

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/165547>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

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

Stefano A. Caria[†] Paolo Falco[‡]

March 14, 2022

Abstract

We conduct a lab-in-the-field 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 entrepreneurs' expectations and we randomly provide information on workers' previous behaviour to study belief updating. We find that entrepreneurs significantly underestimate workers' trustworthiness. This reduces hiring and profit in the experiment. We also find evidence of asymmetric learning: negative signals lower entrepreneurs' expectations while positive signals do not affect them. Experience with hiring in daily life correlates positively with entrepreneurs' expectations and with their willingness to trust. Our evidence corroborates the hypothesis that a low-trust labour market 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 entrepreneurs think workers can be trusted? There is little evidence on the beliefs of entrepreneurs and 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 entrepreneurs overly pessimistic about the trustworthiness of workers. First, trust presents a classic problem of experimentation: an overly pessimistic entrepreneur will hire fewer workers and hence will receive less 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, entrepreneurs 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, entrepreneurs 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, incorrect 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 entrepreneurs' 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 a lab-in-the-field experiment based on a trust game with a real-effort task, played between real entrepreneurs and workers in urban Ghana. Entrepreneurs have the option to trust an anonymous worker by investing an initial endowment to pay his/her wage. If hired, the worker chooses whether to exert effort in order to complete a simple task that requires no particular skill or ability. The entrepreneur's payoff is directly linked to the worker's performance, while the worker's wage, if hired, is fixed. By decoupling the worker's pay from performance and by removing the entrepreneur's ability to monitor and punish low effort, we are able to elicit trust and trustworthiness. Crucially, we elicit entrepreneurs' expectations of workers' trustworthiness and we compare them with actual workers' performance to identify potential misperceptions. We further devise two treatments to study the extent to which expectations are biased against specific worker categories and to investigate how entrepreneurs update their priors as they receive new information.

We find that entrepreneurs significantly underestimate workers' trustworthiness and that low expectations have negative repercussions for hiring and profits. On average, entrepreneurs underestimate workers' trustworthiness by 20 percent and we estimate that closing the expectation gap would increase the probability of hiring in the experiment by approximately 16 percent. The result is robust to controlling for a separate experimental measure of risk-aversion. We also find evidence of a gender gap. Male entrepreneurs expect female workers to be less trustworthy than men, while behaviour in the experiment reveals the opposite.

Furthermore, we show that as entrepreneurs acquire new information they update their

expectations asymmetrically. Negative information on workers' performance has a negative effect on entrepreneurs' beliefs, while positive signals do not change beliefs significantly.

When we explore the correlation between the daily-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 correlation lends further credibility to our measures of trust and expected trustworthiness. Further, it 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 burgeoning 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 hiring decisions.³ Our approach also relates to a growing literature that studies subjects' ability to predict experimental results ([DellaVigna and Pope, 2018](#)).

Furthermore, we provide evidence on why biased expectations may persist. First, the fact that the entrepreneurs who hire fewer workers (in daily-life and in the lab) have beliefs that are further from the truth suggests that insufficient experimentation may sustain an incorrect

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. In our data, there is a positive, but insignificant correlation between age and expected trustworthiness, in line with their conclusion that in low-trust societies only limited experimentation occurs.

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

³This is not to say that entrepreneurs and managers have no information on workers that can be used productively. For example, [Friebel et al. \(2021\)](#) show that, when asked to reduce personnel turnover, store managers are able to reduce staff quits by allocating more time towards HR activities.

perception of the world. Furthermore, we present evidence suggesting that asymmetric learning may entrench over-pessimism. This finding draws 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).⁴

Third, we contribute to a recent literature that studies information frictions in developing countries' labour markets (Abel et al., 2020; Carranza et al., 2020; Abebe et al., 2021; Bassi and Nansamba, 2021). This literature emphasises that employers value worker traits such as trustworthiness and dependability, which are difficult to observe in job interviews and screening tests. Worker-level skill certification is a promising solution for this friction. In a similar vein, Akcigit et al. (2021) and a long-standing literature on the span-of-control argue that the expectation of moral hazard limits firm size in developing countries. Our results advance these literatures by showing that entrepreneurs' perception of average worker trustworthiness may be biased. Thus, the optimal policy response to information frictions may entail a combination of skill certification and information provision about labour market fundamentals.

Furthermore, our results suggest that entrepreneurs may update their beliefs more strongly when they observe a worker under-performing compared to when they observe a worker completing the required task. Certification interventions, on the other hand, have been shown to have positive impacts on employer beliefs (e.g., Bassi and Nansamba (2021)). These interventions provide a very precise signal about worker ability (workers are ranked against each other in terms of their performance in a state-of-the-art personnel selection test), which employers may find harder to ignore. This suggests that belief updating may potentially be more subject to bias when employers observe single instances of behaviour, compared to when they are presented with more precise signals.

⁴The recent literature has also emphasised the possibility of differential updating by the gender of the person whose performance is evaluated (Sarsons, 2019). In our experiment, we can study whether entrepreneurs respond differentially to positive signals that are explicitly about a female worker. We are unable to find an effect, but our statistical power for this test is likely to be limited.

Finally, we contribute to the literature that studies the quality of management in 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 entrepreneurs' beliefs and in the way they assimilate new information, and to show how this affects decision-making.⁵ We also find a strong correlation between our lab-in-the-field measure of trust and daily-life hiring. This showcases the potential for experimental methods in the laboratory to generate variables that have predictive power over real employer decisions, and may help researchers develop better models of employer behaviour in developing countries.

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, the role it plays in the labour market, and its relationship with business growth. In section 3 we outline and discuss 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 on firm hiring. Three pieces of evidence provide support for this view and motivate our analysis.

First, generalised trust among people in Ghana is low. Harmonised cross-country data from the World Values Survey (WVS) shows that out of 60 countries analysed (developed and developing) across all the continents, Ghana has the fourth lowest level of trust (Figure 1).⁶

⁵This relates to the work by Davies and Fafchamps (2017), who run a gift-exchange lab experiment with Ghanaian students and show that subjects fail to screen out partners who have been untrustworthy in the past.

⁶Trust in the World Values Survey is measured by means of a standard question that asks whether the respondent *thinks that others can be trusted*.

Figure 1: Trust around the world



Source: World Values Survey

Second, firms in Ghana make a limited use of extrinsic incentives to motivate their workers, which suggests an important role for trust and trustworthiness in hiring decisions. We document this point using data from the World Management Survey (WMS), which shows that firms in Ghana are close to the bottom of the international distribution of the key indicators capturing the use of worker incentives — “building a high-performance culture through incentives and appraisals” and “removing poor performers” (see Figure A.1 in the Appendix).⁷ This goes hand in hand with the fact that enterprises in Ghana are small on average (Falco et al., 2014; Teal, 2016). Indeed, the WMS data suggests that two strong predictors of the use of worker incentives are (i) firm size (firms in the largest size category have an “incentives and appraisal” score that is .6 standard deviations higher than firms in the smallest size category) and (ii) the country’s GDP (one standard deviation increase in GDP is associated with a .1 standard deviation increase in the “incentives and appraisal” score).

Third, cross-country evidence also shows a positive correlation between entrepreneurs’ willingness to expand their business and aggregate trust levels. The Global Entrepreneur

⁷The WMS includes standardised scores on a wide variety of management practices for 11,000 firms in 35 countries, ranging from the very poor (e.g., Ethiopia), to the very rich (e.g., United States).

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. We merge this data with the trust variable from the World Values Survey, and find that as generalised levels of trust increase, the share of new businesses that expect to hire at least five workers in the coming five years grows.⁸

Overall, this descriptive evidence suggests that Ghana is a setting where trust can play an important role in the labour market. Furthermore, many of the features that we emphasise here — low trust, small average firm size, limited use of performance incentives — are typical of developing countries more broadly. Our results will thus be of interest for the larger set of developing countries that share similar structural characteristics.

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 manipulations we use to study specific mechanisms.

3.1 The players

Each respondent was pre-assigned to one of two experimental roles: *entrepreneur* or *worker* (with real entrepreneurs assigned to the role of experimental entrepreneurs) and randomly allocated to an experimental session. Each session had 10 entrepreneurs and 10 workers in total.⁹ 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 entrepreneurs' and the workers' room respectively). Inside each room,

⁸The relationship is concave and the positive correlation weakens once countries reach high levels of trust (beyond the levels observed in Ghana).

⁹There were few exceptions, when some respondents dropped out and the session was not full. The plan was to have 30 sessions with 10 entrepreneurs and 10 workers, but the average number of entrepreneurs and workers per session was 9.73 and 9.37, i.e. $N = 292$ and $N = 281$, respectively). This did not affect the decisions of players. Entrepreneurs knew that workers played independently and were physically isolated from one another (as described below). Hence, having one or two fewer players in the room did not affect their actions. Similarly, entrepreneurs made their decisions independently and were physically separated from one another. For simplicity, we told entrepreneurs that the workers 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.

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 workers' room were higher to ensure that workers could not see each other at all, minimising peer effects on productivity (Falk and Ichino, 2006; Mas and Moretti, 2009).¹⁰ They also ensured that respondents did not feel observed or monitored by the experimenter. All the features of the workers' 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 C and the full script is available upon request.

Each respondent in the role of entrepreneur 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 . If the worker exerts effort ($e_i = 1$), production is realised. If the worker shirks ($e_i = 0$), 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). Entrepreneurs have full knowledge of this.

When the worker chooses to exert effort and production is successful, the entrepreneur receives profits equal to $p > R$ and the worker receives a wage of W . When the worker shirks and production fails, the entrepreneur does not receive any money, while still having to pay W to the worker. Our chosen payoffs were such that $R = W$ (i.e. a trusting entrepreneur loses

¹⁰Moreover, 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).

the entire endowment when the worker shirks). This type of contractual arrangement, with no monetary penalty for shirking, allows us to isolate trust as a driver of the entrepreneurs' choice. The game tree in Figure A.2 of the Online Appendix outlines the sequence of choices and corresponding payoffs faced by entrepreneurs and workers 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.¹¹

Based on these rules, entrepreneurs were asked to choose between hiring a worker and keeping their endowment. For every entrepreneur that chose to hire, a worker was asked to undertake the real-effort task and could decide whether to exert effort (or shirk).¹² The game was played twice, with a second round announced as a surprise at the end of the first one. Entrepreneurs started the second round with a new endowment equal to the one they had in round 1. The results from the first round (i.e. whether a hired worker exerted effort or not) were not announced to entrepreneurs until the very end of the game to avoid influencing their decision in the second round. Before or after the entrepreneurs' decision to hire in each round (the order was randomised), their expectations of workers' effort 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

¹¹The median of net daily earnings for workers in the reference population was approximately 6 Ghana Cedis according to data from the Oxford University Ghana Household Urban Panel Survey available at the time of the experiment (August 2013). The endowment provided to entrepreneurs in each round of the investment game was 5 Cedis. The endowment for the additional 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.

¹²The match between entrepreneurs and workers was random and entrepreneurs did not know who their worker would be when making their decision. The respondents in the worker room who were not hired remained unemployed and only received a participation fee.

task (e.g. [Fehr et al. \(1993\)](#), [Altmann et al. \(2012\)](#)) as real-effort and chosen effort, although correlated, may diverge substantially. Indeed, [Bruggen and Strobel \(2007\)](#) show that effort is significantly higher in real-effort tasks, compared to chosen-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.¹³

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.¹⁴ 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. Workers were simply left free to use the time at their disposal as they pleased. They could decide to exert effort to complete the task, watch TV, or dedicate to other activities they might have preferred (e.g., texting or browsing the internet with their phones). As explained above, they were constantly shielded from anybody else’s sight. Entrepreneurs were aware of these conditions. Furthermore, it was common knowledge that any participant could leave at any point during the game.

Given the features of the task, we are confident that workers’ success is a valid proxy of their willingness to make an effort.¹⁵ Exerting effort, in turn, is a direct measure of trustworthiness,

¹³The 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.

¹⁴We took care in choosing a show that was widely known and popular across age groups and genders.

¹⁵It 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

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. These measures of trust and trustworthiness align well with those captured by the canonical investment game of [Berg et al. \(1995\)](#), which has been replicated extensively around the world ([Johnson and Mislin, 2011](#)), including in developing countries (e.g. [Barr \(2003\)](#); [Karlan \(2005\)](#); [Ashraf et al. \(2006\)](#)).

