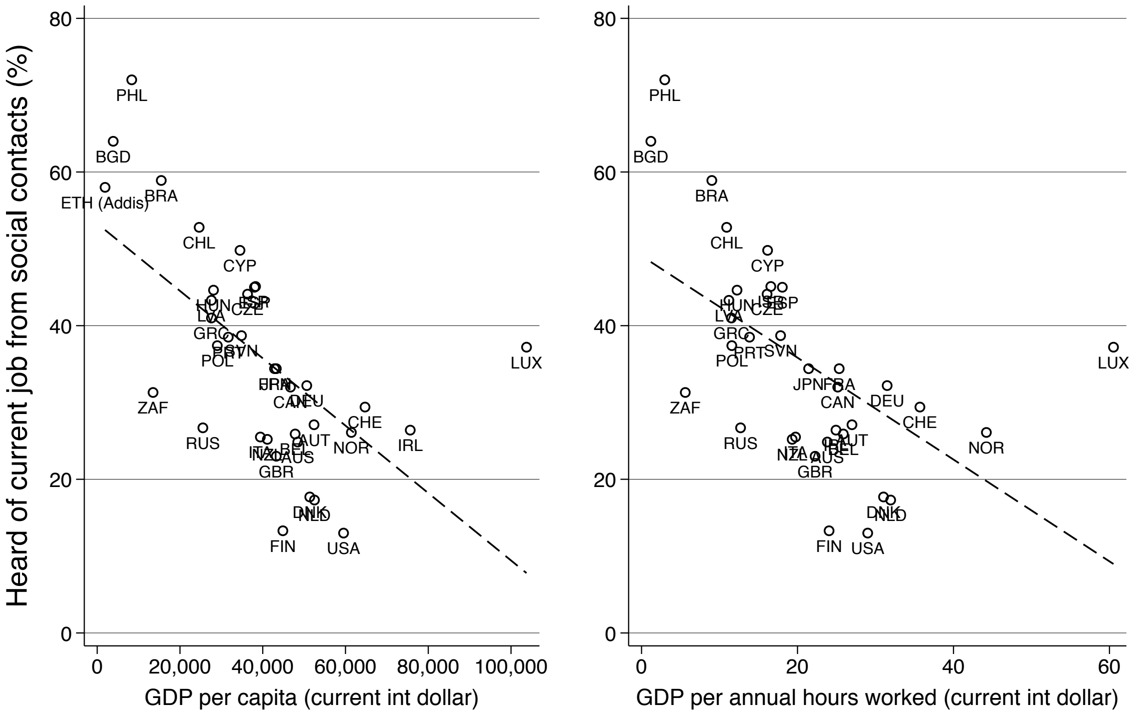

Source: data compiled from ECHP, ISSP, OurWorldInData and own data. Labour productivity is defined as GDP per capita over working hours per year, in 2013

Figure A2: Correlation between finding jobs through social networks and GDP per capita or labour productivity

E.2. Heterogenous effects

In this section, I look at heterogeneous effects, displaying the results on productivity and reciprocity by the degree of the worker making the referrals, as well as the referee.

E.2.1. Productivity by degree

Figure A3 shows that when pooling the production data, there is a small positive relationship between a worker's in-degree centrality and her performance on the job. The relationship does not differ by whether the referral was incentivised for the original referee.[48]

[48]When taking the out-degree centrality instead, the relationship becomes almost precisely zero.

