Do employers trust workers too little?
An experimental study of trust in the labour market*

Stefano A. Caria† Paolo Falco‡

Nov 2018

Abstract

We conduct a field experiment to investigate employers' trust in workers. A sample of real entrepreneurs and workers from urban Ghana are respectively assigned to the roles of employers and employees in the experiment. Employers have the option to hire (trust) an employee, who can in turn choose whether to exert effort (trustworthiness) in a real-effort task with monetary payoffs. We elicit employers and employees' expectations about each other's choices and randomly provide information on previous behaviour. We find that employers significantly underestimate workers' effort. This reduces hiring and profits. We also find that expectations respond to negative signals, but are inelastic with respect to positive news. Finally, we find that experimental behaviour relates directly to real-life choices: employers who hire less in their businesses have more pessimistic expectations and are less likely to trust in the experiment. Our evidence corroborates the hypothesis that an equilibrium with no experimentation and biased beliefs may be self-sustaining.

Keywords: hiring, effort, trust, discrimination, productivity, microenterprise, expectations, learning, experiment, African labour markets.
JEL: J23, J71, O15, C9.

*We are grateful to Abigail Barr, Andrea Bassanini, Iris Bohnet, Helene Bie Lilleør, Federico Cingano, Marcel Fafchamps, Marco Piovesan, Imran Rasul and attendants of the CESS seminar (Oxford), the OECD Applied Economics WiP seminar, the 2014 Royal Economic Society conference, the University of Copenhagen seminars, the 2013 SEEDEC and the Rockwool Foundation seminar (Copenhagen) for their invaluable comments. We are also indebted to Jono Lain, Poorna Mazumdar, Laura Rossouw, Vaclav Tehle and Asmus Zoch for their excellent research assistance, and to the CSAE enumerators in Ghana led by Moses Awoonor-Williams. We are very grateful to the Bill and Melinda Gates Foundation for financial support. All errors are our own.

†University of Oxford; stefano.caria@economics.ox.ac.uk
‡OECD and University of Oxford; paolo.falco@oecd.org

1 Introduction

In many developing countries the economy is dominated by a large number of small firms that employ very few workers, if any (Hsieh and Olken, 2014). These firms are often unable to exploit economies of scale and find it hard to compete in large markets. A host of factors may prevent microenterprises from expanding, but not all of them are well understood. One constraint highlighted in the literature is that entrepreneurs may not trust workers to exert effort on the job (Heath, 2013). Monitoring and rewarding performance is costly and incentive schemes are often resisted by workers, leaving the firm with few options to increase worker productivity (Atkin et al., 2015; Breza et al., 2015). These observations are consistent with the fact that small firms mostly hire through personal connections, which, it has been argued, can be an attempt to reduce moral hazard through peer screening and monitoring (Montgomery, 1991; de Mel et al., 2010; Heath, 2013). They are also consistent with the low take-up observed in a number of wage-subsidy interventions, which indicates that the perceived marginal product of labour is low (Galasso et al., 2004; de Mel et al., 2010; Levinsohn et al., 2012; Groh et al., 2016).

Is low trust in workers justified? There is extremely little evidence to answer this question. However, the recent literature on biased beliefs suggests a number of mechanisms that may make employers overly pessimistic about the trustworthiness of workers. First, trust presents a classic problem of experimentation: a pessimistic manager will not hire any workers and hence will not receive new information to update his or her incorrect priors (Rothschild, 1974; Butler et al., 2009). Further, those managers that receive new information may fail to adjust their beliefs in a rational way (Enke and Zimmermann, 2013; Ambuehl and Li, 2014; Falk and Zimmermann, 2015; Hoffman, 2016). In particular, under loss or disappointment aversion, employers may treat positive signals more cautiously than negative signals, as negative surprises will affect personal utility more than positive surprises (Gill and Prowse, 2012). Biased beliefs of this kind may have important negative economic implications. For example, overly pessimistic expectations of workers' trustworthiness may depress hiring and overall employment.

In this paper, we study whether firm managers are overly pessimistic about the trust-

worthiness of workers. Further, we explore a key mechanism that can lead to pessimism: asymmetric updating in response to new information. To achieve these objectives, we design a trust experiment with a real-effort task, played between *real entrepreneurs* and *real workers* in urban Ghana. Employers have the option to trust an anonymous employee by investing an initial endowment to pay his/her wage. If hired, the employee can choose whether or not to complete a simple task that requires effort but does not require particular skills or ability. The employer's payoff is directly linked to the employee's performance, while the employee's wage, if hired, is fixed. By decoupling the worker's pay from his/her performance and by removing the employer's ability to monitor and punish low effort, we are able to elicit trust and trustworthiness. In addition, we elicit employers' expectations of workes' trustworthiness by means of simple questions requiring no knowledge of probability. We compare these expectations to actual workers' performance to identify potential misperceptions. We further devise two randomized treatments to study the extent to which expectations are biased against specific worker categories and to estimate the elasticity of expectations with respect to new information.

We find that employers significantly underestimate workers' trustworthiness. We also find that employers have biased expectations about the relative trustworthiness of different groups of employees. In particular, male employers expect female workers to be *less* trustworthy then men, while in the experiment female workers are actually *more* trustworthy then men. These low expectations have significant negative repercussions on hiring and cause a loss of profit for employers.

Furthermore, we show evidence of asymmetric updating. Negative information on workers' performance has a stronger negative effect on employers' beliefs than positive signals. This makes undesirable equilibria of low expectations and low trust particularly difficult to escape. When we explore the correlation between *real life* hiring decisions of the entrepreneurs in our experiment and their elicited beliefs, we find that hiring more labour is associated with more positive expectations of workers' trustworthiness and, in turn, with a higher propensity to trust in the game. This corroborates the external validity of our results and gives further evidence in favour of the hypothesis that lack of experimentation can lead

to a self-sustaining equilibrium with overly pessimistic beliefs.

Our results contribute to the literature that studies the quality of management in developing countries (Bloom and Van Reenen, 2007; Bloom et al., 2011; McKenzie and Woodruff, 2015). We advance this literature by using experimental techniques to isolate biases in beliefs and updating. In a related gift-exchange lab experiment with Ghanian students, Davies and Fafchamps (2015) show that subjects fail to screen out partners who have been untrustworthy in the past. Our results support the idea that public policy, in the form of training, information campaigns, and enhanced systems for information sharing, may be useful to correct employers' misperceptions and encourage employment.

We also contribute to the literature that studies expectations and their formation (Enke and Zimmermann, 2013; Falk and Zimmermann, 2015). In a controlled setting, we document that firm mangers hold biased expectations about workers' trustworthiness and that these incorrect assessments are associated with suboptimal decisions in the experiment. In related work, McKenzie et al. (2007) show that migrants underestimate the returns to migration, Jensen (2010) documents that students underestimate the returns to education, and Cruces et al. (2013) demonstrate that individuals hold incorrect beliefs about the income distribution. There is also a small literature that studies subjects' ability to predict experimental results (DellaVigna and Pope, 2016). Our contribution is to show that the forecasts of the actors in our study are inaccurate and that "wisdom-of-the-crowd" effects may be context-specific.

Finally, we show evidence supporting the existence of a mechanism that can contribute to biased expectations – asymmetric updating. The existing literature has highlighted that subjects give more weight to positive feedback about their personal ability and less weight to negative feedback (Mobius et al., 2011). We show that when it comes to forecasting the behaviour of others, subjects assign greater weight to negative information. This asymmetric updating is consistent with asymmetric weighting of gains and losses in the utility function. There is currently much interest in the literature in understanding how the payoffs of actions affect information acquisition and beliefs (Ambuehl, 2015) In future work, we plan to explore

the mechanism we have uncovered with further experimental evidence and a structural model.

The remainder of the paper is structured as follows. In section 2 we motivate our choice of Ghana as the context of our study by providing some stylised facts on trust and on its relationship with business expansion. In section 3 we outline our experimental design. In section 4 we describe our sampling strategy and provide some key descriptive statistics. In section 5 we present our results. Section 6 discusses our findings and our plans for future work.

2 A low-trust country

The experiment was conducted in Ghana, a country that a priori seemed particularly well-suited to study the implications of poor trust for the labour market. Two pieces of descriptive evidence provide support for this view and motivate our analysis.

First, data from the World Value Survey (WVS) shows that trust in Ghana is particularly low. WVS contains harmonised data on trust in 2014 for 60 different countries across all the continents.[1] Out of all the countries analysed, Ghana is the one with the fourth lowest level of trust (Figure 1).

<< Figure 1 here >>

Second, cross-country evidence shows a positive correlation between entrepreneurs' willingness to expand their business and trust in society. The 2014 Global Entrepreneur Monitor (GEM) dataset collects unique cross-country data on entrepreneurs. Crucial for our purposes, the dataset tracks the share of entrepreneurs who have just started a business that expect to hire at least five workers in the coming five years. The relationship between that variable and aggregate levels of trust is positive at the low levels of trust recorded in Ghana.[2] Coupled with the observation that the large majority of business in Ghana are very small and rarely hire any labour (Falco et al., 2014), this descriptive evidence suggests that Ghana constitutes an interesting setting for our study.[3]

[1] Trust is measured by means of a standard question that asks whether the respondent *thinks that others can be trusted*.

[2] The relationship is concave and, at high levels of trust, the positive correlation ceases to exist.

[3] It is important to underline that we did not embark in this exercise with the aim to uncover the main

3 Experimental design

The experiment was conducted in Accra and took place in a central location that was easily reached from different parts of the city.[4] Each respondent was pre-assigned to one of two experimental roles: *employer* or *employee* (crucially, as discussed below, real entrepreneurs were assigned to the role of employers) and randomly allocated to one of 30 experimental sessions. Each session had 10 employers and 10 employees in total.[5] Within each session, the two groups were kept separate and participated in parallel activities. Two separate locations, sufficiently distant from each other, were equipped to host the employers' and the employees' room respectively.[6] Inside each room, respondents sat at separate individual desks, which were equipped with custom-designed wooden screens that ensured respondents' choices in the game could not be observed by their peers. The screens in the employees' room were higher, to ensure that employees could not see each other, minimising peer effects on productivity (see (Falk and Ichino, 2006) and (Mas and Moretti, 2009) for recent evidence on the effect of peers' pressure on workers' effort).[7] They also ensured that respondents did not feel observed or monitored by the experimenter. All the features of the employees' room were outlined to employers to ensure full common knowledge of the experimental setup.

