Sceptical employers:
Experimental evidence on biased beliefs constraining firm growth*

Stefano A. Caria† Paolo Falco‡

June 16, 2020

Abstract

Many firms in developing countries hire no labour. We conduct a novel experiment in Ghana to investigate whether low trust in workers discourages small businesses from hiring. We give real entrepreneurs the option to hire a worker for a simple but tedious task. Shirking attracts no penalty and completion of the task is an indicator of workers' trustworthiness. We elicit employers' expectations and we randomly provide information on workers' previous behaviour to study belief updating. We find that employers significantly underestimate workers' trustworthiness. This reduces hiring and profits. We also find evidence of asymmetric learning: negative signals lower employers' expectations while positive signals do not affect them. Experience with hiring in real life correlates positively with employers' expectations and with their willingness to trust. Our evidence corroborates the hypothesis that a low-trust equilibrium with limited experimentation and biased beliefs can be self-sustaining.

Keywords: hiring, trust, trustworthiness, expectations, learning, effort, productivity, microenterprise, discrimination, experiment, African labour markets.
JEL codes: C91, D22, D84, J23, L25, M51, O12.

*We are grateful to Abigail Barr, Andrea Bassanini, Iris Bohnet, Helene Bie Lilleør, Federico Cingano, Marcel Fafchamps, Luigi Guiso, Marco Piovesan, Imran Rasul, and attendants of the CESS seminar (Oxford), the OECD Applied Economics WiP seminar, the Royal Economic Society conference, the University of Copenhagen seminars, the Symposium on Economic Experiments in Developing Countries (SEEDEC), the Rockwool Foundation seminar (Copenhagen), and the DIW Workshop on Development Economics ("Firms and Productivity") 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.

†Department of Economics, University of Warwick; stefano.caria@warwick.ac.uk
‡Department of Economics, University of Copenhagen; paolo.falco@econ.ku.dk

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 hypothesis that emerges from the literature is that low trust in workers may discourage entrepreneurs from hiring. This may play an important role in much of the developing world, where contracts are harder to enforce (Sánchez de la Sierra, 2016; Ashraf et al., 2019), and schemes to monitor and reward performance can be costly and difficult to introduce (Atkin et al., 2017; Breza et al., 2017). Indeed, firms in many parts of the Global South, including Africa, often make limited use of financial incentives in hiring and retention (Lemos and Scur, 2014; Bloom et al., 2014; Caria, 2019). If managers do not trust workers to exert effort when tight monitoring and incentives are not an option, they may refrain from hiring or may distort their hiring policies to reduce moral hazard. For example, they may rely on network-based hiring, a second-best strategy that can generate significant aggregate costs (Montgomery, 1991; de Mel et al., 2010; Heath, 2018; Chandrasekhar et al., 2020; Hoffman et al., 2018).

Do managers think workers can be trusted? There is extremely little evidence on the beliefs of firm managers in developing countries and we do not know whether existing levels of trust are fully aligned with prevailing levels of trustworthiness. Yet, a growing literature on biased expectations highlights a number of mechanisms that may make employers overly pessimistic about the trustworthiness of workers. First, trust presents a classic problem of experimentation: an overly pessimistic manager will not hire and hence will not receive new information to update incorrect priors (Rothschild, 1974). Recent work by Butler et al. (2016) corroborates this hypothesis by showing that individuals with overly pessimistic beliefs give up profitable opportunities to avoid being cheated. In low-trust societies, these beliefs can become entrenched as parents will pass them on to their children in an effort to protect them from costly mistakes (Guiso et al., 2008). Second, managers that receive new information may fail to adjust their beliefs in a rational way (Enke and Zimmermann, 2017; Ambuehl and Li, 2018; Falk and Zimmermann, 2018; 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). Crucially, negative beliefs of this kind may have important negative economic implications (Bartling et al., 2018). We hypothesise that in a labour market context, overly pessimistic expectations of workers' trustworthiness may depress hiring.

In this paper, we study employers' expectations about workers trustworthiness and a key mechanism that can lead to biased beliefs: asymmetric updating in response to new information. To achieve these objectives, we design an investment game with a real-effort task, played between real entrepreneurs and 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 chooses whether to exert effort in order to complete a simple task that requires no particular skill 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 performance and by removing the employer's ability to monitor and punish low effort, we are able to elicit trust and trustworthiness. Crucially, we elicit employers' expectations of workes' trustworthiness and we compare them with actual workers' performance to identify potential misperceptions. We further devise two randomised treatments to study the extent to which expectations are biased against specific worker categories and to investigate how employers update their priors as they receive new information.

We find that employers significantly underestimate workers' trustworthiness and that low expectations have negative repercussions for hiring and profits. The result is robust to controlling for a separate experimental measure of risk-aversion. We also find evidence of a gender gap. Male employers expect female workers to be less trustworthy than men, while behaviour in the experiment reveals the opposite.

Furthermore, we show that as employers acquire new information they update their expectations asymmetrically. Negative information on workers' performance has a negative effect on employers' beliefs, while positive signals do not change beliefs significantly. When we explore the correlation between the 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 is consistent with the hypothesis that lack of experimentation can lead to a self-sustaining equilibrium with overly pessimistic beliefs. The existence of such low-trust equilibria finds support in the literature (Guiso et al., 2008).[1]

This paper contributes to a burgenoning literature on expectations and their formation (Manski, 2004; Attanasio and Kaufmann, 2017; Attanasio et al., 2019; Delavande and Zafar, 2019; Bursztyn et al., 2020), and to a growing body of evidence on biased beliefs across a range of domains. In related work, McKenzie et al. (2013) 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. In the context of labour markets, existing studies have typically focused on biases in workers' expectations (Alfonsi et al., 2019; Banerjee and Sequeira, 2020), but limited evidence exists on the expectations of firms.[2] We contribute to bridge this gap by showing that entrepreneurs can be biased in their assessment of workers' trustworthiness and this is associated with suboptimal investment decisions. Our approach also relates to a growing literature that studies subjects' ability to predict experimental results (DellaVigna and Pope, 2018).

We also contribute to explore the mechanisms that can lead to biased expectations by providing evidence of asymmetric learning. Our findings draw an interesting parallel with the literature demonstrating that people 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 growing interest in understanding how the payoffs of actions affect information acquisition and beliefs (Ambuehl, 2017).

Finally, we contribute to the literature that studies the quality of management in

[1] Guiso et al. (2008) show that such equilibria can persist across generations, as parents transmit conservative priors to their children in order to protect them from costly mistakes. Their model also predicts that as people age and learn, their beliefs should naturally improve. This is confirmed by our data, where we find a positive correlation between age and expected trustworthiness (results available upon request). The correlation, however, is not very strong, in line with their conclusion that in low-trust societies only limited experimentation occurs.

[2] In a very recent project, we investigate biased beliefs among firm managers in Ethiopia (Abebe et al., 2020).

developing countries (Bloom and Van Reenen, 2007, 2010; Bloom et al., 2012; McKenzie and Woodruff, 2017). We advance this body of work by using experimental techniques to isolate biases in employers' beliefs and in the way they assimilate new information, and to show how this affects decision-making.[3] Our results suggest that information interventions targeting firm managers should take into account biases in belief-updating, as positive and negative information may be absorbed asymmetrically.

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 its relationship with business growth. In section 3 we outline our experimental design. In section 4 we describe our sampling strategy and provide some descriptive statistics. In section 5 we present our results. Section 6 concludes.

2 A low-trust country

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

First, harmonised cross-country data from the World Values Survey (WVS) shows that trust among people in Ghana is particularly low. Out of 60 countries analysed (developed and developing) across all the continents, Ghana is the one with the fourth lowest level of trust (Figure 1).[4]

<< **Figure 1 here** >>

Second, cross-country evidence shows a positive correlation between entrepreneurs' willingness to expand their business and aggregate trust levels. The Global Entrepreneur Monitor (GEM) dataset includes cross-country data on the share of entrepreneurs who have just started a business and expect to hire at least five workers in the coming five years.

[3] This relates to the work by Davies and Fafchamps (2017), who run a gift-exchange lab experiment with Ghanian students and show that subjects fail to screen out partners who have been untrustworthy in the past.

[4] Trust in the World Values Survey is measured by means of a standard question that asks whether the respondent *thinks that others can be trusted*.