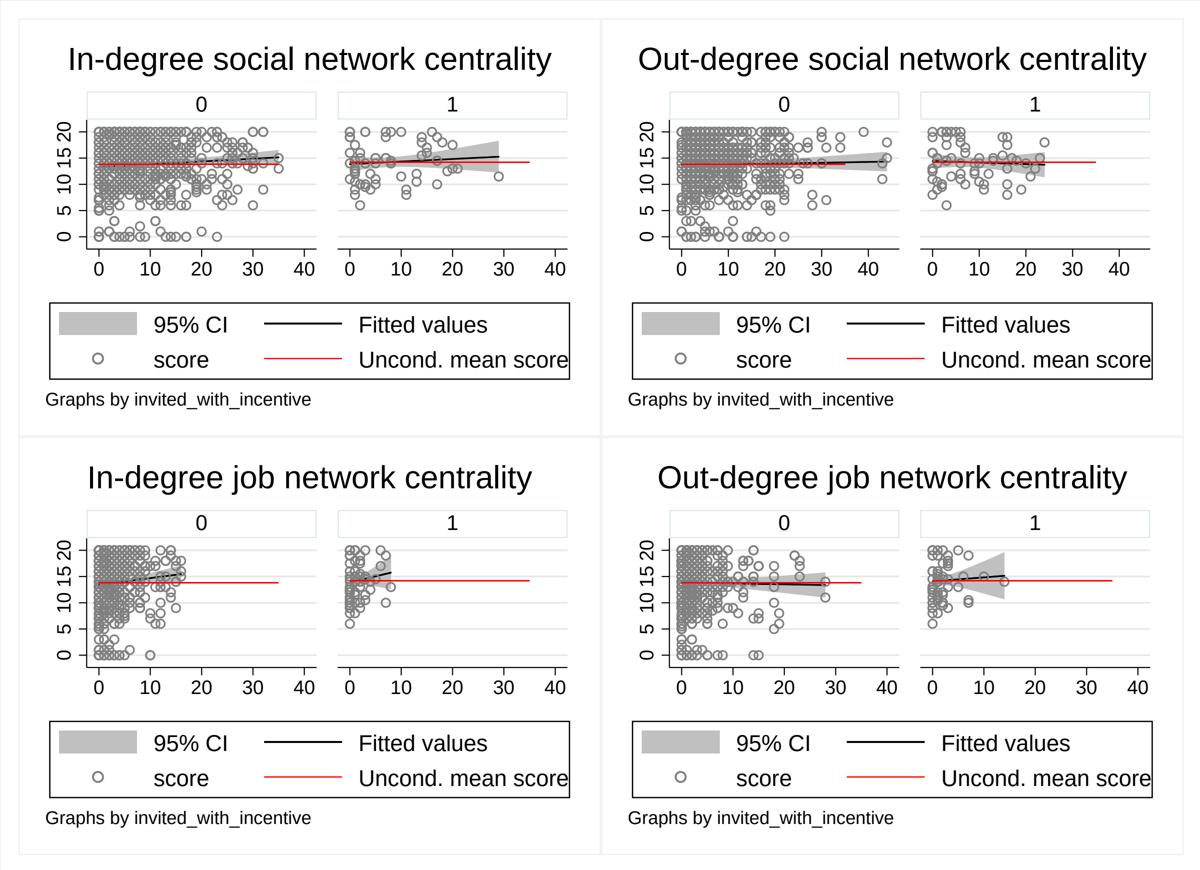

Figure A3: Correlation between productivity and centrality, by whether the referral is incentivised

E.2.2. Reciprocity by degree

Figure A4 displays the relationship between a first round invited worker j's network centrality and her likelihood of inviting i reciprocally to the job, split by whether the first referral from i to j is incentivised or not. I find that under the non-incentivised referral treatment, there is essentially no relationship between centrality and likelihood of reciprocity. Under the incentivised treatment, a similar picture emerges for social network centrality. However, for job networks I find a positive relationship between j's centrality and the likelihood of j reciprocally referring i.

I can interpret this finding with my theoretical model presented in section 2: the increase in P_i due to the incentive treatment decreases the immediate reciprocity concerns ρ_{ji} in the social payoff σ_{ij} function, and instead increases the weight of future benefits κ_{ji}. j's job network centrality now positively predicts her likelihood of engaging in reciprocal referral arrangements. As a consequence, I observe a positive correlation between ρ_{ji} and future benefits κ_{ji}.