3.1 The game

Upon reaching the venue, the two groups were introduced to the structure of the experiment. For the sake of clarity, the following explanation provides a stylised representation of the rules. The instructions respondents received were less schematic and illustrated by means of examples.

cause of low enterprise growth in Ghana. The exercise we conduct is nonetheless important to uncover a potentially important factor that has received little attention in the literature on firm growth. Figure 11 in the appendix D provides a diagramatic representation of how low expectations of employees' trustworthiness on the part of employers may depress hiring.

[4]It was carried out in the summer of 2013 and piloted in the preceding months.

[5]There were few exceptions, where some respondents dropped out and the session was not full.

[6]The distance between the rooms, which were located on two different floors of a large compound, helped us to prevent respondents in different roles from getting to know each other prior to the experiment. Moreover, the regular flow of people in and out of the compound, which hosts a number of different offices and businesses, made it virtually impossible for respondents to infer information about the pool of participants in the other room based on the features of the people they met upon arrival.

[7]Moreover, by preventing workers from knowing which one of their peers was employed in each round, we were able to exclude direct effects of the unemployment rate in each session on effort, as predicted by a model of unemployment as a worker discipline device (Shapiro and Stiglitz, 1984).

Each respondent in the role of employer j owns a firm producing y and in each round of the game has an initial endowment of R, which he/she can choose to keep or invest in a worker's salary. If they decide to hire a worker, production of y occurs through a real effort task (discussed below) by an anonymous worker i (located in the other room), according to the following production function whose only input is worker effort e:

$$y_j = f(e_i) = e_i \quad where \quad e_i = 0, 1 \tag{1}$$

If the worker decides to exert effort, production is realised. If the worker shirks, production is zero. This binary setup is the key to a simple elicitation of employers' expectations and it ensures high levels of understanding in the game. Crucially, the triviality of the real-effort task that was chosen ensures that completion of the jobs is strictly a function of effort (see the next section for a detailed discussion). Employers have full knowledge about that.

When the worker chooses to exert an effort and production is successful (i.e. $e_i = 1 \rightarrow y = 1$), the employer receives a value of $p > R$ and the worker receives a wage of W. When the employee shirks and production fails (i.e. $e_i = 0 \rightarrow y = 0$), the employer does not receive any money, while still having to pay W to the employee. Our chosen payoffs were such that $R = W$ (i.e. a trusting employer loses all his endowment when his/her employee shirks). While this form of contractual arrangement may seem unrealistic, it was important to isolate trust as a driver of employers' choices (discussed below).

Based on these rules, employers were asked to choose between hiring an employee and keeping their endowment. For every employer that chose to hire, an employee in the other room was asked to undertake the real-effort task and could decide whether to exert effort or shirk.[8] The game was played twice, with a second round announced as a surprise at the end of the first one. The results from the first round (i.e. whether a hired worker revealed to be trustworthy or not) were not announced to employers until the very end of the game (to avoid influencing the outcome of their second decision). Before each decision, respondents'

[8]The match between employers and employees was random and employers did not know who their employee would be upon making their decision. The respondents in the employees' room who were not hired remained unemployed and only received a participation fee.

expecations of the other party's behaviour were elicited, as discussed below. They will be central to the analysis in this paper. A complete outline of the experiment protocol is provided in Appendix C.

Before the outcome (i.e. the workers' performance) in the two rounds was revealed, employers took part in a *risk game* and in a *dictator game*. The risk-game follows Gaechter et al. (2010) and it consists of 5 consecutive dichotomous choices between a lottery with two possible outcomes, each carrying a probability of 50%, and a safe option worth a fixed amount. For each of the five choices the respondent was asked whether he/she would like to play the lottery or receive the fixed amount. By progressively decreasing the value of the negative outcome, the lottery becomes increasingly risky, leading respondents to 'switch' to the safe option.[9] Figure 10 in Appendix C shows the visual aid used to explain the lotteries. For comparability with our main trust game, the value of the safe option was equalised to the initial endowment employers could invest to hire a worker. Moreover, in order to mimic the hiring decision as closely as possible, the risk-game was framed in the *loss domain*, with employers receiving an initial endowment they could either *keep* or *risk*, in the same way as they could keep their endowment or hire a worker in the trust game. Our measure of risk-preferences, therefore, is more precisely a measure of 'risk aversion in the loss domain'. In the current draft, we construct this measure by simply counting the number of choices (out of 5) when the respondent preferred the safe option over the risky lottery.[10]

In the dictator game employers received an endowment of 3 Cedis, which they were asked to split between themselves and a randomly chosen anonymous employee.[11] The results of the dictator game are not relevant to the analysis in this draft and will be the subject of future work.

[9]Using a standard strategy method, we asked respondents to make a choice for each of the 5 decisions, knowing that only one of them would be randomly acted out in the end determining their prize.

[10]Given that the large majority of the respondents played according to a monotone 'switching rule' (i.e. once they switched from the lottery to the safe option they never switched back in subsequent choices in the sequence of five), this method is a consistent way of categorising workers. Further checks will be necessary to test the robustness of our results to the exclusion of workers who do not switch monotonously in the choice sequence.

[11]In half of the session the receiving employee would be randomly chosen among those who were *successful* in the task (i.e. exerted high effort) in the task. In the other half, the receiver would be chosen among *unsuccessful* workers (i.e. those who exerted low effort in the task). Employers were informed about this. The split was devised to study the role of reciprocity in altruistic choices.

3.2 A schematic representation of respondents' choices

In order to model the agents' decisions and facilitate the analysis of the game, we assume that each employee has an *unobservable* cost of effort equal to c_i and, for simplicity, he/she attaches an unobservable value of b_i to being trustworthy (e.g. intrinsic motivation). These last two terms are unknown to the employer, who has beliefs on their relative strength. These beliefs play a crucial part in the employers' decision to trust.

The game tree in Figure 2 outlines the sequence of choices and corresponding payoffs faced by employers and employees in each round. Figure 3 shows the same game tree with the chosen monetary payoffs (in Ghana Cedis) replacing the game parameters. The payoffs were large with respect to hourly earnings in the reference population, hence generating strong incentives in the game.[12]

<< **Figure 2 here** >>

<< **Figure 3 here** >>

3.3 The real effort task

We designed a trivial task, such that any worker who was willing to make an effort should succeed (i.e. reveal trustworthy). The task did not rely on any specific know-how or skills, ruling out the influence of human capital, and success was driven by workers' effort (as opposed to ability). This was explicitly and repeatedly explained to the employers, who were fully aware that any worker could attain success if he/she was willing to work steadily without shirking.

We chose a real effort task, as in Fahr and Irlenbusch (2000), van Dijk et al. (2001), Falk and Ichino (2006), as opposed to a chosen effort task (e.g. Fehr et al. (1993), Altmann et al. (2012)) for one fundamental reason. We are interested in the potential mismatch between

[12] In 2012 (the latest available year of data at the time of the experiment) the median of net *daily* earnings for workers in the reference population was approximately 6 Cedis (calculated using Oxford GHUPS data). The endowment given to employers for *each* round of the hiring game *and* for a risk game played at the end was equal to 5 Cedis. A dictator game was also played and it provided an additional endowment of 3 Cedis. The result was a *total endowment of 18 Cedis* for the three games combined. The entire experiment lasted approximately 2 hours.

expected and revealed trustworthiness in a real-effort setting (akin to real employment); and, while shown to be correlated in laboratory settings, real effort and chosen effort may diverge substantially. This may be due, for instance, to the direct utility/disutility of work and effort. Indeed, in a lab-experiment designed to directly compare the outcomes of chosen and real-effort experiments, Bruggen and Strobel (2007) show that effort is significantly higher in real-effort tasks. Our chosen task was the following:

"Starting from a bag with 3 types of beans, your job will be to sort the beans into three smaller bags, each containing only one type, in 10 minutes"

Each bag contained 350 Grams of beans in total (1/3 of each type). This was our best estimate of a *minimum attainable amount under steady effort and no shirking.*[13]

The task was trivial, but it required constant effort and attention. In order to minimise uncertainty about one's optimal effort to succeed, workers were explicitly told in the instructions that by applying a steady level of effort throughout the period, they should be able to complete the task'. Moreover, in order to rule out the possibility that workers may work 'out of boredom', the experiment venue was equipped with a TV screen showing a popular local show and workers had an implicit choice between working on the task and watching TV. This choice was not recorded explicitly (i.e. respondents were not explicitly asked whether they wanted to work or watch TV), as we believe its measurement would have been severely biased.[14]

Given the features of the task, we confidently conclude that workers' success is a valid proxy of their effort. It should be noted, however, that if ability (or other workers' characteristics besides their willingness to make an effort) play a significant role in their rate of success, our estimated level of trustworthiness will be a lower bound for actual trustworthiness. In other words, if some workers tried their best (i.e. were trustworthy), but

[13]The chosen quantity was equal to the *minimum* amount of mixed beans sorted in 10 minutes by a sample of trusted survey staff, *who were instructed to work at a regular pace and were constantly monitored.* These subjects were of mean age close to the survey average, half male and half female. Direct observation of respondents in pilot sessions confirmed that completion of the task was attainable for a wide range of subjects under steady effort.

[14]The respondents' propensity to choose leisure when observed/asked by the experimenter might have been lower than when he/she was behind the surrounding screen and was not observed.

due to lack of skills were unable to complete (and will be mistakenly classified as untrustworthy), overall estimated trustworthiness will be biased downwards. This possibility only strengthens the results in the next sections. Indeed, despite this potential downward bias, we will show high levels of revealed trustworthiness, which are significantly underestimated by employers.

