Hiring Through Referrals in Experimental Markets with Adverse Selection

Aurelie Dariel* Arno Riedl‡ Simon Siegenthaler§

May 31, 2018

Abstract

It is well-known that informational asymmetries can prevent markets from operating efficiently. An important example is the labor market, where employers face uncertainty about the productivity of job candidates. This study explores whether employee referrals can convey information and improve welfare in such environments. In the experiment, each worker has a privately known productivity type. A social network is introduced such that each worker has a social tie to another worker. In a first stage, firms hire workers in a public (posted-offer) market. In a second stage, firms can hire new workers either in the public market or through employee referrals. Referrals occur between workers linked by a social tie and are informative because social ties are more likely between workers of similar productivity. In line with the predictions of our theoretical model, we find that firms take advantage of the opportunity to hire through referral offers, leading to more high productivity hires and increasing efficiency levels compared with a baseline treatment where there are no social ties among workers. We also discuss why adverse selection effects are reinforced in dynamic environments.

JEL Classification: C92, D82, D85, E20

Keywords: Adverse selection, labor market, employee referrals, social networks

*New York University Abu Dhabi. Division of Social Science, New York University Abu Dhabi, P.O. Box 129188, Abu Dhabi, United Arab Emirates, phone: +971 2-628-4000, e-mail: apd5@nyu.edu

‡CESifo, IZA, Netspar, Department of Economics (AE1), School of Business and Economics, Maastricht University, P.O.Box 616, 6200 MD Maastricht, the Netherlands, phone: +31 43-38-84982, e-mail: a.riedl@maastrichtuniversity.nl

§Naveen Jindal School of Management, University of Texas at Dallas, Richardson, TX 75080, USA. Phone: +1 972-883-5871. E-mail: simon.siegenthaler@utdallas.edu.

1 Introduction

Firms frequently fill job vacancies through the social networks of their employees. Up to one half of all open positions in the U.S. and European labor market are allocated based on referrals by friends and relatives (e.g. Ioannides and Loury, 2004; Jackson, 2010a; Pellizzari, 2010; Topa, 2011). Employee referrals are also widely used in developing countries (e.g. Heath, forthcoming), in migrant communities (Munshi, 2003; Beaman, 2011), and are pervasive in different industries ranging from call-centers to high-tech (Burks et al., 2015). Both market sides seem to benefit from employee referrals. Referred workers yield substantially higher profits per worker than non-referred workers, mainly driven by lower turnover and lower recruiting costs, are more likely to be hired, and tend to earn higher wages (e.g. Burks et al., 2015; Brown et al., 2016).

Labor markets are a prominent example where informational asymmetries lead to adverse selection and thus inefficient outcomes (Akerlof, 1970). It is therefore important to study mechanism that can help promote social welfare in such environments. Employee referrals can be one such mechanism. In particular, employee referrals can reveal information about the characteristics of prospective job candidates. The reason is that current employees tend to refer individuals of similar ability (e.g. Rees, 1966; Granovetter, 1985, 1995). Hence, firms have an incentive to search for new workers through the social networks of their current, high-performing employees. More generally, this argument is based on the observation that people tend to interact with others who are like themselves, often referred to as homophily (e.g. McPherson et al., 2001; Currarini et al., 2009; Jackson, 2010b).

This paper provides a test of the hypothesis that employee referrals alleviate adverse selection in labor markets.[1] We study this question experimentally. This allows us to circumvent methodological difficulties related to the estimation of productivity or the degree of homophily in a social network, and to isolate the information value of referrals from other potential benefits such as faster and less expensive job-matching (e.g. Calvó-Armengol and Zenou, 2005; Galenianos, 2014). In our experiment, workers have unobservable productivities, 50% high types and 50% low types. There are two stages. A different set of workers is active in each stage. Firms seek to hire one worker per stage. In the *Baseline* treatment, there are no social ties between workers. The hiring process takes place in a public market lasting for 2 minutes where firms can freely post offers and workers choose if and when to accept. In the *Referral* treatment, we introduce a simple social network: each stage-2 worker has a social tie to a stage-1 worker and, crucially, with a probability of 75% two linked workers are of the same productivity type. The hiring process in stage 1 of the *Referral* treatment is identical to the *Baseline* treatment. In stage 2, however, firms can also make referral offers, that is, they can offer wages directly to the stage-2 that has a social tie to

[1] It should be noted that the use of referrals is not limited to the labor market. In markets for goods and services referrals can take the form of recommendations, e.g. when people refer doctors to their friends or when informing others about trustworthy dealers in second-hand markets.

the previously hired stage-1 worker.

We develop a theoretical model which delivers a number of testable implications. The three most important ones are as follows: Firstly, firms make referral offers only after hiring a high productivity worker in stage 1. Secondly, the existence of a social network leads to a greater proportion of high productivity hires relative to the baseline treatment (in which there are no social ties). And thirdly, despite the fact that referral offers are possible only in the second stage, firms increase their wage offers already in the first stage in order to attract high productivity workers and get access to higher quality social ties. Empirically, the last observation seems particularly challenging as it requires a substantial amount of foresight by the subjects.

Remarkably, all three hypotheses are borne out by the experimental results. Referral hires are common and alleviate inefficiencies arising from informational asymmetries compared with the *Baseline* treatment where firms can hire workers only in the posted-offer market. The positive effect of social ties can be observed in both stages. Our model together with the corroborating evidence generated in a controlled environment thus points to the important role social networks can play in markets impaired by adverse selection. Despite the welfare-enhancing effect of employee referrals, however, overall welfare levels in both treatments fall short of the efficient (second-best) outcome. We show that this is caused by a combination of risk aversion and our dynamic environment in which firms can freely choose when to make offers. Firms in the public market almost always start with low offers and often all low productivity workers are hired before offers are raised above the high productivity workers' reservation wage (this is surprising because in a static environment in which firms make a single offer, the prediction would be to immediately offer a high wage).

There is an extensive theoretical literature on social networks in the labor market. Topics of interest have been job-search through personal contacts (e.g. Mortensen and Vishwanath, 1994; Pissarides, 2000; Topa, 2001; Kugler, 2003), network structure and how outcomes depend on an individual's position in a network (e.g. Boorman, 1975; Zenou, 2013), investment in new network connections to access information about available jobs (e.g., Calvó-Armengol, 2004; Galeotti and Merlino, 2014), job tenure (e.g. Loury, 2006), or how networks perpetuate inequality (e.g. Montgomery, 1991b; Calvo-Armengol and Jackson, 2004). The study that is closest to our paper is Montgomery (1991a) who presents a model explaining why workers who are well-connected fare better than other workers and why firms hiring through referrals tend to earn higher profits. Montgomery assumes that firms are restricted to make a single offer. We choose to develop our own model, designed to capture the key features of the experimental environment where firms can make multiple wage offers over time.[2]

[2]Moreover, in Montgomery (1991a) there is free entry by firms and thus firms earn zero profits. In our model (as in the experiment), there are fewer firms than workers, implying that firms have some market power. Because workers' types are private information, workers can also extract some of the surplus.

Numerous empirical studies find that employee referrals convey information about the quality of the referred workers (e.g. Schmutte, 2014; Pallais and Sands, 2016; Hensvik and Skans, 2016; Dustmann et al., 2015). Thus, the empirical evidence is largely consistent with the assumption we make in the experiment that referrals are informative. On the other hand, Beaman and Magruder (2012) conducted a lab-in-the-field experiment where subjects could refer members of their social networks, emphasizing a tradeoff between referring their own preferred individual or the most qualified individual for the job. Relatedly, Fafchamps and Moradi (2015) show based on recruitment data from the British colonial army in Ghana 1908-1923 that referrals by soldiers did not improve the unobserved quality of recruits. Employee referrals are not infallibly informative across contexts. We contribute to this empirical literature by demonstrating the informational value of referrals in a controlled environment with a clearly identifiable equilibrium. We also observe the counterfactual outcome when there is no social network.

Finally, our study contributes to the growing experimental literature on adverse selection. Kübler et al. (2008) study Spence's education game. Hoppe and Schmitz (2015) examine an adverse selection model where the uninformed party can offer suitable menus of contracts to achieve separation between the different types of the informed party. Mimra and Waibel (2017) vary the degree of market power of the uninformed party, ranging from monopoly to nonexclusive competition. A common finding is that the degree of equilibrium play tends to be high. Siegenthaler (2017) shows that cheap-talk can help alleviate adverse selection in decentralized markets. Bochet and Siegenthaler (2018) examine whether adverse selection can be mitigated through screening in bargaining environments but find that take-it-or-leave-it environments outperform environments with repeated offers. To the best of our knowledge, the present paper is the first experiment studying whether social networks can help promote welfare in markets with adverse selection.[3]

The paper proceeds as follows. In Section 2, we present the experimental design. Section 3 introduces the model and derives a set of hypotheses to guide our data analysis. In Section 4, we present the experimental results. Section 5 concludes.

2 The Experiment

The experiment consists of two treatments. In the *Baseline* treatment there are no social ties among workers. Thus, firms hire workers solely in a public posted-offer market. In the *Referral* treatment social networks are introduced such that each worker has a social tie to another worker. Thus, firms can make referral offers in stage 2 of the experiment. We first present the general

[3]Other experimental studies have focused on moral hazard in labor market settings. Fehr et al. (1993) and Fehr et al. (1997), in line with Akerlof (1982, 1984), developed the gift exchange game to capture contractual incompleteness in labor markets. They show that reciprocal fairness may overcome moral hazard. Brown et al. (2004) show that successful long-term relationships can lead to high effort levels and wage offers. Andreoni (2017) finds that selling goods with a "satisfaction guarantee" helps mitigate problems of moral hazard.

setting and then describe each treatment in more detail.