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 entrepreneurs 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 workers: "*Out of the 10 entrepreneurs 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.¹⁶ The elicitation questions were not independently incentivised.

We took several steps to give participants' the best possible chance of making correct predictions. First, participants were given clear information about the rules and incentives that applied to the other players. Second, enumerators showcased the effort task to the entrepreneurs (who were free to experiment with the task themselves). Entrepreneurs were also aware of the presence of the TV in the employees' room. Third, employers were given accurate information

trustworthiness that are above entrepreneurs' expectations.

¹⁶Under 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.

about how experimental employees were sampled and about the composition of the employee sample, including the median age, gender, and employment status of workers.

3.5 Treatments

In addition to the version of the game described above (which we will call T0), we assigned subjects in randomly selected sessions to one of two additional treatments.

Treatment 1 (T1): Signals of performance. At the end of round 1, entrepreneurs in T1 sessions received a note showing them whether a random worker out of all those hired in the initial pilot sessions completed the task (i.e. proved to be trustworthy) or not. In doing so, the treatment effectively conveyed to entrepreneurs the same information they could gather through their own experience of hiring a worker.¹⁷ Crucially, we provided this information both to entrepreneurs 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 observational data) whereby new information only accrues to entrepreneurs who are willing to experiment.¹⁸ For the sake of comparability, the entrepreneurs who hired in the first round did not know the performance of their worker in round 1 at the time of receiving the signal in round 2 (the outcomes of both rounds were revealed at the end of experiment).

In one half of T1 sessions (T1A), the random signals did not reveal any worker characteristic. In the other half (T1B), they specifically referred to a worker who was female or young (below 25).¹⁹ This additional variation allows us to test whether entrepreneurs' 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.

¹⁷Alternatively, we could have provided entrepreneurs with information on average trustworthiness in previous sessions. We decided against this option in order to maximise the realism of the design. Entrepreneurs presumably form their beliefs by observing specific instances of behaviour and would only exceptionally be presented with summary statistics on a large number of prior experiments.

¹⁸Respondents in the role of workers, on the other hand, were informed between round 1 and round 2 that entrepreneurs had received information on workers' performance (and may have updated their beliefs accordingly). This was intended to test for the potential effect of changing entrepreneurs' expectations on workers' 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 workers in these sessions from the estimation of workers' trustworthiness in round 2 below.

¹⁹Due 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.

Treatment 2 (T2): Changing the composition of the worker pool. In T2 sessions the sample of respondents invited to take part as “workers” 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 (T2W), 80% of invited workers were *women* and, in the other half (T2Y), 80% were *young* (below the age of 25). Entrepreneurs 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 typical gender and age composition of the invited worker pool across the experiment). The rest of the design was identical to T0.

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.²⁰ Figure A.3 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 entrepreneurs 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 entrepreneurs 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) where the respondent preferred the safe option over the risky lottery.²¹

²⁰Using 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.

²¹Given 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 entrepreneurs.

3.7 Discussion of the design

In this final subsection, we discuss two features of our design that merit specific consideration. First, workers in the experiment are not monitored and do not face any pecuniary penalty for low effort. Thus, their decision to complete the task is a measure of their willingness to exert effort when they are not facing an extrinsic incentive to do so. While we acknowledge that in many cases firms are able to motivate workers with monetary rewards or the threat of dismissal, firms in Ghana and in other developing countries are also often likely to be unable to set up efficient incentive schemes (e.g., because of the difficulty of monitoring effort).²² In the previous section, we used cross-country data from the World Management Survey to document this point: compared to the other countries in the data, firms in Ghana score low both on the use of incentives and appraisals, and on removing poor performers. More broadly, the cross-country comparison showed that the limited use of incentives is common among firms in developing countries. To shed additional light on this point, we leverage data from a recent firm survey that we have run in Addis Ababa, another fast growing capital of a Sub-Saharan African country.²³ There are three key results. First, we asked employers whether they experience situations in which monitoring a worker is too costly or difficult, and they simply have to trust their workers to do a good job when unmonitored (Figure A.4 in the Appendix). 88 percent of firms report that this is true either “sometimes” or “often”. Among very small firms with fewer than 3 workers, the proportion goes up to 91 percent. Second, consistent with this, we find that trustworthiness and honesty are the mostly highly ranked worker traits by firm managers among a set that includes practical skills, numeracy, literacy, and other soft skills such as communication and extroversion (Figure A.5 in the Appendix). This result is in

²²It is important to note that even when employers can monitor workers and simply lay off those who prove to be untrustworthy, this is costly for them for at least three reasons. First, under the Labour Act of 2003 (Act 651), a firm wishing to terminate a contract of employment in Ghana must typically give workers one month’s prior notice (reduced for short employment spells) and pay regular wages during the notice period. Second, firing untrustworthy workers may result in litigation if the worker challenges the employer’s decision. Specifically, in the event of a disputed dismissal, a worker can file a complaint with the Labour Commission. If on investigation, the Commission finds that the termination of employment was unfair, it may order the employer to reinstate the worker from the date of termination of employment (either in the same role or in another reasonably suitable role) on the same terms and conditions as before the termination; or order the employer to pay compensation to the worker. If the misconduct cannot be easily proven, as it is often the case, such litigations can be complex. Finally, in addition to being costly for the employer in terms of salary and potential legal costs, untrustworthy workers may do direct damage to the firm by stealing or breaking equipment, goods, intellectual property, and other assets.

²³This survey is described in detail in Appendix B.

line with the findings of [Bassi and Nansamba \(2021\)](#) for Uganda, which have inspired us to include this question in the survey. Third, we asked firms whether, when deciding whether to entrust a worker with a new task, they would prefer to have information about how the worker performed the same task (i) in a previous job when s/he was monitored and incentivised to perform, or (ii) in a previous job when s/he was not incentivised or monitored. 80 percent of the firms reported they would prefer to learn about the worker's performance when the worker was not incentivised (Figure A.6 in the Appendix). Overall, these results confirm our rationale for not including effort incentives in the experimental design: situations where effort cannot be monitored and trust in workers plays an important role are common in the workplace, and the information that is elicited from observing unincentivised performance is valued highly by firms.

Second, the elicitation of beliefs is not 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. Second, even if respondents understood the scoring rule, we were concerned that they could 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. For example, [Gächter and Renner \(2010\)](#) show that incentivised belief elicitation can bias the relationship between beliefs and behaviour in the context of a public good game.²⁴ 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](#)).

²⁴[Gächter and Renner \(2010\)](#) also show that belief-elicitation incentives reduce the variation of reported beliefs, but do not change the mean of the belief distribution. A larger variance in reported beliefs should not be a major problem for us, since our core results are based on a comparison of the mean of the belief distribution against the true empirical counterpart.

4 Sample and descriptive statistics

One of the key objectives of this study is to measure the potential misperceptions of workers' trustworthiness among entrepreneurs in Ghana. To pursue this objective, we follow an established literature that uses lab experiments in the field to measure parameters that are otherwise hard to observe (Harrison and List, 2004). Several papers document the potential of this methodology. For example, Karlan (2005) shows that an experimental measure of trustworthiness predicts credit repayment in Peru; Rustagi et al. (2010) document that communities that exhibit higher levels of conditional cooperation in Ethiopia have better outcomes when managing a commonly-owned natural resource; Finan and Schechter (2012) find that politicians target reciprocal individuals.

To allow for meaningful comparisons between entrepreneurs' expectations and workers' behaviour, we drew a random sample of business owners and workers from the working-age population of Accra, the capital city of Ghana, where the experiment took place. We considered an eligible entrepreneur 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 worker.²⁵

Table 1 describes the main characteristics of our samples of entrepreneurs and workers respectively. Further, Table A.1 in Appendix A shows that entrepreneurs' characteristics are largely balanced between treatment groups.

As expected, given our sampling strategy, entrepreneurs are substantially older than workers, 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. In terms of the businesses that they run, traders are the most common category (41% of the total), followed by manufacturers (including the construction sector) and service-providers (both at

²⁵The sample was drawn from the Ghana Urban household Panel Survey (GUHPS), a representative survey conducted by the Centre for the Study of African Economies in the cities of Accra, Kumasi, Takoradi, and Cape Coast. GUHPS respondents were drawn by stratified random sampling of urban households from the Population and Housing Census. The survey was designed to cover all household members of working age at the time of the interview. To attain the desired sample size, we coupled this strategy with additional sampling based on a simple random-walk rule in random areas of the city: starting from a randomly selected initial location, enumerators were asked to walk in a randomly chosen direction and sample one random adult from every n^{th} household. When they reached a junction, they randomly picked the direction to take.

Table 1: Summary Statistics

Variable	Mean	Std. Dev.
<i>Panel A: Entrepreneurs (N = 292)</i>		
Age	37.582	11.701
Male	0.366	0.483
Married	0.562	0.497
Educ: No Schooling	0.106	0.309
Educ: Primary	0.103	0.304
Educ: Lower secondary	0.312	0.464
Educ: Upper secondary	0.312	0.464
Educ: Tertiary	0.168	0.374
Employed	0.839	0.368
Sect: Trading	0.408	0.492
Sect: Manuf. & Construct.	0.171	0.378
Sect: Services	0.175	0.380
Sect: Unspecified	0.246	0.432
No. Employees	0.787	1.573
<i>Panel B: Workers (N = 281)</i>		
Age	27.954	10.089
Male	0.47	0.5
Married	0.278	0.449
Educ: No Schooling	0.064	0.245
Educ: Primary	0.096	0.295
Educ: Lower secondary	0.263	0.441
Educ: Upper secondary	0.363	0.482
Educ: Tertiary	0.214	0.411
Employed	0.505	0.501

Note: Summary statistics for the sample of entrepreneurs and workers. *Educ:* Tertiary education includes vocational training. *Employed* captures whether the subject currently has a job. The business sector is the one in which the subject last had a business. *No. Employees* is the number of workers the participant was employing at the time when s/he last employed any labour. This allows us to elicit comparable answers from both current and former entrepreneurs. The information on number of employees is available for 263 entrepreneurs since it was obtained from a follow-up survey where 29 respondents could not be re-interviewed.

17% of the total). Average business size is very small: 0.7 worker per business (we plot the full distribution of the number of workers hired in Figure A.7 in the Appendix.)

Workers, on the other hand, are balanced across genders, they are younger than entrepreneurs on average, and substantially less likely to be currently employed. They also have slightly higher educational attainments, most likely a reflection of belonging to a younger generation. If we compare our sample of workers to the overall labour force (using the official Ghana Labour Force Survey for the closest year available, 2015), we find no difference in the gender balance, but our subjects are slightly younger (the average age is 27 in our sample and 33 in the LFS) and less likely to be employed (the employment rate in our sample is 50% while it is 65% in the LFS). It should be stressed, however, that (i) entrepreneurs were informed of the age and employment rate in the sample of workers we selected, and (ii) that the younger age and lower employment rate make our sample closer to the population of *jobseekers*, which is the relevant group that entrepreneurs should consider in their hiring decisions.

5 Results

5.1 Are entrepreneurs' expectations aligned with the revealed trustworthiness of workers?

We begin by addressing the central question in our investigation: *Do entrepreneurs have correct expectations of workers' trustworthiness?* Our experimental design allows us to answer this question directly, by comparing entrepreneurs' elicited expectations of workers' trustworthiness (i.e., the proportion of workers expected to be trustworthy, $E_{jr}[P_r(e = 1)]$, where j denotes the entrepreneur and r denotes the round²⁶) with the revealed trustworthiness of workers (i.e., the proportion of workers who are actually trustworthy, captured by their rate of success in the task, $P_r(e = 1)$):

$$\underbrace{E_{jr}[P_r(e = 1)]}_{\text{Expected trustworthiness in round } r} \quad vs \quad \underbrace{P_r(e = 1)}_{\text{Actual trustworthiness in round } r}$$

²⁶Whenever this is not likely to cause confusion, we will shorten $E_{jr}[P_r(e = 1)]$ with $E_{jr}[P(e = 1)]$.

