

Referral Hiring and Wage Formation in a Market with Adverse Selection*

Aurelie Dariel[‡]

Arno Riedl[§]

Simon Siegenthaler[¶]

April 2020

Abstract

The widespread use of employee referrals raises questions regarding how they affect labor market outcomes: Does referral hiring lead to a more efficient allocation of workers compared to when hiring is possible only on a competitive market? To utilize the social links of their current employees, are employers willing to pay a wage premium? We present a game-theoretic model and provide results from a laboratory experiment to address these questions. In line with the theoretical predictions, employers often hire via employee referrals, which in turn mitigates adverse selection and elevates wages. Importantly, employers anticipate the future value of incumbent high-productivity employees' social links and are willing to take the risk of offering higher wages, even when hiring on the competitive market. We also find that employers' risk aversion and the dynamic nature of the hiring process can help account for the remaining inefficiency in the labor market.

JEL Classification: C92, D82, D85, E20

Keywords: Adverse selection, referral hiring, wage formation, social links

*We would like to thank Olivier Bochet, Nikos Nikiforakis, Jean-Marie Meier and Khai Chiong for useful comments and discussions. We also thank seminar participants at the University of Bern, Maastricht University, Université de Cergy-Pontoise, and the participants at the ESA conference in New-York and the French Experimental Economics Association Meeting in Lyon.

[‡]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 Microeconomics and Public Economics, 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 (corresponding author).

[¶]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

Information asymmetries are a persistent source of inefficiency and have been studied extensively in economics since Kenneth Arrow’s (1963) and George Akerlof’s (1970) seminal works. The premise is that the qualities of goods and services traded in a market are unobservable, which can lead to adverse selection. Labor markets are a prominent example. There, adverse selection leads to a low equilibrium market wage and an inefficient assignment of workers to jobs. However, firms would be willing to pay high wages if they could identify productive workers, which would be optimal for firms and society.¹ Thus, it is important to find mechanisms that can help mitigate inefficiencies in labor markets induced by asymmetric information.

Employee referrals have been suggested as one such mechanism with which employers can ensure better social and economic outcomes (e.g., Granovetter, 1985, 1995; Rees, 1966; Fernandez and Moore, 2000). Empirically it has been shown that workers that are part of the same social network may monitor and help each other more effectively (Kugler, 2003; Castilla, 2005; Heath, 2018). It has also been observed that referred employees benefit firms through lower turnover (Weller et al., 2009) and, in some instances, superior performance (Castilla, 2005; Hensvik and Skans, 2016; Barr et al., 2019), and that they earn higher wages (Burks et al., 2015; Brown et al., 2016). In addition, a number of recent studies suggest that employee referrals convey information about the quality of the referred workers (Yakubovich and Lup, 2006; Schmutte, 2014; Burks et al., 2015; Dustmann et al., 2015; Pallais and Sands, 2016), though this information may be noisy (Fafchamps and Moradi, 2015). Despite this evidence, a number of important questions remain open, specifically regarding the impact of employee referrals on aggregate market performance and the causality of the observed effects.

In this study, we use a laboratory experiment to examine the causal effect of employee referrals and social links on outcomes in a labor market with adverse selection.^{2,3} We contribute to the literature in two ways. First, due to the large scale of many markets, in the field it is usually infeasible to detect the impact of employee referrals on aggregate market performance. Therefore, important previous empirical studies focus on firm-specific outcomes (e.g. Burks et al., 2015; Hensvik and Skans, 2016) or examine how job referrals impact behavior of referrers and the referred (e.g., Beaman and Magruder, 2012; Obukhova and Lan, 2013; Pieper et al., 2019). In contrast, in the laboratory environment everything is observable to the researcher, which admits an analysis of how employee referrals affect *aggregate* market dynamics, market efficiency, and the formation of

¹Greenwald (1986) has shown that adverse selection of workers extends further to the stream of job changers, because firms tend to concentrate on preventing turnover on their better workers.

²Our study is framed in labor market terms but the results and identified mechanisms can be translated to other markets with similar incentives structures.

³It may be argued that behavior in a laboratory labor market is not necessarily representative for behavior in the field. For a critical view on the ecological validity of laboratory experiments, see, e.g., Levitt and List (2007). Recently, however, a number of studies showed that behavior in the laboratory has predictive power for field behavior (Falk et al., 2013; Cooper and Kagel, 2016; Riedl and Smeets, 2017).

market wages. One key result we offer is that the possibility of hiring through referrals results in a wage premium even on the competitive *public* market, where employees are hired via posted offers. This means that incumbent employees benefit when their employers anticipate the use of the employees' social networks in further hiring decisions.

Second, it is generally accepted that employee referrals serve a central function in the hiring process. For instance, about one half of all open positions in the U.S. and European labor market are allocated through referrals by friends and relatives (e.g., Ioannides and Loury, 2004; Jackson, 2010a; Pellizzari, 2010; Topa, 2011).⁴ However, this could be due to a multitude of factors, including the informational value of referrals, but also faster and cheaper job-matching (Calvo-Armengol and Zenou, 2005; Obukhova and Lan, 2013; Galenianos, 2014) or benefits derived from long-term employee attachment to the firm (Castilla, 2005; Pieper et al., 2019). In the field it is difficult to distinguish between these factors and identify causal effects of job referrals, because social links are endogenous and the productivity of employees is difficult to quantify. In some cases, employee referrals may also have unintended consequences such as job segregation (e.g., Rubineau and Fernandez, 2013) and increased inequality (e.g., Calvo-Armengol and Jackson, 2004; Galenianos, 2018) that can affect behavior in unpredictable ways. The laboratory environment permits us full control over the incentives people face and allows us to causally identify the impact of the informational value of referrals.

A central premise of our approach is that workers have social links and that there is a positive correlation in productivity between linked workers. This assumption is supported by a number of important studies showing that people tend to mainly interact with others that are similar to themselves, often referred to as homophily—and which in our setting imply a correlation in productivity among employees with a social link (e.g., Fernandez and Moore, 2000; McPherson et al., 2001; Currarini et al., 2009; Jackson, 2010b; Zeltzer, 2019). We, thus, take the correlation in productivity as given and focus on its *consequences* for aggregate market outcomes, in a setting where efficiency is hindered by substantial information asymmetries.

Our experimental environment corresponds to a standard labor market setting with asymmetric information. The main novelty is that we exogenously vary the availability of employee referrals. The hiring process takes place in an anonymous public market where employers can post offers and workers choose if and when to accept. The number of workers exceeds the number of available jobs. However, workers may nevertheless extract part of the surplus, because they are privately informed about their productivity (low or high). The hiring process takes place in two stages and employers seek to hire at most one worker per stage. The profit of an employer depends on the productivity of the hired workers and the wage the employer pays to these workers. The earnings of the workers are determined by the wage they receive or, if a worker fails to find a job, by his or

⁴Employee referrals are also widely used in developing countries (Heath, 2018), in migrant communities (Munshi, 2003; Beaman, 2011), and are pervasive in different industries ranging from call-centers to high-tech firms (Burks et al., 2015).

her outside option. An outside option could reflect a worker’s benefit from waiting for new jobs openings or from being hired in a different industry. High-productivity workers are assumed to have better outside options than low-productivity workers.

We implement two treatments. In the *Baseline* treatment, there are no social links between workers in stage 1 and stage 2. Hence, knowledge about the productivity of a stage-1 worker does not contain information regarding the productivity of a stage-2 worker. In this treatment, stage 1 and stage 2 can be thought of as two independent markets characterized by the same incentives. In contrast, in the *Referral* treatment, each stage-1 worker has a social link with a stage-2 worker, meaning that the linked workers are of the same productivity type with a high probability. Thus, in this treatment the productivity of a stage-1 worker has informational value regarding the productivity of the stage-2 worker linked to the hired stage-1 worker. Employers can utilize this information by making referral offers to the stage-2 worker. This implies that hiring decisions employers make in stage 1 can directly affect their hiring opportunities in stage 2. Importantly, referral offers do not give exclusive access to a worker: all stage-2 workers can also be hired on the public market.

We develop a theoretical model to derive testable hypotheses for the experiment. The model builds on a well-established theoretical literature on social networks in labor markets and job-search through personal contacts (e.g., Greenwald, 1986; Montgomery, 1991a; Mortensen and Vishwanath, 1994; Pissarides, 2000; Topa, 2001; Kugler, 2003; Obukhova and Lan, 2013; Rubineau and Fernandez, 2013; Galenianos, 2014).⁵ However, existing models typically assume free entry by firms (i.e., the number of firms is endogenous and firms make zero profits), allow for rich and complicated network structures, assume risk neutrality, and look at behavior in the steady state of a market. These features are difficult to implement experimentally. Our model lends itself to an experimental implementation while retaining most of the predictions of the above-cited theoretical literature. Hence, our study may be viewed as an experimental test of some important predictions coming from the theoretical literature on employee referrals in labor markets.