2.1 General Setting

At the beginning of the experiment, each participant is randomly assigned the role of a firm or a worker. There are 4 firms and 12 workers. Roles remain fixed throughout the 15 periods of the experiment. In each period participants interact in a market. Each market has 2 trading stages, where a stage lasts 2 minutes. Firms seek to hire a worker in each stage. The 12 workers are divided into 6 stage-1 workers (active only in stage 1) and 6 stage-2 workers (active only in stage 2). Moreover, workers have two different productivity levels. In particular, in each stage there are 3 low productivity (L-type) and 3 high productivity (H-type) workers. The productivity of a worker and whether they are active in stage 1 or stage 2 is randomly determined at the beginning of each of the 15 periods.

Firms cannot observe workers' types. A low productivity worker produces an output $P_L = 20$. A high productivity worker produces an output of $P_H = 60$. To reduce the risk of negative payoffs, we add a baseline productivity of $B = 20$ to the firms' earnings whenever a contract is concluded.[4] The payoff of a firm in a given stage is then:

$$\pi_F(\theta, w) = \begin{cases} P_\theta - w + 20 & \text{if a worker of type } \theta = \{L, H\} \text{ is hired} \\ 0 & \text{if no worker is hired} \end{cases} \tag{1}$$

where w represents the accepted wage offer. The total payoff of a firm per period equals the sum of payoffs in stage 1 and 2.

Workers with different productivity types have different reservation wages. An L-type worker has a reservation wage of $\lambda_L = 10$. An H-type worker has a reservation wage of $\lambda_H = 30$. The payoff of a worker is thus:

$$\pi_W(\theta, w) = \begin{cases} w & \text{if hired} \\ \lambda_\theta & \text{if not hired and of type } \theta = \{L, H\} \end{cases} \tag{2}$$

Notice how the different reservation wages make it difficult for firms to hire high productivity workers. Hiring such a worker requires an offer of at least 30, but this leads to a low payoff in case a low productivity worker accepts the offer.[5]

[4] As we will see, the baseline productivity guarantees firms a positive payoff when hiring an L-type worker at the H-type workers' reservation wage. Despite this, in the theoretical model we develop below, equilibrium market wages are often below the H-type workers' reservation wage. This leads to adverse selection.

[5] In the experiment workers were paid for both market stages. In the stage in which they were not active, they received the reservation wage. This was done to keep payoffs between firms and workers comparable.

2.2 *Baseline*

In the *Baseline* treatment, participants in the role of a firm make all wage offers in a public, posted-offer market. We will refer to such offers as *public offers*. Firms and workers can observe all public offers. The market lasts for at most 2 minutes or until all firms have hired a worker. Firms are completely free when or how many offers they would like to make. The set of allowed wage offers contains the integers between 0 and 60. Offers have to increase over time. More specifically, we use an improvement rule according to which every new offer has to be higher than the currently highest standing offer. In case there are no more standing offers (because each offer has been accepted by a worker), offers can again start at 0. The first worker who accepts a given offer is hired at the corresponding wage by the firm which made the offer. The remaining workers can still get hired by other firms.

2.3 *Referral*

The *Referral* treatment is identical to the *Baseline* treatment, except that firms have the additional option to hire workers through a simple social network. In stage 1, firms can hire workers only via public offers as in the *Baseline* setting. However, each stage-1 worker has a social tie to one of the stage-2 workers. Social ties are assigned among workers at the beginning of each trading period such that the probability of two linked workers being of the same type is 75%. The social ties affect the hiring opportunities firms have in the stage-2 market. Specifically, firms can make two types of offers in stage 2. On the one hand, they can still make public offers which are observed by everyone. In addition, they can hire a worker through an offer that is only received and observed by the stage-2 worker that has a social tie to the respective firm's stage-1 worker. We will refer to such offers as *referral offers*. Firms can make multiple offers of both types (public and referral offers) and these offers can be made simultaneously. That is, in the experiment firms and workers in stage 2 of the *Referral* treatment observe two panels with offers. In one panel, they can make/accept public offers and in the other panel they can make/accept referral offers made via the social ties. A screen shot of the experimental interface can be found in the Online Appendix.

2.4 Information

In both treatments the participants know the parameters of the experiment. That is, they know that there are 15 periods with 2 trading stages in each period, that there are 4 firms and 6 workers per trading stage, that among the 6 workers there are 3 low and 3 high productivity ones, and what the payoff functions are. In the *Referral* treatment, everyone knows at the start of the experiment that workers who are linked by a social tie have the same productivity level with a probability of 75%. Finally, firms and workers can observe all public offers that are made over the two minutes of a trading stage. A given referral offer can only be seen by a single worker.

2.5 Procedures

The experiment was run at Maastricht University. Participants were recruited via ORSEE (Greiner, 2015). There were 11 sessions, 5 for treatment *Baseline* and 6 for treatment *Referral*. Each session consisted of 16 participants (4 firms and 12 workers). The total number of participants was 176. We used the experiment software developed by Fischbacher (2007). At the beginning of each session participants were randomly assigned to a closed cubicle where they made decisions in privacy. The experiment instructions included a set of control questions to check participants' understanding. The instructions can be found in the Online Appendix. Sessions lasted 100 minutes or less. Payments averaged 19 Euros per participant including a show up fee of 5 Euros.[6]

3 Theoretical Background and Hypotheses

This section derives a set of hypotheses for each experimental treatment. We base the predictions on a model which we believe captures well the basic features of the experimental game described above.

3.1 Model

Consider a market with n_F firms, n_L low productivity workers and n_H high productivity workers. Workers' types are private information. We assume $P_H - \lambda_H > P_L - \lambda_L$, i.e. the gains from trade are larger with H-type workers. There are two stages and each firm can hire at most one worker per stage. In each stage the market opens with firms announcing a finite number of wages. Workers then choose which wage offers they would like to accept (potentially more than one offer). Firms and workers are matched from low to high wages.[7] In particular, market clearing starts at the lowest wage offer, say w_1. Among the workers who chose to accept w_1, one is hired at

[6]Typical for adverse selection settings, in our experiment it is possible that participants receive negative payoffs. This is problematic if a participant goes bankrupt, which is most likely to happen in early periods when participants haven't accrued profits yet and are more likely to make mistakes. We gave participants an additional endowment of 120 experimental points (conversion rate: 1 point = 0.0225 Euro cent) at the start of the experiment to address this issue (we did not have any bankruptcies).

[7]This is in line with the experimental setting where wages have to follow an improvement rule. Notice, however, that the model doesn't mirror the experimental setting one to one. In particular, in the model firms first choose a set of wage offers before workers choose which wages to accept. In the experiment, this process is dynamic and wage offers, rejections, and acceptances occur continuously.

random. Then, the next wage level w_2 is selected resulting in another firm-worker match. This process continues until all firms hired a worker or the highest wage level is reached.[8]

We allow for risk aversion. An agent's utility is given by

$$u(\pi) = 1 - e^{-\sigma\pi} \tag{3}$$

where π represents the payoffs defined in (1) and (2) and $\sigma > 0$ measures the level of (absolute) risk aversion.

3.2 Market Equilibrium in *Baseline*

A *market equilibrium* is reached if firms' wage offers and workers acceptance decisions maximize their respective expected utilities, given everyone else's behavior, and firms' beliefs about the expected quality of workers are correct at all wage levels.[9] We focus on the key predictions and refer to the appendix for a detailed analysis.

In the absence of social ties there is no difference between stage 1 and 2. Equilibria take the following form. Firms randomize between posting the H-type reservation wage $\{\lambda_H\}$ or the set of wages $\{w^*, \lambda_H\}$, where w^* is such that $\lambda_L < w^* < \lambda_H$ and depends on the specific equilibrium as well as on risk aversion.[10] The trade-off between $\{\lambda_H\}$ and $\{w^*, \lambda_H\}$ is apparent: hiring an H-type worker at a wage of λ_H generates large gains from trade but involves the risk of hiring an L-type worker in which case the wage w^* would have been preferable. H-type workers only accept high wages. L-type workers accept both wages w^* and λ_H. The reason they accept w^* is that at a wage of λ_H they face competition from H-type workers, i.e. some workers will remain unemployed. This also explains why firms always include offer λ_H in their wage schedule. A firm that offers $\{w^*, \lambda_H\}$ will hire at the wage of λ_H only if all L-type workers have already left the market.

Figure 1 illustrates the market equilibrium for the experimental parameters. The left-hand figure shows the average accepted offer. The right-hand figure shows the percentage of high productivity workers among all hired workers, where the total number of hires equals 4 (that is, firms always hire some worker in equilibrium). The solid lines in each figure depict the predictions for treatment *Baseline*. Notice that for low levels of risk aversion the equilibrium wage equals $\lambda_H = 30$ and the

[8] The model is inspired by Wilson (1980) who studies markets with adverse selection under different price setting conventions, including the one where firms make wage offers. Wilson assumes that market clearing starts at the highest wage and each firm can only make a single wage offer.

[9] Beliefs are required to be correct also at wages that are not offered in equilibrium. That is, we are interested in equilibrium outcomes that are robust to firms and workers experimenting with different (off-equilibrium) strategies (see Wilson, 1980; Mas-Colell et al., 1995).