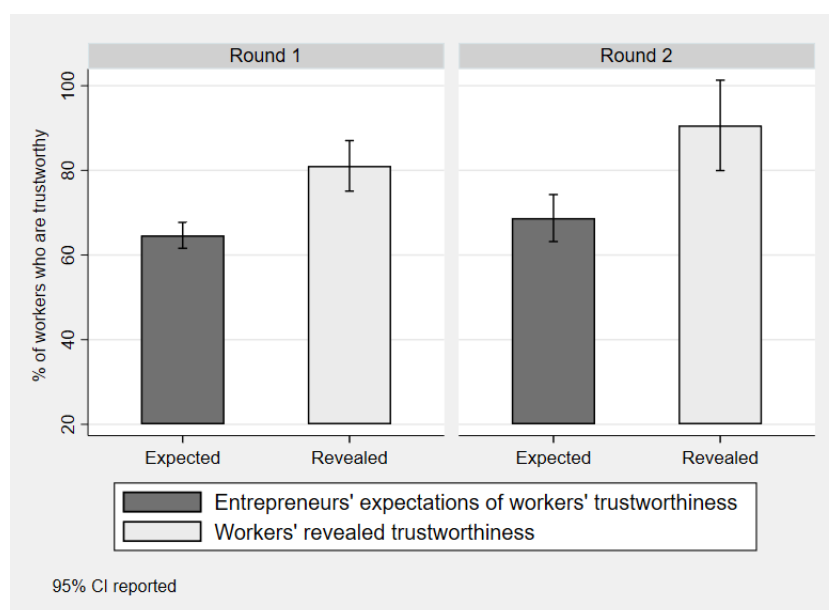
Result 1: Entrepreneurs underestimate workers' trustworthiness.

The results show that the proportion of workers who reveal to be trustworthy is significantly higher than it is expected by entrepreneurs on average (Figure 2). 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). Specifically, in round 1 entrepreneurs' expected trustworthiness is 64.6% on average (95% CI: 61.6, 67.7), while workers' revealed trustworthiness is 81.1% (95% CI: 75.1, 87.03). This is to say that, entrepreneurs expect 6.45 workers out of 10 to be trustworthy, while 8.1 reveal to be trustworthy (an expectation gap of 1.65, or approximately 20% relative to revealed trustworthiness). The gap is slightly larger in round 2.²⁷ Furthermore, plotting the cumulative distribution of entrepreneurs' expectations reveals that the majority of entrepreneurs underestimate the proportion of workers who will be trustworthy (Figure A.9).²⁸

²⁷In round 2, entrepreneurs expected trustworthiness is 68.7% on average (95% CI: 63.2, 74.3) and workers' revealed trustworthiness is 90.6% (95% CI: 79.95, 101.3)

²⁸In Figure A.8 in the Appendix, we plot instead trusting against expected trustworthiness. It is important to note that the level of trusting and expected trustworthiness captured in the experiment are not directly comparable to the trust measure reported in the World Values Survey (WVS) which we discussed in Section 2, for three reasons. First, the WVS records the answers of a representative sample of respondents, while we focus on a sample of entrepreneurs. Second, the WVS captures generic trust in other people, while our experiment measures trust in workers. Third, the questions are not directly comparable. The WVS records binary responses to the following question "Generally speaking, would you say that most people can be trusted or that you need to be very careful in dealing with people", while we ask entrepreneurs whether they would like to hire a worker in the experiment. An entrepreneur could say that "you need to be very careful in dealing with people" in response to the WVS question, while choosing to trust a worker in our experiment, e.g. because the stakes in the experiment are relatively low, or because in the experiment the entrepreneur faces a binary trust decision while in the real world a continuum of options may be available (e.g., an entrepreneur could hire a worker and set up a costly monitoring scheme to avoid being cheated).

Figure 2: Expected and revealed trustworthiness



Note: Percentage of workers who reveal to be trustworthy (i.e., exert effort in the task) as *expected* by the entrepreneurs and *revealed* in the experiment. We exclude T2 sessions and round 2 in T1 sessions (to exclude the effect of the treatments).

When we explore heterogeneity by entrepreneur characteristics, we find that more educated entrepreneurs tend to have higher and less biased expectations about worker trustworthiness, both when we condition on age and gender and when we do not condition on these variables (Table A.2 in the Appendix). This is an interesting finding, which suggests a role for education in counteracting potentially harmful biases. Further work and additional data will however be necessary to study the underlying mechanisms and establish causal links. More educated individuals, for example, may be better able to use the available information to form correct beliefs. Alternatively, there may be an underlying causal effect of cognitive ability (which is positively related to education) on beliefs.²⁹ Finally, entrepreneurs with more education may be running more sophisticated businesses, where the returns to worker trustworthiness are higher. They may thus have stronger incentives to acquire information about worker trustworthiness.³⁰

²⁹In related work, Bassi and Nansamba (2021) find that managers with higher cognitive ability (as measured by a Raven's test) respond more strongly to signals about worker skills.

³⁰Coibion et al. (2018), for example, document that managers' beliefs about inflation tend to be more accurate in businesses that have an incentive to collect better information about prices because they face more competition.

Result 2: Workers correctly predict entrepreneurs' propensity to trust.

By contrast, when we compare workers' elicited expectations of entrepreneurs' propensity to trust with the actual rate of trusting, we find that workers' beliefs are broadly correct (Figure A.10).³¹ Insofar as behaviour in the lab is driven by beliefs developed in the labour market, these two results point to an interesting asymmetry between entrepreneurs' and workers' learning. Workers appear to be aware of entrepreneurs' low expectations and low propensity to trust. On the other hand, the costs and risks of experimentation may be preventing entrepreneurs from forming correct beliefs about workers' trustworthiness.

Furthermore, when we analyse the relationship between workers' trustworthiness and their expectations of entrepreneurs' propensity to trust, we find an interesting correlation (Table A.4). Workers who expect entrepreneurs to be more trusting show higher levels of trustworthiness. This result is consistent with the predictions of models of guilt aversion: workers who expect entrepreneurs to be more trusting experience more guilt if they fail to complete the task. Thus, they are more likely to be trustworthy. Overall, this finding further reassures us that we are eliciting a behavioural trait that correlates in meaningful ways with agents' beliefs.

Result 3: Male entrepreneurs have lower expectations when the workers are predominantly women.

Next, we explore the heterogeneity of entrepreneurs' misperceptions by exploiting the exogenous changes in sample composition we induced in Treatment 2. Recall that at the beginning of each T2 session entrepreneurs were informed that the majority of the workers invited (80%) belonged to a specific category of interest: youth or women (as opposed to T0 and T1 sessions where the sample was balanced). Hence, by comparing first-round entrepreneurs' expectations in T2 to those in T0 and T1, we can identify the effect of workers' characteristics on entrepreneurs' beliefs.³² We do so by estimating the following model, clustering standard

³¹In both rounds, we cannot reject the hypothesis that expected and revealed trust are equal ($p = .58$ and $p = .36$, respectively).

³²By focusing on round 1, we are able to compare T2 with T0 and T1 sessions jointly (since round 1 was identical in the latter two), maximising the available sample size. The same patterns, however, emerge in round 2.

errors at the session level:

$$E_{j1}[P(e = 1)] = \alpha + \beta_W T2W_j + \beta_Y T2Y_j + \gamma X_j + u_j, \quad (1)$$

where: $E_{j1}[P(e = 1)]$ are entrepreneur round 1 expectations, $T2W_j$ is a dummy equal to 1 if j is in T2 with a majority of *female* workers; $T2Y_j$ is a dummy equal to 1 if j is in T2 with majority of *young* (<25) workers; and X_j is a vector of control variables.

When we run the estimation over the entire sample (Table 2, “All”), we find that *entrepreneurs have significantly lower expectations when the majority of workers in the session are women*. This result is driven by male entrepreneurs (Table 2, “Male Employers”), who also display higher expectations when the sample is formed predominantly of younger workers (overall, they still underestimate workers’ trustworthiness but the difference with actual trustworthiness is no longer statistically significant in T2Y). The expectations of female entrepreneurs, on the other hand, are not affected by the composition of the worker pool (Table 2, “Female Employers”).

Table 2: Entrepreneurs’ expectations for different pools of workers

(*Dep Var: Entrepreneur’s expectations of workers’ trustworthiness*)

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

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

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 A.11. It is also consistent with existing evidence from lab-experiments showing that women are more trustworthy (e.g., Buchan et al. (2008)), and hence helps us to corroborate the validity of our approach to measure trustworthiness. Younger workers, on the other hand, are no more (or less) trustworthy than older ones ($p = .34$), as shown in Figure A.12 in the Appendix and in line with previous studies (Sutter and Kocher, 2007). We present a full set of regressions on the correlates of worker trustworthiness in Appendix Table A.3.

5.2 Do expectations matter for trusting?

Having presented our results on entrepreneurs' misperceptions of workers' trustworthiness, we now test whether the expectations we elicited have predictive power on the (incentivised) choice to trust (i.e., hire) 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, clustering standard errors at the session-round level:

$$\text{Trust}_{jr} = \alpha + \beta E_{jr}[P(e = 1)] + \gamma X_j + u_{jr}, \quad (2)$$

where: Trust_{jr} is a dummy equal to 1 if entrepreneur j hires a worker in round r ; $E_{jr}[P(e = 1)]$ are entrepreneur j 's round r expectations of workers' trustworthiness; X_j is a vector of control variables.

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

Table A.8 in the Appendix shows that expected trustworthiness strongly and significantly predicts trusting in all the proposed specifications. A change in entrepreneur'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).

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 entrepreneurs' perceptions of workers' trustworthiness and a second one, which depends on entrepreneurs' willingness to take risks.

By showing that hiring decisions are strongly correlated with entrepreneurs' elicited expectations, the results in this section point to the conclusion that *underestimating workers' trustworthiness leads to sub-optimal hiring and hurts entrepreneurs' profits.*³³

5.3 Does access to information change entrepreneurs' 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 entrepreneurs' expectations? 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, entrepreneurs 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 T1A sessions, the signal carried no worker characteristics. In T1B sessions (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).³⁴ Two entrepreneurs 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 while

³³We 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.

³⁴The distribution of signals was calibrated on the results from a small (N=20) pilot we ran prior to the experiment, where about 60% of the workers completed the task. Due to the limited sample size, this pilot is unsuitable to draw conclusions about average trustworthiness, but that was not our objective. We only used the pilot to motivate our choice to provide positive and negative signals in equal proportions (a slight deviation from the 60% completion rate for the sake of balance).

maintaining sufficient statistical power. In those sessions, all the signals were positive. Since the signals were private knowledge, entrepreneurs could not possibly be aware of this deviation. Given the overall high levels of workers' trustworthiness, we were not worried this minor deception could lead to significant financial losses for entrepreneurs who may have decided to trust as a result of receiving a positive signal.

The treatment was designed to provide entrepreneurs with the same information they could obtain by hiring a worker (i.e. experimentation); and it is important to recall that at the time of receiving the signal the results of the first round had not been revealed.³⁵ This meant that all treated entrepreneurs 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 entrepreneurs 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 assess the impact of information signals on expectations, we first estimate the following model using data from T0 and T1A:

$$E_{j2}[P(e = 1)] = \alpha + \beta_p P_j + \beta_n N_j + \eta E_{j1}[P(e = 1)] + \gamma T1A_j + \delta X_j + u_j, \quad (3)$$

where: $E_{j1}[P(e = 1)]$ and $E_{j2}[P(e = 1)]$ are entrepreneur j 's expectations in rounds 1 and 2; P_j and N_j are dummies indicating whether entrepreneur j received a positive or a negative signal; $T1_j$ is a dummy for being in a T1A session; X_j is a vector of control variables.

We then introduce the data from T1B, and hence add dummies capturing (i) a positive signal on a female worker, (ii) a positive signal on a young (< 25) worker, and (iii) being in a T1B session.

Result 5: Negative signals have a strong downward effect on expectations, while positive signals have no effect.

The results are reported in Table 3. They show that positive signals have no impact on

³⁵As 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.

entrepreneurs' expectations, while negative signals significantly lower them. A negative signal decreases the expectations of an entrepreneur by about 8 percentage points (i.e., a reduction of almost 1 worker out of ten expected to be trustworthy) relative to an entrepreneur who receives no signal. Furthermore, the coefficient on the negative signal is significantly different from the coefficient on the positive signal ($p = 0.03$ in the first specification and $p = 0.01$ in the second specification).