As generalised levels of trust in a country increase, that share grows.[5] Coupled with the observation that the large majority of Ghanaian businesses 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 Experimental design

The experiment is centred around an investment game (Berg et al., 1995) with a real effort task. In this section, we describe the game setup, our methodology to elicit players' expectations, and the experimental manipulations we use to study specific mechanisms.

3.1 The players

Each respondent was pre-assigned to one of two experimental roles: *employer* or *employee* (with real entrepreneurs assigned to the role of employers) and randomly allocated to an experimental session. Each session had 10 employers and 10 employees in total.[6] Within each session, the two groups were kept separate and participated in parallel activities without meeting (two separate locations, sufficiently distant from each other, were equipped to host the employers' and the employees' room respectively). 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 (Falk and Ichino, 2006; Mas and Moretti, 2009).[7]

[5]The relationship is concave and the positive correlation weakens once countries reach high levels of trust (beyond the levels observed in Ghana).

[6]There were few exceptions, when some respondents dropped out and the session was not full. The plan was to have 30 sessions with 10 employers and 10 employees, but the avarage number of employers and employees per session was 9.73 and 9.37, i.e. N = 292 and N = 281, reSymposium on Economic Experiments in Developing Countriesspectively). This did not affect the decisions of players. Employers knew that employees played independently and were physically isolated from one another (as described below). Hence, having one or two players less in the room did not affect their actions. Similarly, employers made their decisions independently and were physically separated. For simplicity, we told employers that the employees in the other room were always ten (and vice-versa). Since the two groups of players were in different locations, respondents had no way to detect the minor deviations that occurred.

[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 on effort, as predicted by a model in which unemployment constitutes a worker discipline device (Shapiro and Stiglitz, 1984).

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.

3.2 The game

Upon reaching the venue, the two groups were introduced to the structure of the experiment. For the sake of conciseness, the following explanation provides a stylised representation of the rules. The instructions respondents received were less schematic and illustrated by means of examples. An outline of the experimental protocol is provided in Appendix B and the full script is available upon request.

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

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

If the worker exerts 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. We operationalise it through a real-effort task that requires no skills and can be completed by any worker who is willing to make an effort (see the next section for a detailed discussion). Employers have full knowledge of this.

When the worker chooses to exert effort and production is successful (i.e. $e_i = 1 \rightarrow y = 1$), the employer receives profits equal to $p > R$ 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 the entire endowment when the employee shirks). This type of contractual arrangement, with no monetary penalty for shirking, allows us to isolate trust as a driver of employers' choices.

The game tree in Figure 2 outlines the sequence of choices and corresponding payoffs

faced by employers and employees in each round. The chosen monetary payoffs (shown in parentheses) are large with respect to earnings in the reference population, hence generating strong incentives in the game.[8]

<< **Figure 2 here** >>

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 was asked to undertake the real-effort task and could decide whether to exert effort (or shirk).[9] 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 exerted effort or not) were not announced to employers until the very end of the game to avoid influencing their decision in the second round. Before or after each decision (the order was randomised), respondents' expectations of the other party's behaviour were elicited, as discussed below. They will be central to the analysis in this paper.

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 to be trustworthy). The task did not rely on any specific know-how or skills, ruling out the influence of human capital. This was explicitly and repeatedly explained to both employers and employees, who were aware that any worker could attain success if s/he 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 expected and revealed trustworthiness in a real-effort setting (akin to real employment); and,

[8]The median of net daily earnings for workers in the reference population was approximately 6 Ghana Cedis according to data from the Oxford University Ghana Househol Urban Panel Survey available at the time of the experiment (August 2013). The endowment provided to employers in each round of the investment game was 5 Cedis. The endowment for the game to elicit risk-preferences (described below) was also equal to 5 Cedis (i.e., the total endowment for the three games amounted to 15 Cedis). The entire experiment lasted approximately 2 hours.

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

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.

In light of these considerations, 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 the *minimum attainable amount under steady effort and no shirking*, based on prior testing.[10]

The task was trivial but 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 they should be able to complete the task in the given time. In order to rule out the possibility that workers may work "out of boredom", the experimental 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.[11] 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.

Given the features of the task, we are confident that workers' success is a valid proxy of their willingness to make an effort.[12] Exerting effort, in turn, is a direct measure of trustworthiness, since it captures workers' willingness to carry out a costly task (for which they have been paid) in a situation where there is no punishment for shirking.

[10]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.

[11]We took care in choosing a show that was widely known and popular across age groups and genders.

[12]It should be noted that if lack of skills did limit workers' success despite their efforts, our estimated level of trustworthiness would be a lower bound estimate of actual trustworthiness. In other words, if some workers tried their best (i.e. were trustworthy), but due to lack of skills were unable to complete (and were mistakenly classified as untrustworthy), overall estimated trustworthiness would be biased downwards. This possibility strengthens the conclusions from the next sections. Indeed, despite this potential downward bias, we will show levels of revealed trustworthiness that are above employers' expectations.

3.4 Eliciting expectations

In each round of the experiment, before or after the hiring decision was made (the order was randomised across sessions to control for potential "order effects"), we asked employers the following question: *"Out of the 10 workers in the other location, one of whom will be assigned to you by chance, how many do you think will complete the task successfully if they are all hired?"* The reverse question was asked to employees: *"Out of the 10 employers in the other location, 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). Having experimental sessions with 10 respondents per room greatly simplified the elicitation strategy, as it allowed us to veer away from questions about probabilities and towards simpler questions about frequencies.[13] To anchor expectations, employers were told that employees were sampled from the working-age population of Accra and they were informed of their average characteristics. Similarly, information about employers was provided to the employees.

The elicitation questions could not be independently incentivised. While we acknowledge the potential limitations of this design feature, a number of important factors led us to the conclusion that introducing such incentives would have been detrimental.

First, 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 on our respondents (whose average levels of numeracy are low). Introducing an additional scoring rule for the belief-elicitation questions appeared to go directly against this principle and we had strong reasons to believe it might increase complexity to the point of lowering overall understanding.

Even if respondents understood the scoring rule, our second concern was that they could

[13]Under the former approach, the question would be "On a scale from 0 to 10, how likely is it that a worker/employer...?" or "On a scale from 0 to 100 percent, what is the probability that a worker/employer...?" Under our approach, the question becomes "Out of 10 workers/employers, how many...?" The latter does not rely on any understanding of probability and should be simpler to interpret, eliciting more accurate answers.

use these questions to hedge against their choices in the trust game, causing distortions in the elicited belief distribution and creating a spurious correlation between beliefs and trusting.[14] Such distortions have been documented in the literature. A methodological study by Gächter and Renner (2010) shows that while monetary incentives may increase the accuracy of beliefs in the context of a public-good game, they may also alter the relationship between beliefs and behaviour (with mixed evidence on the direction of the resulting bias).[15] For this reason, they conclude that if a researcher is interested (as we are) in the relatonship between beliefs and behaviour, one should consider not incentivising belief elicitation. In addition, their results show 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 is that with incentivised questions our results could have been more precise, but they would not have changed qualitatively. The risk, however, would have been to generate distortions in the relationship between beliefs and behaviour.

Third, a 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 (Delavande et al., 2011b). This is in line with recent work by Enke et al. (2020), who set up a large experiment in Kenya to study the effects of incentives on cognitive biases and find little evidence that incentivising respondents improves their performance.

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 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.

[14] A way to resolve this concern and maintain the monetary incentive would have been to devise a randomised 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.

[15] 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.

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 employers (who were assigned to specific "treated" sessions).

Treatment 1 (T1): Information provision. At the end of round 1, employers in T1 sessions 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.[16] 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 rounds were revealed at the end of experiment).

In one half of T1 sessions, the random signals did not reveal any worker characteristic. In the other half, they specifically referred to a worker who was female or young (below 25).[17] This additional variation allows us to test whether employers' expectations are more or less sensitive to new information when the signal is about women and youth, two vulnerable groups with worse labour market outcomes that attract the attention of policy-makers.