We derive the following testable predictions. First, our model predicts that wage offers in the public market are higher when employee referrals are possible (*Referral* treatment) than in the absence of social links (treatment *Baseline*). This is because the likelihood of hiring a high-productivity worker increases with higher wage offers, and employers should anticipate that hiring a high-productivity worker will grant access to a better social network from which employers can draw workers in the subsequent hiring round. Second, if employers succeed in hiring a high-productivity worker in stage 1, they should make use of referral wage offers in stage 2, and workers receiving such offers should accept. This is because the expected productivity of such

⁵The theoretical literature on social networks in labor markets examines a broad range of other important issues such as the impact of network structure on labor market outcomes (e.g. Boorman, 1975; Zenou, 2013), how network connections provide information about available jobs (e.g., Calvo-Armengol, 2004; Galeotti and Merlino, 2014), job tenure (e.g. Loury, 2006), career concerns (Ekinici, 2016), and how social networks may induce inequality (e.g. Montgomery, 1991b; Calvo-Armengol and Jackson, 2004; Galenianos, 2018).

referral workers is higher than the productivity of the average worker on the market. Third, these predictions imply that the availability of employee referrals should increase market efficiency. This effect results from a displacement of low-productivity hires by high-productivity hires, compared with the setting without social links. In particular, employee referrals are not predicted to increase the overall number of hires and, hence, tend to increase inequality among workers.

While these predictions are intuitive, the equilibrium strategies involve mixing and the market equilibrium is characterized by non-trivial trade-offs. Employers trade off the benefits from hiring high-productivity stage-1 workers—i.e., higher profits and access to workers’ social links—with the risk of hiring a low-productivity stage-1 worker at a high wage. Whether offering high wages is a profitable strategy depends on an employer’s expectation about the willingness of low-productivity workers to accept lower wages and on the degree of risk aversion present on both market sides. Workers, on the other hand, trade off the time they wait before accepting a wage offer with the risk of remaining unemployed. Thus, even if social links are by design informative about productivity, whether the presence of employee referrals alleviates adverse selection is not obvious, depends amongst others on the market participants’ expectations, and hence needs empirical validation.

Our results largely support the hypotheses derived from the theoretical model, demonstrating the value of employee referrals in alleviating inefficiencies caused by the asymmetric information present in the market. We find that: (i) the allocation of workers to employers is significantly more efficient in the *Referral* treatment than in the *Baseline* treatment, (ii) hiring through social links is common and mostly concern high-productivity workers, and (iii) employers appear to anticipate this and offer higher stage-1 wages in the *Referral* treatment than in the *Baseline* treatment, the aforementioned risks notwithstanding. Interestingly, the positive effect of social links is partly based on low-productivity workers underestimating the probability of being hired when holding out for high wages.

Despite the positive effect of employee referrals on efficiency, we also observe that it falls short of the second-best efficient outcome, which is defined as the maximum welfare subjects can obtain given the constraint that social links are not fully informative. We show that the failure to reach second-best efficiency is caused by a combination of risk aversion on both sides of the market and the dynamics of the hiring process. Employers almost always start with low wage offers, leading to many early low-productivity hires. The more risky, and socially more efficient, strategy of offering higher wages is adopted only late during market opening time. This suggests that employee referrals are important for mitigating adverse selection in labor markets, but that other mechanisms are needed to further improve efficiency.

The paper proceeds as follows. In Section 2, we present the experimental setting and design. Section 3 introduces the theoretical model and derives three key hypotheses for behavior in the experiment. In Section 4, we present the empirical results and Section 5 concludes.

2 Experiment Design

Our experimental environment is designed to reflect a typical labor market setting with adverse selection. There the existence of asymmetric information implies that it is risky for employers to offer wages that are acceptable to high-productivity workers because they may end up paying a high wage to low-productive workers. Our treatments exogenously vary the availability of employee referrals and are described in more detail below. First, we present the general setting that applies to both treatments.

2.1 General Setting

In each session of the experiment there are 16 participants. At the beginning of a session, the market participants are randomly assigned the role of either a firm or a worker. For clarity, the lab experiment was framed as a labor market in which firms can hire workers. In line with this, in the following we will use the terms firm and worker. In total there are 4 firms and 12 workers that interact in a market. The roles remain fixed throughout a session and each session consists of 15 periods. Each period constitutes a distinct market and the number of firms and workers is the same in all markets.

In each period, subjects interact in a market with two hiring stages. A stage lasts for 2 minutes during which firms seek to hire at most one worker. Firms are free when or how many offers they would like to make and workers can accept any of the standing offers at any time. Firms are active in both stages whereas the 12 workers are divided into 6 stage-1 workers and 6 stage-2 workers. Stage-1 workers are active only in stage-1 and thus can only be hired in this stage; equivalently for stage-2 workers. Whether a worker is active in stage 1 or stage 2 is randomly determined at the beginning of each period.

When making offers, firms have to abide by the improvement rule. That is, every new wage offer has to be higher than the currently highest standing wage offer. When there are no standing offers (which occurs at the beginning of a hiring stage or when all standing offers have been accepted), new offers can again start at the lowest permissible wage. This is a standard rule in experimental markets (e.g., Smith, 1962; Fehr et al., 1993, 1998; Fehr and Falk, 1999; Falk and Szech, 2013). Among others, it rules out the possibility that firms make offers that should never be accepted by workers.

To implement the presence of asymmetric information, workers are further divided into low-productivity (L-type) workers and high-productivity (H-type) workers. Specifically, in each hiring stage there are 3 L-type and 3 H-type workers. The productivity of a worker is randomly determined at the beginning of each period. Importantly, firms know the distribution of worker

types but are informed about their worker’s actual productivity only after the worker has been hired.

Firms’ profits in a stage are determined as follows. A low-productivity worker produces an output $P_L = 20$ and a high-productivity worker produces an output of $P_H = 60$. The set of permissible wage offers is $\{0, 1, \dots, 59, 60\}$. If a firm pays a wage w , her payoff in a given stage is given by

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

where adding the 20 reduces the risk of negative payoffs in the experiment and guarantees firms positive payoffs when hiring an L-type worker at the H-type workers reservation wage (see below). As will be shown in the next section, equilibrium market wages can nevertheless be below the H-type workers’ reservation wage, which gives rise to adverse selection. In a period, the total payoff of a firm is the sum of payoffs in stage 1 and stage 2.

Workers’ payoffs are determined by the wage when hired and their outside option when not hired. It is assumed that workers of different productivity types have different outside options and, thus, different reservation wages. Specifically, an L-type worker has a reservation wage of $\lambda_L = 10$ and an H-type worker has a reservation wage of $\lambda_H = 30$. In a stage, the payoff of a worker is given by

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

The different reservation wages imply that it is risky for firms to try to hire high-productivity workers. Hiring such workers requires a wage offer of at least 30, which entails the risk of a low payoff in case a low-productivity worker accepts the offer.

Workers are active either in stage 1 or stage 2, whereas firms are active in both stages. Therefore, to keep payoffs between firms and workers comparable, workers receive their reservation wage in the stage in which they are not active.

2.2 Treatments

We implemented two treatments that differ in the possibility of hiring workers via referrals. In the *Baseline* treatment, the productivity of any given stage-1 worker is independent of the productivity of any given stage-2 worker. In both hiring stages, firms make wage offers in an anonymous public posted-offer market. We will refer to such offers as *public offers*. All firms and all workers active in a stage can observe all public offers. Workers not active in a stage do not

observe any offers. A worker that accepts an offer is hired at the corresponding wage by the firm which made the offer, and this worker-firm pair leaves the market. The remaining workers can still get hired by the other firms on the market. In both stages, the market is open for at most 2 minutes or until all firms have hired a worker.

The *Referral* treatment differs from *Baseline* in the availability of employee referrals. In particular, social links connect the productivity of stage-1 workers to the productivity of stage-2 workers, such that two linked workers are of the same type with a probability of 75%. Social links are randomly assigned among workers at the beginning of each trading period and each stage-1 worker is uniquely linked to one stage-2 worker. The first stage of the *Referral* treatment is exactly the same as in the *Baseline* treatment and all hires take place via public offers. However, social links affect the hiring opportunities of firms in the stage-2 market.