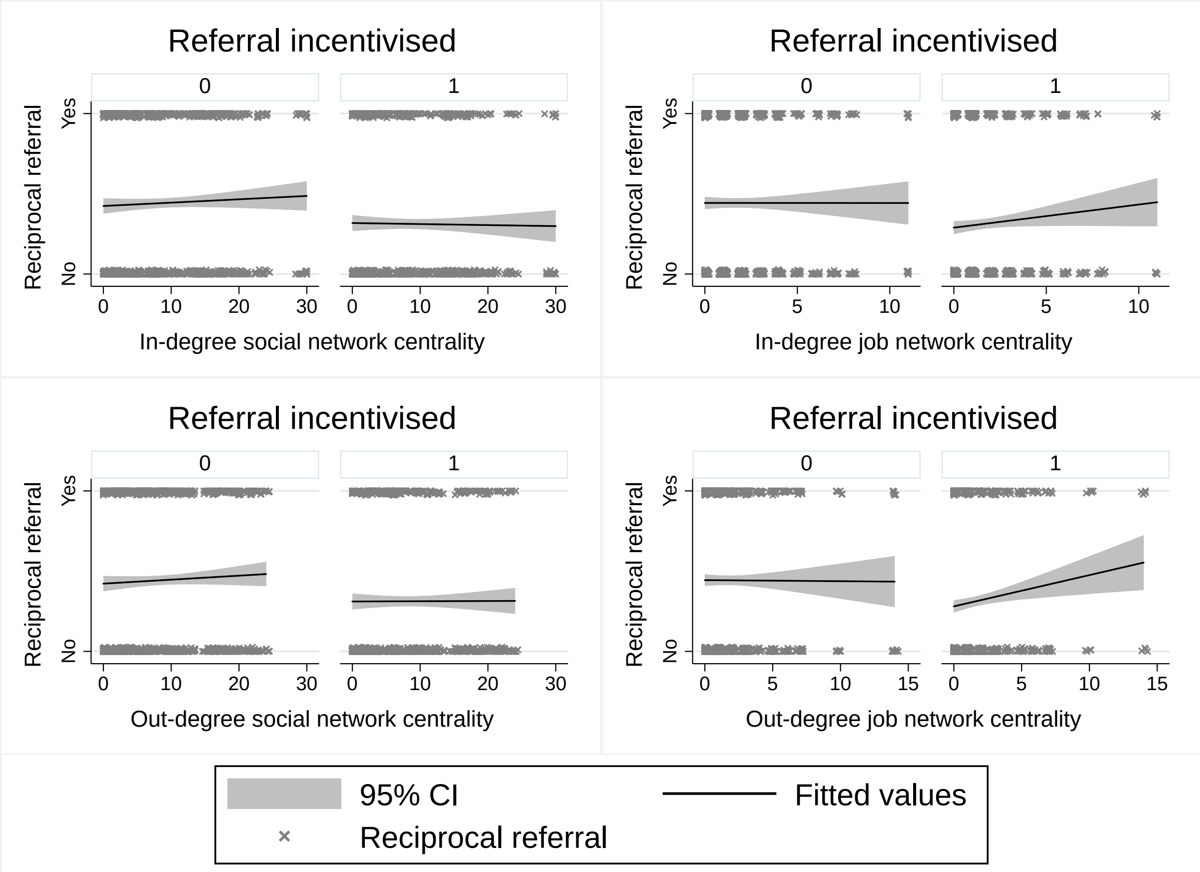

Figure A4: Correlation between reciprocity and centrality, by whether the referral is incentivised

E.3. Correlation between centrality measures and getting a referral

The second panel of table E5 shows results for a regression of a binary indicator of whether individual i referred individual j to the job on a variety of job network centrality measures. The positive effects found in general social networks for in- and out-degree centrality persist, with in-degree dominating again. If individual j is in the job information network of one more person, her chances of being referred to the job increase by almost 6%. Interestingly, when all network centrality measures enter the regression simultaneously, the individual clustering coefficient enters as a strong negative predictor of the chances of receiving a job referral.[49] An intuitive explanation of the local clustering coefficient is how close j's neighbours (*excluding* j) are to being a clique, or a complete graph (Watts and Strogatz, 1998), meaning that all neighbours N_i are also interconnected among themselves and not just through j. This brings more nuance to the previous finding: if j is a job information contact of individual i, j's chances of being referred to the job increase. However, if individual i is also connected

[49]The local clustering coefficient of node j in the network is defined as the share of network contacts N_i of j who are directly connected among themselves.