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. 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. In Appendix Table A.6, we also show that we obtain a similar set of results if we drop entrepreneurs with expectations above 70%, who may have limited scope for upward adjustment.³⁶

Finally, in Table A.5 in the Appendix, we explore the role of entrepreneurs' characteristics. When we split the sample into younger and older entrepreneurs (below and above median age), we find that the aggregate effects we detected are driven by younger (i.e., less experienced) entrepreneurs. This is a highly plausible result. The negative signals we provide are likely to carry a higher weight for less experienced entrepreneurs who have more limited knowledge about the labor market and hence noisier beliefs. On the other hand, when we split the sample by the level of education of the entrepreneur, the results are generally noisier. Among more educated entrepreneurs, the coefficients for the positive and negative signals have the opposite

³⁶Similarly, we obtain coefficients that are consistent with asymmetric updating if we (i) remove entrepreneurs with expectations above 70% or below 30%; (ii) remove entrepreneurs with expectations above 60% or below 40% (Table A.7 in the Appendix). In the regressions reported in both Table A.6 and Table A.7, as the sample size drops, statistical precision naturally decreases, and we become unable to reject equality of the coefficients capturing the impacts of negative and positive signals (in Table A.6 the p-values are 0.23 and 0.13; in Table A.7, they are .40 and .20). However, while the coefficients on the negative signal dummy become if anything more negative, the coefficients on the positive signal dummy move in the wrong direction (and remain statistically insignificant). Thus, overall, these results provide evidence against the hypothesis that ceiling or floor effects drive the asymmetric updating results reported in Table 3.

sign and the same magnitude (we can reject equality, $p = .02$). However, neither coefficient is significantly different from zero. Thus, while these results suggest that asymmetric updating may be driven by the less educated entrepreneurs, we highlight that this conclusion is tentative given the large statistical uncertainty on these effects.

Overall, our evidence shows that entrepreneurs update their beliefs asymmetrically when they receive new information: they are sensitive to negative signals, but they are not affected by positive ones. 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 entrepreneurs' beliefs are hard to change, potentially leading to a vicious cycle of low worker performance and low entrepreneur expectations.

We acknowledge that the limited sample size at our disposal constitutes a potential limitation of the analysis. Thus, to further probe this finding, we run a new survey experiment on belief updating in the context of the survey of firm managers in Addis Ababa, which we introduced in Section 3 and in Appendix B. In this new experiment, we ask firm managers to consider a hypothetical scenario where they have to hire an unknown worker to complete a task that they cannot monitor. Managers are first asked to report the number of workers that they think would complete the task under such conditions out of 10. We then ask half of the managers to imagine that the worker they hired completed the task successfully, and the other half to imagine that the worker failed to complete the task. Finally, managers in the first (second) group are asked whether this experience would leave them more (less) confident or equally confident about the probability that the next random worker they may need to hire for the same task would complete it successfully. In this case, too, we find clear evidence that subjects update asymmetrically: 53 percent of the managers who are exposed to the positive scenario report that they would feel "much more confident", while 68 percent of the managers who are exposed to the negative scenario report that they would feel "much less confident" (a difference that is significant at the 5 percent level). We report these results in Table A.10 in the

Table 3: Expectations updating
(Dep Var: Entrepreneur's expectations)

	Basic	All
E[Trustw, R1]	.649 (.054)***	.682 (.050)***
Negative Signal	-.855 (.396)**	-.776 (.379)**
Positive Signal	-.045 (.455)	.070 (.432)
Positive Signal + Female worker		.175 (.561)
Positive Signal + Young worker (< 25)		-.503 (.627)
T1A: basic signals	.552 (.473)	.479 (.458)
T1B: signals with gender/age info		-.200 (.486)
Const.	2.523 (.619)***	2.422 (.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 entrepreneurs' 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. Entrepreneurs' expectations are measured on a scale from 1 to 10 (number of workers the entrepreneur 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 entrepreneur received a basic signal that carried no information about the gender or age of the worker. *Positive Signal + Female worker* and *Positive Signal + Young worker* are equal to 1 if the entrepreneur received a signal that was associated to a female or a young worker. *T1A: basic signals* is a dummy equal to 1 if the entrepreneur was in a session with basic signals (no info about gender or age of the worker). *T1B: signals with gender/age info* is equal to 1 if the entrepreneur was in a session with signals that also carried information about gender and age (these session effects can be identified because 2 entrepreneurs per session received no signal). Entrepreneur Char.s \equiv Age, Gender, Education, Marital status, Employment status. Session Char.s \equiv Time of session and Timing of expectation elicitation (before or after the hiring decision). *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session).

Appendix.³⁷ This replication helps to strengthen our confidence in the existence of asymmetric updating on worker trustworthiness.

Finally, it is worth remarking that our objective was to provide entrepreneurs with the same information they could obtain from hiring a single worker. This is because we were

³⁷The result is robust to excluding managers who expect more than 70% of workers to be trustworthy prior to being exposed to the positive or negative scenario, and it holds when we restrict the sample to small firms (< 5 employees) despite a drop in precision due to the smaller sample size (Table A.11). Similarly, the result is robust to dropping managers both at the bottom and at the top of the distribution of expected trustworthiness: though statistical precision decreases as sample size shrinks, the coefficient remains stable (Table A.12). We reproduce the distribution of managers' expectations prior to the treatment in Figure A.13.

interested in replicating the learning process that occurs through experimentation. Of course, one can easily think of more informative interventions that may have a stronger impact on expectations (e.g., revealing *average productivity* over the entire sample). If one were to design such an 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 entrepreneurs 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 of equation (2) are reported in Table 4. The first two columns show the results from a naïve OLS regression of trusting as a function of round 2 expectations and individual characteristics, as in Table A.10 above.³⁸ In the third specification, we instrument the expectations variable using the signals received by employers and we find that its effect grows considerably in magnitude, while remaining highly significant. Increasing entrepreneurs' 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 almost 16 percentage points.³⁹ There are at least three potential explanations for the fact that the IV estimate is larger than the OLS estimate. First, some omitted confounders may be affecting entrepreneurs' propensity to trust and their expectations of workers' trustworthiness in opposite directions. For example, under betrayal aversion (Bohnet et al., 2008), entrepreneurs who decide to trust may contemporaneously reduce their beliefs of workers' trustworthiness to minimise the possibility of feeling let down and betrayed. Second, measurement error may be generating attenuation bias. Third, under the LATE interpretation, the IV coefficient captures the treatment effect for the group of employers that revise their beliefs in response to a new signal about workers' trustworthiness. This group may make hiring decisions that are particularly responsive to beliefs, since it is likely to be composed of more sophisticated individuals and of individuals who have noisier priors and hence a stronger need for new

³⁸We restrict the analysis to round 2 because exogenous signals in Treatment 1 were only provided after the first round. We include round 1 expectations as a control.

³⁹Table A.9 in the Appendix shows the reduced form impacts of the signals on trusting.

information.

In terms of the magnitude of these effects, the IV estimates imply that closing the expectation gap of the average entrepreneur would increase the likelihood that the entrepreneur hires a worker in the experiment by approximately .26 ($.158 \times 1.65$), i.e., 26 percentage points.⁴⁰ Alternatively we can consider the more conservative expectation gap between the upper bound of the CI on entrepreneurs' beliefs and the lower bound of the CI on revealed trustworthiness reported in Section 5.1. Closing this gap would have a positive impact on the probability of hiring of .11 ($.158 \times 0.7$), i.e., 11 percentage points. This is an increase of 16% compared to the proportion of entrepreneurs who hire in the experiment (69.7% in round 1).⁴¹ Finally, the OLS estimates also imply economically meaningful impacts: closing the expectation gap for the average entrepreneur would increase hiring by about .05 (0.032×1.65), i.e., 5 percentage points, equivalent to approximately 7.5 percent compared to observed hiring in the experiment. These estimates suggest that a policy intervention aiming to improve entrepreneurs' expectations of workers' trustworthiness may have a strong impact on their willingness to hire.

One important caveat to this analysis relates to the strength of our instrument. The first-stage F-statistic of the IV model is 1.22. Weak instruments can cause two potential problems in our setting. The first problem is that weak instruments can bias the IV estimator in the direction of the OLS estimates. Since the OLS estimates are closer to zero than the IV estimates, this suggests that our IV estimates may be a lower bound on the true impact of expectations on hiring. This point is corroborated by two additional pieces of evidence. First, when we use a limited-information maximum likelihood estimator, as recommended by Angrist and Pischke (2008) since this estimator is median unbiased even when instruments are weak, the coefficient on expectations almost doubles (0.28 vs. 0.16). Second, when we drop the sessions with heterogeneous signals, where our instrument is mechanically weaker since these sessions only included the less impactful positive signal, the IV estimate grows to .21 (and the first-stage F-statistic grows to 1.64). The second problem is that under weak instruments, the unadjusted standard errors can be too small. We therefore conduct an additional statistical test that provides

⁴⁰We multiply the IV coefficient by the difference between average entrepreneurs' beliefs and workers revealed trustworthiness (equal to 1.65 workers out of 10, as reported in Section 5.1).

⁴¹The proportion is slightly lower in round 2 (64.4%).

weak-instrument-robust inference (i.e., the Anderson-Rubin Wald test). This test rejects the null that expected trustworthiness does not have an effect on hiring at the 5 percent level ($p = .03$), corroborating our conclusions. In sum, the result of this analysis suggests that weak instruments are unlikely to be a problem for our findings. However, as we cannot rule this out conclusively, we flag this as a potential caveat.

Table 4: The impact of raising expectations on trust

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

	OLS	OLS	IV
E[Trustw, R2]	.036 (.011)***	.032 (.012)***	.158 (.064)**
Risk Aversion		-.059 (.017)***	-.057 (.022)***
Const.	.434 (.086)***	.617 (.150)***	-.158 (.428)
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 entrepreneurs' decision to trust in round 2. *OLS1* and *OLS2* replicate the naïve regressions in Table A.8 using only round 2 observations. *IV* uses the signals received by employers in T1A and T1B (as well as dummies for being in a T1A and in a T1B session) to instrument expectations. Entrepreneurs' expectations (E[Trustw, R2]) are measured on a scale from 1 to 10 (number of workers the entrepreneur thinks will reveal to be trustworthy out of 10 in round 2). Entrepreneur Char.s \equiv Age, Gender, Education, Marital status, Employment status. Session Char.s \equiv Time of session and Timing of expectation elicitation (before or after the hiring decision). *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session).

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

In this final section, we explore the relationship between the results of our experiment and the daily-life hiring patterns of entrepreneurs who took part in the game. There are two reasons why we should expect a positive correlation between expected trustworthiness and daily-life hiring. First, more optimistic entrepreneurs should expect higher effort from their workers (or a lower cost of incentives) and hence higher returns from an additional hire. Thus, more optimistic entrepreneurs should have a higher propensity to hire in daily life. Second, hiring outside the experiment should help entrepreneurs correct their overly pessimistic beliefs on

worker trustworthiness. Thus, entrepreneurs who hire more in daily life should become more optimistic about worker trustworthiness.

Thanks to a brief survey carried out three months after the experiment, we were able to obtain information on past hiring (specifically, the number of employees the entrepreneur was employing at the time when s/he last employed any labour).⁴² The majority (70%) of the entrepreneurs do not hire labour, 18% hire between 1 and 2 workers, nearly 10% hire between 3 and 5 workers, and only less than 2% hire more than 5 employees (Figure A.7 in the Appendix).

Result 7: Entrepreneurs who hire more in their daily-life businesses have more positive expectations and are more likely to trust workers in the experiment.

Table 5 shows that the number of employees an entrepreneur hires in his/her businesses correlates positively with both entrepreneurs' expectations of employees' trustworthiness and with entrepreneurs' propensity to trust in the experimental setting.⁴³ In terms of magnitudes, we find that an increase in the number of workers by one standard deviation (1.5 workers), is associated with a 4 percent increase in expected trustworthiness and a 4.5 percentage points (7%) increase in trusting. A more sizeable increase of 5 workers is associated with a 13 percent increase in expected trustworthiness and a 13 percentage points (22%) increase in trusting.

One potential concern with the estimated relationship between hiring in daily life and trusting in the game 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

⁴²It 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.