[10] Notice that we assume $n_L + n_H \geq n_F > n_L$, as in the experiment. If $n_F \leq n_L$ firms are essentially monopsonists and offer either $\{\lambda_L\}$ or $\{\lambda_H\}$. If $n_F > n_L + n_H$, competition between firms drives offers above λ_H such that firms' expected profit is zero.

Figure 1: Equilibrium Outcomes

(a) Average Wage in Public Market

(b) Percentage of H-Type Hires

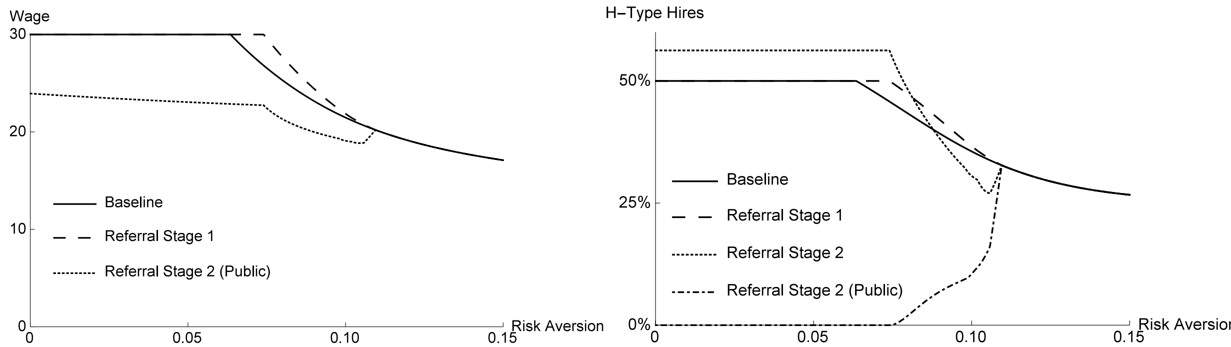

Theoretical predictions for the experimental parameters $n_F = 4$, $n_L = n_H = 3$, P_L=20, $P_H = 60$, $\lambda_L = 10$, $\lambda_H = 30$, $B = 20$, and different levels of risk aversion. Figure (a): average accepted offers in public market. Figure (b): percentage of H-type workers among all hires (the total number of hired workers equals 4). For stage 2 of the *Referral* treatment the dotted line shows the percentage of H-type workers hired combined for public and referral offers; the dash-dotted line shows the corresponding percentage for the public market only.

percentage of hired H-type workers is 50%. All 6 workers accept a wage of 30 and on average 2 workers of each productivity type get hired. As agents become more risk averse firms start randomizing between $\{w^*, \lambda_H\}$ and $\{\lambda_H\}$ exerting downward pressure on the average wage level. Two effects are at play: first, firms become more reluctant to make high offers and second, L-type workers' willingness to accept low wages increases. The latter happens because the risk of not getting hired plays a larger role in L-type workers' utility calculations. As a result, the percentage of hired H-type workers falls towards 25% for higher levels of risk aversion implying a stronger adverse selection effect. Notice that from a social welfare perspective we would like to see many H-type hires, because the gains from trade with an H-type equal $60 - 30 = 30$ while for an L-type they are only $20 - 10 = 10$.[11]

3.3 Market Equilibrium in *Referral*

We next ask whether the possibility to make referral offers alleviates adverse selection. In stage 1, the hiring process happens in a public market, as in the *Baseline* model. In stage 2, firms can still hire in the public market but importantly they can also hire through referral offers using the social ties of the hired period-1 workers (with a 75% probability that workers linked by a social tie have the same productivity type).

The dashed lines in figure 1 depict the equilibrium outcomes in stage 1 of the *Referral* treatment.

[11]In the Online Appendix we show that the Walrasian equilibrium outcome exhibits inefficiencies much as the equilibrium of our model (at equilibrium 3 L-type workers and 1 H-type worker are hired, while efficiency requires 3 H-type hires and 1 L-type hire).

We can see that average wages tend to be higher in stage 1 of treatment *Referral* than in *Baseline* and as a result more H-type workers are hired. Firms are willing to make higher offers in *Referral* because hiring an H-type worker has the additional benefit of getting access to a valuable social tie that can be used in stage 2. At equilibrium, up to a risk parameter of about 0.1, firms that have hired an H-type worker in stage 1 always use referral offers of 30 in stage 2 in order to hire another H-type with a probability of 75%. For higher levels of risk aversion, there is a sharp decrease in the use of social ties until the differences between the treatments disappears. Then, even a 75% probability of hiring an H-type worker is too risky and firms focus on hiring L-type workers in the public market.

The dotted lines in figure 1 depict outcomes in stage 2. We can see that due to the referral offers the percentage of H-type hires is higher in the *Referral* treatment for most levels of risk aversion. There is however an offsetting effect in that the referral hires reduce the number of H-type workers active in the public markets, illustrated by the lower wages in figure 1a. In other words, most or all H-type hires in stage 2 of treatment *Referral* occur through referral offers. The overall effect of referrals on average productivity is thus ambiguous, although for most levels of risk aversion the productivity gain due to referral hires dominates the negative effect in the public market.

We summarize the discussion:

Hypothesis 1 (Public Market Wages): Wages in the stage-1 public market tend to be higher in *Referral* than in *Baseline*. In contrast, wages in the stage-2 public market tend to be lower in *Referral* than in *Baseline*.

Hypothesis 2 (Referral Offers): Firms are more likely to use referral offers in stage 2 if they hired an H-type rather than an L-type worker in stage 1. Referral offers are equal to the H-type workers' reservation wage of 30.

Hypothesis 3 (Hires): The percentage of H-type hires in stage 1 and stage 2 is higher in *Referral* than in *Baseline*. This is because (i) hiring a stage-1 H-type worker in *Referral* provides access to a social tie which (ii) can then be used to hire an H-type worker in stage 2 with a large probability.

4 Results

The discussion of the results is organized as follows. We first discuss offers and accepted wages in section 4.1, then examine employment rates and efficiency in section 4.2, and finally take a more detailed look at the strategies adopted on the individual level and the role of risk aversion in section 4.3. The experiment consisted of 15 periods with a stage-1 and a stage-2 market in each period. In the following we focus on the final 10 periods (periods 6 - 15) where behavior is

Figure 2: Wage Offers

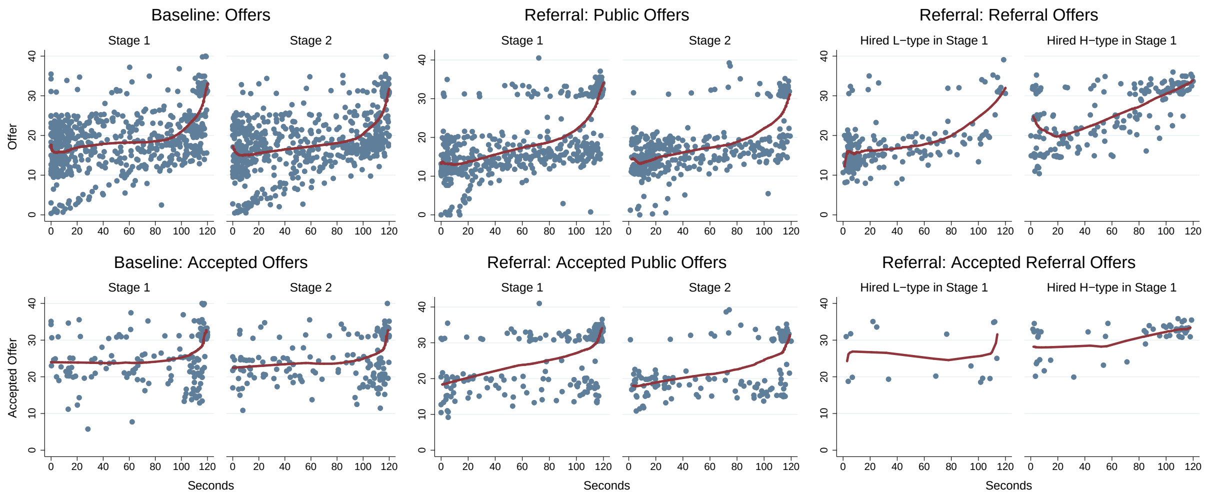

The first row shows all offers (accepted or not) over time by treatment, stage and public versus referral offers in *Referral*. The second row shows accepted offers (wages). Graphs include smoothed values from locally weighted regressions.

more stable and can thus be interpreted as equilibrium play.[12]

4.1 Offers and Wages

We start with a general look at firms' wage offers over the 2 minutes of a market stage. The first row in figure 2 shows all offers for the two treatments and stages. There is a clear pattern which we summarize in the following result.

Result 1: *In the public markets of both treatments most firms engage in screening behavior. They start by offering wages targeted at L-types, usually around 20, followed by a sharp increase in offers above the H-type workers' reservation wage of 30 during the final 20 seconds of a market.*

Looking at the second row of figure 2 shows that while acceptances occur continuously throughout the 120 seconds and some L-type workers accept offers below 30 there is also a concentration of hires in the last 20 seconds. Figure 2 is meant to convey the general offer patterns observed in the experiment. We next provide a more formal analysis starting with the behavior in stage 1.