Table E5: Pooled cross-sections in general and job networks: More central individuals are referred to the job

General networks	Referral$_{ij}$	Referral$_{ij}$	Referral$_{ij}$	Referral$_{ij}$	Referral$_{ij}$	Referral$_{ij}$	Referral$_{ij}$	Referral$_{ij}$	Referral$_{ij}$
Outdegree centrality$_j$	0.0100***		0.00410						0.00593*
	(4.62)		(1.42)						(1.96)
Indegree centrality$_j$		0.0120***	0.00920***						0.0101***
		(5.40)	(3.09)						(3.15)
Between centrality$_j$				0.0000304					-0.000588**
				(0.13)					(-2.16)
Out-Katz centrality$_j$					-0.000691				0.00119
					(-0.40)				(0.66)
In-Katz centrality$_j$						-0.00152			-0.000473
						(-0.93)			(-0.28)
Clustering centrality$_j$							0.140**		-0.00628
							(2.38)		(-0.09)
Eigenvector centrality$_j$								0.454	
								(1.14)	
Constant	0.346***	0.327***	0.316***	0.431***	0.448***	0.464***	0.375***	0.482***	0.309***
	(13.63)	(12.60)	(11.69)	(20.38)	(10.89)	(12.34)	(12.67)	(8.62)	(4.73)
N	836	836	836	836	836	836	836	406	836
Job networks									
Outdegree centrality$_j$	0.0123***		0.00295					0.00382	
	(3.13)		(0.57)					(0.66)	
Indegree centrality$_j$		0.0208***	0.0183***					0.0151**	
		(4.13)	(2.73)					(1.99)	
Between centrality$_j$				0.000585*				-0.000304	
				(1.88)				(-0.83)	
Out-Katz centrality$_j$					-0.00701***			-0.00485**	
					(-4.16)			(-2.52)	
In-Katz centrality$_j$						-0.00523***		-0.00290	
						(-3.63)		(-1.56)	
Clustering centrality$_j$							-0.0218	-0.163***	
							(-0.44)	(-2.94)	
Constant	0.397***	0.373***	0.372***	0.420***	0.712***	0.640***	0.439***	0.734***	
	(19.25)	(16.74)	(16.55)	(22.85)	(10.29)	(10.74)	(20.66)	(6.73)	
N	836	836	836	836	836	836	836	836	

Notes: OLS regression of a binary indicator indicating whether individual i referred individual j to the job on individual j's centrality in general and job networks. The regressions are run on the pooled cross-sectional data (i.e. one observation per individual). * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

to j's contacts k and l, j's chances decrease *ceteris paribus*. In other words, if j and i are a connected pair of job information partners, j has higher chances of being invited to a job than in a triplet where i, j and k are all connected (complete graph). This underlines the rivalrous aspects of job referrals: existing job information networks predict who gets the referral. In a complete clique of job information sharers, the referral can go to multiple individuals, whereas in pairs of individuals, the potential recipient is only one person – unless the referee links up with a new node.