Specifically, in stage 2, firms can make two types of offers. First, they can make public offers which are observed by all firms and all stage-2 workers. Second, each firm that has hired a worker in stage 1, can try to hire a worker through an offer that is only received and observed by the stage-2 worker that has a social link with that firm's stage-1 worker. We will refer to such offers as *referral offers*. Referral offers can be changed in any way and do not have to abide by the improvement rule. Firms can make multiple offers of both types (public and referral) simultaneously during the 2 minutes the stage-2 market is open.⁶ Importantly, workers that receive referral offers can still accept wage offers in the public market. Thus, unless all firms only make referral offers, the firms still compete for workers. Referral offers have the advantage that, due to the social link, there is less noisy information regarding the stage-2 worker's productivity.⁷

2.3 Information

In both treatments, participants are informed about all relevant parameters of the experiment. That is, they know that there are (a) 15 periods with two trading stages in each period, (b) 4 firms and 6 workers per trading stage, (c) 3 low- and 3 high-productivity types among the 6 workers, and they also know (d) the payoff functions of both firms and workers. In the *Referral* treatment, participants know from the start of the experiment that stage-1 and stage-2 workers that are connected by a social link are of the same productivity type with a probability of 75%.

⁶In the experiment, this was implemented via two panels showing standing offers that firms and workers could observe. In one panel, firms could make public offers and all workers could accept. In the other panel, firms could make referral offers to their linked stage-2 worker. Since each such worker is linked to exactly one firm, a given worker could only observe the referral offers of this one firm. A screen shot of the experimental interface can be found in the online appendix.

⁷From the perspective of stage-2 workers, while they do not know the true productivity of the stage-1 worker they are linked to, receiving a referral offer is informative about said productivity and hence may affect a stage-2 worker's expectation about future wage offer.

Finally, firms and workers active in a stage can observe all public offers made in that stage. In contrast, referral offers can only be seen by the worker that receives the offer.⁸

2.4 Procedures

The experiment was run in the BEElab at Maastricht University, using experiment software z-Tree (Fischbacher, 2007) and participants were recruited from the BEElab participant pool via ORSEE (Greiner, 2015). Altogether, 11 sessions were conducted, 5 for treatment *Baseline* and 6 for treatment *Referral*, with 16 participants—4 firms and 12 workers—in each session, amounting to a total of 176 participants. At the beginning of a session each participant was randomly assigned to a cubicle where they made decisions in private. The experiment instructions included a set of control questions to check participants’ understanding.⁹ Each session lasted 100 minutes or less and payments averaged 19 Euros per participant including a show up fee of 5 Euros.¹⁰

3 Theoretical Background and Hypotheses

We present a model to derive a set of hypotheses for behavior in the experiment. As mentioned in the introduction, our main predictions are in line with previous theoretical studies on job referrals in labor markets. Our experiment may therefore be seen as a general test of (a subset of) the theoretical insights established in this literature. The main reason for providing a new model is to better capture the features of the experimental setting. Two important differences to the existing theoretical literature are that we allow for risk aversion and that we use a market clearing rule that is standard in market experiments.

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. Using the notation introduced in Section 2.1, the produc-

⁸The design element that referral offers can not be seen by other firms is based on Greenwald (1986), who argues that for firms it is difficult to observe wage offers that workers receive from their current employer; for instance, because wage offers from current employers can involve credible future promises and because current employers can often directly respond to outside offers, arguments which arguably also apply to referral offers. In a similar vein, Waldman (1984) states that “firms other than the employer are frequently at an informational disadvantage when it comes to experienced workers, but this is somewhat attenuated by the fact that such firms can use an individual’s job assignment as an imprecise signal of the individual’s ability.”

⁹The instructions can be found in the online appendix.

¹⁰Typical for adverse selection settings, in our experiment it is possible that participants receive negative pay-offs. This could be problematic if a participant goes bankrupt, which most likely happens in early periods when participants have not yet accumulated earnings. To deal with this potential issue, participants received an initial endowment of 120 experimental points (conversion rate: 1 point = 2.25 Euro cent). No participant went bankrupt.

tivity of a worker of type θ is P_θ and the reservation wage is λ_θ . 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 the wage offers they would like to accept. A worker can accept none, one, or several of the offers made by the firms. Given these decisions, firms and workers are then matched from low to high wages. In particular, market clearing starts at the lowest accepted wage offer, say w_1 . Among the workers that chose to accept w_1 , one is hired at random. Then, the next highest accepted wage offer w_2 is selected resulting in another firm-worker match. This process continues until all firms hired a worker or the highest accepted wage offer is reached.

The assumption that wages clear from below is in line with the improvement rule typically used in market experiments (see Section 2.1). Recall that the improvement rule says that if a firm chooses to make a new wage offer, it must be higher than the most recent non-accepted wage offer. Importantly, wage clearing from below does *not* require firms to offer low wages—indeed, we will see that in equilibrium some firms post a single wage offer equal to the productivity level of the high-type worker.

It should be noted that the model does not mirror the experimental setting one-to-one. In the experiment, wage offers and acceptances occur dynamically to allow for a realistic scenario. In the model, for tractability, we assume that firms first choose a set of wage offers, then workers choose which wages to accept. However, the fact that markets clear from below implicitly allows for dynamic hiring. To see this, consider a firm that chooses to make two offers, a low offer and a high offer. The firm will first compete for the pool of workers that are willing to accept low offers and, if the firm is not able to hire such a worker, it can still compete for the remaining workers that demand higher offers.¹¹

Regarding the agents' utility function, we allow for risk aversion for the following main reasons. First, in experiments, subjects are usually found to be risk averse. Second, it is likely that risk attitudes are also important in the field. There, workers face the risk of unemployment and recruiters, while not directly incurring profit losses from a bad hire, can bear substantial costs associated with unsatisfactory job performance. We assume that 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), respectively, and $\sigma > 0$ measures the degree

¹¹In the online appendix, we show that the Walrasian equilibrium outcome and the predictions from a model in which wages clear from above exhibit inefficiencies due to adverse selection that are similar to those predicted by our model.

of absolute risk aversion. For tractability, we assume that risk preferences are identical across agents.¹²

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 the behavior of everyone else, and firms' beliefs about the expected quality of workers are correct at all wage levels.¹³ In the following, we focus on the key predictions and refer to the appendix for a detailed analysis.

Recall that in the *Baseline* treatment stages 1 and 2 are identical, because there are no social links. The equilibria in each stage 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.¹⁴ 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 and, thus, some workers will remain unemployed. This explains why it is always optimal for a firm that offers w^* to also include offer λ_H in its wage schedule: because w^* is accepted by L-type workers and market clearing occurs from low to high wages, at equilibrium a firm that offers $\{w^*, \lambda_H\}$ will hire at a wage of λ_H only if all L-type workers have already left the market. In other words, including wage offer w^* in the set of offered wages guarantees a positive profit.

Figure 1 provides an illustration of the market equilibrium for different levels of agents' risk aversion. The figure is based on the parameters that we used in the experiment but the qualitative predictions are robust to variations in these parameters. Figure 1(a) shows the average accepted offer and Figure 1(b) shows the percentage of high-productivity workers among all hired workers, where the total number of hires is four (in equilibrium each firm hires a worker). The solid lines depict the predictions for treatment *Baseline*. Notice that for low levels of risk aversion

¹²The model can also accommodate situations in which firms and workers have different risk preferences. In that case, the qualitative predictions for the effect of employee referrals are similar although the adverse selection effect is mitigated when firms are less risk averse or risk-neutral. In the online appendix, we discuss the case of risk-neutral employers. The main effect is that competition between firms for high-productivity workers is fiercer, as firms are less concerned about the downside of hiring a low-productivity worker at a high wage. Interestingly, however, the availability of employee referrals is still predicted to improve the average productivity of the hired workers.

¹³Beliefs 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).

¹⁴Notice 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

Notes: Theoretical predictions for the experimental parameters $n_F = 4$, $n_L = n_H = 3$, $P_L = 20$, $P_H = 60$, $\lambda_L = 10$, $\lambda_H = 30$, $B = 20$ (added productivity parameter), and different levels of risk aversion. Figure (a): average accepted wage offers in the public market. Figure (b): percentage of H-type workers among all hires. 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.

the equilibrium wage equals $\lambda_H = 30$ and the percentage of hired H-type workers is 50%. All six workers accept a wage of 30 and on average two workers of each productivity type are hired. As agents become more risk averse, firms start to randomize between $\{w^*, \lambda_H\}$ and $\{\lambda_H\}$, which depresses 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 strong adverse selection effect and ex-post social welfare loss. From a social welfare perspective, all (three) H-type workers should be hired, as the gains from trade with an H-type equal 30 whereas for an L-type they are only 10.