Exerting effort, in this context, represents a measure of trustworthiness since it captures workers' willingness to carry out a costly task, for which they have been paid, in a situation when there is no punishment for shirking. We believe this is a relevant factor in hiring decisions, as it applies to many situations in which employers cannot closely monitor the conduct of their employees. Of course, it can be argued that pure trustworthiness of this kind may not be the primary driver of the decisions made by workers, who are instead motivated by tangible rewards and punishments that typically do exist in the labour market. We do not aim to dispute that argument. However, in light of the well-researched links between trust and economic development, we believe that our mechanism of interest deserves rigorous investigation. To attain that objective, we need a clear setup where pure trustworthiness is salient.

3.4 Eliciting expectations

The focus of this paper is on employers' expectations of workers' trustworthiness and on employees' expectations of workers' trusting. In this section we outline our procedure for eliciting these expectations.

In each round of the experiment, before or after the hiring decision was made (the order was randomized across sessions to control for potential 'order effects'), we asked employers the following question: *"Out of the 10 workers in the other room, one of whom will be assigned to you by chance, how many do you think will complete the task successfully?"* The reverse question was asked to employees: *"Out of the 10 employers in the other room, to one of whom you will be assigned by chance, how many do you think will choose to hire a worker?"* Answers to these questions were gathered by means of visual aids consisting of 10 tokens (small yellow plastic disks) that respondents were asked to distribute between a blue and a red circle printed on a game card (respectively representing success and failure). Our

design was inspired by a methodological study by Delavande et al. (2011a), and by the prior work of Manski (2004) and Attanasio and Kaufmann (2009). Moreover, having experimental sessions with exactly 10 respondents per room allowed us to simplify the elicitation strategy used in much existing work, veering away from questions about 'probabilities' and towards simpler questions about 'frequencies'.[15]

The elicitation questions could not be independently incentivised. While we acknowledge the potential limitations of this design feature, a number of important factors prevented us from introducing such incentives. The first concern was that respondents could use these questions to hedge against their choices in the trust game, causing severe distortions to the elicited belief distribution and creating a spurious correlation between beliefs and trusting.[16]

Second, we strongly wished to minimise complexity, especially since the results of pilot sessions had clearly indicated that we should attempt to contain the cognitive burden placed upon our respondents (whose average levels of numeracy are low). Introducing an additional scoring rule for the elicitation questions appeared to go directly against this principle and we had strong reasons to believe it might have increased complexity to the point of lowering overall understanding.

Third, a recent review of the literature on belief-elicitation in developing countries concludes that non-incentivised questions are effective in capturing expectations over a wide range of outcomes and that no conclusive evidence exists to indicate the superiority of incentivised tasks (Delavande et al. (2011b)). In fact, there may be reasons to believe that monetary incentives attributed by means of a scoring rule may distort the elicited distribution depending on the rule that is applied. A methodological study by Gächter and Renner (2010) with subject from developed countries (UK and Germany) in the context

[15]Under the former, more common, approach, a typical question would be "On a scale from 1 to 10, how likely is it that a worker/employer...?". Under our approach, the question of interest becomes "Out of 10 workers/employers in the other room, how many...? The latter does not rely on employees grasping basic probability concepts and it should hence deliver more accurate answers.

[16]A way to resolve this concern and maintain the monetary incentive would have been to devise a randomized rule whereby *either* the elicitation question *or* the trust game would determine people's payoffs at the end of the game. Piloting the experiment, however, clearly showed that adding additional features to the current design would have resulted in a significant loss of understanding.

of a public good game shows that while monetary incentives may increase the accuracy of beliefs, they may also alter the relationship between beliefs and behaviour (with mixed evidence on the direction of the resulting bias[17]). For this reason they conclude that if a researcher is interested (as we are) in the relatonship between beliefs and behaviour, belief elicitation should *not* be incentivised. In addition, looking at their results in detail, it appears that the increased precision of beliefs resulting from the incentive is symmetric around the mean (zero) of the error distribution. In other words, by introducing the monetary incentive respondents are equally *less* likely to *over*-estimate and *under*-estimate the object of interest. The implication for our work is that by incentivising the elicitation we could have tightened our estimate of expected trustworthiness, but we wouldn't have changed its central tendency. Our results could have thus been even stronger (both statistically and in magnitude, as imprecision may have led to classic attenuation bias in our regressions), but they would not have changed qualitatively. This improvement, however, would have come at the cost of hitting a binding constraint on the level of complexity we could allow without compromising understanding, and the incentive might have caused spurious behavioural changes. This is a risk we chose to avoid at the design stage.

Finally, and perhaps most importantly, upon testing our experiment in pilot sessions, we observed that *unincentivised expectations strongly predicted behavour in the incentivised trust game* (a finding that is strongly confirmed by the results in the next sections), lending support to the conclusion that people truthfully revealed their beliefs despite the elicitation itself was not incentivised.

3.5 Treatments

In addition to the basic version of the game (control), we conducted two variants (treatments) of the experiment over random sub-samples of the employer population (who were assigned to specific 'treated' sessions). This section describes the main features of the two treatments.

Treatment 1 (T1): Information provision. At the end of round 1, employers in T1 sessions

[17] While Gächter and Renner (2010) show that monetary incentives *increase* respondents' contributions in a public good game, they report a result by Croson (2000) showing the opposite.

received a note showing them whether a random employee out of all those hired in previous sessions completed the task (i.e. proved to be trustworthy) or not. In doing so, the treatment effectively conveyed to employers the same information they could gather through their own experience of hiring a worker. Crucially, we provided this information both to employers who had been willing to hire in round 1 and to those who had not, overcoming a fundamental problem of endogenous selection (which occurs in the real world) whereby new information only accrues to employers who are willing to experiment. For the sake of comparability, the employers who hired in the first round did not know the performance of their worker at the time of receiving the signal (the outcomes of both the first and the second round real-effort tasks were revealed at the end of experiment).[18] In one half of T1 sessions the random signals did not reveal any worker characteristics. In the other half, they revealed if the worker whose performance was being reported was male/female or young/old. This additional variation was meant to test whether employers' expectations are more or less elastic towards women and young workers (two groups that are considered at risk of having poorer labour market outcomes across the world (OECD, 2015)).

Treatment 2 (T2): Changing the composition of the workers' pool. In T2 sessions the sample of respondents invited to take part as 'employees' was drawn so that 80% of them would belong to one of two categories of interest. In half of T2 sessions, 80% of invited workers were *women* and, in the other half, 80% were *young* (below the age of 25).[19] Employers in these sessions were informed of the peculiar sample composition at the beginning of the game (while respondents in all the other non-T2 sessions were informed of the average gender and age composition of the invited worker pool across the experiment). The rest of the design was identical to control sessions.

[18] In addition, respondents in the role of workers were informed between round 1 and round 2 that some employers had received positive information on workers' performance (and may have updated their beliefs accordingly). This was intended to test for the potential effect of changing employers' expectations on employees' effort. This additional feature is not relevant to the analysis in the current draft and the performance of treated employees in round 2 of these sessions will be excluded from the subsequent discussion.

[19] While the composition of the 'invited' workers' pool was designed to reflect these proportions, the sample of workers who 'attended' the session may have differed. This was due to respondents changing their availability in unpredictable ways, forcing the research team to accommodate their needs by re-allocating them to different sessions. It was impossible, under these circumstances, to strictly maintain the designated proportions. However, since employers were separated from and unable to see the workers' pool, and since we do not believe employers had any reason to think the pool of attendees should be 'systematically different' from the pool of invited workers, this issue should not affect our results.

4 Sampling and descriptive statistics

4.1 Sample

One of the key objectives of this study is to measure the potential misperceptions of workers' trustworthiness among employers in Ghana. While we acknowledge the potential limitations of a comparison between laboratory outcomes and real-world decisions, the best way to achieve this goal was to assign *real entrepreneurs* and *real workers* to the experimental roles of employers and employees respectively. It seems reasonable to assume that respondents' behaviour in the lab experiment is based on the priors they develop in their actual labour market experience; and in order to give our experiment the best chances to elicit those priors, the most appealing strategy was to assortatively match respondents to their roles based on real-life experience (a similar approach is used by Barr and Zeitlin (2011) who study the incentives of primary school providers in Uganda by means of a lab experiment). Here is how we pursued this goal in practice. Using data from the *Ghana Household Urban Panel Survey (GHUPS)*, a representative survey of the urban Ghanaian population conducted by the Centre for the Study of African Economies (University of Oxford), we extracted a random sample of workers *who owned a business in at least one of the 2 most recent survey waves (2010, 2012)*. We assigned those respondents to the role of *employers*. The sample of *employees*, on the other hand, was randomly drawn from the general working-age population.[20]

4.2 Descriptive statistics

Table 1 and 2 describe the main characteristics of our samples of employers and employees respectively.

<< **Table 1 here** >>

<< **Table 2 here** >>

As expected, given our sampling strategy, employers are significantly older, more likely to be married and, given the structure of the Ghanaian economy (where female participation in

[20]Since the subset of eligible GHUPS respondents in Accra was not sufficient to attain the desired sample size, we recruited additional respondents by means of visits to random households in areas of Accra that were near but did not overlap with GHUPS enumeration areas. These additional participants were assigned to the two experimetnal roles using the same criteria described above.

entrepreneurial activities is high), the majority of them are women. Employees, on the other hand, are balanced across gender, they are younger on average, and significantly more likely to be currently unemployed. Figure 4 plots the distribution of educational attainments for the two groups. It shows that employees have slightly higher educational attainments, most likely the reflection of belonging to a younger generation.

<< **Figure 4 here** >>

Finally, Table 3 tests the balance of covariates across the different treatment groups to which employers were assigned. It shows that random assignment largely succeeded in achieving balance. In order to improve upon the precision of our estimates, the analysis in the next section will control for the most important employer characteristics.