Result 2: *In stage 1, firms in Referral are more likely to offer wages above 30 than firms in Baseline. In line with this, the percentage of L-type workers earning a wage above 30 is larger in Referral than in Baseline.*

[12]In the Online Appendix we show that the differences in behavior between *Baseline* and *Referral* are stable or, if anything, become more pronounced over time, indicating that the treatment differences are robust to experience and learning.

Table 1: Key Summary Statistics

Treatment:	**Baseline**				**Referral**							
	Stage 1		Stage 2		Stage 1		Stage 2 Overall		Stage 2 Public Offers		Stage 2 Referral Offers	
Type	H	L	H	L	H	L	H	L	H	L	H	L
Offers	19.3		19.3		19.1		19.6		17.8		21.5	
Offers $\geq$ 30	14%		13%		26%		21%		16%		26%	
Wages	30.9	24.3	31.4	23.6	31.1	24.0	32.0	21.6	31.2	20.8	32.6	25.3
Wages $\geq$ 30	88%	30%	92%	23%	92%	44%	95%	23%	92%	22%	100%	28%
Hires (No.)	0.9	2.7	0.8	2.8	1.3	2.4	1.3	2.4	0.6	2	0.7	0.4
Hires (%)	25%	75%	22%	78%	35%	65%	34%	66%	23%	77%	61%	39%
Efficiency	53%		52%		63%		62%		–		–	

Data from periods ≥ 6. By row: average offers (accepted or rejected), % of offers above 30, average accepted wages, % of accepted wages above 30, average number of workers hired, % of H versus L-types hires, and efficiency levels (% of realized gains from trade relative to the first-best outcome where 3 H-type and 1 L-type workers are hired).

Table 2: Regressions on Offers and Wages in Public Markets

	Stage 1				Stage 2			
Dep. Var:	Offer $\geq$ 30		L-type Wage $\geq$ 30		Offer $\geq$ 30		L-type Wage $\geq$ 30	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Referral	0.12***	0.00	0.13*	-0.04	0.03	-0.01	-0.01	-0.04
	(0.03)	(0.03)	(0.07)	(0.16)	(0.02)	(0.02)	(0.07)	(0.05)
Late (> 60 sec)		0.22***		0.17		0.21***		0.23***
		(0.06)		(0.16)		(0.05)		(0.04)
Referral × Late		0.22***		0.28		0.15**		0.10
		(0.08)		(0.19)		(0.07)		(0.10)
Constant	0.01	0.06**	0.20*	0.20	0.09***	0.05	0.18	0.05
	(0.04)	(0.28)	(0.12)	(0.16)	(0.02))	(0.04)	(0.15)	(0.08)
Wald Test[(a)]		$p = 0.002$		$p < 0.001$		$p = 0.035$		$p = 0.511$
Observations	1197	1197	277	277	1047	1047	260	260

Linear random effects models with standard errors clustered on 11 experimental sessions in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Bootstrap standard errors yield identical results. Period dummies included in all regressions. (a) Wald chi-squared test for the hypothesis $Referral + Referral \times Late = 0$. The reference group is treatment *Baseline* in models (1), (3), (5), and (7), and the first 60 seconds of a market stage in treatment *Baseline* in the remaining models.

The first two rows of table 1 present the average wage offers and the percentage of offers above 30. Notice that average offers are similar across the two treatments. More importantly, the percentage of offers exceeding the reservation wage of the H-type worker is larger in *Referral* than in *Baseline*, 26% versus 14% in stage 1 (Mann-Whitney U, $p = 0.017$). Median accepted wages in stage 1 are also significantly larger in treatment *Referral* than in *Baseline* (Mann-Whitney U, $p = 0.039$). To understand why there is no difference in terms of average wage offers between the treatments, note that there is range of wages between X and 30 that firms should never offer. Rather than offering a wage above X but below 30, firms prefer to offer 30 to get a chance at hiring an H-type. The maximum offer X below 30 is lower in *Referral* than in *Baseline*: L-type workers in *Referral* are willing to accept lower wages because some firms hire via referral offers and thus the firm-worker ratio in the public market decreases relative to the *Baseline* setting. That offers below 30 are lower in *Baseline* can indeed be observed in the experiment (see figure 2). So, despite the larger number of offers above 30 in *Referral*, the fact that offers below 30 are lower than in *Baseline* implies that the average offer is similar.

Further insights are given in table 2 where we present random effects regressions. The dependent variable in models (1) and (2) is the percentage of offers above 30 in stage 1. The explanatory variables are the treatment and a dummy *Late* equaling one if an offer happened in the second half of a market (i.e. after 60 seconds). In stage 1, offers above 30 are significantly more likely in *Referral* and the effect is exclusively due to behavior in the second half of a market—the hypothesis $Referral + Referral \times Late = 0$ is rejected by a Wald chi-squared test with a p-value of 0.002. In line with these observations, the percentage of L-type workers earning a wage above 30 is 44% in *Referral* and 30% in *Baseline*, see row 4 in table 1. The higher likelihood of observing high offers in *Referral* imply less success at screening L-type workers (in stage 1). Regressions (3) and (4) in table 2 confirm that the effect is significant, again mainly due to behavior at seconds 61 to 120 ($p < 0.001$).

Result 3: *In stage 2, firms in Referral are more likely to offer wages above 30 than firms in Baseline. However, the percentage of L-type workers earning a wage above 30 is not significantly different in Referral and Baseline.*

A look at table 1 shows that the percentage of offers exceeding the reservation wage of the H-type worker is 21% in *Referral* and 13% in *Baseline* (Mann-Whitney U, $p = 0.017$). Regressions (5) and (6) in table 2 reveal that the difference is mainly due to the behavior in the second half of a market ($p = 0.035$). In contrast to stage 1, however, screening success in the public market does not differ across treatments in stage 2. The percentage of L-type workers earning a wage above 30 is 22% in *Referral* and 23% in *Baseline*, and there is no significant difference either overall or when focussing on the second half of the market, see models (7) and (8) in table 2.

Remarkably, the observation that the percentage of L-type workers hired at a wage above 30 is larger in *Referral* than *Baseline* in stage 1 but not in stage 2 is fully in line with the theoretical predictions of our model. In stage 1 the reason firms are more likely to make offers exceeding 30 is that hiring an H-type worker provides access to a valuable social tie. Thus, less screening is required before high offers become preferable for firms. In stage 2, increased competition among L-type workers in the public market increases their willingness to accept low offers (since some firms hire workers through referral offers and don't participate in the public market). Firms thus find it easier to screen L-type workers.

Results 2 and 3 confirm our first hypothesis on public market wages. The results also show our design succeeded in creating an environment where information asymmetries lead to low wages, impeding trading opportunities for high productivity workers. We next take a look at referral offers to see if treatment *Referral* was able to mitigate adverse selection.

Result 4: *In stage 2 of treatment Referral, firms mainly use referral offers after hiring an H-type worker in stage 1. Most referral offers exceed the H-type workers' reservation wage of 30.*

The right-most panels in figure 2 show referral offers in stage 2 of treatment *Referral* in the first row and the corresponding accepted wages in the second row. In the second row, we can see that firms that hired an L-type worker in stage 1 only rarely hire workers through referral offers, while there are more referral hires for firms that hired an H-type worker in stage 1 (and a large majority of them at wages above 30). In line with this, table 1 shows that referral offers tend to be higher than public offers. To confirm that these effects are significant, we ran random effects regressions (cluster-robust standard errors). The results are reported in the Online Appendix. We summarize the results: Firms are on average 35.3% points ($p < 0.001$) more likely to make referral offers if they previously hired an H-type worker in stage 1. Further, on the intensive margin, referral offers of firms that hired an H-type in stage 1 are on average 7.72 points ($p < 0.001$) higher and 26.6% points ($p < 0.001$) more likely to exceed 30 than referral offers of firms that hired an L-type worker in stage 1. Overall, firms that hired an H-type worker in stage 1 are 48% points more likely to hire an H-type worker in stage 2 than other firms ($p < 0.001$). These findings are in line with our second hypothesis.

4.2 Hires and Efficiency

We are now ready to state our main result about the ability of referrals to increase the number of H-type hires and promote efficiency. All p-values in this section are based on non-parametric Mann-Whitney U tests with session averages as independent observations.

Result 5: *The number and percentage of H-type workers hired in both stage 1 and 2 of treatment Referral is higher than in Baseline. As a result, overall efficiency is higher in the Referral treatment.*

Table 1 presents the average number of H and L-type hires per market stage and the percentage of

Figure 3: Firm Profits

Histogram of firm profits separated by whether accepted wages exceed H-type workers' reservation wage of 30.

H versus L-type hires for each treatment and stage (third and second row from below). Summing up H and L-type hires shows that in both treatments and stages the total number of hired workers is close to the maximum of 4. The total number of hires is not significantly different between *Referral* and *Baseline* ($p > 0.404$). However, in *Referral* more H-type workers were hired than in *Baseline* ($p = 0.042$ in stage 1, $p = 0.016$ in stage 2, $p = 0.021$ overall). The opposite holds for the number of L-type hires ($p = 0.098$ in stage 1, $p = 0.020$ in stage 2, $p = 0.053$ overall). As a result, in *Referral* the percentage of H-type workers among all hires is 35% in stage 1 and 34% in stage 2, while in *Baseline* the corresponding percentages are only 25% in stage 1 and 22% in stage 2 ($p = 0.044$ in stage 1, $p = 0.006$ in stage 2, $p = 0.017$ overall).