3.3 Market Equilibrium in *Referral*

We next discuss whether the opportunity to make referral offers alleviates adverse selection. In stage 1 of the *Referral* treatment, the hiring process happens on a public market, as in the *Baseline* treatment. In stage 2, firms can still hire through public offers but, importantly, they can also hire through referral offers using the social links of the hired stage-1 worker. Recall, that with a chance of 75% the linked stage-2 worker is of the same productivity type as the stage-1 worker.

The dashed lines in Figure 1 depict the equilibrium outcomes in stage 1 of the *Referral* treatment. It shows that for some levels of risk aversion, average wages in stage 1 are higher in the *Referral* treatment than in the *Baseline* treatment. Firms are willing to make higher wage offers in the

former treatment, because hiring an H-type worker has the additional benefit of providing access to a valuable social link that can be used in stage 2. As a result, compared to the *Baseline* treatment, more H-type workers are hired.

In stage 2 of the *Referral* treatment, up to a risk coefficient of about 0.1, firms that have hired an H-type worker in stage 1 always make referral offers of 30 in stage 2. For higher levels of risk aversion, there is a sharp decrease in the use of social links until the differences between the treatments disappear. For higher risk aversion, even a 75% probability of hiring an H-type worker is too risky and firms focus on hiring L-type workers through public offers. The dotted line in Figure 1(b) shows that the percentage of H-type hires in stage 2 is higher in the *Referral* treatment than in the *Baseline* treatment, as long as the degree of risk aversion is not too large. Thus, the possibility of using referrals is predicted to mitigate the adverse selection problem.¹⁵ Interestingly, referral hires in stage 2 lead to a “lemons” effect in the pool of workers that can be hired via public offers. This is illustrated by the dotted line in Figure 1(a), showing relatively low average wages in the public market of stage 2 in the *Referral* treatment.

The theoretical predictions can be summarized as follows:

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

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

Hypothesis 3 (Hires): The percentage of H-type hires is higher in *Referral* than in *Baseline*. This holds in both stages.

In the *Referral* treatment hiring an H-type worker in stage 1 provides access to a social link, which can be used to hire an H-type worker in stage 2 with a higher probability. Therefore, social links are predicted to promote H-type hires not just in stage 2 but also in stage 1.

4 Empirical Results

We will first discuss offers and accepted wages in Section 4.1, then examine hires 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. Recall that 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

¹⁵Below, the empirical results will reveal that in the experiment the constant absolute risk coefficient for firms is around 0.05, which is in the range where the possibility of making referral offers promotes hires of H-type workers.

periods (periods 6 - 15) to account for learning effects and behavior that is more stable than in earlier periods.¹⁶

4.1 Offers and Wages

We start with a look at firms’ wage offers over the 120 seconds a market stage was (maximally) open. Figure 2 conveys the general pattern in wages offered and accepted. The upper row in the figure shows all offers in the two treatments and both stages. There is a clear pattern visible which we summarize in the following result.

Result 1. *In the public markets of both treatments firms start by offering wages that are clearly below 30, thus targeting L-type workers. Only towards the end of a trading stage, wage offers increase to the H-type workers’ reservation wage of 30 and above. These observations are consistent with firms being engaged in screening behavior.*

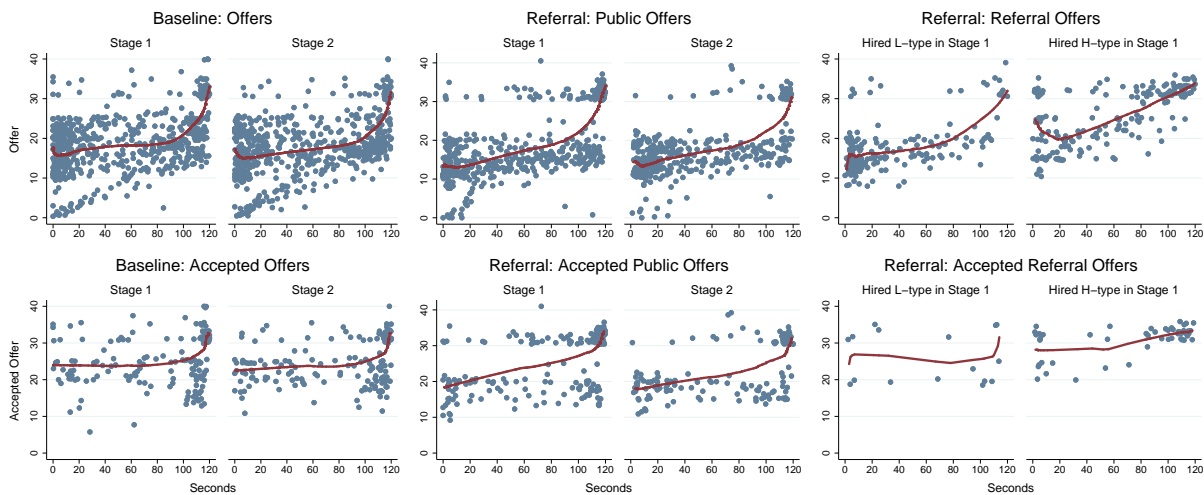


Figure 2: Wage Offers and Accepted Wages

Notes: The first row shows all offers (accepted or not) over time within a trading stage 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.

The lower panels of Figure 2 show accepted wage offers (i.e., hires) and provide two eye-catching observations. First, acceptances of offers below 30 occur throughout the 120 seconds a stage is open, with some concentration at the end of the trading stage. These acceptances are presumably by L-type workers. Second, acceptances of offers higher than 30 occur predominantly towards the end of a trading stage, mainly in the last 20 seconds.

¹⁶In the online appendix we show that the differences in behavior between *Baseline* and *Referral* remain over time, becoming more pronounced in later periods.

Table 1: Summary of Wage Offers, Accepted Wages, Hires and Efficiency

Treatment:	Baseline				Referral								
	Stage 1		Stage 2		Stage 1		Stage 2		Stage 2		Stage 2		
Type	H	L	H	L	H	L	H	L	Overall	Public Offers	Referral Offers	H	L
Offers	19.3		19.3		19.1		19.6		17.8		21.5		
Offers ≥ 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 ≥ 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%		-		-		

Notes: Data from periods ≥ 6 . “H” (“L”) stands for H-(L)-type worker. Variables by row: average offers; % of offers greater or equal to 30; average wages (accepted offers) of H and L, respectively; % of wages (accepted offers) greater or equal to 30 of H and L, respectively; average number of H and L hired; % of H and L hired; efficiency levels as % of realized gains from trade relative to the first-best outcome where three H-type and one L-type workers are hired.

In what follows we provide a quantitative analysis of offers and wages, starting with the behavior in stage 1. We first state results and thereafter provide the empirical support.

Result 2. *In stage 1, firms in the Referral treatment are more likely to offer wages greater or equal to 30 than firms in the Baseline treatment. Consistent with this, the percentage of L-type workers receiving a wage greater or equal to 30 is higher in the Referral treatment than in the Baseline treatment.*

Table 1 provides an overview of average wage offers (‘Offers’), accepted wage offers (‘Wages’), frequency of hires of H- and L-type workers (‘Hires’), and achieved efficiency levels (‘Efficiency’). The second row of the table presents the percentage of offers greater or equal to 30, which in stage 1 is higher in *Referral* than in *Baseline* (26% versus 14%, Mann-Whitney, $p = 0.017$).¹⁷ Consistent with this, the median accepted wage offer in *Referral* is 31, while it is only 25 in *Baseline* (Mann-Whitney, $p = 0.039$). This difference in stage-1 wage offers between treatments is corroborated by random effects regression analyses presented in Table 2, which additionally provide insights into the timing of offers. In regression models (1) and (2), the dependent variable is binary, indicating whether an offer is greater or equal to 30 in stage 1. The explanatory variables are a dummy for the *Referral* treatment and—in model (2)—a dummy *Late* equaling one for offers made in the second half of the market opening time (i.e., after 60 seconds) as well as an interaction between the treatment and the timing of the offer.¹⁸

¹⁷Unless stated otherwise, all non-parametric tests are two-sided and use session averages as the unit of observation with 5 observations in *Baseline* and 6 observations in *Referral*.