⁴³We focus on round 1 to exclude the impacts of our information treatment in round 2. Table A.13 in the Appendix shows that the result is robust to using the log of the number of workers hired. On the other hand, we are unable to find a significant correlation when we use a dummy for whether the entrepreneur has one or more workers in daily life (Table A.14 in the Appendix). One potential reason for this is that some of the businesses that do not have any employees may face other constraints to firm growth (e.g. technological constraints), which weaken the relationship between expected trustworthiness and hiring. Consistent with this hypothesis, we re-run the analysis by first excluding entrepreneurs who cannot read and write and then by focusing on entrepreneurs with tertiary education. These are ways to identify entrepreneurs with stronger growth potential. For these samples, we are again able to document a positive and significant correlation between expectations and the dummy for having one or more workers in daily life. However, given the limited sample size, this is only suggestive evidence and further work will be necessary to investigate the underlying mechanisms.

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).⁴⁴

While the results in this section only constitute correlational evidence of the relationship between our experiment and actual hiring, they are consistent with the idea that biased beliefs may derive from lack of experience with hiring labour and may therefore be self-sustaining.

Table 5: The relationship between daily-life hiring and trust in the experiment

(Dep Var.s: *Entrepreneur's expectations & Whether the entrepreneur trusts/hires a worker*)

	E[Trustw]	Trust	E[Trustw]	Trust
Num. Workers Employed	0.168*** (0.0596)	0.0298** (0.0131)	0.125* (0.0671)	0.0271** (0.0126)
Const	6.453*** (0.129)	0.661*** (0.0436)	5.145*** (0.688)	0.562*** (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 daily-life businesses and: (i) entrepreneurs' expectations of workers' trustworthiness (E[Trustw]), and (ii) whether the entrepreneur hires in the game (P[Hire]). Entrepreneurs' expectations (E[Trustw]) are measured on a scale from 1 to 10 (number of workers the entrepreneur thinks will reveal to be trustworthy out of 10). Analysis limited to round 1 to exclude treatment effects in round 2. *Num. Workers Employed* indicates the number of workers employed by the entrepreneur at the time when s/he last employed workers (the information is missing for 29 entrepreneurs, hence the lower sample size compared to previous tables). Entrepreneur Char.s \equiv Age, Gender, Education, Marital status, Employment status; Session Char.s \equiv Time of session and Timing of expectation elicitation (before/after hiring decision). *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session).

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 entrepreneurs and workers 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

⁴⁴The controls included in the analysis, such as the respondent's educational background, are meant to further attenuate some of these concerns.

and decrease employment. From an empirical standpoint, little is known about the extent to which entrepreneurs' expectations and willingness to trust a worker are a realistic reflection of workers' behaviour. This study is an attempt to determine whether entrepreneurs' 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 a lab experiment consisting of a trust game with a real-effort task, played between real entrepreneurs and workers in urban Ghana. By eliciting entrepreneurs' 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 treatments to study the underlying learning mechanism and the extent to which expectations are biased against vulnerable groups.

We find that entrepreneurs significantly underestimate workers' trustworthiness, while workers' correctly predict entrepreneurs' propensity to (under)trust them. This points to the conclusion that the costs and risks of experimentation may lead entrepreneurs to an undesirable equilibrium of pessimistic expectations and sub-optimal trusting. When we correlate daily-life hiring decisions by the entrepreneurs with the elicited beliefs in the game, we find descriptive support for this hypothesis: entrepreneurs who have experimented less with hiring in the past have lower expectations of workers' 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 entrepreneurs' 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 entrepreneur's career may have sustained scarring effects on his/her propensity to hire later on. Our results suggest that successful information interventions targeting entrepreneurs and firm managers should take into account both biases in beliefs and in the updating process.

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*.
- Abebe, G., S. Caria, M. Fafchamps, P. Falco, S. Franklin, and S. Quinn (2021). Anonymity or distance? job search and labor market exclusion in a growing african city. *Review of Economic Studies*.
- Abel, M., R. Burger, and P. Piraino (2020). The value of reference letters: Experimental evidence from south africa. *American Economic Journal: Applied Economics* 12(3), 40–71.
- Akcigit, U., H. Alp, and M. Peters (2021). Lack of selection and limits to delegation: firm dynamics in developing countries. *American Economic Review* 111(1), 231–75.
- 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.
- Angrist, J. D. and J.-S. Pischke (2008). *Mostly Harmless Econometrics*. Princeton university press.

- 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. (2003, 07). Trust and expected trustworthiness: experimental evidence from zimbabwean villages. *Economic Journal* 113(489), 614–630.
- Bartling, B., E. Fehr, D. B. Huffman, and N. Netzer (2018, October). The Causal Effect of Trust. *IZA Discussion Papers* (11917).
- Bassi, V. and A. Nansamba (2021). Screening and signaling non-cognitive skills: Experimental evidence from uganda. *Working Paper*.
- 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., F. Greig, B. Herrmann, and R. Zeckhauser (2008). Betrayal aversion: Evidence from brazil, china, oman, switzerland, turkey, and the united states. *American Economic Review* 98(1), 294–310.
- 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.
- Buchan, N. R., R. T. Croson, and S. Solnick (2008). Trust and gender: An examination of behavior and beliefs in the investment game. *Journal of Economic Behavior & Organization* 68(3-4), 466–476.
- Bursztyl, 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.
- Carranza, E., R. Garlick, K. Orkin, and N. Rankin (2020). Job search and hiring with two-sided limited information about workseekers? skills. *Working Paper*.
- 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.
- Coibion, O., Y. Gorodnichenko, and S. Kumar (2018). How do firms form their expectations? new survey evidence. *American Economic Review* 108(9), 2671–2713.
- 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.

- Della Vigna, S. and D. Pope (2018). Predicting experimental results: Who knows what? *Journal of Political Economy* 126(6), 2410–2456.
- 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.
- Finan, F. and L. Schechter (2012). Vote-buying and reciprocity. *Econometrica* 80(2), 863–881.
- Friebel, G., M. Heinz, and N. Zubanov (2021). Middle managers, personnel turnover, and performance: A long-term field experiment in a retail chain. *Management Science*.
- Gächter, S. and E. Renner (2010, September). The effects of (incentivized) belief elicitation in public goods experiments. *Experimental Economics* 13(3), 364–377.
- Gächter, 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.

- Harrison, G. W. and J. A. List (2004). Field experiments. *Journal of Economic literature* 42(4), 1009–1055.
- 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.
- Johnson, N. D. and A. A. Mislin (2011). Trust games: A meta-analysis. *Journal of economic psychology* 32(5), 865–889.
- Karlan, D. S. (2005). Using experimental economics to measure social capital and predict financial decisions. *American Economic Review* 95(5), 1688–1699.
- 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 Silanes, 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.
- Rustagi, D., S. Engel, and M. Kosfeld (2010). Conditional cooperation and costly monitoring explain success in forest commons management. *science* 330(6006), 961–965.
- Sánchez de la Sierra, R. (2016). Ethnic State Capacity and Contract Enforceability. *PEDL Research Papers*.
- Sarsons, H. (2019). Interpreting signals in the labor market: evidence from medical referrals. *Working Paper*.
- Shapiro, C. and J. E. Stiglitz (1984, June). Equilibrium unemployment as a worker discipline device. *American Economic Review* 74(3), 433–44.
- Sutter, M. and M. G. Kocher (2007). Trust and trustworthiness across different age groups. *Games and Economic behavior* 59(2), 364–382.
- Teal, F. (2016). Firm size, employment and value added in african manufacturing firms: Why ghana needs its 1%. *Journal of African Economies*.
- 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.

ONLINE APPENDIX


A Additional figures and tables

Figure A.1: Use of performance incentives across the world

(a) Building a high-performance culture through incentives and appraisals

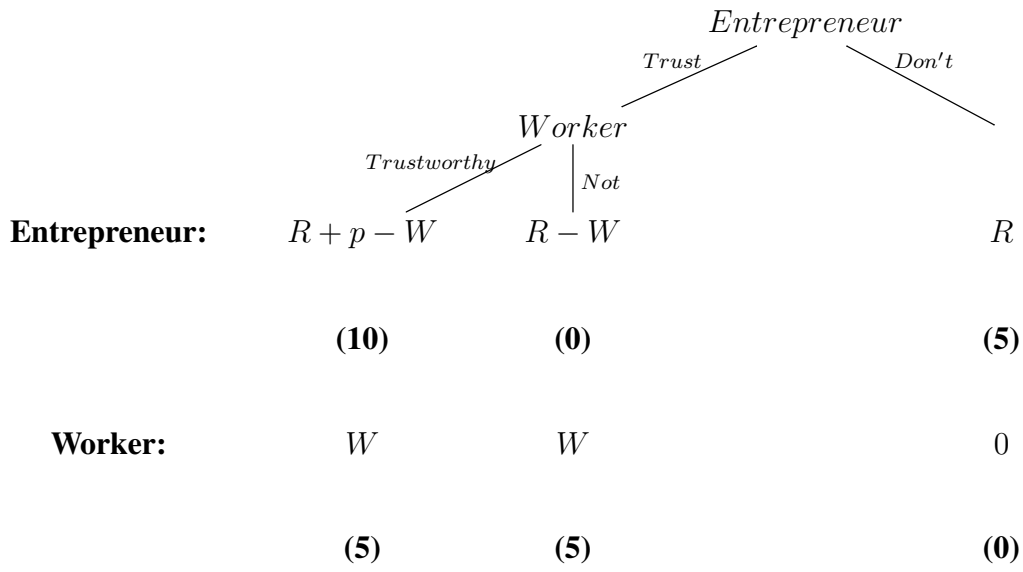


(b) Removing poor performers and making room for talent











Source: World Management Survey (WMS): www.worldmanagementsurvey.org. Note: Panel (a) shows an index (1-5) measuring the extent to which "firms use a systematic approach to identify good and bad performers, and reward them proportionately". A firm attains the lowest score in this dimension if "it rewards its workers equally irrespective of performance level". Panel (b) shows an index (1-5) measuring "how well the organisation is able to deal with underperformers". A firm attains the lowest score in this dimension if "poor performers are rarely removed from their position". WMS collects data on individual firms and we plot the average by country.

Figure A.2: The game tree



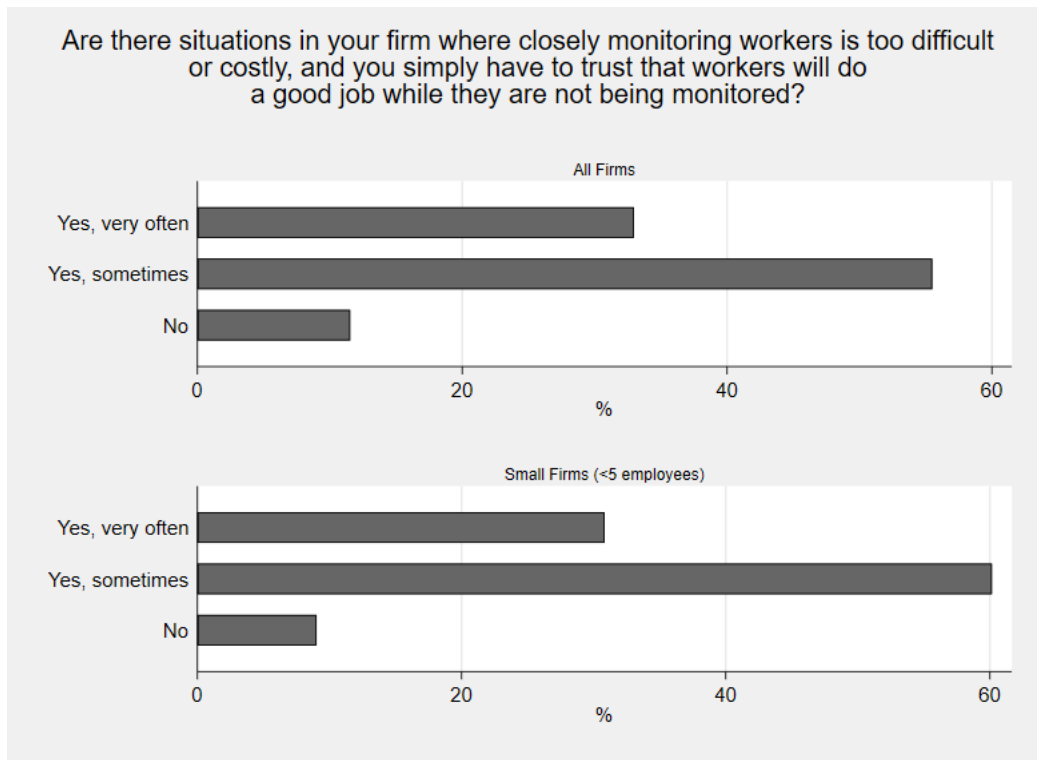
Note:: The numbers in parentheses indicate the monetary values of the payoffs (in Ghana Cedis).