<< **Table 3 here** >>

5 Results

5.1 Are employers' expectations realistic?

We begin by addressing the central question in our investigation: *Do employers have 'realistic' expectations of workers' trustworthiness?* Our experimental design allows us to answer this question directly, by comparing employers' elicited expectations to the revealed trustworthiness of workers (captured by their rate of success in the task).

$$H_0 : \underbrace{\{E_{j,t}[e_{i,R=1,2}] - Pr[e_{i,R=1,2} = 1]\}}_{\text{(Exp. - Actual Trustworthiness in Round 1,2)}} = 0 \tag{2}$$

***Result 1:* Employers significantly underestimate workers trustworthiness.**

Figure 5 shows our main result very clearly. In both rounds of the game the proportion of workers who reveal to be trustworthy by carrying out the task correctly is significantly higher than it is expected by employers on average. Moreover, the cumulative distribution of employers' expectations, plotted in Figure 6, reveals that the vast majority of employers underestimate employees' trustworthiness.

<< **Figure 5 here** >>

<< **Figure 6 here** >>

***Result 2:* Workers correctly predict employers' trust.**

Second, we compare employees' elicited expectations of employers' propensity to trust with the actual rate of trusting. The results are shown in Figure 7. Interestingly, they reveal that workers' average expectations correctly predict employers' trusting, since the share of employers who choose to hire is not significantly different from workers' predicted proportion.

<< **Figure 7 here** >>

Insofar as behaviour in the lab is driven by beliefs developed in the labour market, these two results point to an interesting asymmetry between employers' and employees' learning. Employees appear to be aware of employers' low expectations and low propensity to trust. On the other hand, the costs and risks of experimentation may be preventing employers from forming correct beliefs about workers' productivity. Further theoretical work will be required to place a structure on the process leading to this asymmetry.

***Result 3:* Male employers have lower expectations of female employees.**

Next, we explore the heterogeneity of employers' misperceptions by exploiting the exogenous changes in sample composition we induced in Treatment 2. Recall that at the beginning of each T2 session, employers were informed that the majority of the workers invited (80%) belonged to a specific category of interest: youth or women (as opposed to Control and T1 sessions where the sample was balanced). Hence, by comparing employers' expectations between T2 and the rest we can identify the effect of workers' characteristics on employers' beliefs. We do so by estimating the following model:

$$E_j[e_i = 1] = \alpha + \beta_W T_{2,W,j} + \beta_Y T_{2,Y,j} + \gamma X_j + u_j \tag{3}$$

where:

$T_{2,W,j} = 1$[Employer j is in T2 with majority of FEMALE workers]

$T_{2,Y,j} = 1$[Employer j is in T2 with majority of YOUNG workers]

X_j is a vector of control variables.

The results, reported in Table 4, are restricted to round 1 to conserve space. Moreover, by focusing on round 1, we are able to compare T2 to Control *and* T1 sessions jointly (since

round 1 was identical in the latter two), maximising the available sample size.[21]

When we run the estimation over the entire sample (col 1), we find that *employers have significantly lower expectations when the majority of workers in the session are women.* The result is not very strong (only statistically significant at the 10% level). However, when we split the estimation by the gender of the employer, we find a much clearer effect. The expectations of male employers are significantly lower when the sample of workers is predominantly female. The expectations of female employers, on the other hand, are not affected by the worker pool composition (col 3). The result is particularly striking when confronted with revealed trustworthiness by worker gender. Figure 8 shows that female workers are, in fact, significantly more trustworthy than men.

<< Table 4 here >>

<< Figure 8 here >>

The results in Table 4 also suggest that male employers have a positive expectations bias towards young workers, which is absent among female employers. This bias, however, is not reflected in revealed employees' trustworthiness, which shows no differential between older and younger workers (see Figure 9).

<< Figure 9 here >>

5.2 Do expectations matter for trusting?

Having presented our main results on employers' misperceptions of employees trustworthiness, we now test whether the expectations we elicited have any predictive power on the (incentivised) choice to trust a worker or not. This test follows in the footsteps of Ashraf et al. (2006), who find a significant effect of expected trustworthiness on trusting in a monetary trust game. We estimate the following model of trust, pooling observations from the two rounds:[22]

$$Pr(H_{j,R} = 1) = \alpha + \beta E_{j,R}[e_i = 1] + \gamma X_j + u_{j,R} \tag{4}$$

[21] The same patters, however, emerge in round 2.

[22] And accordingly clustering the standard errors at the specific session-round level.

where:

$Pr(H_{j,R} = 1) \equiv$ Probability of Trusting (Hiring) in round R.

$E_{j,R}[e_i = 1] \equiv$ Employer j's expectations in round R.

X_j is a vector of control variables.

Result 4: Employers' expectations are a strong predictor of trusting.

Table 5 shows that expected trustworthiness strongly and significantly affects the probability of trusting in all the proposed specifications. Increasing an employers' expectations by 10 percentage points increases the probability of trusting by nearly 3 percentage points on average. The estimated relationship is in line with the conclusions of Ashraf et al. (2006). We also find that our elicited measure of risk-preferences significantly predicts trust *independently* of expectations. This is an important result for the interpretation of our findings, since it shows that we are clearly identifying two different effects: one related to employers' perceptions of the moments of the trustworthiness distribution and a second one, which depends on the curvature of their utility function in the loss-domain. It should be noted, once again, that the loss-domain is the most relevant one for the choice at hand, since employers who choose to hire stand to 'lose' an endowment they could otherwise keep for themselves.

<< Table 5 here >>

Most importantly, by showing that hiring decisions are strongly correlated with employers' elicited expectations, the results in this section suggest that *underestimating workers' trustworthiness may lead to sub-optimal hiring and hurt employers' profits.*

5.3 Does access to information change expectations?

Next, we set out to answer the most policy-relevant questions in our analysis: *Can access to information change employers' expectationos? And how large is the impact of changing expectations on hiring?* For this part of the investigation we exploit the design features of Treatment 1. Recall that between the first and the second round of the experiment, employers in T1 received a random (and private) signal, which informed them of whether a random worker out of those who had been hired in previous sessions had completed the task. In half of T1 sessions, the signal carried no worker characteristics. In the other half, it

revealed whether the randomly drawn worker was male or female, and whether he/she was old or young (above or below 25).[23] The treatment was designed to provide employers with *the same information they could obtain by hiring a worker* (i.e. experimentation). Also recall that at the time of receiving the signal the results of the first round had not been revealed to the employers.[24] This meant that *all* employers had the same information prior to treatment and they all received a random signal, whether or not they had hired in round 1. We could therefore overcome, by experimental design, a classic problem of endogeneity whereby employers who are less prone to trusting are also the ones who are less likely to receive information from the market (and may be the ones whose expectations arehardest to change).

In order to gauge the impact of information signals on expectations, we estimate the following model:

$$E_{j,2}[e_{i,2}=1] = \alpha + \eta E_{j,2}[e_{i,2}=1] + \sum_{\kappa} \beta_{\kappa} S_{T1,j,\kappa} + \gamma T1_j + \delta X_j + \nu_j \tag{5}$$

where:

- $E_{j,R}[e_{i,2}=1] \equiv$ Employer j' expectations in each round, R = (1,2).
- $S_{T1,j,\kappa} = 1[Signal\ received\ by\ j\ in\ T1 = \kappa],\ where\ \kappa = Pos\ or\ Neg\ Sig.$
- $T1_j = 1[j\ was\ in\ session\ T1]$.[25]

[23]This added feature was designed to explore potential heterogeneity in the elasticity of beliefs with respect to new information. The distribution of signals we used was calibrated on the results of a pilot we ran prior to the main experiment, which revealed a level of trustworthiness among employees of around 60%. In order to have a balanced distribution of signals, we deviated slightly from that proportion and provided 50% of the time a positive signal and a negative one in the remaining 50% of the cases. Two employers in each session received 'no signal', to test for the sheer effect of being in a treated session. When we further cross-cut the analysis by gender and age (for half of T1), it was impossible to provide a distribution of signals that reflected the real distribution while maintaining sufficient statistical power to identify heterogenous effects. In those sessions, all the signals were positive. Since the signals were private knowledge, employers could not possibly be aware of this deviation. Moreover, given the overall high levels of employees' trustworthiness, we were not worried this minor deception could lead to significant financial losses for employers who decided to trust as a result of receiving a positive signal (which, anyway, only happened in extremely rare cases).

[24]As explained above, they were only revealed at the very end of the game, to avoid any wealth effects in round 2 caused by knowledge of having won (or lost) in round 1.

[25]Positive signals were further split into (Positive + Female worker) and (Positive + Young worker) as explained in a previous note. The coefficient on the session dummy is identified thanks to the fact that some workers in each session do not receive a signal. This allows us to test for the sheer effect of taking part in a treated session.

Result 5: ***Negative signals have a stronger (downward) effect on expectations than positive ones.***

The results are reported in Table 6. They show that positive signals have no impact on employers' expectations, while negative signals significantly lower them. A negative signal appears to decrease the expectations of an employer by 7-8 percentage points relative to one that receives no signal.[26] In column 1 we pool all the positive signals together and estimate a single average impact of positive information. Column 2, instead, splits the effects of positive signals by whether they carried information about a worker's gender or age, but finds no significant heterogeneity along these dimensions. In column 3 we drop sessions with heterogeneous signals and focus on those where the signal carried no information about worker characteristics.[27] Despite a slight drop in the significance level due to the considerable reduction in sample size, the main result remains the same.

Overall, our evidence suggests that workers may update their beliefs asymmetrically when they receive new information: they are more sensitive to negative signals than to positive ones. The reduced level of significance in some specifications warrants some caution, but, considering the small sample sizes available, we can at least consider our results prima facie evidence of an important mechanism.[28] This is an important finding, since it indicates that the scarring effects of early disappointments in an entrepreneur's career may be difficult to overcome and that undesirable equilibria of low expectations and low trust may be hard to exit.[29]

<< **Table 6 here** >>

Finally, we are aware that the amount of information provided in Treatment 1 is rather limited. Our design was driven by the idea of providing employers with the same information

[26] Thanks to our design, we are able to compare treated individuals to untreated ones who are in the same sessions but receive no signal. This allows us to control for the potential spurious effect that being in a treated session might have even when subjects receive no information.