We measure efficiency as realized gains from trade divided by maximum possible gains from trade. The latter correspond to $n_H(P_H - \lambda_H) + (n_F - n_H)(P_L - \lambda_L) = 3 * (60 - 30) + 1 * (20 - 10) = 100$ and require 3 H-type and 1 L-type hires. The last row in table 1 shows the realized efficiency levels. The availability of referral offers increased efficiency by 10% points ($p = 0.028$). Two channels are at work. In stage 1 efficiency is increased by firms raising wages in anticipation of the benefit of H-type workers' social ties. In stage 2 efficiency is increased directly by firms hiring H-type workers through referral offers. Taken together, these observations confirm hypothesis 3 and constitute clear evidence for the value of social networks in alleviating adverse selection.

4.3 Individual Level Analysis: Risk Aversion and Profits

This section takes a more detailed look at participants' strategies. Given the offer patterns observed in figure 2, it is useful to focus on two strategies for firms. One strategy is to target L-type workers, offering a wage around 20 even towards the end of the 2 minutes of a market stage. Following our model, we would expect that this strategy is primarily used by risk averse firms. The second strategy a firm could follow is to attempt at trade with H-type workers. This is risky, because a firm may end up paying a wage above 30 to a low productivity worker.

Figure 3 depicts the distribution of firm profits for the two strategies. If the accepted offer was below the H-type workers' reservation wage of 30 (transparent bars) profits are centered around 20. Hiring at wages above 30 (black bars) either leads to profits above 45 or, if an L-type worker is hired, to profits below 10. The probability of hiring an H-type when the wage is above 30 is 51% in *Baseline* and 54% in *Referral*. One way to look at the data is thus to say that participants in the role of a firm essentially choose between a certain, intermediate profit or a gamble between low and high profits.

The regressions in table 3 help us understand better the parameters dictating this gamble. Model (1) shows that in *Baseline* the average firm profit is 7.25 points higher when the accepted offer exceeds 30 than when it is below 30. This represents a risk premium. Model (2) adds as an explanatory variable a dummy for the worker's productivity type. For offers above 30, firms on average earn 12.74 points less than firms making low offers if the hired worker is an L-type and 26.76 (adding the coefficients of *Offer* $\geq$ 30 and *H-type*) more if the hired worker is an H-type. Accounting for the constant, this means that firms could either pick a certain payoff of 19.97 (when making offers around 20) or a 50-50 bet between receiving a payoff of 7.23 and 46.73. Indifference between the certain payoff and the gamble in *Baseline* implies a CARA coefficient of 0.044 (or a CRRA coefficient of 0.89).[13,14] Columns (3) and (4) tell a similar story for the *Referral* treatment, the main differences being that offering high wages is particularly profitable for referral offers.

Workers face a different trade-off than firms. In each market (with 4 firms and 6 workers) at least 2 workers will not be hired. L-type workers can always choose to accept a wage below 30. The alternative strategy is to hold out for higher wages, risking to remain unemployed. Models (5) to (8) in table 3 provide information about the payoff consequences of both strategies. Models (5) and (7) show that in the public markets there is a risk premium of 5.2 in *Baseline* and 3.97 in *Referral*, that is, holding out for a high offer promises a larger expected profit. Model (6) tells us that accepting an offer below 30 yields on average a payoff of 21.14 (note that workers who accept low wages are almost always hired). Holding out for high offers either results in a payoff of 32.82 if hired or a payoff of 10.14 if not hired. The latter is just the L-type workers' reservation wage. More specifically, the probability to get hired for an L-type that didn't accept offers below 30 was 74%. Hence, the gamble for an L-type worker is to earn 32.82 with probability 0.74 and earn 10.14 with probability 0.26 or choose the safe option and almost certainly earn 21.14. The implied risk coefficients are 0.098 for CARA and 1.58 for CRRA. Similar conclusions apply to the *Referral* treatment.

[13] The risk coefficients are broadly in line with the literature on risk elicitation, see e.g. Holt and Laury (2002)'s high-stakes treatment and Dave et al. (2010).

[14] In the Online Appendix we show that in line with the above results subjects that are more likely to target H-types on average earn more than subjects who tend to make only low offers.

Table 3: Regressions on Firm and L-type Worker Profits

Dep. Var:	Firm Profit				L-type Worker Profit			
	Baseline		Referral		Baseline		Referral	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Offer $\geq$ 30	7.25**	-12.74***	7.08***	-14.89***	5.20*	11.68***	3.97***	14.77***
	(3.32)	(1.49)	(1.33)	(0.66)	(2.75)	(1.64)	(1.13)	(0.85)
H-type		39.50***		40.04***				
		(0.42)		(0.14)				
Not Hired						-22.68***		-22.36***
						(0.93)		(0.33)
Referral Offer (RO)			-6.47***	-4.30***			5.74***	4.58***
			(1.65)	(0.60)			(1.21)	(0.78)
Offer $\geq$ 30 $\times$ RO			17.19***	4.01***			7.01***	-3.13**
			(1.47)	(0.65)			(1.49)	(1.34)
Constant	20.88***	19.97***	24.38***	22.10***	20.88***	21.14***	18.77***	18.51***
	(2.80)	(1.42)	(2.42)	(0.63)	(0.69)	(1.03)	(1.80)	(0.58)
Observations	349	349	429	429	299	299	356	356
Subjects (Sessions)	20 (5)	20 (5)	24 (6)	24 (6)	60 (5)	60 (5)	72 (6)	72 (6)

Linear mixed effects models with subject and session random intercept; standard errors clustered on sessions in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Period and stage dummies included in all regressions. The reference groups for *Baseline* are offers below 30 in models (1) and (5), offers below 30 with an L-type worker in model (2), and offers below 30 if hired in model (6). The reference group for *Referral* are public offers below 30 in models (3) and (7), public offers below 30 with an L-type worker in model (4), and public offers below 30 if hired in model (8).

Figure 4: Availability of Workers and Hires Above H-Type Workers' Reservation Wage

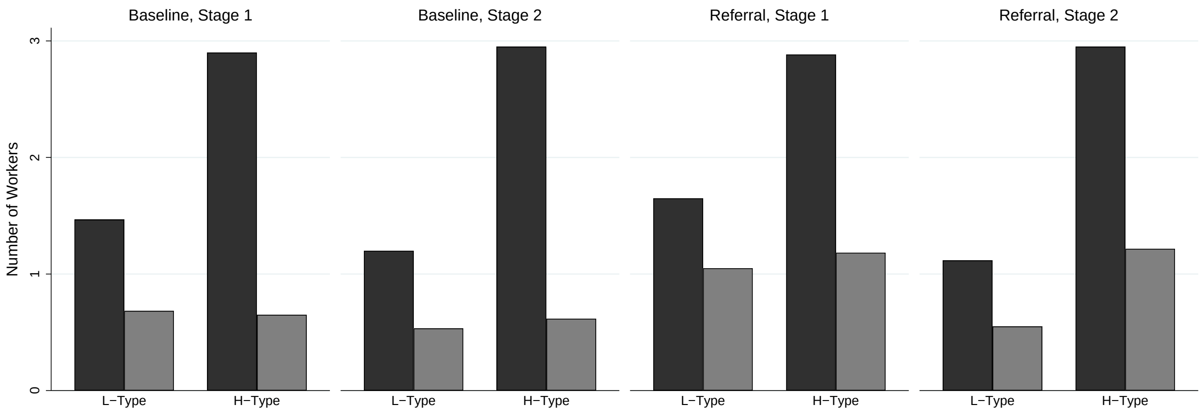

Black bars show the average number of workers (separated by treatment and productivity type) who are still available for hire when wage offers are at or above H-type workers' reservation wage of 30. Grey bars show the average number of workers who are hired at wages at or above 30.

The degree of risk aversion seems to be much higher for workers than for firms. For example, the derived CARA coefficient is 0.044 for firms and 0.098 for workers. We offer two explanations for this phenomenon. First, the nature of the gamble differs between firms and workers. Specifically, a worker may dislike the outcome in which she is not hired beyond the immediate payoff consequences. Such a "joy of being employed" would be reminiscent of the "joy of winning" in the auction literature (e.g. Dohmen et al., 2011).

Our second explanation has to do with the fact that L-type workers are much more likely to accept a wage above 30 than H-type workers. The black bars in figure 4 show the number of workers who are still available to be hired when wage offers have reached the H-type workers' reservation wage of 30. Notice that at such wages all three H-type workers are still available to be hired, but more than half of the L-type workers have already accepted a lower wage. The grey bars show the number of workers who are eventually hired at a wage above 30. Interestingly, in treatment *Baseline* the number of L-type and H-type workers hired at such wages is the same, despite the fact that there are many more H-type workers available for hire. Intuitively, H-type workers can always consume their outside option of 30, while L-type workers are much more eager to accept offers around 30. In other words, L-type workers tend to be much quicker than H-type workers in accepting offers of, say, 31 or 32. If L-type workers don't anticipate this effect, they may underestimate the probability of being hired when holding out for a high wage, in which case we would overestimate their risk parameters. Panel 3 in figure 4 shows that similar conclusions hold for stage 1 of treatment *Referral*. In panel 4 (stage 2 of *Referral*), the number of H-type workers hired at wages above 30 is larger than the number of L-type workers, which arises because of the referral offers targeted primarily at H-type workers.