Figure A.3: Risk-game lotteries

	 <p>50%</p>	 <p>50%</p>
<p>1</p>	 <p>LOSE 1 CEDI</p>	 <p>GAIN 5 CEDI</p>
<p>2</p>	 <p>LOSE 2 CEDI</p>	 <p>GAIN 5 CEDI</p>
<p>3</p>	 <p>LOSE 3 CEDI</p>	 <p>GAIN 5 CEDI</p>
<p>4</p>	 <p>LOSE 4 CEDI</p>	 <p>GAIN 5 CEDI</p>
<p>5</p>	 <p>LOSE 5 CEDI</p>	 <p>GAIN 5 CEDI</p>

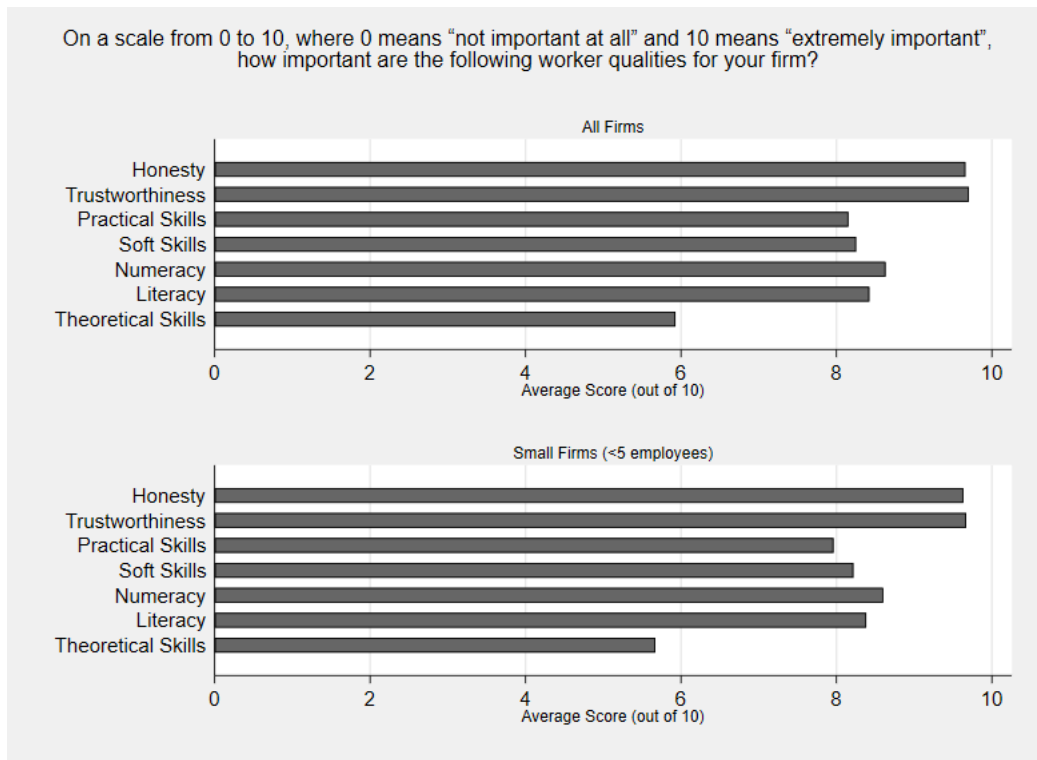
Note: Visual aid for the risk-preference elicitation game.

Figure A.4: Monitoring in the Workplace



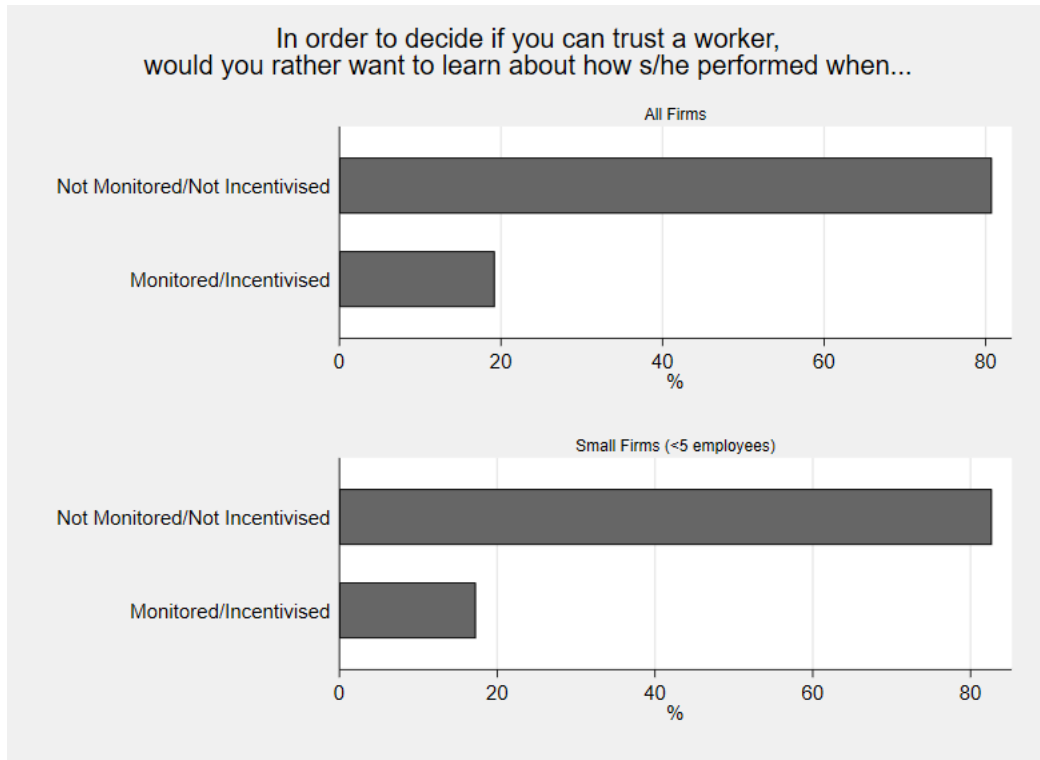
Note: The figure is based on a survey of 182 employers conducted in Addis Ababa to explore the replicability of our results and further test mechanisms (see Appendix B).

Figure A.5: Most important worker traits according to employers



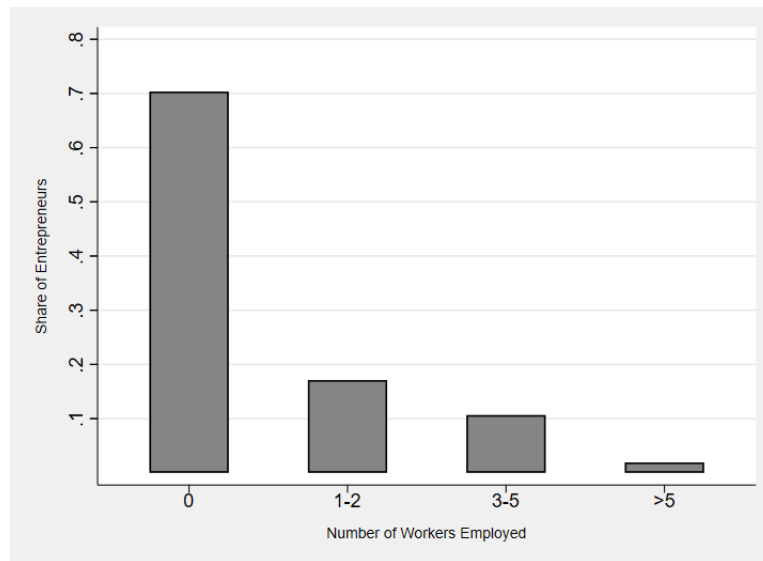
Note: The figure is based on a survey of 182 employers conducted in Addis Ababa to explore the replicability of our results and further test mechanisms (see Appendix B).

Figure A.6: Most valuable information to judge workers' trustworthiness



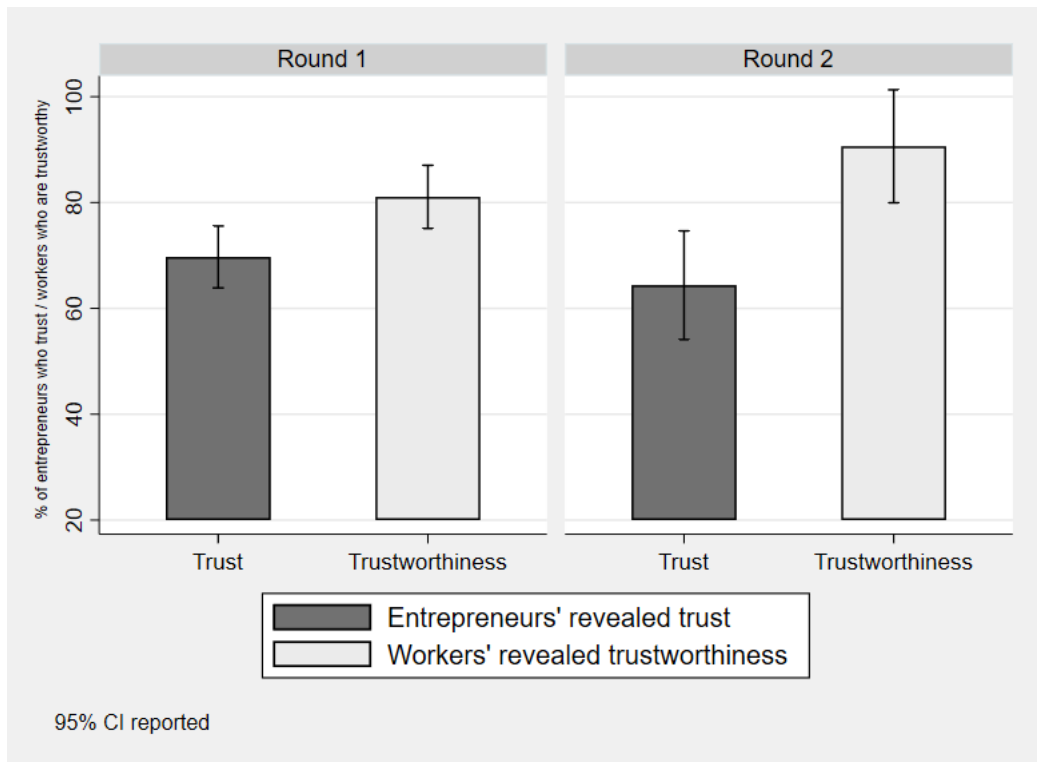
Note: The figure is based on a survey of 182 employers conducted in Addis Ababa to explore the replicability of our results and further test mechanisms (see Appendix B). The full question read as follows: *In order to decide whether you can trust a worker to carry out a very simple task that anybody can do but requires effort, which one of the following pieces of information would you find more convincing? 1. Knowing that when the worker was hired in the past for the same task he/she did the job well even if he/she was not monitored by a supervisor and could not be punished if he/she did not work; 2. Knowing that when the worker was hired in the past for the same task he/she did the job well, but was constantly monitored by a supervisor who could punish him/her for low effort.*

Figure A.7: Number of workers employed by entrepreneurs in their businesses




Note: Distribution of workers employed by participating entrepreneurs in their daily-life businesses. In order to elicit comparable answers from both current and former entrepreneurs, we ask for the number of workers the participant was employing at the time when s/he last employed any labour.

Figure A.8: Trust and Trustworthiness



Note: Percentage of entrepreneurs who trust workers (i.e., choose to hire) and percentage of workers who reveal to be trustworthy (i.e., exert effort in the task). Excluding T2 sessions and round 2 in T1 sessions (to exclude the effect of the treatments).

Figure A.9: Cumulative distribution of entrepreneurs' expectations




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

Figure A.10: Expected and revealed trust




Note: Percentage of entrepreneurs who trust a worker (i.e., choose to hire) as *expected* by the workers and *revealed* in the experiment. We exclude T2 sessions and round 2 in T1 sessions (to exclude the effect of the treatments).