[27] Those are the sessions that allow us to identify the impact of negative signals.

[28] One reason for not finding a stronger result may be, of course, that the new information generated by the signal about a single worker is relatively weak. Nonetheless, we detect an effect. We discuss the motivation behind our treatment design below.

[29] It also suggest that workers may have a weak incentive to signal their trustworthiness in the labour market if employers' beliefs are hard to change, potentially leading to a vicious cycle of low employee performance and low employer expectations. Further theoretical work will be necessary to characterise the mechanisms at play and the nature of such equilibria.

they could obtain from hiring their first worker (i.e. one instance of experimentation). One can easily think of more powerful interventions that may have a stronger impact on workers' expectations. For instance, one may provide information on *average* worker performance, which would be much more informative than a single observation and could have a stronger impact on expectations.

However, before designing a more powerful treatment, it would be useful to know *what is the impact of changing expectations on the rate of hiring/trusting*. The analysis in this section can help us gauge the magnitude of this impact, since we can achieve identification by using the random assignment of signals in Treatment 1 as an instrument for expectations.

Result 6: ***Raising expected trustworthiness has a strong impact on trusting***. The estimates from our IV estimation are reported in Table 7. Column 1 and 2 show the results from a naïve OLS regression of the probability of trusting as a function of expectations and individual characteristics.[30] In column 3, we instrument the expectations variable by assignment to T1 and we find that its effect grows considerably in magnitude, while remaining highly significant. The results show that increasing expected trustworthiness by 10 percentage points would increase the probability of trusting a worker by more than the same amount. These estimates suggest that a policy intervention aiming to improve employers' expectations of workers' trustworthiness may have a strong impact on their willingness to hire.

<< Table 7 here >>

5.4 Do entrepreneurs who hire more labour in *real life* have more positive beliefs?

In this final section, we explore the relationship between the results of our experiment and the *real-life* hiring patterns of entrepreneurs who took part in the game. In particular, we are interested to uncover whether those who hire employees outside the experiment have different beliefs and a different propensity to trust workers. If the hypothesis that lack of experimentation with employees prevents entrepreneurs from forming unbiased expectations

[30] We restrict the analysis to round 2 because exogenous T1-signals were only provided after the first round.

about workers' trustworthiness is correct, we should find that hiring outside the experiment helps employers to overcome their negative biases and, therefore, correlates positively with their expectations and with their propensity to trust in the game. Thanks to a brief survey carried out shortly after the experiment, we were able to obtain information on the last instance when each respondent hired someone in his or her business, and we asked for the precise number of employees he/she had at that time.[31]

Result 7: ***Employers who hire more in their real-life businesses have more positive expectations and are more likely to trust workers in the experiment.***
Table 8 shows that the number of employees entrepreneurs hire in their businesses correlates positively and significantly with both employers' expectations of employees' trustworthiness and with their propensity to trust in the experimental setting.[32] The results are robust to adding a number of control variables, including education, marital status, and current employment status, which are the best available controls for potential wealth effects (i.e. the possibility that being more affluent may drive trusting in the game and be correlated with past hiring). With regard to this concern, however, it is important to underline that past hiring not only correlates with greater trusting in the experiment, but also, crucially, with higher expectations. The elicitation of beliefs, which was not incentivised for reasons discussed above, is less likely to be directly affected by respondents' economic conditions (since it has no direct bearing on the monetary payoff of the game).[33] The result is even stronger when confining the sample to literate subjects, who had a better understanding of the follow-up questionnaire (and hence could provide a more precise account of their past hiring).[34]

[31]Unfortunately, 29 respondents could not be re-interviewed. This explains the reduction in sample size. To maximise precision and reduce costs, the survey was extremely short and only had a handful of questions on hiring that could be directly linked to the experimental setting. That limits the possibility of carrying out a more refined analysis. We focus on the last instance when the entrepreneur hired some labour to capture the effect of past learning. Ideally, we would have liked to reconstruct the entire history of past hirings by the entrepreneur, but that exercise proved to be too demanding for respondents.

[32]We focus on round 1 of the experiment, when employers display biased expectations, as shown in the previous sections.

[33]Of course, a respondent's economic condition may influence his/her beliefs in other ways). The controls used in the analysis, most importantly the worker's educational background, are meant to address some of these concerns within the strict limits imposed by a small sample.

[34]Such additional results are available upon request.

While this only constitutes descriptive evidence of the relationship between our experiment and actual hiring, the result is consistent with the idea that biased beliefs may derive from lack of experience with hiring labour and may therefore be self-sustaining.

<< **Table 8 here** >>

6 Conclusions

A growing literature documents the crucial role of trust in the process of development (Knack and Keefer, 1997; La Porta et al., 1997; Algan and Cahuc, 2010; Bohnet et al., 2010). Little emphasis, however, has been placed on the role that trust between employers and employees plays in the labour market. From a theoretical point of view, low expectations of workers' trustworthiness may depress hiring and decrease employment. From an empirical standpoint, little is known about the extent to which employers' expectations and willingness to trust an employee are a realistic reflection of workers's behaviour. This study is an attempt to determine whether employers' expectations are aligned with workers' actual trustworthiness, and to what extent such expectations can be changed through new information, stimulating higher trust.

We pursue this goal by means of an original lab experiment consisting of a binary trust-game with a real-effort task, played between *real entrepreneurs* and *real workers* in urban Ghana. In our design, entrepreneurs have the option to trust an anonymous employee by investing an initial capital endowment to pay his/her wage. If hired, the employee can choose whether or not to complete a trivial task that requires no skills or ability (and whose completion, is therefore, solely a function of effort). The employer's payoff is directly linked to the employee's performance, while the employee's wage is fixed. By decoupling the worker's pay from his/her performance and by removing the employer's ability to monitor and punish low effort, we are able to elicit trust and trustworthiness. Then, by eliciting employers' expectations of workers' trustworthiness and comparing them to actual workers' performance, we are able to identify potential misperceptions leading to sub-optimal hiring. We further devise two randomized treatments to study the extent to which expectations are biased against specific worker categories (women and youth), and to estimate the elasticity

of expectations with respect to new information.

Our main finding is that *employers significantly underestimate workers' trustworthiness, while workers' correctly predict employers' propensity to (under)trust them.* This points to the conclusion that the costs and risks of experimentation may lead employers to an undesirable equilibrium of pessimistic expectations and sub-optimal trusting. When we correlate real life hiring decisions by the entrepreneurs with the elicited beliefs in the game, we find descriptive support for this hypothesis: employers who have experimented less with hiring in the past have lower expectations of employees' trustworthiness and are less likely to trust in the game. We also find evidence of asymmetric updating: negative signals (i.e. signals of workers' untrustworthiness) have a stronger (downward) impact on employers' expectations than positive ones. This important asymmetry may crucially sustain an undesirable equilibrium of low expectations and low trust. It also indicates that negative experiences early in an employer's career may have sustained scarring effects on his/her propensity to hire later on. Finally, we estimate that raising employers' expectations would have a strong impact on their propensity to trust. Our results support the idea that public policy, in the form of training, information campaigns, and enhanced systems of information sharing, may be useful to correct employers' misperceptions and encourage hiring.

In future work, we plan to explore the mechanism we have uncovered with further experimental evidence and a structural model. More generally, we are planning to pursue a research agenda that will aim to understand the role of biased beliefs in labour market choices. For instance, we are currently designing a research project to investigate expectations mismatches among platform workers.

References

Algan, Y. and P. Cahuc (2010, December). Inherited trust and growth. *American Economic Review 100*(5), 2060–92.

Altmann, S., A. Falk, and M. Wibral (2012). Promotions and incentives: The case of multistage elimination tournaments. *Journal of Labor Economics 30*(1), 149 – 174.

Ambuehl, S. (2015). An offer you can?t refuse. *Unpublished manuscript, Stanford University*.

Ambuehl, S. and S. Li (2014). Belief updating and the demand for information. *Available at SSRN 2461904*.

Ashraf, N., I. Bohnet, and N. Piankov (2006, September). Decomposing trust and trustworthiness. *Experimental Economics 9*(3), 193–208.

Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2015). Organizational barriers to technology adoption: Evidence from soccer-ball producers in pakistan. *NBER Working Paper*.

Attanasio, O. and K. Kaufmann (2009, July). Educational choices, subjective expectations, and credit constraints. NBER Working Papers 15087, National Bureau of Economic Research, Inc.

Barr, A. and A. Zeitlin (2011). Conflict of interest as a barrier to local accountability. *CSAE Working Paper Series 2011-13*.

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2011). Does management matter? evidence from india. Technical report, National Bureau of Economic Research.

Bloom, N. and J. Van Reenen (2007). Measuring and explaining management practices across firms and countries. *Quarterly Journal of Economics*.

Bohnet, I., B. Herrmann, and R. Zeckhauser (2010). Trust and the reference points for trustworthiness in gulf and western countries. *The Quarterly Journal of Economics 125*(2), 811–828.

Breza, E., S. Kaur, and Y. Shamdasani (2015). The morale effects of pay inequality. *Working paper*.

Bruggen, A. and M. Strobel (2007, August). Real effort versus chosen effort in experiments. *Economics Letters 96*(2), 232–236.

Butler, J., P. Giuliano, and L. Guiso (2009). The right amount of trust. *NBER Working Paper*.

Croson, R. T. A. (2000, March). Thinking like a game theorist: factors affecting the frequency of equilibrium play. *Journal of Economic Behavior & Organization 41*(3), 299–314.

Cruces, G., R. Perez-Truglia, and M. Tetaz (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics 98*, 100–112.

Davies, E. and M. Fafchamps (2015). When No Bad Deed Goes Punished: A Relational Contracting Experiment in Ghana. *CSAE Working Paper*.

de Mel, S., D. McKenzie, and C. Woodruff (2010, May). Wage subsidies for microenterprises. *American Economic Review 100*(2), 614–18.