F. Job network maps

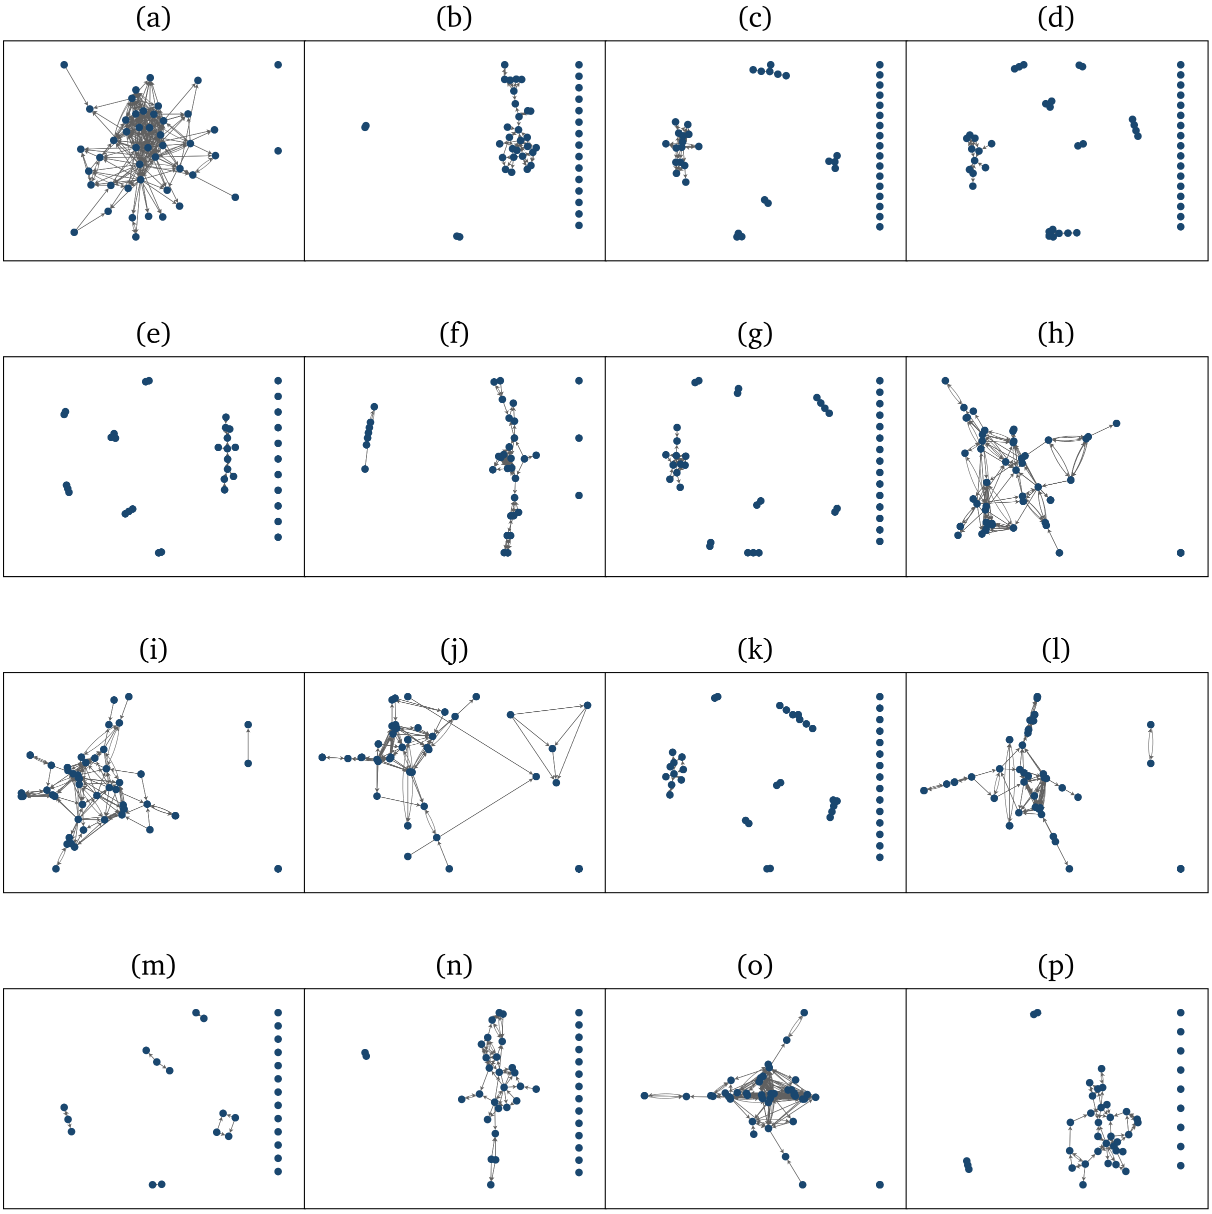

Figure A5: Overview of directed job referral network graphs at baseline

G. Referral sheets

Referral sheet -- 003-05 Nefas Silk Lafto -- Session 1/2					
ያንተ/ች ስም:					
Tick referral:	Name	Nickname	Number 1	Number 2	በአካባቢዉ የሚኖሩ ጓደኞች ብዛት
	[redacted]				6
	[redacted]				8
	[redacted]				16
	[redacted]				5
	[redacted]				11
	[redacted]				12
	[redacted]				4
	[redacted]				
	[redacted]				
	[redacted]				13
	[redacted]				13
	[redacted]				14
	[redacted]				
	[redacted]				1
	[redacted]				11
	[redacted]				7
	[redacted]				10
	[redacted]				6
	[redacted]				15
	[redacted]				17
	[redacted]				6
	[redacted]				3
	[redacted]				18
	[redacted]				2
	[redacted]				11
	[redacted]				3
	[redacted]				6
	[redacted]				0
	[redacted]				0
	[redacted]				10
	[redacted]				2
	[redacted]				3
	[redacted]				9
	[redacted]				9
	[redacted]				12
	[redacted]				
	[redacted]				11
	[redacted]				9
	[redacted]				11
	[redacted]				6

Figure A6: Exemplary home referral list

<u>Other referral sheet -- 003-05 Nefas Silk Lafto</u>

Below is a list of unemployed young people from a neighborhood similar to the one you live in. Please pick one person who you would like to make a referral to the same job you did today. We will implement this referral in reality (with a 1/15 probability) and give the other person your name, phone number and location, so that he or she can get back to you to thank you.