Figure A.11: Revealed trustworthiness by gender



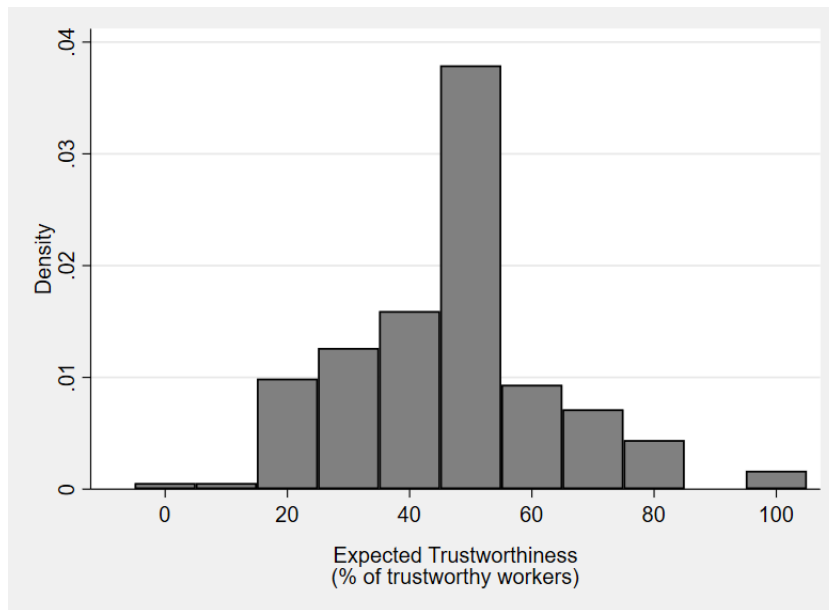
Note: Percentage of workers who reveal to be trustworthy (i.e., exert effort in the task) by gender. Round 1 only.

Figure A.12: Revealed trustworthiness by age



Note: Percentage of workers who reveal to be trustworthy (i.e., exert effort in the task) by age. Round 1 only.

Figure A.13: Distribution of Expected Trustworthiness among firms in Addis Ababa



Note: Distribution of entrepreneurs' expectations of workers trustworthiness (i.e., expected percentage of workers who would successfully complete a basic task if unmonitored). The figure is based on a survey of 182 employers conducted in Addis Ababa to explore the replicability of our results and further test mechanisms (see Appendix B).

Table A.1: Covariates balance across treatments

Variable	T1	T0,T2	Equality	T2	T0,T1	Equality
	Mean		<i>p</i> -value	Mean		<i>p</i> -value
Male	.35	.37	.69	.37	.36	.87
Age	38.98	36.89	.15	37	39.18	.16
Education (years)	9.68	9.68	.99	9.81	9.31	.37
Married	.63	.53	.1	.54	.63	.17
Employed	.84	.84	.9	.86	.79	.22
Sect: Trading	.4	.41	.89	.41	.41	.95
Sect: Manuf. & Constr.	.19	.16	.65	.16	.21	.35
Sect: Services	.13	.19	.2	.19	.13	.21
Sect: Unspecified	.28	.23	.38	.24	.26	.81
No. Employees	.93	.72	.31	.74	.91	.43
Obs.*	195	97	292	78	214	292

Note: Tests of balance in the characteristics of entrepreneurs between treatment arms. T1 is compared with T0 + T2, while T2 is compared with T0 + T1. Education is defined as number of years of education completed. Note that, if we use dummies for specific levels of education attainment, we find that completion of primary and lower secondary education are balanced, while there are some imbalances in higher secondary and tertiary. We thus include education attainment dummies in all specifications that include entrepreneur controls. *As reported in Table 1, *No. Employees* is only available for 263 out of 292 subjects.

Table A.2: Heterogeneity in entrepreneurs' expectations

(Dep Var: Expectations of workers' trustworthiness)

	(1)	(2)	(3)	(4)
Age	.016 (.012)			.018 (.012)
Female		-.272 (.342)		-.116 (.335)
Low Education			-.617 (.266)**	-.641 (.264)**
Const.	5.934 (.442)***	6.710 (.293)***	6.667 (.152)***	6.071 (.503)***
Obs.	292	292	292	292

Note: The table shows how entrepreneurs' expectations of workers' trustworthiness vary with a range of entrepreneur characteristics. The dependent variable is the number of workers (out of ten) who the entrepreneur thinks will complete the task if hired. *Low Education* is defined as having at most grade 6 (primary school). The analysis focuses on round 1 in order to exclude the influence of the treatment introduced in round 2. *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session).

Table A.3: Heterogeneity in workers' trustworthiness

(Dep Var: Worker's trustworthiness)

	(1)	(2)	(3)	(4)	(5)
Age (ys)	-.001 (.002)				-.002 (.003)
Female		.229 (.061)***			.215 (.060)***
No School			.061 (.094)		.022 (.093)
Primary			-.055 (.118)		-.032 (.131)
Lower Secondary			.054 (.061)		.036 (.059)
Tertiary			-.084 (.096)		-.070 (.108)
Unemployed				-.045 (.051)	-.044 (.052)
Const.	.834 (.067)***	.686 (.060)***	.805 (.049)***	.820 (.043)***	.759 (.086)***
Obs.	208	208	208	208	208

Note: The table shows how workers' trustworthiness (i.e., their likelihood of completing the task) varies with a range of individual characteristics. The dependent variable is a dummy that takes value 1 if the worker completes the task. The omitted category for education is *Upper Secondary* (the most frequent). The number of observations is lower than the total number of workers in the sample because a fraction of (randomly selected) workers does not get hired by entrepreneurs in the game and, therefore, does not perform the task. The analysis focuses on round 1 in order to exclude the influence of the treatments introduced in round 2. *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session).

Table A.4: Trustworthiness and workers' expectations of entrepreneurs' propensity to trust

(Dep Var: Worker's trustworthiness)

	OLS1	OLS2	Probit1	Probit2	Probit-MfX
E[Entrepreneurs' Trust]	.017 (.009)*	.019 (.009)**	.072 (.034)**	.088 (.037)**	.018
Const.	.736 (.080)***	.922 (.091)***	.565 (.288)**	1.461 (.398)***	
Worker Chars.	No	Yes	No	Yes	Yes
Session Chars.	No	Yes	No	Yes	Yes
Obs.	413	413	413	413	413

Note: The table shows the relationship between the probability that a worker is trustworthy and his/her expectations of entrepreneurs' propensity to trust workers. Workers' expectations are measured on a scale from 1 to 10 (number of entrepreneurs the worker thinks will reveal to be trusting, i.e., will trust a worker, out of 10). The estimation pools observations from round 1 and round 2. Worker Char.s \equiv Age, Gender, Education, Marital status, Employment status. Session Char.s \equiv Time of session and Timing of expectation elicitation (before or after the hiring decision). Probit-Mfx \equiv Marginal effects from Probit2. *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session + Round).

Table A.5: Heterogeneity in expectation updating

(Dep Var: Entrepreneur's expectations)

	All	Age < 37	Age ≥ 37	Edu < 12 ys	Edu ≥ 12 ys
E[Trustw, R1]	.682 (.050)***	.727 (.076)***	.645 (.071)***	.637 (.100)***	.725 (.072)***
Negative Signal	-.776 (.379)**	-1.397 (.493)***	.305 (.721)	-.382 (.745)	-.847 (.552)
Positive Signal	.070 (.432)	-.205 (.543)	.682 (.674)	-.488 (.563)	.963 (.695)
Positive + Female/Young Worker	.070 (.432)	.735 (.869)	-.996 (.623)	-.613 (.589)	.636 (.778)
T1A: basic signals	.479 (.458)	1.042 (.628)*	-.654 (.598)	.646 (.630)	.448 (.575)
T1B: signals with gender/age info	-.200 (.486)	-.527 (.866)	-.091 (.414)	.166 (.411)	-.440 (.758)
Const.	2.422 (.564)***	3.461 (1.083)***	4.002 (.926)***	2.541 (.664)***	1.666 (.713)**
Entrepreneur Char.s	Yes	Yes	Yes	Yes	Yes
Session Char.s	Yes	Yes	Yes	Yes	Yes
Obs.	292	145	147	152	140

Note: The table shows the effect of receiving positive and negative signals of workers' trustworthiness on entrepreneurs' expectations in round 2 (controlling for expectations in round 1). It replicates the analysis in Column 2 of Table 3 (reported in Column 1 here) and then splits the sample by age (above/below median) and education (above/below median). Median age in the employer sample is 37. The median level of education is 12 years (upper secondary education). The signals (Treatment 1) consisted of revealing whether a random worker had been trustworthy in previous rounds of the experiment. Entrepreneurs' expectations are measured on a scale from 1 to 10 (number of workers the entrepreneur thinks will reveal to be trustworthy out of 10). The dummies *Negative Signal* and *Positive Signal* are equal to 1 if the entrepreneur received a basic signal that carried no information about the gender or age of the worker. The dummy *Positive + Female/Young Worker* combines *Positive Signal + Female worker* and *Positive Signal + Young worker* (see Table 3 for additional details). This choice of aggregation is driven by the fact that, upon splitting the sample, the number of observations in each sub-treatment arm becomes very small. As in Table 3, *T1A: basic signals* is a dummy equal to 1 if the entrepreneur was in a session with basic signals (no info about gender or age of the worker). *T1B: signals with gender/age info* is equal to 1 if the entrepreneur was in a session with signals that also carried information about gender and age (these session effects can be identified because 2 entrepreneurs per session received no signal). Entrepreneur Char.s \equiv Age, Gender, Education, Marital status, Employment status. Session Char.s \equiv Time of session and Timing of expectation elicitation (before or after the hiring decision). *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session).

Table A.6: Expectation updating (Excl. subjects with high initial expectations)

(Dep Var: Entrepreneur's expectations)

	Basic	Basic Restr. Sample	All	All Restr. Sample
E[Trustw, R1]	.649 (.054)***	.574 (.130)***	.682 (.050)***	.698 (.099)***
Negative Signal	-.855 (.396)**	-.942 (.572)*	-.776 (.379)**	-.952 (.555)*
Positive Signal	-.045 (.455)	-.325 (.650)	.070 (.432)	-.215 (.638)
Positive Signal + Female worker			.175 (.561)	-.589 (.568)
Positive Signal + Young worker (< 25)			-.503 (.627)	-.791 (.581)
T1A: basic signals	.552 (.473)	.513 (.634)	.479 (.458)	.482 (.623)
T1B: signals with gender/age info			-.200 (.486)	.036 (.499)
Const.	2.523 (.619)***	2.933 (1.083)***	2.422 (.564)***	2.214 (.801)***
Entrepreneur Char.s	Yes	Yes	Yes	Yes
Session Char.s	Yes	Yes	Yes	Yes
Obs.	195	124	292	183

Note: The table reports the results of repeating the analysis in Table 3 after dropping subjects with high expectations (i.e., those with expectations of workers' trustworthiness above 70% in round 1). To facilitate comparisons, Column 1 and 3 report the estimates shown in Table 3. Column 2 and 4 report the results of estimating the same specifications on the restricted sample. *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%.

Table A.7: Expectation updating (Further sample restrictions)

(Dep Var: Entrepreneur's expectations)

	All	$30\% \geq E[Trustw, R1] \leq 70\%$	$40\% \geq E[Trustw, R1] \leq 60\%$
E[Trustw, R1]	.649 (.054)***	.575 (.221)***	.375 (.341)
Negative Signal	-.855 (.396)**	-1.294 (.718)*	-1.715 (.976)*
Positive Signal	-.045 (.455)	-.864 (.619)	-.884 (.950)
T1A: basic signals	.552 (.473)	.843 (.722)	.818 (.989)
Const.	2.523 (.619)***	3.968 (1.404)***	5.039 (1.883)***
Entrepreneur Char.s	Yes	Yes	Yes
Session Char.s	Yes	Yes	Yes
Obs.	195	112	81

Note: The table reports the results of repeating the analysis in Column 1 of Table 3 after restricting the analysis to different sub-samples. Column (1) reports the main results in Table 3; Column (2) removes subjects with expectations in round 1 below 30% and above 70%; Column (3) removes subjects with expectations in round 1 below 40% and above 60%. Confidence: *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%.