Treatment 2 (T2): Changing the composition of the worker 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 the same two vulnerable groups used in T1. In half of T2 sessions, 80% of invited workers were women and, in the other half, 80% were young (below the age of

[16] Respondents in the role of employees, on the other hand, were informed between round 1 and round 2 that employers had received 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 for the analysis in this paper. In order to exclude the potential impacts of this design feature, we drop employees in these sessions from the estimation of workers' trustworthiness in round 2 below.

[17] Due to concerns of statistical power, the signals that carried this additional information were all positive (i.e., they all indicated that the worker had completed the task). This choice is discussed upon presenting the results.

25). 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.

3.6 Eliciting risk-aversion

At the end of the experiment, but before performance and payoffs were revealed, we also conducted a separate game to elicit employers' risk aversion. The 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 s/he 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.[18] Figure 8 in Appendix A shows the visual aid used to explain the lotteries. For comparability with our main investment game, the value of the safe option was equal to the initial endowment employers could invest to hire a worker. Moreover, in order to be as consistent as possible with the investment decision, 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 investment game. Our measure of risk preferences, therefore, is more precisely a measure of "risk aversion in the loss domain". 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.[19]

4 Sample and descriptive statistics

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

[18]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.

[19]Given that the large majority of the respondents played according to a monotonic 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 simple method is a consistent way of categorising workers.

of a comparison between laboratory outcomes and real-world decisions, the best way to achieve this goal was to assign real entrepreneurs and workers to the experimental roles of employers and employees respectively. It seems reasonable to assume that respondents' behaviour in the lab is based on the priors they develop in the labour market; and in order to give our experiment the best chance 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). We pursued this objective by drawing a random sample of business owners and employees from the working-age population of Accra, the capital city of Ghana, where the experiment took place. We considered eligible for the role of employer any adult between the age of 18 and 65 who currently owned a business or had owned one over the past three years. All adults in the same age range from the general population were eligible for the role of employee.

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 entrepreneurial activities is high), the majority of them are women. Employees, on the other hand, are balanced across genders, they are younger on average, and significantly more likely to be currently unemployed. They also have slightly higher educational attainments, most likely a reflection of belonging to a younger generation. Table 8 in Appendix A shows that covariates are largely balanced between treatment groups.

5 Results

5.1 Are employers' expectations aligned with the revealed trustworthiness of employees?

We begin by addressing the central question in our investigation: *Do employers have correct expectations of workers' trustworthiness?* Our experimental design allows us to answer this question directly, by comparing employers' elicited expectations of employees' trustworthiness (i.e., the proportion of employees expected to be trustworthy, $E_j[P(e_i = 1)]$) with the revealed trustworthiness of employees (i.e., the proportion of employees who are actually trustworthy, captured by their rate of success in the task, $P(e_i = 1)$) in each round ($R = 1, 2$):

$$\underbrace{E_j[P(e_{i,R=1,2} = 1)]}_{\text{Expected trustworthiness in R1,2}} \quad vs \quad \underbrace{P(e_{i,R=1,2} = 1).}_{\text{Actual trustworthiness in R1,2}}$$

***Result 1:* Employers underestimate employees' trustworthiness.**

The results show that the proportion of workers who reveal to be trustworthy is significantly higher than it is expected by employers on average (Figure 3). This is true in round 1 as well as round 2 of the game ($p < .001$ in both upon testing equality between expected and revealed trustworthiness). Furthermore, plotting the cumulative distribution of employers' expectations reveals that the majority of employers underestimate the proportion of employees who will be trustworthy (Figure 4).

<< **Figure 3 here** >>

<< **Figure 4 here** >>

***Result 2:* Employees correctly predict employers' propensity to trust.**

By contrast, when we compare employees' elicited expectations of employers' propensity to trust with the actual rate of trusting, we find that workers' beliefs are broadly correct (Figure 5).[20] Insofar as behaviour in the lab is driven by beliefs developed in the labour market, these

[20] In both rounds, we cannot reject the hypothesis that expected and revealed trust are equal ($p = .58$ and $p = .36$, respectively).

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' trustworthiness.

<< **Figure 5 here** >>

Result 3: ***Male employers have lower expectations when the employees are predominantly women.***

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:

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

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 (<25)* workers];
X_j is a vector of control variables.

When we run the estimation over the entire sample (Table 3, "All"), we find that *employers have significantly lower expectations when the majority of workers in the session are women.*[21] 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

[21]The results are restricted to round 1 to conserve space. Moreover, by focusing on round 1, we are able to compare T2 with Control *and* T1 sessions jointly (since round 1 was identical in the latter two), maximising the available sample size.[22]

is predominantly female. Male employers also have a positive expectations bias towards young workers, which is absent among female employers. The expectations of female employers, on the other hand, are not affected by the composition of the worker pool.

<< **Table 3 here** >>

The result by gender is particularly striking when confronted with revealed trustworthiness, since female workers are, in fact, significantly more trustworthy than men ($p < .001$ upon testing equality), as shown in Figure 6. Younger workers, on the other hand, are no more (or less) trustworthy than older ones ($p = .34$), as shown in Figure 7.

<< **Figure 6 here** >>

<< **Figure 7 here** >>

5.2 Do expectations matter for trusting?

Having presented our results on employers' misperceptions of employees' trustworthiness, we now test whether the expectations we elicited have 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:[23]

$$P(H_{j,R} = 1) = \alpha + \beta E_{j,R}[P(e_i = 1)] + \gamma X_j + u_{j,R}, \tag{3}$$

where:

$P(H_{j,R} = 1) \equiv$ Probability of Trusting (Hiring) in round R;
$E_{j,R}[P(e_i = 1)] \equiv$ Employer j's expectations in round R;
X_j is a vector of control variables.

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

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

Table 4 shows that expected trustworthiness strongly and significantly predicts trusting in all the proposed specifications. A change in employer's expectations by 1 (i.e., one additional worker out of 10 expected to be trustworthy, or a 10 percentage point increase) raises the probability of trusting by 3 percentage points on average. The estimated relationship is in line with the conclusions of Ashraf et al. (2006).

<< Table 4 here >>

We also find that our elicited measure of risk-preferences significantly predicts trusting, over and above the effect of expectations. This is an important robustness check, which reveals evidence of the two separate channels one would expect to be at play: an effect related to employers' perceptions of the moments of the trustworthiness distribution and a second one, which depends on the curvature of their utility function.[24]

By showing that hiring decisions are strongly correlated with employers' elicited expectations, the results in this section point to the conclusion that *underestimating workers' trustworthiness leads to sub-optimal hiring and hurts employers' profits.*[25]

5.3 Does access to information change employers' expectations?

Next, we set out to explore the learning mechanism that underlies expectation formation. In doing so, we will 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 (*positive signal*) or not (*negative signal*). Two control subjects in each session received *no signal*. In

[24] As discussed above, our measure of risk-aversion is well-suited to draw this distinction, since it is framed in the loss domain like the hiring choice employers have to make (which entails risking an endowment they own).

[25] We also tested whether including a measure of altruistic preferences changes the estimated relationship between expectations and trust, and it does not (the results are not included for conciseness, but available upon request). The measure of altruistic preferences was obtained from a *dictator game* played after we conducted the main experiment and the elicitation of risk-preferences, but before the payoffs from the main experiment were revealed.

half of T1 sessions, the signal carried no worker characteristics. In the other half (where we could only provide positive signals due to limited statistical power), it revealed whether the randomly drawn worker was male or female, and whether s/he was old or young (above or below 25).[26]

The treatment was designed to provide employers with the same information they could obtain by hiring a worker (i.e. experimentation); and it is important ro recall that at the time of receiving the signal the results of the first round had not been revealed.[27] This meant that all treated employers had the same information prior to treatment and they received a 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 are hardest to change).

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

$$E_{j,2}[P(e_{i,2}=1)] = \alpha + \eta E_{j,1}[P(e_{i,1}=1)] + \sum_{\kappa} \beta_{\kappa} S_{T1,j,\kappa} + \gamma T1_j + \delta X_j + u_j, \qquad (4)$$

where:

$E_{j,R}[P(e_{i,2}=1)] \equiv$ Employer j' expectations in round $R = (1,2)$;

$S_{T1,j,\kappa} = 1[\text{Signal received by } j \text{ in } T1 = \kappa]$, where $\kappa =$ Positive or Negative signal;[28]

$T1_j = 1[j$ was in Treatment 1];[29]

X_j is a vector of control variables.