Delavande, A., X. Giné, and D. McKenzie (2011a, 04). Eliciting probabilistic expectations with visual aids in developing countries: how sensitive are answers to variations in elicitation design? *Journal of Applied Econometrics 26*(3), 479–497.

Delavande, A., X. Giné, and D. McKenzie (2011b, March). Measuring subjective expectations in developing countries: A critical review and new evidence. *Journal of Development Economics 94*(2), 151–163.

DellaVigna, S. and D. Pope (2016). Predicting experimental results: who knows what? *NBER Working Paper*.

Enke, B. and F. Zimmermann (2013). Correlation neglect in belief formation.

Fahr, R. and B. Irlenbusch (2000, March). Fairness as a constraint on trust in reciprocity: earned property rights in a reciprocal exchange experiment. *Economics Letters 66*(3), 275–282.

Falco, A., A. Kerr, P. Paci, and B. Rijkers (2014). *Working toward Better Pay: Earnings Dynamics in Ghana and Tanzania*. The World Bank.

Falk, A. and A. Ichino (2006, January). Clean evidence on peer effects. *Journal of Labor Economics 24*(1), 39–58.

Falk, A. and F. Zimmermann (2015). Information processing and commitment.

Fehr, E., G. Kirchsteiger, and A. Riedl (1993, May). Does fairness prevent market clearing? an experimental investigation. *The Quarterly Journal of Economics 108*(2), 437–59.

Gächter, S. and E. Renner (2010, September). The effects of (incentivized) belief elicitation in public goods experiments. *Experimental Economics 13*(3), 364–377.

Gaechter, S., E. J. Johnson, and A. Herrmann (2010, November). Individual-level loss aversion in riskless and risky choices. (2010-20).

Galasso, E., M. Ravallion, and A. Salvia (2004, October). Assisting the transition from workfare to work: A randomized experiment. *Industrial and Labor Relations Review 58*(1), 128–142.

Gill, D. and V. Prowse (2012). A structural analysis of disappointment aversion in a real effort competition. *American Economic Review 102*(1), 469–503.

Groh, M., N. Krishnan, D. McKenzie, and T. Vishwanath (2016). Do wage subsidies provide a stepping-stone to employment for recent college graduates? evidence from a randomized experiment in jordan. *Review of Economics and Statistics* (00).

Heath, R. (2013). Why do firms hire using referrals? evidence from bangladeshi garment factories. *University of Washington, Unpublished Paper*.

Hoffman, M. (2016). How is information valued? evidence from framed field experiments. *Economic Journal*.

Hsieh, C.-T. and B. A. Olken (2014). The Missing ?Missing Middle? *The Journal of Economic Perspectives 28*(3), 89–108.

Jensen, R. (2010, May). The (perceived) returns to education and the demand for schooling. *The Quarterly Journal of Economics 125*(2), 515–548.

Knack, S. and P. Keefer (1997, November). Does social capital have an economic payoff? a cross-country investigation. *The Quarterly Journal of Economics 112*(4), 1251–88.

La Porta, R., F. L. de Silane, A. Shleifer, and R. W. Vishny (1997, May). Trust in large organizations. *American Economic Review 87*(2), 333–38.

Levinsohn, J., N. Rankin, G. Roberts, and V. Schoer (2012). Wage subsidies and youth employment in south africa. evidence from a randomised control trial. *CSAE Annual Conference 2012*.

Manski, C. F. (2004, 09). Measuring expectations. *Econometrica 72*(5), 1329–1376.

Mas, A. and E. Moretti (2009, March). Peers at work. *American Economic Review 99*(1), 112–45.

McKenzie, D., J. Gibson, and S. Stillman (2007, March). A land of milk and honey with streets paved with gold : do emigrants have over-optimistic expectations about incomes abroad ? Policy Research Working Paper Series 4141, The World Bank.

McKenzie, D. and C. Woodruff (2015). Business practices in small firms in developing countries. *NBER Working Paper*.

Mobius, M. M., M. Niederle, P. Niehaus, and T. S. Rosenblat (2011). Managing self-confidence: Theory and experimental evidence. *NBER Working Paper*.

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

OECD (2015). *OECD Employment Outlook 2015*.

Rothschild, M. (1974, October). A two-armed bandit theory of market pricing. *Journal of Economic Theory 9*(2), 185–202.

Shapiro, C. and J. E. Stiglitz (1984, June). Equilibrium unemployment as a worker discipline device. *American Economic Review 74*(3), 433–44.

van Dijk, F., J. Sonnemans, and F. van Winden (2001, February). Incentive systems in a real effort experiment. *European Economic Review 45*(2), 187–214.

A Figures

Figure 1: Trust around the world (WVS 2014 data)

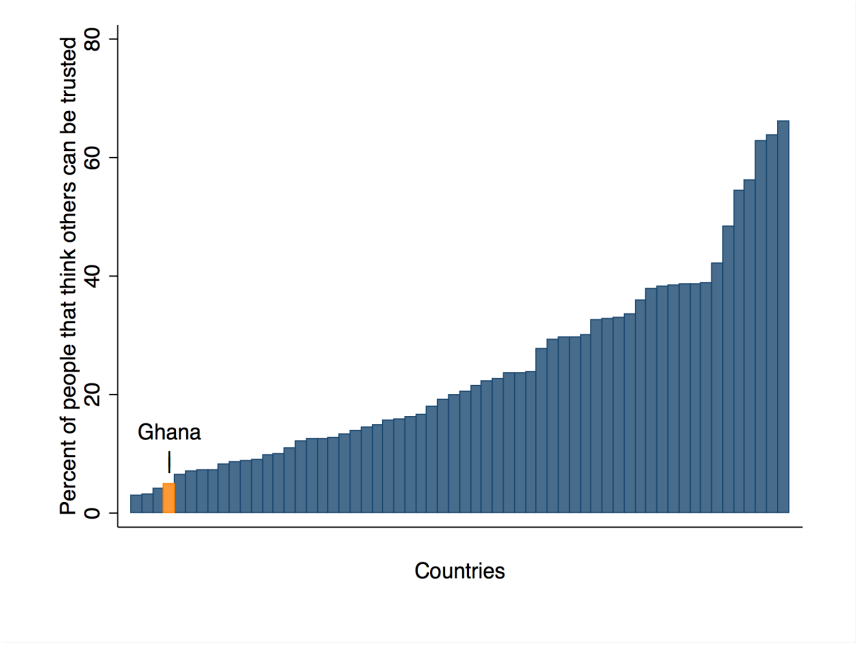

Figure 2: The Game Tree

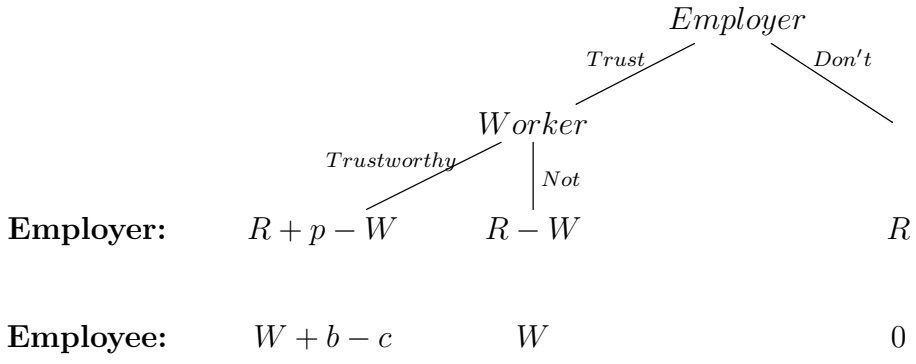

Figure 3: The Game Tree with Monetary Values (Ghana Cedis)