¹⁸For easier interpretation of the estimates, in particular of the interaction terms, the table reports results from linear random effects regressions. Probit random effects regression yield qualitatively identical results.

Table 2: Regression Analysis of Offers and Wages

Dep. Var:	Stage 1				Stage 2			
	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*** (0.03)	0.00 (0.03)	0.13* (0.07)	-0.04 (0.16)	0.08*** (0.02)	0.03 (0.02)	-0.01 (0.07)	-0.04 (0.05)
Late (> 60 sec)		0.22*** (0.06)		0.17 (0.16)		0.21*** (0.05)		0.23*** (0.04)
Referral \times Late		0.22*** (0.08)		0.28 (0.19)		0.19*** (0.06)		0.10 (0.10)
Constant	0.01 (0.04)	0.06** (0.28)	0.20* (0.12)	0.20 (0.16)	0.03 (0.03)	0.03 (0.04)	0.18 (0.15)	0.05 (0.08)
Wald Test ^(a)		$p = 0.002$		$p < 0.001$		$p < 0.001$		$p = 0.511$
Observations	1197	1197	277	277	1411	1411	260	260

Notes: Linear random effects models with cluster-robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Bootstrap standard errors yield qualitatively identical results. Period dummies are included in all regressions. (a) Wald chi-squared tests for the hypothesis ‘Referral + Referral \times Late’ = 0. The reference group in models (1), (3), (5) and (7) is the *Baseline* treatment, whereas in the remaining models it is the first 60 seconds of a market stage in *Baseline*. For stage 2, regressions (5) and (6) include public and referral offers, whereas regressions (7) and (8) only include wages in the public market.

Regression model (1) shows that in stage 1, offers of at least 30 are significantly more likely in *Referral* than in *Baseline* ($p < 0.001$) and model (2) indicates that this effect is exclusively due to offers made in the second half of the market opening time ($Referral + Referral \times Late = 0$, Wald chi-squared test, $p = 0.002$). Consistent with these observations, the percentage of L-type workers earning a wage greater or equal to 30 is higher in *Referral* than in *Baseline* (44% vs. 30%, see row ‘Wages \geq 30’ in Table 1). To see if this difference is statistically significant, models (3) and (4) in Table 2 replicate models (1) and (2) for accepted wages of L-type workers only. Model (3) shows that the percentage of wages of at least 30 accepted by L-type workers is indeed (marginally) significantly higher in *Referral* than in *Baseline* ($p = 0.06$). This effect is due to behavior in the second half of the market opening time in stage 1, where it is also significant (Wald chi-squared test, $p < 0.001$).

Given these differences between treatments, it seems surprising that the average offers in stage 1 are similar in the two treatments (see row ‘Offers’ in Table 1). To understand this, we refer to Figure 2. There, it can be seen from the stage-1 panels ‘Baseline: Offers’ and ‘Referral: Public Offers’ that in the *Baseline* treatment there are almost no offers between 25 and 30 but plenty of offers between 20 and 25, whereas in the *Referral* treatment there are almost no offers over the entire range of 20 to 30. Thus, despite the larger percentage of offers at and above 30 and the higher median offer in *Referral*, it happens that the average offers in both treatments are similar, because—conditional on offers being below 30—offers in *Baseline* are higher than in *Referral*.

This observation is consistent with the predictions of our model. The presence of social links in *Referral* increases firms’ expected profits when trying to hire H-type workers with wage offers of at least 30. To compensate for the higher expected profits, for firms still willing to hire L-type workers, L-type hires must occur at lower wages. We now turn to public offers in stage 2.

Result 3. *In stage 2, average wages in the public market are lower in the Referral treatment than in the Baseline treatment. However, the percentage of L-type workers earning a wage greater or equal to 30 is not significantly different between treatments.*

The average public offer in stage 2 of *Baseline* is 19.3, while it is 17.8 in *Referral* (see ‘Offers’ in Table 1). The corresponding average wage in the public market is 25.4 in *Baseline* and 23.2 in *Referral*, and this difference is significant (Mann-Whitney, $p = 0.044$). Focusing only on L-type workers, row ‘Wages’ in Table 1 shows that the average wage of 23.6 in the *Baseline* treatment exceeds the average wage of 20.8 of such workers hired on the public market in the *Referral* treatment. Also this difference is significant (Mann-Whitney, $p = 0.044$). The reason for the lower wages in the *Referral* treatment is that the competition between L-type workers is intensified in stage-2 public markets because some firms hire workers through referral offers and do not participate in the public market. The pressure on wages in the public stage-2 market of the *Referral* treatment does not imply that firms are more successful at screening L-type workers: the percentage of L-type workers hired at a wage of at least 30 on the public market is similar across treatments (23% versus 22%). Regression models (7) and (8) in Table 2 show that there is no significant difference in these percentages between treatments either overall or when focusing on the second half of the market opening time.

Overall, the results up to here show that information asymmetries lead to relatively low wages on public (competitive) markets, impeding trading opportunities for high-productivity workers. Also, Results 2 and 3 largely support our Hypothesis 1 on public market wages. Specifically, the prediction that the existence of referral hiring in stage 2 gives an incentive to firms to offer high wages already in the competitive stage 1 is borne out by the data. We next examine referral offers and explore if they have the predicted mitigating effect on adverse selection.

Result 4. *In stage 2 of the Referral treatment, firms mainly use referral offers after they hired an H-type worker in stage 1, and most referral offers are greater or equal to the H-type workers’ reservation wage of 30.*

The right-most panels in Figure 2 show referral offers and accepted wages in stage 2 of the *Referral* treatment, split by whether the worker hired in stage 1 was of the L- or H-type. The lower panel shows that firms which hired an L-type worker in stage 1 only rarely hire workers through referral offers. In contrast, there are many more referral hires of firms which hired an H-type worker in stage 1, and a large majority of them are hired at wages that equal 30 or are slightly greater. Consistent with this, Table 1 shows that referral offers tend to be higher than public offers in

Table 3: *Referral* Treatment—Likelihood of Referral Offers, Offer Levels, and Likelihood of H-type Hires in Stage 2

Dep. Var.	(1) Referral Offer (Yes/No)	(2) Offer Level	(3) Offer Level \geq 30	(4) H-type Hire in Stage 2
Hired H-type in Stage 1	0.353*** (0.0595)	0.235 (1.291)	0.0665 (0.0723)	0.157 (0.124)
Referral Offer (Yes/No)		0.155 (0.425)	-0.0399 (0.0223)	0.0396 (0.169)
Hired H-type in Stage 1 \times Referral Offer (Yes/No)		7.727*** (1.128)	0.266*** (0.0611)	0.323 (0.279)
Constant	0.352*** (0.0830)	18.66*** (0.826)	0.179*** (0.0445)	0.225 (0.126)
Observations	735	735	735	204

Notes: Linear random effects models with cluster-robust standard errors in parentheses, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Bootstrap standard errors yield qualitatively identical results. Dependent variables: ‘Referral Offer’ = 1 if ‘yes’, ‘Offer Level’ equals offered wage level, ‘Offer Level \geq 30’ = 1 if yes, ‘H-type Hire in Stage 2’ = 1 if yes. All independent variables are dummy variables equaling 1 if the statement in table is true. Period dummies are included in all regressions. For model (4), a Wald chi-squared test is used for the hypothesis ‘Hired H-type in Stage 1 + Hired H-type in Stage 1 \times Referral Offer’ = 0 is rejected at $p < 0.001$.

stage 2 of the *Referral* treatment as well as the *Baseline* treatment (21.5 versus 17.8 and 19.3, respectively).

To test for statistical significance we conduct random effects regressions which are reported in Table 3. Specification (1) shows that firms are 35.3 percentage points ($p < 0.001$) more likely to make referral offers when they hired an H-type worker in stage 1. Specifications (2) and (3), respectively, show that referral offers of firms that hired an H-type in stage 1 are on average 7.72 points higher ($p < 0.001$) and 26.6 percentage points more likely to exceed 30 ($p < 0.001$) than referral offers of firms that hired an L-type worker in stage 1. Together these results support Hypothesis 2.

Moreover, specification (4) indicates that firms which have hired an H-type worker in stage 1 and make a referral offer in stage 2 are $15.7 + 32.3 = 48$ percentage points more likely to hire an H-type worker in stage 2 ($p < 0.001$). Further, combining public and referral offers, firms are more likely to offer wages above 30 in the *Referral* treatment than in the *Baseline* treatment. This can be seen from Table 1 (‘Offers \geq 30’) which shows that in stage 2 the percentage of public and referral offers exceeding the reservation wage of the H-type worker is 21% in *Referral* and only 13% in *Baseline* (Mann-Whitney, $p = 0.017$). Regressions (5) and (6) in Table 2 reveal that this difference is due to the behavior in the second half of the market opening time ($p < 0.001$).