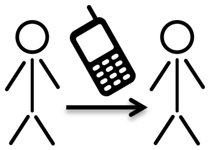

ያንተ/ች ስም:				
Tick referral:	ስም	በአካባቢህ/ሽ የሚኖሩ ጓደኞች ብዛት	እድሜ	ጾታ
	00808I00301	9	17	ሴት
	00808G00201	0	28	ሴት
	00808D00902	21	19	ሴት
	00808D00601	0	26	ወንድ
	00808D01001	1	27	ሴት
	00808G00602	26	23	ሴት
	00808D00701	4	21	ሴት
	00808I01001	24	28	ሴት
	00808I00401	9	24	ወንድ
	00808I00804	16	25	ወንድ
	00808G00401	0	23	ወንድ
	00808I00803	14	25	ወንድ
	00808I00901	19	26	ሴት
	00808D00501	21	25	ሴት
	00808D01101	14	27	ሴት
	00808D00801	15	23	ወንድ
	00808I00801	15	24	ወንድ
	00808D00101	20	29	ወንድ
	00808I00201	2	27	ወንድ
	00808I01002	0	19	ሴት
	00808I00501	29	22	ወንድ
	00808D00901	17	25	ሴት
	00808I00602	27	25	ሴት
	00808I00101	2	25	ሴት
	00808I00701	1	25	ሴት
	00808D00401	28	20	ወንድ
	00808G00301	1	19	ወንድ
	00808D00301	20	19	ወንድ
	00808I00601	0	27	ሴት
	00808D00201	1	27	ወንድ
	00808G00601	30	26	ወንድ
	00808G00501	2	19	ወንድ
	00808I00802	19	26	ወንድ
	00808G00101	10	22	ወንድ
	00808D01201	3	21	ሴት

Figure A7: Exemplary first ‘other’ referral list: Open

<u>Other anon referral sheet -- 003-05 Nefas Silk Lafto</u>

Below is a list of unemployed young people from a neighborhood similar to the one you live in. Please pick one person who you would like to make a referral to the same job you did today. We will implement this referral in reality (with a 1/15 probability), but we will <u>not</u> give the other person your phone number or any other personal details.

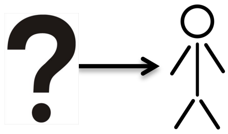

ያንተ/ች ስም:				
Tick referral:	**ስም**	**በአካባቢህ/ሽ የሚኖሩ ጓደኞች ብዛት**	**እድሜ**	**ጾታ**
	00808E00401	19	29	ሴት
	00808B00201	16	27	ሴት
	00808E01302	3	25	ወንድ
	00808B00901	18	21	ወንድ
	00808F00702	19	25	ወንድ
	00808E00402	9	21	ሴት
	00808B00401	3	24	ሴት
	00808B00601	0	28	ወንድ
	00808B00501	3	20	ወንድ
	00808E01101	27	23	ወንድ
	00808F00901	2	19	ወንድ
	00808E01502	31	20	ወንድ
	00808E01201	15	21	ወንድ
	00808E01401	7	24	ሴት
	00808E00501	10	24	ሴት
	00808E01202	23	20	ወንድ
	00808F00801	14	18	ወንድ
	00808E00801	25	19	ወንድ
	00808E01501	1	25	ሴት
	00808F00902	23	25	ወንድ
	00808E00101	6	21	ሴት
	00808E01001	14	28	ወንድ
	00808B00101	4	28	ወንድ
	00808F00701	2	24	ወንድ
	00808E00601	1	28	ሴት
	00808B00801	11	20	ሴት
	00808E00901	12	29	ወንድ
	00808B00701	0	18	ሴት
	00808E00301	24	24	ሴት
	00808E00602	6	28	ሴት
	00808B00301	29	26	ሴት
	00808F00601	8	21	ሴት
	00808E01301	1	25	ወንድ
	00808E00701	13	19	ሴት
	00808E00201	21	20	ሴት

Figure A8: Exemplary second ‘other’ referral list: Anonymous