5 Conclusion

We investigate in a laboratory experiment whether firms use social networks of their employees to alleviate information asymmetry in labor markets. Employee referrals are informative about the productivity of prospective workers because individuals that are part of the same social network tend to have similar abilities. We manipulate the existence of social networks and therefore the availability of hiring via employee referrals in two different treatments. In the *Referral* treatment there are social ties between workers that firms can use to make referrals offer (a referral offer is received only by a worker that has a tie to the firm's existent employee). In our *Baseline* treatment there are no social ties among workers and thus firms can only hire in a public market.

The experimental results provide strong support for our theoretical hypotheses. We find that social networks alleviate adverse selection effects. More high productivity workers are hired and efficiency is higher in our *Referral* treatment than in the *Baseline* where there is no social network. Remarkably, we find that social networks lead to higher wages not only for workers hired via referrals but may do so also in the public market. The reason is that higher wages increase the probability of hiring high productivity workers, which will allow firms to access these workers' valuable social ties in the future.

On a more general level, our study contributes to the question of how to promote efficiency in markets impaired by asymmetric information. Our results are encouraging, as participants effectively use networks to improve market outcomes. But we also observe that in most of our markets initial offers were low, resulting in many low productivity hires and suggesting that in dynamic environments adverse selection is a particularly hard problem to overcome. Montgomery (1991b) and Calvo-Armengol and Jackson (2004) present network models where individuals who are not well-connected are at a disadvantage and thus networks can perpetuate inequality. Our positive evaluation of referral offers would need to be revised if their informational value would come at the cost of an increase in inequality. Such a view would be premature, however. Social groups that are less well-connected (e.g., first-time job seekers or immigrants) are often the same groups about whom there exists little objective information. Employee referrals, as one source of credible information, should thus be particularly effective for such groups. Exploring which of the two effects dominates is a fruitful avenue for future research.

References

Akerlof, George A., "The market for "lemons": Quality uncertainty and the market mechanism," *The Quarterly Journal of Economics*, 1970, pp. 488–500.

Akerlof, George A, "Labor contracts as partial gift exchange," *The Quarterly Journal of Economics*, 1982, *97* (4), 543–569.

_ , "Gift exchange and efficiency-wage theory: Four views," *The American Economic Review*, 1984, *74* (2), 79–83.

Andreoni, James, "Satisfaction Guaranteed: When Moral Hazard meets Moral Preferences," *American Economic Journal: Microeconomics*, 2017, *74* (5), 1365–1384.

Beaman, Lori A., "Social networks and the dynamics of labour market outcomes: Evidence from refugees resettled in the US," *The Review of Economic Studies*, 2011, *79* (1), 128–161.

Beaman, Lori and Jeremy Magruder, "Who gets the job referral? Evidence from a social networks experiment," *The American Economic Review*, 2012, *102* (7), 3574–3593.

Bochet, Olivier and Simon Siegenthaler, "Better Later than Never? An Experiment on Bargaining under Adverse Selection," *International Economic Review*, 2018, *59* (2), 947–972.

Boorman, Scott A, "A combinatiorial optimization model for transmission of job information through contact networks," *The Bell Journal of Economics*, 1975, pp. 216–249.

Brown, Martin, Armin Falk, and Ernst Fehr, "Relational contracts and the nature of market interactions," *Econometrica*, 2004, *72* (3), 747–780.

Brown, Meta, Elizabeth Setren, and Giorgio Topa, "Do informal referrals lead to better matches? Evidence from a firms employee referral system," *Journal of Labor Economics*, 2016, *34* (1), 161–209.

Burks, Stephen V., Bo Cowgill, Mitchell Hoffman, and Michael Housman, "The value of hiring through employee referrals," *The Quarterly Journal of Economics*, 2015, *130* (2), 805–839.

Calvó-Armengol, Antoni, "Job contact networks," *Journal of Economic Theory*, 2004, *115* (1), 191–206.

Calvo-Armengol, Antoni and Matthew O Jackson, "The effects of social networks on employment and inequality," *The American Economic Review*, 2004, *94* (3), 426–454.

Calvó-Armengol, Antoni and Yves Zenou, "Job matching, social network and word-of-mouth communication," *Journal of Urban Economics*, 2005, *57* (3), 500–522.

Currarini, Sergio, Matthew O. Jackson, and Paolo Pin, "An economic model of friendship: Homophily, minorities, and segregation," *Econometrica*, 2009, *77* (4), 1003–1045.

Dave, Chetan, Catherine C. Eckel, Cathleen A. Johnson, and Christian Rojas, "Eliciting risk preferences: When is simple better?," *Journal of Risk and Uncertainty*, 2010, *41* (3), 219–243.

Dohmen, Thomas, Armin Falk, Klaus Fliessbach, Uwe Sunde, and Bernd Weber, "Relative versus absolute income, joy of winning, and gender: Brain imaging evidence," *Journal of Public Economics*, 2011, *95* (3-4), 279–285.

Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker, "Referral-based job search networks," *The Review of Economic Studies*, 2015, *83* (2), 514–546.

Fafchamps, Marcel and Alexander Moradi, "Referral and job performance: Evidence from the Ghana colonial army," *Economic Development and Cultural Change*, 2015, *63* (4), 715–751.

Fehr, Ernst, Georg Kirchsteiger, and Arno Riedl, "Does fairness prevent market clearing? An experimental investigation," *The Quarterly Journal of Economics*, 1993, *108* (2), 437–459.

_ , Simon Gächter, and Georg Kirchsteiger, "Reciprocity as a contract enforcement device: Experimental evidence," *Econometrica: Journal of the Econometric Society*, 1997, pp. 833–860.

Fischbacher, Urs, "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 2007, *10* (2), 171–178.

Galenianos, Manolis, "Hiring through referrals," *Journal of Economic Theory*, 2014, *152*, 304–323.

Galeotti, Andrea and Luca Paolo Merlino, "Endogenous job contact networks," *International Economic Review*, 2014, *55* (4), 1201–1226.

Granovetter, Mark, "Economic action and social structure: The problem of embeddedness," *American Journal of Sociology*, 1985, *91* (3), 481–510.

_ , *Getting a job: A study of contacts and careers*, The University of Chicago Press, 1995.

Greiner, Ben, "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association*, 2015, *1* (1), 114–125.

Heath, Rachel, "Why do firms hire using referrals? Evidence from Bangladeshi Garment Factories," *Journal of Political Economy*, forthcoming.

Hensvik, Lena and Oskar Nordström Skans, "Social networks, employee selection, and labor market outcomes," *Journal of Labor Economics*, 2016, *34* (4), 825–867.

Holt, Charles A. and Susan K. Laury, "Risk aversion and incentive effects," *The American Economic Review*, 2002, *92* (5), 1644.

Hoppe, Eva I. and Patrick W. Schmitz, "Do sellers offer menus of contracts to separate buyer types? An experimental test of adverse selection theory," *Games and Economic Behavior*, 2015, *89*, 17–33.

Ioannides, Yannis M. and Linda D. Loury, "Job information networks, neighborhood effects, and inequality," *Journal of Economic Literature*, 2004, *42* (4), 1056–1093.

Jackson, Matthew O., "An overview of social networks and economic applications," in "The Handbook of Social Economics," Vol. 1, North Holland Press Amsterdam, 2010, pp. 511–585.

_ , *Social and economic networks*, Princeton University Press, 2010.

Kübler, Dorothea, Wieland Müller, and Hans-Theo Normann, "Job-market signaling and screening: An experimental comparison," *Games and Economic Behavior*, 2008, *64* (1), 219–236.

Kugler, Adriana D., "Employee referrals and efficiency wages," *Labour Economics*, 2003, *10* (5), 531–556.

Loury, Linda Datcher, "Some contacts are more equal than others: Informal networks, job tenure, and wages," *Journal of Labor Economics*, 2006, *24* (2), 299–318.

Mas-Colell, Andreu, Michael Dennis Whinston, Jerry R Green et al., *Microeconomic theory*, Vol. 1, Oxford university press New York, 1995.

McPherson, Miller, Lynn Smith-Lovin, and James M Cook, "Birds of a feather: Homophily in social networks," *Annual Review of Sociology*, 2001, *27* (1), 415–444.

Mimra, Wanda and Christian Waibel, "(Non)Exclusive Contracting Under Adverse Selection: An Experiment," 2017. Working paper available at SSRN: https://ssrn.com/abstract=3037097.

Montgomery, James D., "Social networks and labor-market outcomes: Toward an economic analysis," *The American Economic Review*, 1991, *81* (5), 1408–1418.

_ , *Social networks and persistent inequality in the labor market*, Center for Urban Affairs and Policy Research, 1991.

Mortensen, Dale T and Tara Vishwanath, "Personal contacts and earnings: It is who you know!," *Labour Economics*, 1994, *1* (2), 187–201.

Munshi, Kaivan, "Networks in the modern economy: Mexican migrants in the US labor market," *The Quarterly Journal of Economics*, 2003, *118* (2), 549–599.

Pallais, Amanda and Emily Glassberg Sands, "Why the referential treatment? Evidence from field experiments on referrals," *Journal of Political Economy*, 2016, *124* (6), 1793–1828.

Pellizzari, Michele, "Do friends and relatives really help in getting a good job?," *Industrial Labor & Relations Review*, 2010, *63* (3), 494–510.

Pissarides, Christopher A, *Equilibrium unemployment theory*, MIT press, 2000.

Rees, Albert, "Information networks in labor markets," *The American Economic Review*, 1966, *56* (1/2), 559–566.