Table A.8: Trust and expected trustworthiness

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

	OLS1	OLS2	Probit1	Probit2	Probit-Mfx
E[Trustw]	.030 (.008)***	.027 (.008)***	.083 (.023)***	.080 (.025)***	.027
Risk Aversion		-.048 (.011)***		-.148 (.036)***	-.051
Round 2	-.016 (.052)	-.015 (.048)	-.043 (.150)	-.038 (.146)	-.013
Const.	.494 (.066)***	.593 (.107)***	-.042 (.177)	.267 (.317)	
Entrepreneur 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 entrepreneur trusts (hires) and his/her expectations of workers' trustworthiness. The estimation pools observations from round 1 and round 2. Entrepreneurs' expectations are measured on a scale from 1 to 10 (number of workers the entrepreneur thinks will reveal to be trustworthy out of 10). Entrepreneur Char.s \equiv Age, Gender, Education, Marital status, Employment status. Session Char.s \equiv Time of session and Timing of expectation elicitation (before or after the hiring decision). Probit-Mfx \equiv Marginal effects from Probit2. *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session + Round).

Table A.9: The direct effect of the signals of performance on trust

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

	Basic	All
Negative Signal	-.199 (.071)***	-.200 (.065)***
Positive Signal	-.075 (.083)	-.074 (.073)
Positive Signal + Female worker		-.116 (.105)
Pos Sig + Young worker (<25)		.039 (.121)
T1A: basic signals	.222 (.097)**	.232 (.094)**
T1B: signals with gender/age info		.149 (.115)
Const.	.697 (.204)***	.707 (.167)***
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 entrepreneurs' propensity to trust (hire) in round 2. The signals (Treatment 1) consisted of revealing whether a random worker had been trustworthy in previous rounds of the experiment. *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 entrepreneur received a basic signal that carried no information about the gender or age of the worker. *Positive Signal + Female worker* and *Positive Signal + Young worker* are equal to 1 if the entrepreneur 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 entrepreneur was in a session with signals that also carried information about gender and age (these session effects can be identified because 2 entrepreneurs per session received no signal). Entrepreneur Char.s \equiv Age, Gender, Education, Marital status, Employment status, Risk Aversion. Session Char.s \equiv Time of session and Timing of expectation elicitation (before or after the hiring decision). *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session).

Table A.10: Expectation updating following a negative vs positive experience with a hired worker

	All Firms		Small Firms (< 5 employees)	
	Update	Strongly Update	Update	Strongly Update
Negative Experience	.059 (.035)*	.147 (.072)**	.103 (.041)**	.121 (.085)
Const.	.909 (.025)***	.534 (.052)***	.883 (.030)***	.550 (.063)***
Obs.	182	182	133	133

Note: The figure is based on a survey of 182 employers conducted in Addis Ababa to explore the replicability of our results and further test mechanisms (see Appendix B). We present a random half of respondents with a positive scenario in which a hired worker performs a task satisfactorily without monitoring. We present the other half with a negative scenario in which a hired worker does not perform the same task satisfactorily (under the same conditions). We then ask the first group whether they would subsequently feel equally confident (no update), more confident (update), much more confident (strongly update) that a new worker hired for the same task under the same conditions would perform satisfactorily following that first experience. Conversely, we ask the second group whether they would subsequently feel equally confident (no update), less confident (update), or much less confident (strongly update) that a new worker hired for the same task would perform satisfactorily following that first experience. The table shows the results of analysing the likelihood that respondents who are faced with the negative scenario update (column 1 and 3) or strongly update (column 2 and 4) relative to those who are faced with the positive scenario. In column 1 and 3, the dependent variable is a dummy that takes value 1 if respondents update or strongly update. In column 2 and 4, the dependent variable is a dummy that takes value 1 if respondents strongly update. *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%.

Table A.11: Expectation updating following a negative vs positive experience with a hired worker (Dropping subjects with high initial expectations (>70%))

	All Firms		Small Firms (< 5 employees)	
	Update	Strongly Update	Update	Strongly Update
Negative Experience	.065 (.037)*	.195 (.074)***	.115 (.044)***	.171 (.088)*
Const.	.901 (.027)***	.494 (.054)***	.870 (.033)***	.500 (.066)***
Obs.	171	171	124	124

Note: The figure is based on a survey of 182 employers conducted in Addis Ababa to explore the replicability of our results and further test mechanisms (see Appendix B). We present a random half of respondents with a positive scenario in which a hired worker performs a task satisfactorily without monitoring. We present the other half with a negative scenario in which a hired worker does not perform the same task satisfactorily (under the same conditions). We then ask the first group whether they would subsequently feel equally confident (no update), more confident (update), much more confident (strongly update) that a new worker hired for the same task under the same conditions would perform satisfactorily following that first experience. Conversely, we ask the second group whether they would subsequently feel equally confident (no update), less confident (update), or much less confident (strongly update) that a new worker hired for the same task would perform satisfactorily following that first experience. The table shows the results of analysing the likelihood that respondents who are faced with the negative scenario update (column 1 and 3) or strongly update (column 2 and 4) relative to those who are faced with the positive scenario. In column 1 and 3, the dependent variable is a dummy that takes value 1 if respondents update or strongly update. In column 2 and 4, the dependent variable is a dummy that takes value 1 if respondents strongly update. We exclude firms that report baseline expectations of workers trustworthiness above 70% (in a previous question, prior to be presented with the scenarios), since they have limited room for updating their expectations upwards (we apply the same robustness check to our main results from the core experiment). *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%.

Table A.12: Expectations updating following a negative vs positive experience with a hired worker (Dropping subjects with high and low initial expectations)

(Dep Var: Strongly update expectations)

	$20\% \geq E[Trw.] \leq 80\%$	$30\% \geq E[Trw.] \leq 70\%$	$40\% \geq E[Trw.] \leq 60\%$
Negative Experience	.147 (.073)**	.124 (.080)	.127 (.092)
Const.	.523 (.053)***	.543 (.058)***	.529 (.069)***
Obs.	177	151	115

Note: The figure is based on a survey of 182 employers conducted in Addis Ababa to explore the replicability of our results and further test mechanisms (see Appendix B). We report results from the same specification presented in column 2 of table A.11, estimated over a restricted sample that excludes observations at the tails of the distribution of expected trustworthiness, where there may be limited room for updating. Column headings report the cutoffs on the distribution of expected trustworthiness used to define the restricted samples. *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%.

Table A.13: The relationship between daily-life hiring and trust in the experiment (LOGS)

(Dep Var.s: *Entrepreneur's expectations & Whether the entrepreneur trusts/hires a worker*)

	E[Trustw]	Trust	E[Trustw]	Trust
Log Num. Workers Employed	1.169*** (0.247)	0.127* (0.0664)	0.887** (0.364)	0.105 (0.0807)
Const.	5.664*** (0.334)	0.620*** (0.0938)	2.031** (0.778)	0.543 (0.321)
Entrepreneur Char.s	No	No	Yes	Yes
Session Char.s	No	No	Yes	Yes
Obs.	78	78	78	78

Note: The table shows the relationship between hiring in daily-life businesses and: (i) entrepreneurs' expectations of workers' trustworthiness (E[Trustw]), and (ii) whether the entrepreneur trusts/hires in the game (Trust). It replicates the analysis in Table 5, but it uses the natural logarithm of the number of workers employed as the regressor. *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session).

Table A.14: The relationship between daily-life hiring (binary indicator) and expectations in the experiment

(Dep Var.s: Entrepreneur's expectations)

	Full Sample	Excl. Illiterate Subjects	Tertiary Educ. Only
Dummy: Employs Workers	-0.0305 (0.268)	0.536* (0.315)	1.349* (0.674)
Const.	6.595*** (0.137)	6.408*** (0.188)	6.207*** (0.543)
Obs.	263	183	38

Note: The table shows the relationship between hiring in daily-life businesses and entrepreneurs' expectations of workers' trustworthiness. It replicates the analysis in Column 1 of Table 5, but it uses a dummy capturing whether the entrepreneur employs any labour as the main regressor. Illiterate subjects are those who cannot read and write based on the information collected in the short survey conducted at the end of the experiment. *Confidence:* *** \leftrightarrow 99%, ** \leftrightarrow 95%, * \leftrightarrow 90%. Robust standard errors in parentheses (cluster. by Session).

B The additional firm survey

As an extension to our original experiment, we conducted a survey of 182 firms in Addis Ababa, Ethiopia.⁴⁵ This survey has two objectives. First, it gives us an opportunity to study the external validity of our key result on asymmetric updating in a similar context – the urban labour market of another growing capital city in Sub-Saharan Africa. Second, it enables us to collect unique descriptive evidence that sheds light on a number of subtle issues related to the role played by expected trustworthiness in hiring.⁴⁶

The sampling frame of this new survey consisted of all firms registered at the Addis Ababa Trade and Industry Bureau. The business are located in Addis Ababa. The sample was stratified by legal status: either (i) sole proprietor or (ii) other (which included: share company, PLC, and other categories).

Given the second objective of the survey, which revolved around understanding how firms monitor their workers, how highly they value different worker traits, and what information they consider most relevant to judge a worker’s quality upon hiring, we restricted the sample to firms that had at least one paid worker (either in the current or in the previous calendar year).

Table B.1 below reports some key summary statistics for the sample. Firm size is generally quite small with a few large outliers (for this reason, we truncate the distribution at the 99th percentile). Mean firm size is 15, but almost 70 percent of the firms have fewer than 5 workers, and 30 percent of the firms have fewer than 2 workers. In 87 percent of cases, the respondent is both the owner and the main manager of the firm. Finally, the majority of firms are in the service sector (70 percent), while 8 percent of firms are in the manufacturing sector and 8 percent of firms are in the construction sector. This is in line with the sectoral breakdown of our sample for Ghana, where the sum of Trading and Services, which would jointly fall under Services in our Ethiopian sample, accounts for approximately 77 percent percent of the businesses that we are able to classify based on the available data, while manufacturing and construction account for approximately 23 percent.

Table B.1: Summary statistics for the firm sample in Ethiopia

	Mean	N
Firm size*	14.98	180
Firm size < 5	0.68	182
Firm size < 2	0.30	182
Respondent is owner and manager	0.87	182
Sector: Services	0.72	182
Sector: Manufacturing	0.08	182
Sector: Construction	0.08	182

Note: The table shows summary statistics for a sample of 182 firms surveyed in Addis Ababa to further test mechanisms and investigate the external validity of our results. *The mean of *Firm Size* is calculated after excluding two outliers.

⁴⁵The survey took place between January and February 2021.

⁴⁶The choice of the location for the new survey was driven by the first objective. Though we also considered conducting additional surveying in Ghana, this was made impossible by the onset of the COVID-19 pandemic.

C Experimental protocol

T0

Round 1

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

[Note: 2 and 3 randomly inverted in half of the sessions]

4. *Worker's* expectations of entrepreneurs' trust elicited.
5. *Worker* chooses effort (not revealed to entrepreneur 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. *Entrepreneur's* expectations (re)elicited.
2. *Entrepreneur* chooses hire/not hire a (new) random worker.

[Note: 1 and 2 randomly inverted in half of the sessions]

3. *Worker's* expectations (re)elicited.
4. *Worker* 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: Signals of performance

Round 1

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

[Note: 2 and 3 randomly inverted in half of the sessions]

4. *Worker*'s expectations of entrepreneurs' trust elicited.
5. *Worker* chooses effort (not revealed to entrepreneur 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. *Entrepreneur informed of trustworthiness of random worker from previous sessions.*
2. *Entrepreneur*'s expectations (re)elicited.
3. *Entrepreneur* chooses hire/not hire a (new) random worker.

[Note: 2 and 3 randomly inverted in half of the sessions]

4. *Worker*'s expectations (re)elicited.
5. *Worker* 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. *Entrepreneur informed that 80% of workers belong to a specific category (i.e., young or female).*

2. *Entrepreneur's expectations of workers' trustworthiness elicited.*

3. *Entrepreneur chooses hire/not hire a random and anonymous worker.*

[Note: 2 and 3 randomly inverted in half of the sessions]

4. *Worker's expectations of entrepreneurs' trust elicited.*

5. *Worker chooses effort (not revealed to entrepreneur 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. *Entrepreneur's expectations (re)elicited.*

2. *Entrepreneur chooses hire/not hire a (new) random worker.*

[Note: 1 and 2 randomly inverted in half of the sessions]

3. *Worker's expectations (re)elicited.*

4. *Worker 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.