4.2 Hires and Efficiency

We explore now the main research question whether the possibility of hiring via referrals helps to overcome adverse selection, i.e., if it increases the number of H-type hires and thus promotes efficiency. As before, we state first the result followed by statistical support.

Result 5. *Both in stage 1 and 2, the number and percentage of H-type workers hired is higher in the Referral treatment than in the Baseline treatment. Consequently, overall efficiency is also higher in the Referral treatment than in the Baseline treatment.*

In Table 1 the rows labeled ‘Hires (No.)’ and ‘Hires (%)’ present the average number and the percentage of H and L-type hires for each treatment and market stage. First, we note that the sum of H- and L-type hires in both treatments and both stages is with 3.6 to 3.7 close to the maximum of 4 hires. In both stages, the total number of hires is not significantly different between the *Referral* and the *Baseline* treatment (Mann-Whitney, $p > 0.404$).¹⁹ Importantly, however, significantly more H-type workers are hired in *Referral* 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). Consequently, in *Referral*, the percentage of H-type workers among all hires is with 35% in stage 1 and 34% in stage 2 significantly larger than in *Baseline* with corresponding percentages of 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).

To test for treatment differences in achieved efficiency levels, we define efficiency as the realized gains from trade divided by the maximally possible gains from trade. The latter is achieved with three H-type hires and one L-type hire, and is given by $3 * (60 - 30) + 1 * (20 - 10) = 100$. In Table 1, the row labeled ‘Efficiency’ reports the realized efficiency levels for both hiring stages in both treatments. It shows that the availability of referral offers increases efficiency by 10 percentage points ($p = 0.028$). In the two stages two different channels are responsible for the increased efficiency in the *Referral* treatment. In stage 1, efficiency is increased by firms raising wages in anticipation of the stage 2 benefit of hiring an H-type worker in stage 1. 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 links in alleviating adverse selection.

As predicted by the theoretical literature on job referrals, the increase in market efficiency comes at the cost of an increase in inequality between L-type and H-type workers. Table 1 shows that L-type workers’ average wage level does not differ much between *Baseline* and *Referral* (average wages vary only between 21.6 and 24.3). The main negative effect on L-type workers comes from the more efficient assignment of workers to firms. Specifically, the number of L-type hires per

¹⁹All p -values reported in this sub-section are two-sided and coming from non-parametric Mann-Whitney U tests with session averages as unit of observation.

market stage drops from 2.75 to 2.4, whereas the number of H-type hires increases from 0.85 to 1.3. Thus, job referrals prevent some of the L-type workers from finding a job, as these jobs are taken by H-type workers. Unemployed L-type workers earn their reservation wage of 10 instead of the average L-type wage of 24, thereby experiencing a significant drop in earnings. At the same time, H-type workers do not benefit much, because their wage tends to be around 31, which is not much higher than the reservation wage of 30 they receive when not being hired. Together this implies that the main benefactors of employee referrals are the firms.

4.3 Individual Level Analysis: Risk Aversion and Profits

Having established statistical support for our model-based hypotheses, we now take a more detailed look at individual strategies. We pool the data from stage 1 and 2 and first look into firms' strategies followed by an exploration of workers' strategies.

The patterns of offers observed in Figure 2 suggests two main strategies firms may apply when making their offers. First, firms may target L-type workers by offering a wage below 30, even towards the end of the market opening time. According to our model, such a strategy should primarily be used by risk averse firms. Second, firms may attempt to hire H-type workers by offering wages of at least 30. This is a more risky strategy as it may lead to hires of L-type workers at relatively high wages. Consequently such a strategy should be used by relatively less risk averse firms.

The difference in riskiness between these two strategies is confirmed by the data. Figure 3 depicts the distribution of firms' profits for both strategies. When the accepted wage offer is below the H-type workers' reservation wage of 30, profits are approximately uni-modal and centered around 20 (transparent bars). Hiring at wages at or above 30 leads to a bi-modal distribution of profits, with profits mostly being either above 45 or below 10 (black bars). The frequency of hiring an H-type worker when the wage is at or above 30 is 51% in *Baseline* and 54% in the public market of the *Referral* treatment. Thus, by using one of the described strategies, participants in the role of firms essentially choose between a relatively certain, intermediate profit and a gamble with a low and a high profits outcome.

Models (1)–(4), in the regression analyses reported in Table 4 look at how firms' different strategies impact their profits. In Model (1), the coefficient of 'Acc. Offer ≥ 30 '—which indicates accepted offers of at least 30—shows that in *Baseline* average firm profit is 7.25 points higher for such offers compared to when they are below 30. Thus, firms receive a risk premium for making such high wage offers. Model (2) adds as an explanatory variable a dummy for the hired worker's productivity type. The coefficient of 'Acc. Offer ≥ 30 ' then shows that firms hiring L-type workers at wages at or above 30 earn on average 12.74 points less than firms hiring at lower wages. At the same time, firms hiring H-type workers at the higher wages earn on average 26.76




Figure 3: Firm Profits

Notes: Histogram of firm profits separated by whether the accepted wage exceed an H-type worker’s reservation wage of 30. Data includes both stages except in the third histogram as referral offers are only made in stage 2 of treatment *Referral*.

($= -12.74 + 39.50$) more than firms hiring at lower wages. Accounting for the constant, this implies that firms could either go for a certain payoff of 19.97 (when making offers around 20) or a 50-50 bet between receiving a payoff of 7.23 ($= 19.97 - 12.74$) or of 46.73 ($= 19.97 - 12.74 + 39.50$). It follows that in the *Baseline* treatment, indifference between the certain payoff and the gamble implies a CARA coefficient of 0.044 (or a CRRA coefficient of 0.89). These average risk coefficients are broadly in line with the literature on risk elicitation (see, e.g., the high-stake treatment in Holt and Laury, 2002, and Dave et al., 2010).²⁰ Models (3) and (4) report the regressions for the *Referral* treatment. Results are very similar with the additional result that offering high wages is particularly profitable when offers are made through referrals.

When deciding on their acceptance strategy, 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 either choose to accept a wage below 30 or to hold out for higher wages with a higher risk of being not hired. Models (5)–(8) in Table 4 provide information on the payoff consequences of both strategies. Models (5) and (7) show that in the public markets there is a risk premium of 5.20 (in *Baseline*) and 3.97 (in *Referral*), when holding out for a wage offer of at least 30. Model (6) adds the variable ‘Not Hired’ accounting for the effect of remaining unemployed. The constant in this model shows that accepting an offer below 30 yields on average a payoff of 21.14. This payoff is almost certain, because workers that accept low offers are almost always hired. The model also indicates that holding out for wage offers of at least 30 results in either a payoff of 32.82 ($= 21.14 + 11.68$) if hired or a payoff of 10.14 ($= 21.14 + 11.68 - 22.68$) if not hired. The latter reflects the L-type workers’ reservation wage. The probability that an L-type worker that does not accept offers below 30 is eventually hired amounts to 73%. Therefore, for an L-type worker,

²⁰Consistent with these results, in the online appendix we show that participants that are more likely to target H-type workers, on average, earn more than participants that tend to make only low offers.

Table 4: Firm and L-type Worker Profits in Dependence of Wage Offers and Hirings

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

Notes: Linear mixed effects models with subject and session random intercept; cluster-robust standard errors in parentheses, * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Bootstrap standard errors yield qualitatively identical results. Data include both stages. ‘Acc. Offer ≥ 30 ’ is a dummy variable that equals 1 if the accepted offer is not smaller than 30; ‘H-type’ is a dummy variable equal to 1 if the hired worker is of the H-type; ‘Referral Offer’ is a dummy variable that is equal to 1 if the accepted offer in stage 2 is a referral offer; ‘Not Hired’ is a dummy variable that is equal to 1 if a L-type worker is not hired in a stage and period. The reference groups for *Baseline* are accepted offers below 30 in models (1) and (5), and offers below 30 accepted by an L-type worker in models (2) and (6). The reference group for *Referral* are accepted offers below 30 in the public market models (3) and (7), and offers below 30 made in the public market and accepted by an L-type worker in models (4) and (8). Period and stage dummies are included in all regressions.

holding out for a high wage offer is equivalent to a gamble which earns 32.82 with probability 0.73 and 10.14 with probability 0.27, while accepting a low wage offer gives the option of earning 21.14 almost certainly. From this, the implied risk coefficients of L-type workers are 0.094 for CARA and 1.64 for CRRA. Similar conclusions can be drawn for the *Referral* treatment.