Schmutte, Ian M., "Job referral networks and the determination of earnings in local labor markets," *Journal of Labor Economics*, 2014, *33* (1), 1–32.

Siegenthaler, Simon, "Meet the lemons: An experiment on how cheap-talk overcomes adverse selection in decentralized markets," *Games and Economic Behavior*, 2017, *102*, 147–161.

Topa, Giorgio, "Social interactions, local spillovers and unemployment," *The Review of Economic Studies*, 2001, *68* (2), 261–295.

_ , "Labor markets and referrals," in "Handbook of Social Economics," Vol. 1, Elsevier, 2011, pp. 1193–1221.

Wilson, Charles, "The nature of equilibrium in markets with adverse selection," *The Bell Journal of Economics*, 1980, pp. 108–130.

Zenou, Yves, "Social interactions and the labor market," *Revue d'économie politique*, 2013, *123* (3), 307–331.

A Equilibrium Characterization

This appendix derives the market equilibria. We will construct the symmetric equilibrium. We also focus on the market structure relevant to the experiment where $n_S \equiv n_L + n_H \geq n_F > n_L$. Moreover, $P_H - \lambda_H > P_L - \lambda_L$. Both the derivation of asymmetric equilibria and for equilibria in other market structures follow reasoning analogous the discussion below. Let $\hat{n}_F \equiv n_F - 1$ and similarly for workers.

A.1 Equilibrium in *Baseline*

It is sufficient to discuss equilibrium behavior in stage 1 (identical predictions apply to stage 2). Notice that $n_F \leq n_S$ implies that the highest offer is λ_H. Moreover, all firms always include λ_H in their set of wage offers. We will show below that at equilibrium L-type workers accept offers $w^* < \lambda_H$ with probability (w.p.) 1. Thus, when offering a wage $w^* < \lambda_H$, it is always profitable to also offer λ_H, as the latter offer will only be accepted if no L-type workers are left in the market.

Consider next the following equation for L-type workers:

$$u_L(w) = \frac{n_F - q}{n_S - q} u_L(\lambda_H) + \frac{n_S - n_F}{n_S - q} u_L(\lambda_L). \tag{A.1}$$

If $q = \hat{n}_L$, the wage $\underline{w}$ solving (A.1) represents a lower bound for acceptable wages. For any lower wage, an L-type worker would prefers to wait for the offer λ_H, even if $q = \hat{n}_L$ firms and other L-type workers have left the market and the probability to be hired at a wage of λ_H is only $(n_F - \hat{n}_L)/(n_S - \hat{n}_L)$. Similarly, the wage $\bar{w}_1$ solving (A.1) for $q = 0$ yield the wage level above which L-type workers accept with probability 1.

On the firms' side, an upper bound for wage offers below λ_H is reached if they would prefer offering λ_H immediately rather than hiring an L-type worker at a wage that exceed $\bar{w}_2$:

$$u_F(P_L + B - \bar{w}_2) = \Psi(\lambda_H). \tag{A.2}$$

where

$$\Psi(\lambda_H) = \frac{n_L}{n_S} u_F(P_L + B - \lambda_H) + \frac{n_H}{n_S} u_F(P_H + B - \lambda_H). \tag{A.3}$$

denotes a firm's expected utility when offering λ_H and all workers are still in the market.

Suppose that at equilibrium firms offer $\{\lambda_H\}$ with probability β^* and $\{w^*, \lambda_H\}$ with probability $1 - \beta^*$. We claim that $w^* < \lambda_H$ can be supported as an equilibrium wage only if $w^* \in [\underline{w}, \max(\bar{w}_1, \bar{w}_2)]$. We prove the claim:

- From (A.1) it follows directly that wages below $\underline{w}$ are always rejected. Further, we must have $w^* \leq \bar{w}_2$ for if not (A.2) implies that firms strictly prefer to offer $\{\lambda_H\}$. Finally, we must have $w^* \leq \bar{w}_1$. If not, firms would have an incentive to offer a lower wage. If $n_{FL} \leq n_L$, a firm is guaranteed to hire an L-type workers with an any offer $w > \bar{w}_1$. If $n_{FL} > n_L$ firms may not hire an L-type worker (if other firms offer more), but in this case the firm prefers to hire at the wage λ_H as it implies that they hire an H-type w.p. 1 (and earn strictly more due to the higher gains from trade). Notice that L-type workers must accept $w^* < \lambda_H$ with probability 0 or 1. If they accept with a positive probability less than 1, firms could slightly lower the wage to $w^* - \epsilon$, knowing that L-type workers would accept such an offer in return for a strictly higher probability to get hired if $n_{FL} \leq n_L$ (as wages are cleared from below) or they would hire an H-type at λ_H if $n_{FL} > n_L$.

There is more than one possible equilibrium wage level below λ_H. To see this, let $w_{FL} \leq \bar{w}_1$ be the solution to (A.1) for $q = \min(\hat{n}_L, n_{FL})$, where n_{FL} is the number of firms offering $\{w^*, \lambda_H\}$. If $w^* < w_{FL}$, L-type workers reject w^* given the realization of n_{FL}. If $w^* \in [w_{FL}, \bar{w}_1]$, L-type workers face a coordination problem: if $\hat{n}_L$ L-type workers accept w^* the remaining worker want to accept as well, but similarly if $\hat{n}_L$ L-type workers reject w^* so does the remaining one. The reason is that the risk of not being hired increases in $\hat{n}_L$, see (A.1). The threshold wage level at which L-type workers switch from accepting w^* to rejecting $w^* - \epsilon$ can thus be anywhere in $[w_{FL}, \bar{w}_1]$. At the threshold, firms also don't deviate to higher offers, because at equilibrium they are indifferent between offering $\{\lambda_H\}$ and $\{w^*, \lambda_H\}$, i.e. they don't want to hire an L-type worker at wage larger than w^*. Hence, the supportable wage levels can go as low as $w^* = \underline{w}$ (although the latter can be an equilibrium only if $\bar{w}_1 \leq \bar{w}_2$ holds, i.e. firms are not too eager to screen L-types).

Offering $\{\lambda_H\}$ yields an expected utility of

$$U_F(\{\lambda_H\}; \hat{n}_{FL}) = \frac{n_L - \min(\hat{n}_{FL}, n_L)}{n_S - \min(\hat{n}_{FL}, n_L)} u_F(P_L + B - \lambda_H) + \frac{n_H}{n_S - \min(\hat{n}_{FL}, n_L)} u_F(P_H + B - \lambda_H). \tag{A.4}$$

Offering $\{w^*, \lambda_H\}$ yields an expected utility of

$$U_F(\{w^*, \lambda_H\}; \hat{n}_{FL}) = \Psi(\lambda_H) \tag{A.5}$$

if $w^* < w_{FL}$ and of

$$U_F(\{w^*, \lambda_H\}; \hat{n}_{FL}) = \frac{\min(n_L, n_{FL})}{n_{FL}} u_F(P_L + B - w^*) + \left(1 - \frac{\min(n_L, n_{FL})}{n_{FL}}\right) u_F(P_H + B - \lambda_H) \tag{A.6}$$

if $w^* \geq w_{FL}$. The equilibrium probability β^* of offering $\{\lambda_H\}$ renders firms indifferent between $\{\lambda_H\}$ and $\{w^*, \lambda_H\}$:

$$\sum_{i=0}^{\hat{n}_F} (1 - \beta^*)^i (\beta^*)^{\hat{n}_F - i} \binom{\hat{n}_F}{i} \left[U_F(\{\lambda_H\}; i) - U_F(\{w^*, \lambda_H\}; i)\right] = 0. \tag{A.7}$$

The solutions to (A.4) - (A.7) for $w^* = \underline{w}$ and $w^* = \bar{w}_1$ give the relevant bounds for the minimum and maximum number of hired L-type workers. The equilibrium reported in Figure 1 is for $w^* = \bar{w}_1$. Note that if expression (A.7) exceeds 0 even for $\beta = 1$, the equilibrium value is $\beta^* = 1$ (the opposite case never occurs if $n_F > n_L$ and the gains from trade a larger with H-type workers).

A.2 Equilibrium in *Referral*

We assume that firms and workers in stage 2 can observe the number of firms hiring through referral offers. While in the experiment individuals weren't explicitly informed about referral hires they could observe other firms' behavior in the public market, i.e., a low activity in the public market indicates that many firms intend to or have already hired a worker through a referral offer. Denote the homophily parameter by $\alpha > \max(1/2, n_H/n_S)$. Without the condition on α the model wouldn't make sense as L-type workers would be more likely than H-type workers to have a social tie to an H-type.

We first show that firms that hired an L-type worker in stage 1 don't make a referral offer in stage 2.

Lemma A.1. Firms that have hired a stage-1 L-type worker do not benefit from the option to make referral offer in stage 2, i.e., their expected utility when hiring in the public market is larger than when hiring through referral offers.