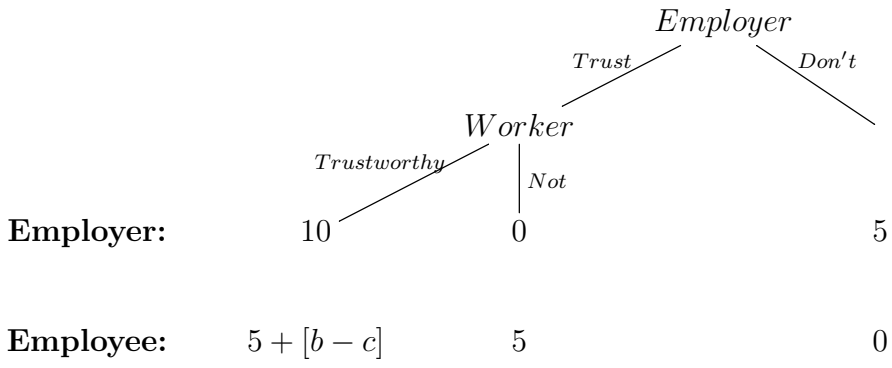

Figure 4: Education

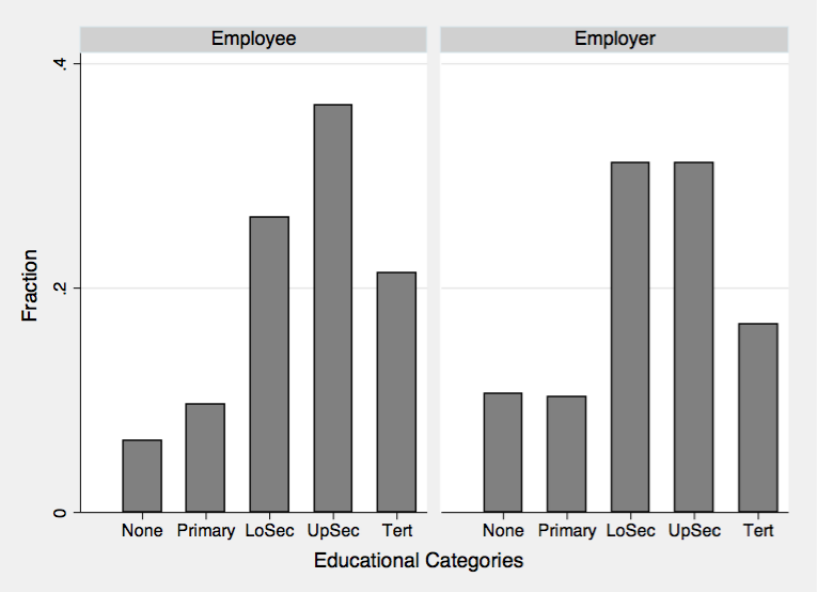

Figure 5: Expected and Revealed Trustworthiness

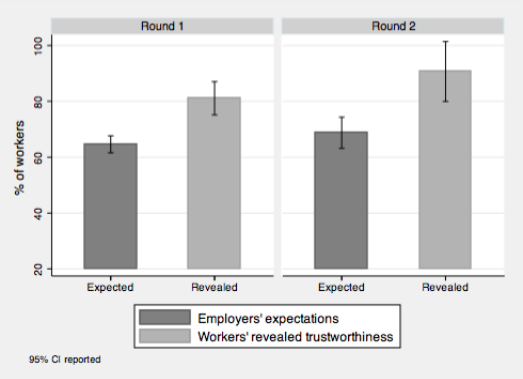

Figure 6: Cumulative distribution of Employers' expectations

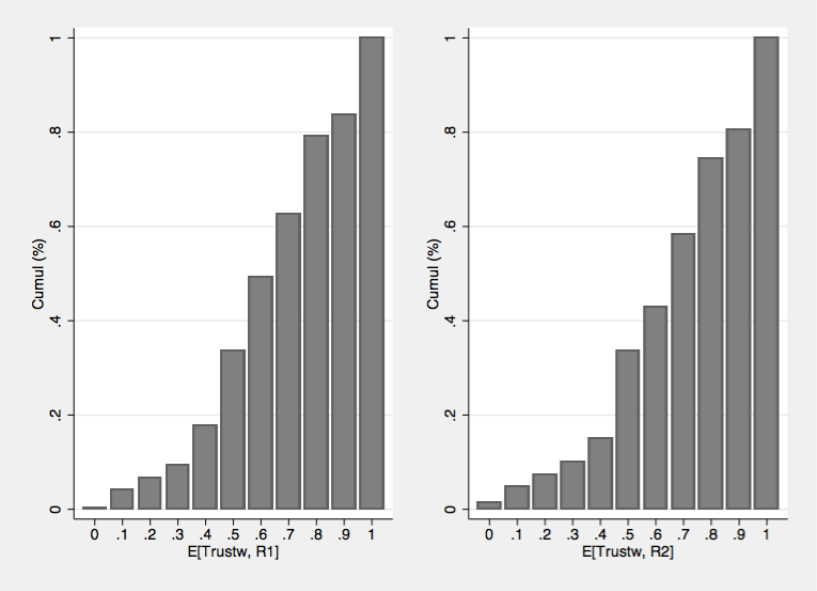

Figure 7: Expected and Revealed Trusting

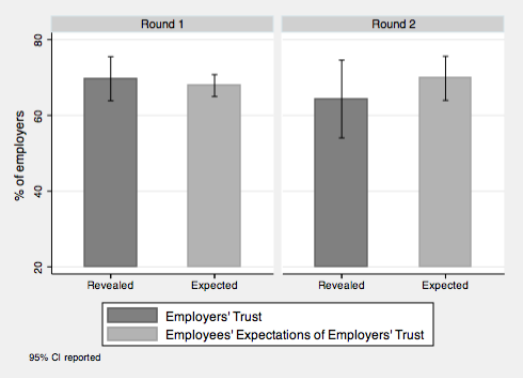

Figure 8: Revealed Trustworthiness by Gender

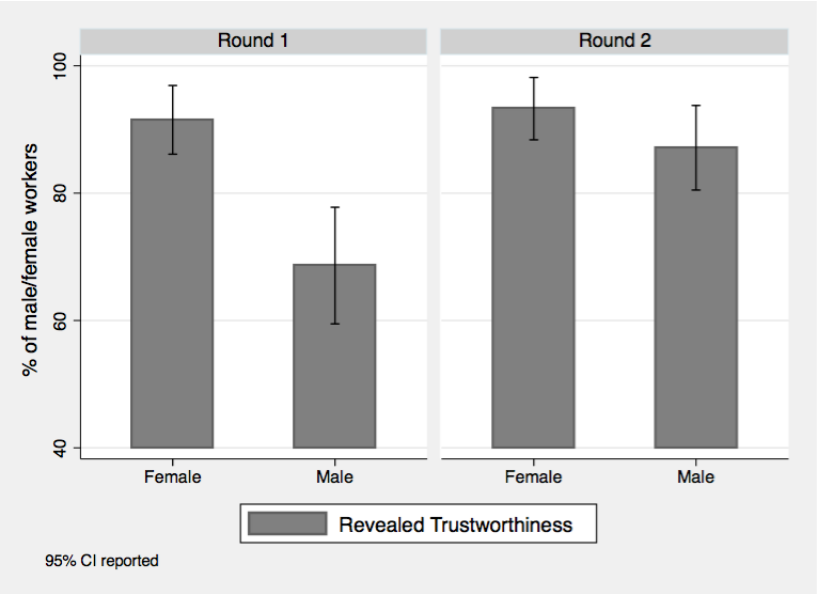

Figure 9: Revealed Trustworthiness by Age

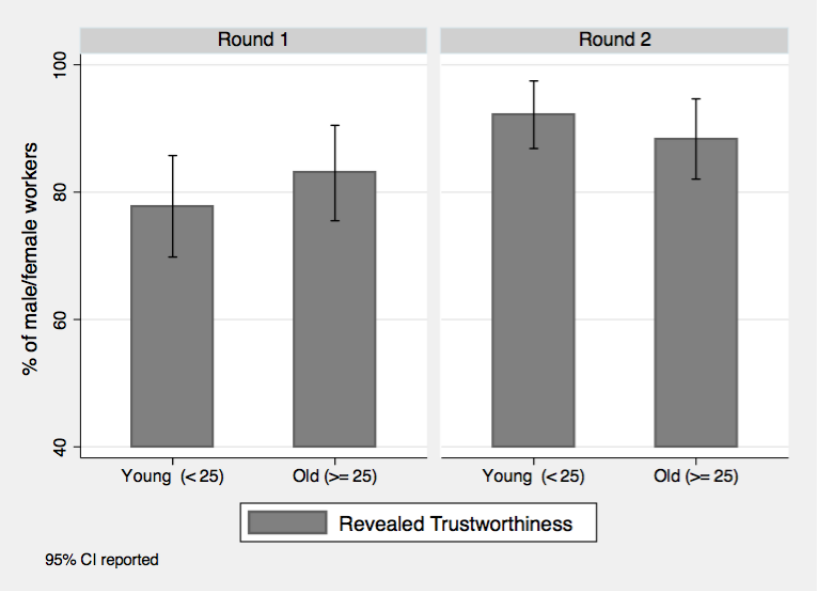

B Tables

Table 1: Employers

Variable	Mean	Std. Dev.
Age	37.582	11.701
Married	0.562	0.497
NoSchool	0.106	0.309
Primary	0.103	0.304
LoSecondary	0.312	0.464
UpSecondary	0.312	0.464
Tertiary	0.168	0.374
Male	0.366	0.483
Unemployed	0.106	0.309
N	292	

Table 2: Employees

Variable	Mean	Std. Dev.
Age	27.954	10.089
Married	0.278	0.449
NoSchool	0.064	0.245
Primary	0.096	0.295
LoSecondary	0.263	0.441
UpSecondary	0.363	0.482
Tertiary	0.214	0.411
Male	0.47	0.5
Unemployed	0.356	0.48
N	281	

Table 3: Covariates Balance across treatments

	T1	T0,T2	Diff	T2	T0,T1	Diff
Variable	Mean		t-value	Mean		t-value
Male	.36	.37	.23	.39	.35	-0.63
Age	36.7	38.0	0.9	37.1	37.8	0.53
Primary	.08	.11	.84	.10	.10	-.014
Lower Secondary	.27	.34	1.21	.33	.30	-.47
Upper Secondary	.33	.30	-.39	.40	.27	-2.37**
Tertiary	.15	.18	.48	.12	.19	1.42
Married	.49	.60	1.76*	.57	.56	-0.13
Unemployed	.06	.13	1.77*	.11	.10	-0.28
Student/Apprentice	.06	.05	-.34	.08	.04	-1.46
Trader	.37	.26	-1.95*	.23	.33	1.79*
OtherSelf	.13	.14	.15	.10	.15	1.19
N	98	194	292	97	195	292

Note: Confidence: *** ↔ *99%*, ** ↔ *95%*, * ↔ *90%*

Table 4: Employers' Expectations towards different worker-types (Dep Var: $E_{R=1}$[Trustw])

	All	MaleEnt	FemalEnt
Young Work.s	.179	1.412	-.368
	(.577)	(.539)***	(.731)
Women Work.s	-.656	-1.621	-.100
	(.340)*	(.633)**	(.368)
Age	.011	.026	.010
	(.017)	(.024)	(.021)
Male	.120		
	(.314)		
Education (ys)	.051	.046	.059
	(.028)*	(.064)	(.029)**
Married	.352	.036	.209
	(.397)	(.744)	(.364)
Unemployed	.468	.728	.323
	(.332)	(1.050)	(.513)
Session Char.s	Yes	Yes	Yes
Const.	5.308	4.542	5.724
	(.665)***	(1.037)***	(.821)***
Obs.	292	107	185

Note: Confidence: *** ↔ *99%*, ** ↔ *95%*, * ↔ *90%.; Robust standard errors in parentheses (cluster. by Session); Session Char.s = Time and Timing of Expect.s Q.*

Table 5: Trust and Expected Trustworthiness (Dep Var: Pr[Hire])

	OLS1	OLS2	Probit1	Probit2	Probit-Mfx
E[Trustw]	.030 (.008)***	.027 (.008)***	.083 (.023)***	.077 (.024)***	.027
Loss Aversion		-.048 (.012)***		-.142 (.037)***	-.049
Age		.001 (.002)		.003 (.006)	.001
Male		.027 (.035)		.087 (.106)	.030
Education (ys)		.004 (.006)		.011 (.017)	.004
Married		-.029 (.040)		-.085 (.119)	-.030
Unemployed		.042 (.057)		.124 (.177)	.042
Exp Q b/f choice		-.111 (.051)**		-.330 (.151)**	-.116
Round 2	-.016 (.052)	-.015 (.051)	-.043 (.150)	-.038 (.152)	-.013
Const.	.494 (.066)***	.594 (.125)***	-.042 (.177)	.279 (.358)	
Obs.	584	584	584	584	584

Note: Confidence: *** ↔ 99%, ** ↔ 95%, * ↔ 90%.; Robust standard errors in parentheses (cluster. by Session+Round). E[Trustw] is measured on a scale from 1 to 10.