[26]The distribution of signals we provided was calibrated on the results from a pilot we ran prior to the experiment, which revealed a level of trustworthiness among employees of around 60%. In order to have a balanced distribution, we deviated slightly and provided the two signals in equal proportions (50%). Two employers in each session received no signal to test for the sheer effect of being in a treated session. When we added the gender and age of the worker, it was impossible to provide a distribution of signals that reflected the real distribution in previous sessions 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. 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).

[27]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.

[28]Positive signals were further split into "Positive + Female worker" and "Positive + Young worker" in one half of T1 sessions, as explained above.

[29]The coefficient on this dummy (γ) is identified thanks to the fact that two workers in each T1 session did 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 strong downward effect on expectations, while positive signals have no effect.***

The results are reported in Table 5. They show that positive signals have no impact on employers' expectations, while negative signals significantly lower them. A negative signal decreases the expectations of an employer by about 8 percentage points (i.e., a reduction of almost 1 worker out of ten expected to be trustworthy) relative to someone who receives no signal. In the first specification ("Basic"), we focus on sessions where the signal carried no information about worker characteristics and only indicated whether the worker was/was not trustworthy.[30] In the second specification ("All"), we include all sessions and we are able to identify separate effects for positive signals that carried information about a worker's gender and age (namely, indicating that the worker was female or below 25). However, we find no significant heterogeneity along these dimensions. In both specifications, only the effect of negative signals is significant and does not change considerably between the two columns.

Overall, our evidence shows that employers update their beliefs asymmetrically when they receive new information: they are sensitive to negative signals, but they are not affected by positive ones.[31] 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 escape. It also suggests 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.

<< **Table 5 here** >>

Finally, it is worth remarking that our objective was to provide employers with the same information they could obtain from hiring their first worker (i.e. one instance of experimentation). This is the most relevant approach in a context where the vast majority of micro-entrepreneurs hire no labour (as in most developing countries). Of course, one

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

[31] We also test whether this may be due to the fact that the distribution of expectations is naturally truncated at 100% and the median is around 70%, mechanically leaving more space for downward adjustments than for upward adjustments. When we restrict the analysis to employers who have expectations closer to 50% (i.e., we drop employers with expectations above 70% in the first round), the results do not change, though the drop in sample size reduces precision.

can easily think of more informative interventions that may have a stronger impact on expectations (e.g., revealing *average productivity* over the entire sample), but our interest was to study a learning process that cannot typically count on such detailed information and relies on experimentation.

If one were to design a more powerful information treatment, however, it would be useful to know *what is the causal effect of changing expectations on the rate of hiring/trusting*. The final piece of analysis in this section will answer this question. Our strategy is to use the random signals provided to workers in Treatment 1 to instrument expectations and uncover their causal impact on trust.

***Result 6:* Exogenously raising expectations has a strong impact on trusting**. The results from the instrumental variable estimation are reported in Table 6. The first two columns show the results from a naïve OLS regression of the probability of trusting as a function of expectations and individual characteristics, as in Table 5 above.[32] In the third specification, we instrument the expectations variable by assignment to Treatment 1 and we find that its effect grows considerably in magnitude, while remaining highly significant. Increasing employers' expectations by 1 (i.e., one extra worker out of 10 who is expected to be trustworthy, or an increase of 10 percentage points), increases the probability of trusting by amost 15 percentage points. 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 6 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. If the hypothesis that lack of experimentation with employees prevents entrepreneurs from forming unbiased expectations about workers' trustworthiness is correct, we should find that hiring outside the

[32] We restrict the analysis to round 2 because exogenous signals in Treatment 1 were only provided after the first round.

experiment helps employers to overcome their negative biases and correlates positively with their expectations (and with their propensity to trust). Thanks to a brief survey carried out shortly after the experiment, we were able to obtain information on past hiring (specifically, the number of employees the entrepreneur was hiring at the time when s/he last hired any labour).[33]

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 7 shows that the number of employees an entrepreneur hires in his/her businesses correlates positively with both employers' expectations of employees' trustworthiness and with employers' propensity to trust in the experimental setting.[34] The result is stronger when confining the sample to subjects with higher literacy, who had a better understanding of the follow-up questionnaire (and could thus provide a more precise account of their past hiring).[35] One potential concern with this correlation is that it may be partly driven by the fact that richer entrepreneurs with larger businesses can afford to trust more. We acknowledge this possibility, but it is worth underlining that while the incentivised choice to invest may be directly influenced by a respondent's economic conditions, the unincentivised question about expectations had no direct bearing on the monetary payoff from the game (and is less likely to be directly driven by wealth effects).[36]

While this only constitutes correlational 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 7 here >>

[33]It was not possible to interview 29 respondents. This explains the reduction in sample size. We focus on the last instance of hiring since we are interested in the effect of past experimentation. The variable is equal to 0 if the entrepreneur never hired any labour. 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.

[34]We focus on round 1 to exclude the impacts of our information treatment in round 2.

[35]The additional estimates are not included for conciseness, but available upon request.

[36]A respondent's economic conditions may influence his/her beliefs in other ways (e.g., through education). The controls used in the analysis, most importantly the respondent's educational background, are meant to address some of these concerns within the strict limits imposed by a small sample. We acknowledge the limitations of this approach and the analysis in this section is meant to be exploratory.

6 Conclusions

An established 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, despite the fact that a salient feature of labour markets in developing countries is the large numbers of micro-enterprises that hire little or no labour. 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 an investment game with a real-effort task, played between real entrepreneurs and workers in urban Ghana. By eliciting employers' expectations of workers' trustworthiness and comparing them with actual workers' performance, we are able to identify potential misperceptions leading to sub-optimal hiring. We also devise two randomised treatments to study the underlying learning mechanism and the extent to which expectations are biased against vulnerable groups.

We find 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 strong (downward) impact on employers' expectations. Positive signals have no effect. This asymmetry may 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. Our results suggest that successful information interventions targeting firm managers should take into account both biases in beliefs *and in the updating process*. Further work is needed to understand how misperceptions can be corrected when new information is assimilated asymmetrically.

References

Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin, and S. Quinn (2020). Matching frictions and distorted beliefs: evidence from a job fair experiment in Africa. *Unpublished manuscript*.

Alfonsi, L., O. Bandiera, V. Bassi, R. Burgess, I. Rasul, M. Sulaiman, and A. Vitali (2019). Tackling Youth Unemployment: Evidence from a Labour Market Experiment in Uganda. *PEDL Research Paper*.

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. (2017). An offer you can't refuse? incentives change how we inform ourselves and what we believe. *CESifo Working Paper Series* (6296).

Ambuehl, S. and S. Li (2018). Belief updating and the demand for information. *Games and Economic Behavior 109*, 21 – 39.

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

Ashraf, N., A. Delfino, and E. L. Glaeser (2019, October). Rule of law and female entrepreneurship. Working Paper 26366, National Bureau of Economic Research.

Atkin, D., A. Chaudhry, S. Chaudry, A. K. Khandelwal, and E. Verhoogen (2017, 03). Organizational barriers to technology adoption: Evidence from soccer-ball producers in pakistan. *The Quarterly Journal of Economics 132*(3), 1101–1164.

Attanasio, O., T. Boneva, and C. Rauh (2019, January). Parental beliefs about returns to different types of investments in school children. Working Paper 25513, National Bureau of Economic Research.

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

Attanasio, O. P. and K. M. Kaufmann (2017). Education choices and returns on the labor and marriage markets: Evidence from data on subjective expectations. *Journal of Economic Behavior & Organization 140*, 35–55.

Banerjee, A. and S. Sequeira (2020). Spatial mismatches and imperfect information in the job search. CEPR Discussion Papers DP14414.

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

Bartling, B., E. Fehr, D. B. Huffman, and N. Netzer (2018, October). The Causal Effect of Trust. *IZA Discussion Papers* (11917).

Berg, J., J. Dickhaut, and K. McCabe (1995). Trust, reciprocity, and social history. *Games and economic behavior 10*(1), 122–142.

Bloom, N., B. Eifert, A. Mahajan, D. McKenzie, and J. Roberts (2012, 11). Does Management Matter? Evidence from India. *The Quarterly Journal of Economics 128*(1), 1–51.