Tying the implied risk aversion parameters back to the model, we can explain why too many L-type workers are hired relative to the second-best efficient outcome.²¹ In particular, for the derived CARA parameters between 0.044 and 0.094, Figure 1 shows that the average wage predicted by the model is below 30. This implies that in equilibrium firms and L-type workers with these risk preferences start by offering and accepting, respectively, low wages and only once a fraction of the L-type workers has accepted does the model predict wage offers ≥ 30 , in line with the data.

Interestingly, the implied degree of risk aversion appears to be substantially higher for L-type workers than for firms. A possible explanation for this might be that L-type workers underestimate

²¹We refer to second-best efficiency here as the relevant benchmark because in environments with adverse selection first-best outcomes can typically not be supported in equilibrium by any mechanism (Samuelson, 1984).

the likelihood of being hired when holding out for a wage offer ≥ 30 (the H-type workers reservation wage). In that case, the risk aversion estimated from L-type workers acceptance behavior would be overestimated. To see if this holds, in Figure 4 we look at the number of workers of each type that are still on the market when wage offers reach the H-type workers reservation wage and the chance for each type to be hired at such wage offers. The black bars depict the number of workers that are still on the market when wage offers reach 30, and the gray bars show the number of workers that are eventually hired at a wage greater or equal to 30. The former show that all three H-type workers are still on the market and can be hired, whereas about two-thirds of the L-type workers have already left the market because they accepted a lower wage. At the same time, the gray bars show that the likelihood of being hired is fairly equal for H- and L-type workers, despite the fact that there are many more H-type workers available for hire.




Figure 4: Supply and Hires of Workers at (and above) H-Type Worker’s Reservation Wage

Notes: Black bars show the average number of workers (separated by treatments, stages, and productivity type) that are active on the market when wage offers reach the H-type workers’ reservation wage of 30. Gray bars show the average number of workers that are hired at wages at or above 30.

An intuitive reason for this is that for wages close to H-type workers’ reservation wage the gains from accepting an offer are much larger for L-type workers than for H-type workers. Thus, L-type workers tend to be much quicker than H-type workers in accepting offers of 30 or slightly above. If L-type workers do not fully anticipate this possibility they may underestimate the probability of being hired when holding out for a high wage, which indeed could explain the estimated relatively high risk aversion of these workers.²² For completeness we note that in stage 2 of the *Referral* treatment (panel ‘Referral, Stage 2’ in the figure) the number of H-type workers hired at wages of at least 30 is larger than the number of L-type workers. This happens because, due to the social link between workers, in this case referral offers are targeted primarily at H-type workers.

²²Another more speculative explanation relates to the difference in the nature of the gamble firms and workers are facing on the market. Specifically, a worker may derive disutility from the outcome where she is not hired, beyond the immediate payoff consequences. Such a “displeasure of being unemployed” is reminiscent of the opposite idea of the “joy of winning” in the auction literature (e.g. Dohmen et al., 2011).

5 Conclusion

We have investigated whether employers, by utilizing the social links of their workers, succeed in alleviating the negative effects of information asymmetry in labor markets. In our experimental and theoretical setup, we mirror the empirical observation that employee referrals (which operate through social links) can often be informative about the productivity of prospective workers. Our model predicts important effects of employee referrals for aggregate market outcomes, including higher average productivity and wage premiums internalizing the benefits of employees' social links. Our empirical results provide strong support for the theoretical hypotheses. In particular, we document significant treatment effects on efficiency and wages when comparing experimental markets featuring employee referrals to markets without employee referrals.

In addition, two empirical results stand out. First, we find that employee referrals lead to higher wages not only for new workers that are hired via referral offers, but also for incumbent workers that have been hired in the competitive market with publicly posted offers. While this finding is predicted by the model—firms are willing to raise wage offers in the competitive market in order to increase the probability of hiring high-productivity workers, which are more likely to have a high-productivity social link—it requires taking into account non-trivial trade offs on the part of the firms.

Second, we find that, despite the benefits of social links, market efficiency still falls short of the second-best outcome. This can be explained by risk aversion, inducing employers to predominantly start out by making low wage offers, which only low-productivity employees are willing to accept. In addition, our data indicate that low-productivity workers' willingness to accept such offers is 'too high', in the sense that this behavior can only be explained by substantially higher levels of risk aversion than those observed for individuals in the role of employers. This is likely due to a novel empirical regularity that we document for high wage offers, and which is insufficiently anticipated by low-productivity workers. Specifically, low-productivity workers are over-proportionally hired at wages that exceed the high-productivity workers' reservation wage (i.e., low-productivity workers are hired at such wages substantially more often than their empirical frequency on the market when such offers are made). The reason for this result lies in the market dynamics where low-productivity workers are quick at accepting such offers, whereas high-productivity workers are less eager to do so, because such wage offers are usually only slightly higher than their reservation wage (i.e., the opportunity cost of delaying acceptance is lower for high-productivity types).

Our study has implications that go beyond the labor market setting. It broadly contributes to the question of how to promote efficiency in markets impaired by asymmetric information, e.g., doctor-patient relationships or second-hand markets. From a market design perspective, our results highlight the benefits of accounting for and enabling the use of the social networks.

The behavioral effects we uncover also play an important role, because they point to potential issues that are specific to dynamic (as opposed to static) markets: risk aversion leading to low opening offers and thus low efficiency levels in dynamic markets and the failure of individuals to internalize a correlation between productivity and the eagerness of an individual to accept offers quickly.

Our study also contributes to a research agenda in experimental economics investigating competitive markets in the presence of incomplete information,²³ and in particular to a number of studies exploring different tools that can alleviate informational inefficiencies. A common finding in this literature is that the degree of equilibrium play tends to be high. That is, people often recognize and utilize the strategic opportunities offered by mechanisms such as quality signaling (Miller and Plott, 1985; Kübler et al., 2008), screening contracts (Cabrales et al., 2011; Hoppe and Schmitz, 2015; Mimra and Waibel, 2017), cheap-talk (Siegenthaler, 2017), or bargaining (Bochet and Siegenthaler, 2018, 2019).²⁴ Our results demonstrate that the same applies in the context of employee referrals. Given the prevalence of referral hiring in labor markets, our results contribute to closing a gap in this literature.

Montgomery (1991b) and Calvo-Armengol and Jackson (2004) present theoretical network models where individuals that are not well-connected are at a disadvantage and thus networks can perpetuate inequality. Indeed, recent evidence suggests that recruitment via job networks may disadvantage qualified women (Beaman et al., 2018). In line with these studies, our results show some increase in inequality among workers when job referrals are possible. At the same time it stands to reason that social groups that are not well-connected (e.g., first-time job seekers or immigrants) are also groups of which often little objective information is available. Employee referrals, as a source of credible information, could thus be particularly effective for such groups and may even mitigate labor market discrimination. Research for a better understanding of the mechanisms that mitigate or exacerbate discrimination in referral networks could be an interesting future research line.

²³Seminal contributions include Smith (1962, 1965, 1982), Plott and Smith (1978), Fehr et al. (1993) and Fehr and Falk (1999).

²⁴Other 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 showing that reciprocal fairness may overcome moral hazard. Bartling et al. (2012) and Brown et al. (2004) examine the difference between trust and control-based contracts on labor market outcomes, and Schram et al. (2010) compare how firms choose between hiring on a competitive market versus hiring through bilateral negotiations. Andreoni (2017) finds that selling goods with a “satisfaction guarantee” helps mitigate problems of moral hazard.