Table 6: Expectations Updating (Dep Var: E[Trustw, R2])

	(1)	(2)	(3)
E[Trustw, R1]	.680 (.051)***	.679 (.052)***	.639 (.053)***
Neg Sig * Homog T1	-.760 (.331)**	-.682 (.380)*	-.782 (.413)*
Pos Sig	.030 (.339)	.162 (.419)	.032 (.471)
Pos Sig (F) * Heter T1		.196 (.539)	
Pos Sig (Y) * Heter T1		-.397 (.591)	
Homog T1	.542 (.387)	.468 (.459)	.609 (.485)
Heter T1	-.257 (.374)	-.189 (.465)	
Const.	2.303 (.590)***	2.296 (.592)***	2.221 (.673)***
Personal Char.s	Yes	Yes	Yes
Session Char.s	Yes	Yes	Yes
Obs.	292	292	195

Note: Confidence: *** ↔ 99%, ** ↔ 95%, * ↔ 90%.; Robust standard errors in parentheses (cluster. by Session); Col 1 includes all T1 sessions, all positive signals are pooled together; Col 2 includes all T1 sessions, different positive signals are separated; Col 3 includes only T1 sessions with homogenous signals (i.e. where no info on worker age and gender were attached to the signals).

Table 7: The impact of raising expectations on trust (Dep Var: Pr[Hire])

	R2	R2C	R2IV
E[Trustw, R2]	.036	.034	.171
	(.011)***	(.012)***	(.065)***
Age		-.0005	-.002
		(.002)	(.002)
Male		.053	.042
		(.045)	(.056)
Education (ys)		-.001	-.003
		(.007)	(.008)
Married		-.030	-.101
		(.056)	(.062)
Unemployed		.090	.013
		(.077)	(.122)
Exp Q b/f choice		-.123	-.062
		(.076)	(.082)
Const.	.434	.529	-.292
	(.086)***	(.156)***	(.423)
Obs.	292	292	292

Note: Confidence: *** ↔ 99%, ** ↔ 95%, * ↔ 90%.; Robust standard errors in parentheses (cluster. by Session). E[Trustw, R2] is measured on a scale from 1 to 10.

Table 8: The relationship between real-life hiring and trust in the experiment

	(1)	(2)	(3)	(4)
	E[Trustw, R1]	Pr[Hire, R1]	E[Trustw, R1]	Pr[Hire, R1]
Num. Employees last hired	0.168***	0.0298**	0.128*	0.0280**
	(0.0596)	(0.0131)	(0.0665)	(0.0127)
Male			-0.0127	0.00149
			(0.337)	(0.0482)
Education (ys)			0.0495	0.00993
			(0.0302)	(0.0101)
Age			0.0195	-0.000174
			(0.0166)	(0.00329)
Married			0.146	0.0177
			(0.386)	(0.0629)
Unemployed			0.429	0.0108
			(0.375)	(0.0887)
Observations	263	263	263	263

* $p < .1$, ** $p < .05$, *** $p < .01$

Note: Confidence: *** ↔ 99%, ** ↔ 95%, * ↔ 90%.; Robust standard errors in parentheses (cluster. by Session).

C Experiment Protocol

C.1 Control

Round 1

1. j informed of average worker char.s (gender, age, etc.) to anchor expectations.
2. j's expectations of workers trustworthiness elicited.
3. j chooses hire/not hire a random and anonymous i.
4. i chooses effort (not revealed to j until the end).

Round 2 (announced as a 'surprise' at the end of R1)

1. j's expectations (re)elicited

   [they cannot be assumed equal to R1 (e.g. due to expected learning or fatigue)]
2. j chooses hire/not hire a (new) random i.
3. i chooses effort.

Post-Exper:

i.Risk Game

ii.Dictator Game

iii.End Questionnaire,

iv. Payoffs Revealed and Prizes Distributed

C.2 Treatment 1

Round 1

1. j informed of average worker char.s (gender, age, etc.) to anchor expectations.
2. j's expectations of workers trustworthiness elicited.
3. j chooses hire/not hire a random and anonymous i.
4. i chooses effort (not revealed to j).

Round 2 (announced as a 'surprise' at the end of R1)

1. ***T1: Inform j of the trustworthiness of a random i from prev. sessions.***
2. j's (new) expectations elicited.

3. j chooses hire/not hire a (new) random i.
4. i chooses effort.

Post-Exper:

i.Risk Game

ii.Dictator Game

iii.End Questionnaire,

iv. Payoffs Revealed and Prizes Distributed

C.3 Treatment 2

Round 1

1. **T2: j informed that 80% of workers invited to the session belong to a specific category (i.e. young, female)**.
2. j's expectations of workers trustworthiness elicited.
3. j chooses hire/not hire a random and anonymous i.
4. i chooses effort (not revealed to j)

Round 2 (announced as a 'surprise' at the end of R1)

1. j's expectations (re)elicited

   [they cannot be assumed equal to R1 (e.g. due to expected learning or fatigue)]
2. j chooses hire/not hire a (new) random i.
3. i chooses effort.

Post-Exper:

i.Risk Game

ii.Dictator Game

iii.End Questionnaire,

iv. Payoffs Revealed and Prizes Distributed

Figure 10: Risk Game Lotteries

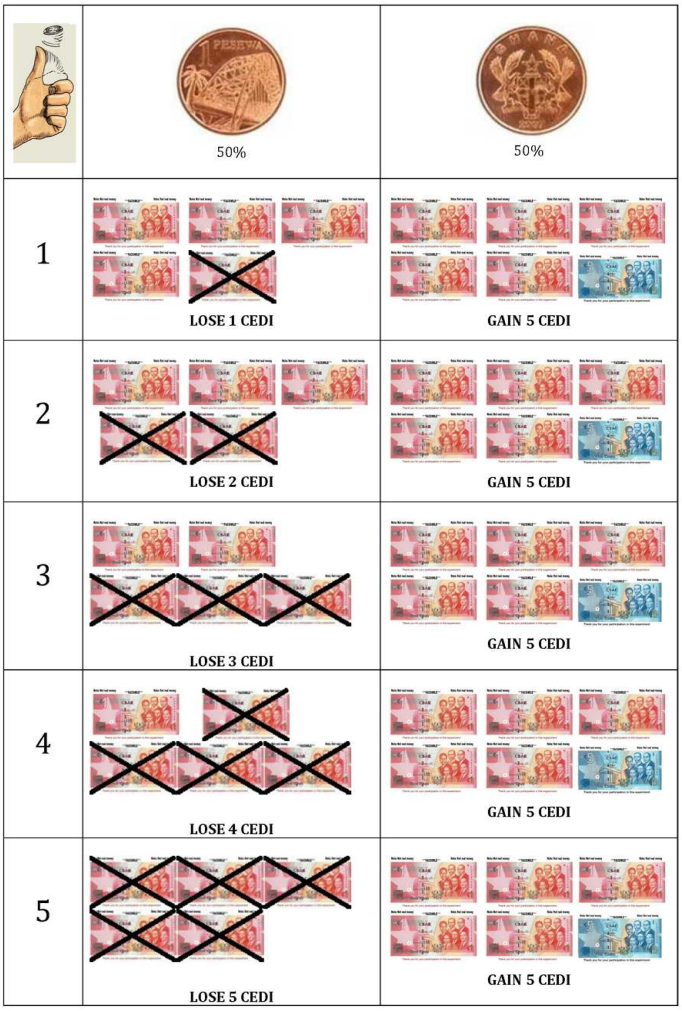

D Additional Figures and Tables

Figure 11: The hiring decision of a micro-entrepreneur

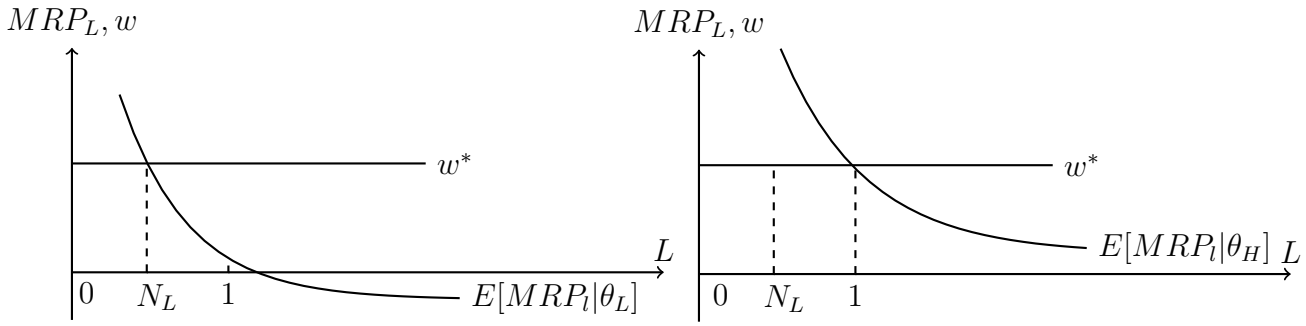

(a) Low Expected Trustworthiness

(b) High Expected Trustworthiness