Bloom, N., R. Lemos, R. Sadun, D. Scur, and J. Van Reenen (2014, 08). JEEA-FBBVA Lecture 2013: The New Empirical Economics of Management. *Journal of the European Economic Association 12*(4), 835–876.

Bloom, N. and J. Van Reenen (2007). Measuring and explaining management practices across firms and countries. *The quarterly journal of Economics 122*(4), 1351–1408.

Bloom, N. and J. Van Reenen (2010, March). Why do management practices differ across firms and countries? *Journal of Economic Perspectives 24*(1), 203–24.

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 (2017, 10). The morale effects of pay inequality. *The Quarterly Journal of Economics 133*(2), 611–663.

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

Bursztyn, L., A. Rao, C. Roth, and D. Yanagizawa-Drott (2020). Misinformation during a pandemic. *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2020-44).

Butler, J. V., P. Giuliano, and L. Guiso (2016). The right amount of trust. *Journal of the European Economic Association 14*(5), 1155–1180.

Caria, S. (2019). Industrialization on a knife's edge : Productivity, labor costs and the rise of manufacturing in ethiopia. Policy Research Working Paper 8980, The World Bank.

Chandrasekhar, A. G., M. Morten, and A. Peter (2020, February). Network-based hiring: Local benefits; global costs. Working Paper 26806, National Bureau of Economic Research.

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 (2017, February). When No Bad Deed Goes Punished: Relational Contracting in Ghana versus the UK. *NBER Working Papers* (23123).

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.

Delavande, A. and B. Zafar (2019). University choice: the role of expected earnings, nonpecuniary outcomes, and financial constraints. *Journal of Political Economy 127*(5), 2343–2393.

DellaVigna, S. and D. Pope (2018). Predicting experimental results: Who knows what? *Journal of Political Economy 126*(6), 2410–2456.

Enke, B., U. Gneezy, B. Hall, D. Martin, V. Nelidov, T. Offerman, and J. van de Ven (2020). Cognitive biases: Mistakes or missing stakes? *CESifo Working Paper Series* (8168).

Enke, B. and F. Zimmermann (2017, 12). Correlation Neglect in Belief Formation. *The Review of Economic Studies 86*(1), 313–332.

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 (2018). Information processing and commitment. *The Economic Journal 128*(613), 1983–2002.

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. *CeDEx Discussion Paper No.2010-20*.

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

Guiso, L., P. Sapienza, and L. Zingales (2008, 05). Social Capital as Good Culture. *Journal of the European Economic Association 6*(2-3), 295–320.

Heath, R. (2018). Why do firms hire using referrals? evidence from bangladeshi garment factories. *Journal of Political Economy 126*(4), 1691–1746.

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

Hoffman, M., L. B. Kahn, and D. Li (2018). Discretion in hiring. *The Quarterly Journal of Economics 133*(2), 765–800.

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.

Lemos, R. and D. Scur (2014). Management practices in manufacturing in mozambique. *IGC Working Paper F-36108-MOZ-1*.

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 (2013). A land of milk and honey with streets paved with gold: Do emigrants have over-optimistic expectations about incomes abroad? *Journal of Development Economics 102*, 116 – 127. Migration and Development.

McKenzie, D. and C. Woodruff (2017). Business practices in small firms in developing countries. *Management Science 63*(9), 2967–2981.

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. *American Economic Review 81*(5), 1408–1418.

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

Sánchez de la Sierra, R. (2016). Ethnic State Capacity and Contract Enforceability. *PEDL Research Papers*.

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.

Figures

Figure 1: Trust around the world

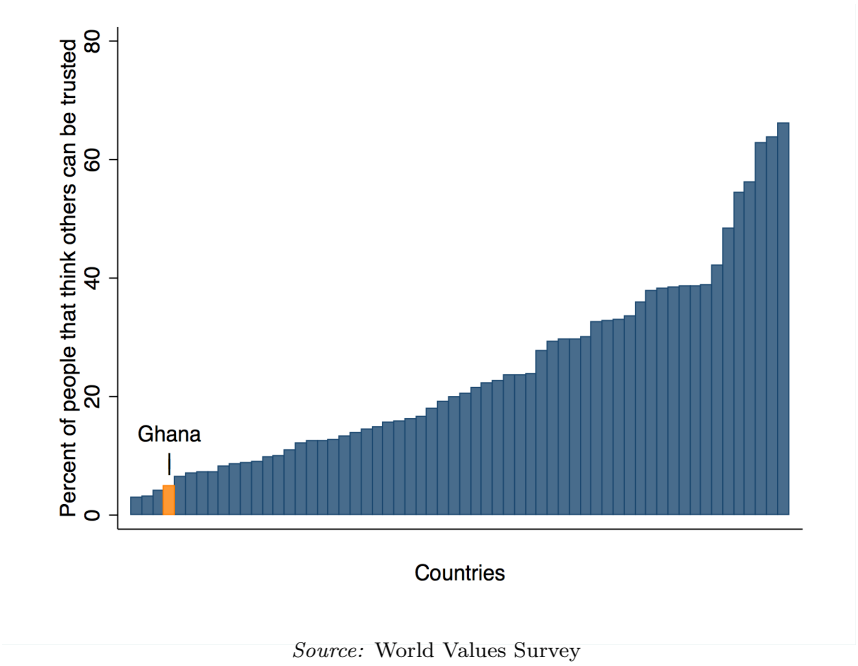

Source: World Values Survey

Figure 2: The game tree

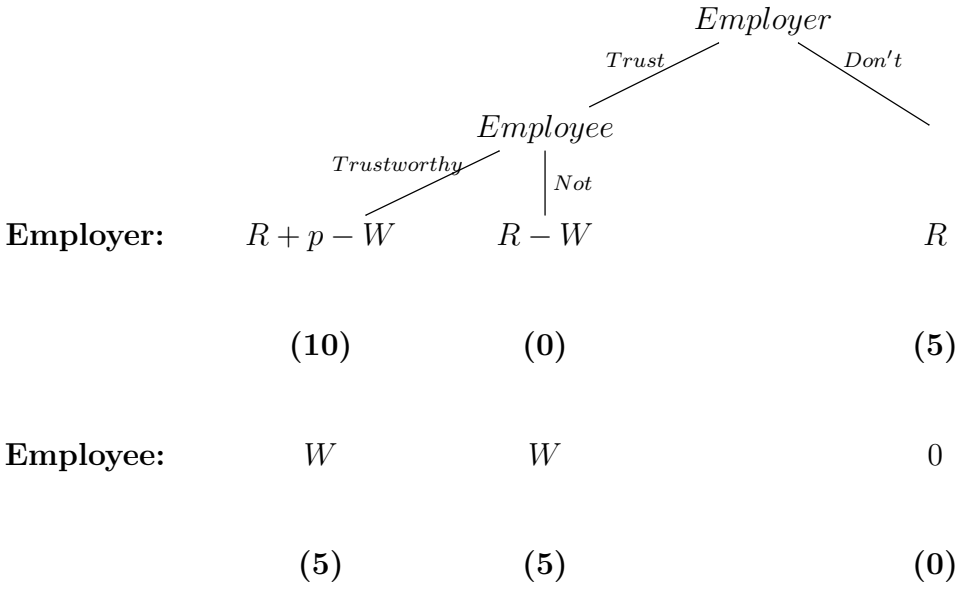

Note:: The numbers in parentheses indicate the monetary values of the payoffs (in Ghana Cedis).