Proof: Suppose that only firms with an H-type stage-1 worker make a referral offers and in stage 1 all firms offered $\{\lambda_H\}$. Then the expected fraction of H-type workers active in the stage-2 public market is at a minimum and equals

$$\frac{n_H - \frac{n_H}{n_S} n_F \alpha}{n_S - \frac{n_H}{n_S} n_F} > 1 - \alpha \tag{A.8}$$

where $\frac{n_H}{n_S} n_F$ is the expected number firms with a stage-1 H-type hire and the inequality follows from plugging in $\alpha > \max(1/2, n_H/n_S)$. Hence the expected fraction of H-type workers active in the stage-2 public market is strictly higher than $1-\alpha$. The latter is the probability to hire an H-type worker through a referral offer for a firm with a stage-1 L-type worker. It follows that for a firm with a stage-1 L-type worker the referral offer of $\{\lambda_H\}$ (or higher) is strictly dominated by the same offer in the public market. Referral offers below λ_H are inconsequential: they are rejected by H-type workers and L-type workers willing to accept such offers will do so in the public market as well. □

When workers choose which wages to accept in the stage-2 public market, they are aware of the number of firms f and the number of workers s still active in the market. Similarly to (A.1), the wage surely accepted by L-type workers $\bar{w}_1^2(f, s)$ follows from solving

$$u_L(\bar{w}_1^2(f, s)) = \frac{f}{s} u_L(\lambda_H) + \frac{s - f}{s} u_L(\lambda_L). \tag{A.9}$$

The maximum wage $\bar{w}_2^2$ below λ_H a firm is willing to offer, similarly to (A.3), follows from

$$u_F(P_L + B - \bar{w}_2^2) = \Psi(\lambda_H, f, s) \tag{A.10}$$

where, letting $\bar{f} = n_F - f$ be the number of firms that hired through a referral offer,

$$\Psi(\lambda_H, f, s) = \sum_{i=0}^{\bar{f}} \alpha^i (1-\alpha)^{\bar{f}-i} \binom{\bar{f}}{i} \left(\frac{n_L - (\bar{f} - i)}{s} u_F(P_L + B - \lambda_H) + \frac{n_H - i}{s} u_F(P_H + B - \lambda_H) \right). \tag{A.11}$$

As in the *Baseline* treatment, equilibrium offers below λ_H satisfy $w^{2,*} \leq \max(\bar{w}_1^2(f,s), \bar{w}_2^2(f,s))$ and firms mix between $\{\lambda_H\}$ and $\{w^{2,*}, \lambda_H\}$. For simplicity, let us focus on $w^{2,*} = \bar{w}_1^2(f,s)$ (equilibria for the other possible wage levels are derived analogous to the *Baseline*). Let $fl \leq f$ be the number of firms offering $\{w^{2,*}, \lambda_H\}$. The expected utility *conditional* on l, h, and fl when offering $\{\lambda_H\}$ is

$$U_F(\{\lambda_H\}; l, h, \hat{fl}) = \frac{l - \min(\hat{fl}, l)}{s - \min(\hat{fl}, l)} u_F(P_L + B - \lambda_H) + \frac{h}{s - \min(\hat{fl}, l)} u_F(P_H + B - \lambda_H). \tag{A.12}$$

Offering $\{w^{2,*}, \lambda_H\}$ yields an expected utility of

$$U_F(\{w^{2,*}, \lambda_H\}; l, h, \hat{fl}) = \frac{\min(l, fl)}{fl} u_F(P_L + B - w^{2,*}) + \left(1 - \frac{\min(fl, l)}{fl}\right) \left[\frac{h}{s - l} u_F(P_H + B - \lambda_H)\right]. \tag{A.13}$$

The equilibrium value $\beta^{2,*}$ with which firms in the public stage-2 market choose to offer $\{\lambda_H\}$ follows from solving

$$\sum_{i=0}^{\hat{f}} (1 - \beta^{2,*})^i (\beta^{2,*})^{\hat{f} - j} \binom{\hat{f}}{i} \sum_{j=0}^{\bar{f}} \alpha^j (1 - \alpha)^{\bar{f} - j} \binom{\bar{f}}{j} \Big[U_F(\{\lambda_H\}; l, h, \hat{fl}) - U_F(\{w^{2,*}, \lambda_H\}; l, h, \hat{fl})\Big] = 0. \tag{A.14}$$

If (A.14) exceeds 0 even for $\beta^{2,*} = 1$, the equilibrium value is $\beta^{2,*} = 1$ and vice versa for $\beta^{2,*} = 0$. This fully characterizes behavior in the public market of stage 2 in the *Referral* treatment.

The next question is whether firms that hired an H-type worker in stage 1 will make a referral offer. We denote the probability with which firms make such an offer by γ. The expected utility when hiring through a referral offer is

$$\Psi_r(\lambda_H) = (1 - \alpha) u_F(P_L + B - \lambda_H) + \alpha u_F(P_H + B - \lambda_H). \tag{A.15}$$

The expected utility when offering in the stage-2 public market depends on $\beta^{2,*}$, which we derived in (A.14), on the number y of other firms that have hired at a wage of λ_H in stage 1 (this is observed), and on the probability γ with which such firms make referral offers. The expected utility is given by

$$U_F^2(\gamma, y) = \sum_{q=0}^{\min(\hat{n_H}, y)} \frac{\left(\mathbb{1}_{q=0} + \prod_{j=1}^{q-1} (\hat{n_L} - j)\right) \left(\mathbb{1}_{q=y} + \prod_{j=0}^{y-q-1} (n_L - j)\right)}{\prod_{j=0}^{y-1} (\hat{n_S} - j)} \binom{y}{q} \sum_{i=0}^{q} \gamma^i (1 - \gamma)^{q-i} \binom{q}{i} U_F^2(n_F - i) \tag{A.16}$$

where the first term cycles through the probabilities that in stage 1 $q = 0$ to $q = \min(\hat{n_H}, y)$ other firms have hired an H-type worker (from the perspective of a firm that hired such a worker), the second terms determines the number of referral hires i given q, and $U_F^2(n_F - i)$ is the expected utility in the stage-2 public market if the number of active firms is $f = n_F - i$; we omit writing out the latter, because it can be found using the same procedure as in (A.12) - (A.14). Notice that $\gamma^* = 0$ if $\Psi_r(\lambda_H) < U_F^2(0, y)$, $\gamma^* = 1$ if $\Psi_r(\lambda_H) \geq U_F^2(0, y)$, and $\gamma^* \in (0, 1)$ solving $\Psi_r(\lambda_H) = U_F^2(0, y)$ otherwise.

With this in hand, we can now determine the behavior in stage 1. Denote the equilibrium expected utility in the public market of stage 2 conditional on y by $U_F^{2,\text{public}}(y)$. Similarly, denote the expected utility in stage 2 when attempting at hiring an H-type worker in stage 1 by $U_F^{2,\text{referral}}(y)$. The wage surely accepted by L-type workers $\bar{w}_1^1$ is exactly the same as in the Baseline, see (A.1). The maximum wage $\bar{w}_2^1 < \lambda_H$ a firm is willing to offer is different than in the *Baseline*, because of the valuable social ties that come with hiring an H-type worker. It solves

$$u_F(P_L + B - \bar{w}_2^1) + U_F^{2,\text{public}}(\hat{n_F}) = \Psi(\lambda_H) + U_F^{2,\text{referral}}(\hat{n_F}) \tag{A.17}$$

where

$$U_F^{2,\text{referral}}(\hat{n_F}) = \mathbb{1}_{\Psi_r(\lambda_H)\geq U_F^{2,\text{public}}(\hat{n_F})} \left(\frac{n_L}{n_S} U_F^{2,\text{public}}(\hat{n_F}) + \frac{n_H}{n_S} \Psi_r(\lambda_H) \right) + \mathbb{1}_{\Psi_r(\lambda_H)< U_F^{2,\text{public}}(\hat{n_F})} U_F^{2,\text{public}}(\hat{n_F}). \tag{A.18}$$

The left-hand side of (A.17) is the sum of expected utilities over both stages when making low offers in both stages and all other firms offer $\{\lambda_H\}$ in stage 1 and follow the behavior derived above in stage 2. The right-hand side is the corresponding sum of expected utilities when making a high offer only in stage 1, hoping to hire a referral worker in stage 2.[15]

As in the *Baseline*, firms will mix between $\{\lambda_H\}$ and $\{\bar{w}_1^1, \lambda_H\}$ if $\bar{w}_1^1 \geq \bar{w}_2^1$. Let $\beta^{1,*}$ be the probability that firms offer $\{\lambda_H\}$. It is found by solving

$$\sum_{i=0}^{\hat{n}_F} (1-\beta^{1,*})^i (\beta^{1,*})^{\hat{n}_F - i} \binom{\hat{n}_F}{i} \left[U_F(\{\lambda_H\}; i) + U_F^{2,\text{referral}}(\hat{n}_F - i) - \left(U_F(\{\bar{w}_1^1, \lambda_H\}; i) + U_F^{2,\text{public}}(\hat{n}_F - i) \right) \right] = 0 \tag{A.19}$$

where $U_F(\{\lambda_H\}; i)$ and $U_F(\{\bar{w}_1^1, \lambda_H\}$ have been derived in (A.4) and (A.6), respectively. If $\bar{w}_1^1 < \bar{w}_2^1$, $\beta^{1,*} = 1$. Notice that because $U_F^{2,\text{referral}}(\hat{n}_F - i) \geq U_F^{2,\text{public}}(\hat{n}_F - i)$, the probability $1 - \beta^{1,*}$ to observe offers below λ_H is smaller in stage 1 of the *Referral* treatment than in the *Baseline* treatment. This completes the construction of the market equilibrium.

[15] The wage $\bar{w}_2^1$ is reached if all other firms offer $\{\lambda_H\}$, because this makes offering low more attractive in stage 1 and it also reduces the possible benefits from offering low in stage 2 because hiring an L-type becomes more likely (recall that when offering low a firm always includes a high offer as well, hoping that others will hire all L-types first).