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,” *American Economic Review*, 1984, *74* (2), 79–83.
- and, Emilio J Fernandez Roberto M and Paul Moore**, “Social capital at work: Networks and employment at a phone center,” *American Journal of Sociology*, 2000, *105* (5), 1288–1356.
- Andreoni, James**, “Satisfaction Guaranteed: When Moral Hazard meets Moral Preferences,” *American Economic Journal: Microeconomics*, 2017, *74* (5), 1365–1384.
- Arrow, Kenneth J.**, “Uncertainty and the welfare economics of medical care,” *American Economic Review*, 1963, *53* (5), 941–973.
- Barr, Tavis, Raicho Bojilov, and Lalith Munasinghe**, “Referrals and Search Efficiency: Who Learns What and When?,” *Journal of Labor Economics*, 2019, *37* (4), forthcoming.
- Bartling, Björn, Ernst Fehr, and Klaus M Schmidt**, “Screening, competition, and job design: Economic origins of good jobs,” *American Economic Review*, 2012, *102* (2), 834–64.
- 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.
- , **Niall Keleher, and Jeremy Magruder**, “Do job networks disadvantage women? Evidence from a recruitment experiment in Malawi,” *Journal of Labor Economics*, 2018, *36* (1), 121–157.
- 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.
- and – , “Price Transparency and Competition in the Market for Lemons: Experimental Evidence,” *Available at SSRN 2940110*, 2019.
- Boorman, Scott A.**, “A combinatorial optimization model for transmission of job information through contact networks,” *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.
- Cabrales, Antonio, Gary Charness, and Marie Claire Villeval**, “Hidden information, bargaining power, and efficiency: an experiment,” *Experimental Economics*, 2011, *14* (2), 133–159.
- 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,” *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.
- Castilla, Emilio J**, “Social networks and employee performance in a call center,” *American Journal of Sociology*, 2005, *110* (5), 1243–1283.
- Cooper, David J. and John H. Kagel**, “Other regarding preferences: A selective survey of experimental results,” *The Handbook of Experimental Economics*, 2016, *2*, 217–289.
- 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.
- Ekinci, Emre**, “Employee referrals as a screening device,” *The RAND Journal of Economics*, 2016, *47* (3), 688–708.
- 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.

- Falk, Armin and Nora Szech**, “Morals and markets,” *Science*, 2013, *340* (6133), 707–711.
- , **Stephan Meier, and Christian Zehnder**, “Do lab experiments misrepresent social preferences? The case of self-selected student samples,” *Journal of the European Economic Association*, 2013, *11* (4), 839–852.
- Fehr, Ernst and Armin Falk**, “Wage rigidity in a competitive incomplete contract market,” *Journal of Political Economy*, 1999, *107* (1), 106–134.
- , **Georg Kirchsteiger, and Arno Riedl**, “Does fairness prevent market clearing? An experimental investigation,” *The Quarterly Journal of Economics*, 1993, *108* (2), 437–459.
- , – , and – , “Gift exchange and reciprocity in competitive experimental markets,” *European Economic Review*, 1998, *42* (1), 1–34.
- , **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.
- , “Referral networks and inequality,” *Available at SSRN 2768083*, 2018.
- 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.
- Greenwald, Bruce C**, “Adverse Selection in the Labour Market,” *The Review of Economic Studies*, 1986, *53* (3), 325–347.
- 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*, 2018, *126* (4), 1691–1746.
- 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,” *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.
- Levitt, Steven D and John A List**, “What do laboratory experiments measuring social preferences reveal about the real world?,” *Journal of Economic Perspectives*, 2007, *21* (2), 153–174.
- 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.
- Miller, Ross M and Charles R Plott**, “Product quality signaling in experimental markets,” *Econometrica*, 1985, pp. 837–872.
- 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,” *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.
- Obukhova, Elena and George Lan**, “Do job seekers benefit from contacts? A direct test with contemporaneous searches,” *Management Science*, 2013, *59* (10), 2204–2216.
- 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.
- Pieper, Jenna R, Charlie O Trevor, Ingo Weller, and Dennis Duchon**, “Referral hire presence implications for referrer turnover and job performance,” *Journal of Management*, 2019, *45* (5), 1858–1888.
- Pissarides, Christopher A**, *Equilibrium unemployment theory*, MIT press, 2000.
- Plott, Charles R and Vernon L Smith**, “An experimental examination of two exchange institutions,” *The Review of Economic Studies*, 1978, *45* (1), 133–153.
- Rees, Albert**, “Information networks in labor markets,” *The American Economic Review*, 1966, *56* (1/2), 559–566.
- Riedl, Arno and Paul Smeets**, “Why do investors hold socially responsible mutual funds?,” *Journal of Finance*, 2017, *72* (6), 2505–2550.
- Rubineau, Brian and Roberto M Fernandez**, “Missing links: Referrer behavior and job segregation,” *Management Science*, 2013, *59* (11), 2470–2489.
- Samuelson, William**, “Bargaining under asymmetric information,” *Econometrica*, 1984, pp. 995–1005.
- Schmutte, Ian M.**, “Job referral networks and the determination of earnings in local labor markets,” *Journal of Labor Economics*, 2014, *33* (1), 1–32.
- Schram, Arthur, Jordi Brandts, and Klarita Gërxhani**, “Information, bilateral negotiations, and worker recruitment,” *European Economic Review*, 2010, *54* (8), 1035–1058.
- 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.
- Smith, Vernon L.**, “An experimental study of competitive market behavior,” *Journal of Political Economy*, 1962, *70* (2), 111–137.

- Smith, Vernon L**, “Experimental auction markets and the Walrasian hypothesis,” *Journal of Political Economy*, 1965, 73 (4), 387–393.
- , “Microeconomic systems as an experimental science,” *American Economic Review*, 1982, 72 (5), 923–955.
- 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.
- Waldman, Michael**, “Job assignments, signalling, and efficiency,” *The RAND Journal of Economics*, 1984, 15 (2), 255–267.
- Weller, Ingo, Brooks C Holtom, Wenzel Matiaske, and Thomas Mellewig**, “Level and time effects of recruitment sources on early voluntary turnover.,” *Journal of Applied Psychology*, 2009, 94 (5), 1146.
- Wilson, Charles**, “The nature of equilibrium in markets with adverse selection,” *The Bell Journal of Economics*, 1980, pp. 108–130.
- Yakubovich, Valery and Daniela Lup**, “Stages of the recruitment process and the referrers performance effect,” *Organization Science*, 2006, 17 (6), 710–723.
- Zeltzer, Dan**, “Gender Homophily in Referral Networks: Consequences for the Medicare Physician Earnings Gap,” *American Economic Journal: Applied Economics (forthcoming)*, 2019.
- Zenou, Yves**, “Social interactions and the labor market,” *Revue d’économie politique*, 2013, 123 (3), 307–331.

Appendix

A Equilibrium Characterization

This appendix derives the market equilibria. We will construct the symmetric equilibria. 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, $P_H - \lambda_H > P_L - \lambda_L$. 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 since the stages are independent). 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 for firms 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:

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

If $q = \hat{n}_L$, the wage w solving (A.1) represents a lower bound for acceptable wages. For any lower wage, an L-type worker 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 $q = 0$ yields 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 prefer offering λ_H immediately rather than hiring an L-type worker at a wage that exceeds \bar{w}_2 :

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

where

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

denotes a firm's expected utility when offering λ_H and all workers are still in the market. The parameter B is a baseline productivity, set to $B = 20$ in the experiment.

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 $w^* \in [\underline{w}, \min(\bar{w}_1, \bar{w}_2)]$. We prove the claim:

- From (A.1) it follows directly that wages below 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. Let n_{FL} be the number of firms offering $\{w^*, \lambda_H\}$. If $n_{FL} \leq n_L$, a firm is guaranteed to hire an L-type workers with 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 in case $n_{FL} > n_L$.

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 $q = \min(\hat{n}_L, n_{FL})$. 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 the number of other L-type workers that accept w^* , 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 $w^* = w$.

Now, offering $\{\lambda_H\}$ yields an expected utility of

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

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

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

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

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

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

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

The solutions to (A.4) - (A.7) for $w^* = w$ 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 are 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 link 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 offers 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

$$\frac{n_H - \frac{n_H}{n_S} n_F \alpha}{n_S - \frac{n_H}{n_S} n_F} > 1 - \alpha \quad (\text{A.8})$$

where $\frac{n_H}{n_S} n_F$ is the expected number of 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. \square

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

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

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

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

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

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

As in the *Baseline* treatment, 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

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

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

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

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

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

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

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

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

$$U_F^2(\gamma, y) = \sum_{q=0}^{\min(\hat{n}_H, y)} \frac{\left(\mathbb{1}_{q=0} + \prod_{j=1}^{q-1} (\hat{n}_L - j) \right) \left(\mathbb{1}_{q=y} + \prod_{j=0}^{y-q-1} (n_L - j) \right)}{\prod_{j=0}^{y-1} (\hat{n}_S - j)} \binom{y}{q} \sum_{i=0}^q \gamma^i (1-\gamma)^{q-i} \binom{q}{i} U_F^2(n_F - i) \quad (\text{A.16})$$

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 term

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, as it is 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 links that come with hiring an H-type worker. It solves

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

where

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

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 only a high offer in stage 1, hoping to hire a referral worker in stage 2.²⁵

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

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

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 symmetric market equilibrium.

²⁵The 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).