Figure 3: Expected and revealed trustworthiness

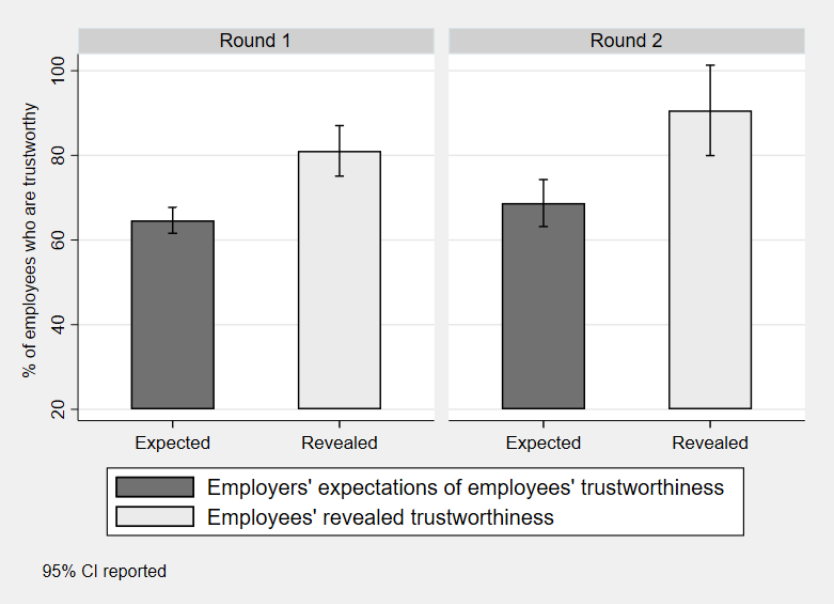

Note: Percentage of employees who reveal to be trustworthy (i.e., exert effort in the task) as *expected* by the employers and *revealed* in the experiment. Control sessions only.

Figure 4: Cumulative distribution of employers' expectations

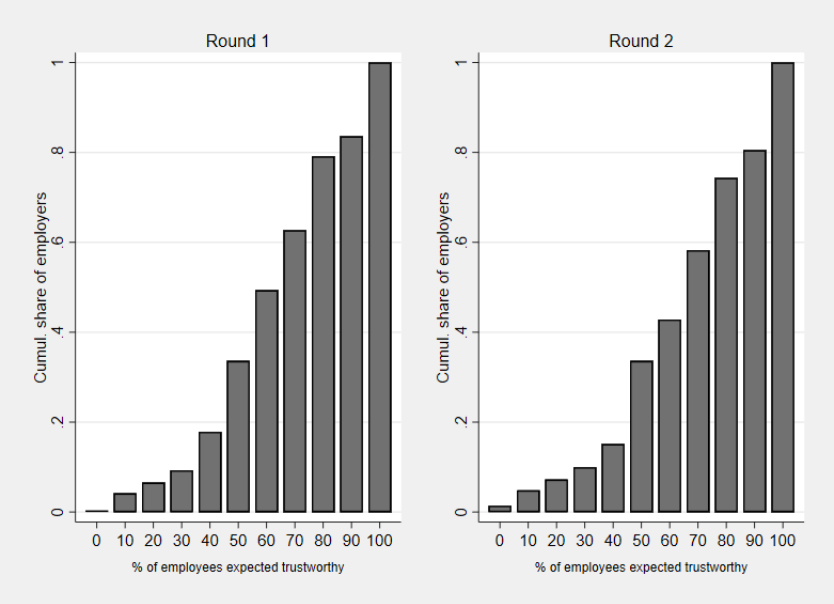

Note:: Cumulative distribution of the percentage of employees expected to be trustworthy by employers (i.e., expected to exert effort in the task).

Figure 5: Expected and revealed trust

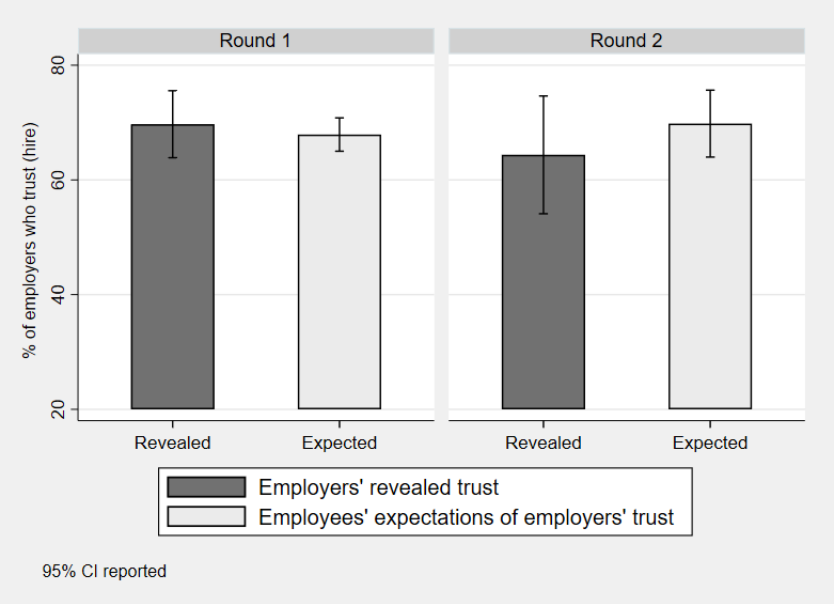

Note: Percentage of employers who trust workers (i.e., choose to hire) as *expected* by the employees and *revealed* in the experiment. Control sessions only.

Figure 6: Revealed trustworthiness by gender

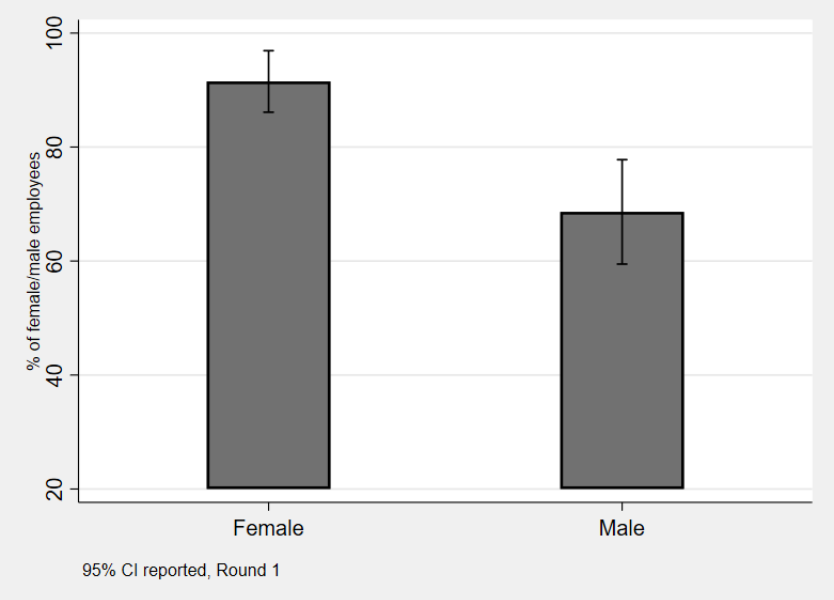

Note: Percentage of employees who reveal to be trustworthy (i.e., exert effort in the task) by gender.

Figure 7: Revealed trustworthiness by age

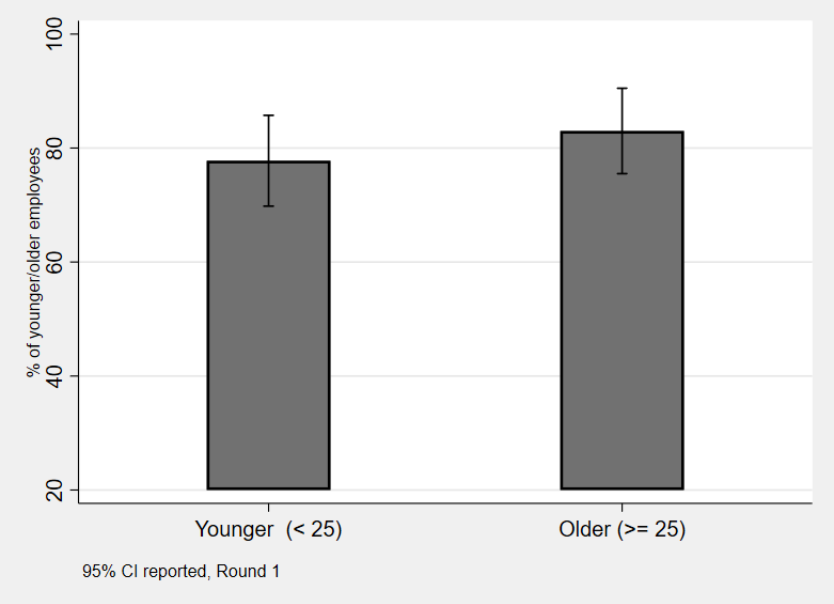

Note: Percentage of employees who reveal to be trustworthy (i.e., exert effort in the task) by age.

Tables

Table 1: Employers

Variable	Mean	Std. Dev.
Age	37.582	11.701
Male	0.366	0.483
Married	0.562	0.497
No Schooling	0.106	0.309
Primary	0.103	0.304
Lower secondary	0.312	0.464
Upper secondary	0.312	0.464
Tertiary	0.168	0.374
Unemployed	0.106	0.309
N	292	

Note: Summary statistics for the sample of *employers*.

Table 2: Employees

Variable	Mean	Std. Dev.
Age	27.954	10.089
Male	0.47	0.5
Married	0.278	0.449
No Schooling	0.064	0.245
Primary	0.096	0.295
Lower secondary	0.263	0.441
Upper secondary	0.363	0.482
Tertiary	0.214	0.411
Unemployed	0.356	0.48
N	281	

Note: Summary statistics for the sample of *employees*.

Table 3: How employers' expectations change with variations in the worker pool

(Dep Var: Employer's expectations of workers' trustworthiness)

	All	Male Employers	Female Employers
Young workers (< 25)	.169	1.492	-.417
	(.594)	(.600)**	(.765)
Female workers	-.589	-1.480	-.055
	(.327)*	(.623)**	(.370)
Const.	5.284	4.573	5.666
	(.587)***	(.876)***	(.762)***
Employer Char.s	Yes	Yes	Yes
Session Char.s	Yes	Yes	Yes
Obs.	292	107	185

Note: The table shows how employers' expectations of workers' trustworthiness change with the composition of the worker pool (Treatment 2). *Young workers* is a dummy equal to 1 if the employer was in a session where 80% of employees were below 25; *Female workers* is a dummy equal to 1 if the employer was in a session where 80% of employees were female. Employers' expectations are measured on a scale from 1 to 10 (i.e., number of workers the employer thinks will reveal to be trustworthy out of 10). Employer Char.s ≡ Age, Gender, Education, Marital status, Employment status. Session Char.s ≡ Time of session and Timing of expectation elicitation (before or after the hiring decision). *Confidence:* *** ↔ 99%, ** ↔ 95%, * ↔ 90%. Robust standard errors in parentheses (cluster. by Session).

Table 4: Trust and expected trustworthiness

(Dep Var: Whether the employer trusts/hires a worker)

	OLS1	OLS2	Probit1	Probit2	Probit-Mfx
E[Trustw]	.030	.027	.083	.080	.027
	(.008)***	(.008)***	(.023)***	(.025)***	
Risk Aversion		-.048		-.148	-.051
		(.011)***		(.036)***	
Round 2	-.016	-.015	-.043	-.038	-.013
	(.052)	(.048)	(.150)	(.146)	
Const.	.494	.593	-.042	.267	
	(.066)***	(.107)***	(.177)	(.317)	
Employer Char.s	No	Yes	No	Yes	Yes
Session Char.s	No	Yes	No	Yes	Yes
Obs.	584	584	584	584	584

Note: The table shows the relationship between the probability that an employer trusts (hires) and his/her expectations of workers' trustworthiness. The estimation pools observations from round 1 and round 2. Employers' expectations (E[Trustw]) are measured on a scale from 1 to 10 (number of workers the employer thinks will reveal to be trustworthy out of 10). Employer Char.s ≡ Age, Gender, Education, Marital status, Employment status. Session Char.s ≡ Time of session and Timing ofexpectation elicitation (before or after the hiring decision). Probit-Mfx ≡ Marginal effects from Probit2. *Confidence:* *** ↔ 99%, ** ↔ 95%, * ↔ 90%. Robust standard errors in parentheses (cluster. by Session + Round).

Table 5: Expectation updating

(Dep Var: Employer's expectations)

	Basic	All
E[Trustw, R1]	.649	.682
	(.054)***	(.050)***
Negative Signal	-.855	-.776
	(.396)**	(.379)**
Positive Signal	-.045	.070
	(.455)	(.432)
Positive Signal + Female employee		.175
		(.561)
Positive Signal + Young employee (< 25)		-.503
		(.627)
T1: basic signals	.552	.479
	(.473)	(.458)
T1: signals with gender/age info		-.200
		(.486)
Const.	2.523	2.422
	(.619)***	(.564)***
Entrepreneur Char.s	Yes	Yes
Session Char.s	Yes	Yes
Obs.	195	292

Note: The table shows the effect of receiving positive and negative signals of workers' trustworthiness on employers' expectations in round 2 (controlling for expectations in round 1). The signals (Treatment 1) consisted of revealing whether a random worker had been trustworthy in previous rounds of the experiment. Employers' expectations are measured on a scale from 1 to 10 (number of workers the employer thinks will reveal to be trustworthy out of 10). *Basic* includes T1 sessions where the signals only carried information about the worker's trustworthiness. *All* includes all T1 sessions (i.e., also those where the signal indicated that the worker was young or female to test for heterogeneous effects). The dummies *Negative Signal* and *Positive Signal* are equal to 1 if the employer received a basic signal that carried no information about the gender or age of the worker. *Positive Signal + Female employee* and *Positive Signal + Young employee* are equal to 1 if the employer received a signal that was associated to a female or a young worker. *T1: basic signals* is a dummy equal to 1 if the employer was in a session with basic signals (no info about gender or age of the worker). *T1: signals with gender/age info* is equal to 1 if the employer was in a session with signals that also carried information about gender and age (these session effects can be identified because 2 employers per session received no signal). Employer Char.s ≡ Age, Gender, Education, Marital status, Employment status. Session Char.s ≡ Time of session and Timing of expectation elicitation (before or after the hiring decision). *Confidence:* *** ↔ 99%, ** ↔ 95%, * ↔ 90%. Robust standard errors in parentheses (cluster. by Session).

Table 6: The impact of raising expectations on trust

(Dep Var: Whether the employer trusts/hires a worker)

	OLS	OLS	IV
E[Trustw, R2]	.036	.032	.147
	(.011)***	(.012)***	(.074)**
Risk Aversion		-.059	-.057
		(.016)***	(.022)***
Const.	.434	.627	-.066
	(.086)***	(.155)***	(.480)
Entrepreneur Char.s	No	Yes	Yes
Session Char.s	No	Yes	Yes
Obs.	292	292	292

Note: The table shows the impact of raising expectations on employers' decision to trust in round 2. *OLS1* and *OLS2* replicate the naïve regressions in Table 4 for round 2 only. *IV* uses assignment to T1 (where employers received random signals of employees' trustworthiness just before round 2) as an instrument for expectations. Employers' expectations (E[Trustw, R2]) are measured on a scale from 1 to 10 (number of workers the employer thinks will reveal to be trustworthy out of 10 in round 2). Employer Char.s ≡ Age, Gender, Education, Marital status, Employment status. Session Char.s ≡ Time of session and Timing of expectation elicitation (before or after the hiring decision). *Confidence:* *** ↔ 99%, ** ↔ 95%, * ↔ 90%. Robust standard errors in parentheses (cluster. by Session).

Table 7: The relationship between real-life hiring and trust in the experiment

(Dep Var.s: Employer's expectations & Whether the employer trusts/hires a worker)

	E[Trustw]	P[Hire]	E[Trustw]	P[Hire]
Num. Employees hired	0.168***	0.0298**	0.125*	0.0271**
	(0.0596)	(0.0131)	(0.0671)	(0.0126)
Const	6.453***	0.661***	5.145***	0.562***
	(0.129)	(0.0436)	(0.688)	(0.176)
Entrepreneur Char.s	No	No	Yes	Yes
Session Char.s	No	No	Yes	Yes
Obs.	263	263	263	263

Note: The table shows the relationship between hiring in real-life businesses and: (i) employers' expectations of workers' trustowrthiness (E[Trustw]), and (ii) whether the employer hires in the game (P[Hire]). Employers' expectations (E[Trustw]) are measured on a scale from 1 to 10 (number of workers the employer thinks will reveal to be trustworthy out of 10). Analysis limited to round 1 to exclude treatment effects in round 2. *Num. Employees hired* indicates the number of employees hired by the entrepreneur at the time when s/he last hired workers (the information is missing for 29 entrepreneurs, hence the lower sample size compared to previous tables). Employer Char.s ≡ Age, Gender, Education, Marital status, Employment status; Session Char.s ≡ Time of session and Timing of expectation elicitation (before/after hiring decision). *Confidence:* *** ↔ 99%, ** ↔ 95%, * ↔ 90%. Robust standard errors in parentheses (cluster. by Session).

APPENDIX

A Additional figures and tables

Figure 8: Risk-game lotteries

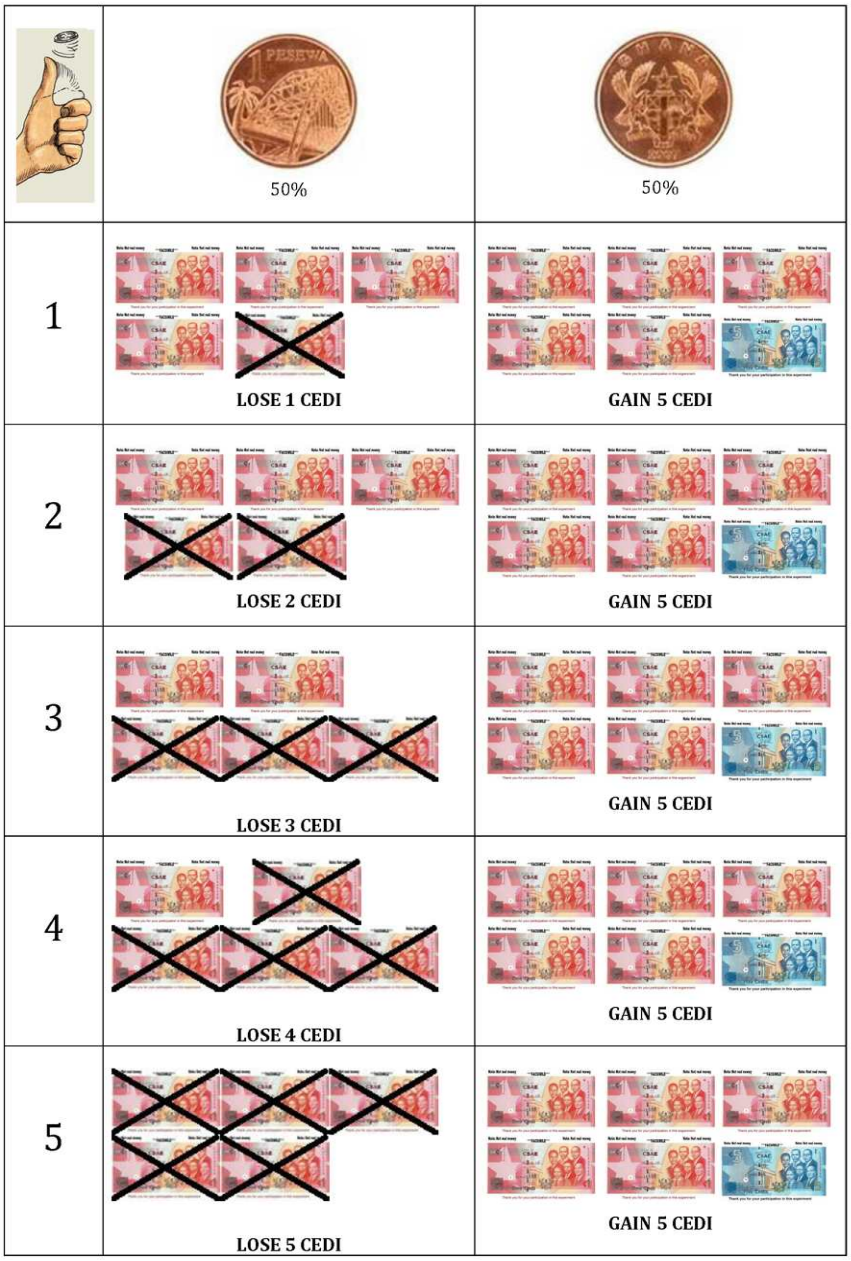

Note: Visual aid for the risk-preference elicitation game.

Table 8: Covariates balance across treatments

	T1	T0,T2	Equality	T2	T0,T1	Equality
Variable	Mean		p-value	Mean		p-value
Age	36.7	38.0	.37	37.1	37.8	.59
Primary	.08	.11	.40	.10	.10	.99
Lower secondary	.27	.34	.23	.33	.30	.64
Upper secondary	.33	.30	.70	.40	.27	.02
Tertiary	.15	.18	.63	.12	.19	.16
Married	.49	.60	.08	.57	.56	.90
Male	.36	.37	.82	.39	.35	.53
Unemployed	.06	.82	.08	.11	.10	.78
Obs.	98	194	292	97	195	292

Note: Tests of balance in the characteristics of employers between treatment arms. T0 ≡ Control sessions; T1 ≡ Information provision; T2 ≡ Changing the composition of the worker pool. T1 is compared with T0 + T2 while T2 is compared with T0 + T1, as those are the relevant comparisons for the analysis (i.e., our design is such that we can estimate the effects of Treatment 1 by comparing subjects in T1 with those in T0 + T2 sessions, and we can estimate the effects of Treatment 2 by comparing subjects in T2 with those in T0 + T1 sessions). Primary - Tertiary indicate the highest educational attainment of the participant.

B Experimental protocol

Control

Round 1

1. *Employer* informed of average worker char.s (gender, age, etc.) to anchor expectations.
2. *Employer*'s expectations of workers' trustworthiness elicited.
3. *Employer* chooses hire/not hire a random and anonymous employee.

   [Note: 2 and 3 randomly inverted in half of the sessions]
4. *Employee*'s expectations of employers' trust elicited.
5. *Employee* chooses effort (not revealed to employer until the end of round 2).

   [Note: 4 and 5 randomly inverted in half of the sessions]

Round 2 (announced as a surprise at the end of round 1)

1. *Employer*'s expectations (re)elicited.
2. *Employer* chooses hire/not hire a (new) random employee.

   [Note: 1 and 2 randomly inverted in half of the sessions]
3. *Employee*'s expectations (re)elicited.
4. *Employee* chooses effort.

   [Note: 3 and 4 randomly inverted in half of the sessions]

Post-game:

1. Additional questions (including elicitation of risk preferences)
2. Payoffs revealed and prizes distributed.

Treatment 1: Information provision

Round 1

1. *Employer* informed of average worker char.s (gender, age, etc.) to anchor expectations.
2. *Employer*'s expectations of workers' trustworthiness elicited.
3. *Employer* chooses hire/not hire a random and anonymous employee.

   [Note: 2 and 3 randomly inverted in half of the sessions]
4. *Employee*'s expectations of employers' trust elicited.
5. *Employee* chooses effort (not revealed to employer until the end of round 2).

   [Note: 4 and 5 randomly inverted in half of the sessions]

Round 2 (announced as a surprise at the end of round 1)

1. ***Employer informed of trustworthiness of random employee from previous sessions.***
2. *Employer*'s expectations (re)elicited.
3. *Employer* chooses hire/not hire a (new) random employee.

   [Note: 2 and 3 randomly inverted in half of the sessions]
4. *Employee*'s expectations (re)elicited.
5. *Employee* chooses effort.

   [Note: 4 and 5 randomly inverted in half of the sessions]

Post-game:

1. Additional questions (including elicitation of risk preferences)
2. Payoffs revealed and prizes distributed.

Treatment 2: Changing the composition of the worker pool

Round 1

1. ***Employer informed that 80% of workers belong to a specific category (i.e., young or female)***.
2. *Employer*'s expectations of workers' trustworthiness elicited.
3. *Employer* chooses hire/not hire a random and anonymous employee.

   [Note: 2 and 3 randomly inverted in half of the sessions]
4. *Employee*'s expectations of employers' trust elicited.
5. *Employee* chooses effort (not revealed to employer until the end of round 2).

   [Note: 4 and 5 randomly inverted in half of the sessions]

Round 2 (announced as a surprise at the end of round 1)

1. *Employer*'s expectations (re)elicited.
2. *Employer* chooses hire/not hire a (new) random employee.

   [Note: 1 and 2 randomly inverted in half of the sessions]
3. *Employee*'s expectations (re)elicited.
4. *Employee* chooses effort.

   [Note: 3 and 4 randomly inverted in half of the sessions]

Post-game:

1. Additional questions (including elicitation of risk preferences)
2. Payoffs revealed and prizes distributed.