THE
ECONOMIC
JOURNAL

The Economic Journal, **132** (*February*), 471–511 https://doi.org/10.1093/ej/ueab071
Advance Access Publication Date: 1 October 2021

SCREENING AND SIGNALLING NON-COGNITIVE SKILLS: EXPERIMENTAL EVIDENCE FROM UGANDA*

Vittorio Bassi and Aisha Nansamba

We study how employers and job seekers respond to credible information on skills that are difficult to observe, and how this affects matching in the labour market. We experimentally vary whether certificates on workers' non-cognitive skills are disclosed to both sides of the market during job interviews between young workers and small firms in Uganda. The certificates cause workers to increase their labour market expectations, while high-ability managers revise their assessments of the workers' skills upwards. The reaction in terms of beliefs leads to an increase in positive assortative matching and to higher earnings for workers, conditional on employment.

Labour productivity and wages remain far lower in developing countries (Hall and Jones, 1999; Caselli, 2005; Bloom *et al.*, 2010). This is particularly true in Africa, where the number of workers earning less than \$3.10 per day is increasing by almost four million a year (ILO, 2020). Understanding which factors contribute to keeping productivity and wages low in such contexts is thus of primary importance not only to foster business growth, but also to raise incomes and improve standards of living.

Theoretical models of the labour market highlight how the information available to both workers and firms plays a key role in determining the efficiency of the job matching process, and hence labour productivity and wages (Jovanovic, 1979; Chade and Eeckhout, 2017): difficulties in *screening* workers can prevent firms from selecting the right employees; at the same time, difficulties in *signalling* skills to employers can impact the ability of workers to match with the right jobs, or even their ex ante decision to acquire human capital (Spence, 1973).

Little is known on how information frictions on skills affect matching in urban labour markets in low-income countries, despite the potential importance of this question (McKenzie, 2017). On the one hand, information frictions may have less severe consequences in developing countries,

* Corresponding author: Vittorio Bassi, Department of Economics, University of Southern California, Los Angeles, CA, USA. Email: vbassi@usc.edu

This paper was received on 23 January 2020 and accepted on 6 September 2021. The Editor was Barbara Petrongolo.

The data and codes for this paper are available on the Journal repository. They were checked for their ability to reproduce the results presented in the paper. Given the highly demanding nature of the algorithms, the reproducibility checks were run on a simplified version of the code, which is also available on the Journal repository. The replication package for this paper is available at the following address: https://doi.org/10.5281/zenodo.5253572.

We are grateful to Imran Rasul, Richard Blundell and Costas Meghir for their advice, and to everyone at BRAC Uganda who made this project possible. We also thank Martin Abel, Ach Adhvaryu, Jie Bai, Dan Bennett, Nick Bloom, Michael Callen, Lorenzo Casaburi, Alan Crawford, Erika Deserranno, Francesco Drago, Christian Dustmann, Jan Eeckhout, Constanca Esteves-Sorenson, Andrea Ichino, Michele Fioretti, Rachel Heath, Sarojini Hirshleifer, Matt Kahn, Asim Khwaja, Leon Matagi, David McKenzie, Monica Morlacco, Paul Niehaus, Emily Nix, Anant Nyshadham, Paulina Oliva, Tommaso Porzio, Nick Ryan, Simone Schaner, Uta Schönberg, John Strauss, Miri Stryjan, Anna Vitali, Alessandra Voena, Jeff Weaver and numerous seminar and conference participants for valuable comments. This document is an output from a project funded by the UK Department for International Development (DFID) and the Institute for the Study of Labor (IZA) for the benefit of developing countries. The views expressed are not necessarily those of DFID or IZA. AEA RCT Registry number: AEARCTR-0001005. A previous draft was circulated under the title 'Information Frictions in the Labor Market: Evidence from a Field Experiment in Uganda'. Vittorio Bassi acknowledges funding from IFS and ESRC. All errors are our own.

where the production process is likely simpler and so the scope for inefficiencies related to mismatch may be lower. On the other hand, information problems could be exacerbated in more informal labour markets, due to weak institutions and limited rule of law (Bloom *et al.*, 2010; 2012; 2013); for instance, if managers cannot prosecute workers in the case of misconduct, this makes screening on reliability or trustworthiness at recruitment even more important.

This paper studies how lack of information on the skills of workers at recruitment affects matching in a developing country. To do so, we design a field experiment in the Ugandan labour market that has two key components: (*i*) a *matching* component, whereby firms and job seekers are matched together for real job interviews, and (*ii*) a *signalling* component, by introducing experimental variation in whether a credible signal on skills that are difficult to observe is shown to both sides of the market during the recruitment process, through the provision of certificates.

Our main contribution is to show how the credible signal affects the allocation of labour. On the firm side, we identify whether managers revise their beliefs on the skills of workers; on the worker side, we study whether the certificates impact the workers' perceptions about what they can achieve in the labour market, such as their expected earnings. We then show how the updating of beliefs of workers and managers translates into reduced-form impacts on the allocation of labour, overall employment and worker earnings in the two years post intervention.

Our sample includes young workers fresh out of vocational training and looking for jobs, as well as small and medium enterprises (SMEs) looking for workers. Young workers can be particularly affected by difficulties in signalling their skills, given their lack of work experience. Similarly, SMEs do not have access to sophisticated screening technologies, and so might be less able to screen workers, as compared to larger firms.

We focus the information revelation on non-cognitive or 'soft' skills, which have been shown to have high labour market returns in both high- and low-income countries, but which are hard to observe by nature.[1] In our context, managers report difficulties in observing the soft skills of job applicants as among their most important concerns. We focus on five specific soft skills that our manager survey reveals to be important but are difficult to observe: communication skills, willingness to help others, trustworthiness, creativity and attendance. We assess our sample of workers on these soft skills, using a combination of teacher surveys, incentivised trust games and psychometric scales. We then schedule and observe over 1,200 real job interviews between our sample of SMEs and job seekers. In a randomly selected half of these interviews, certificates on the workers' soft skills are revealed to both the worker and the firm, while in the other half, neither the worker nor the firm get to observe the certificate. We collect information on the result of each job interview, and track the sample of firms and workers for two years.

Our first set of results relate to whether managers and workers respond to the certificates by updating their beliefs. We find that the information strengthens the correlation between workers' skills and managers' assessment of their skills. This effect is driven by workers with relatively higher skills: while managers revise upwards their beliefs for workers in the upper part of the skill distribution, we find little evidence of negative updating for workers further down the distribution. We explain this by documenting that: (*i*) our sample of workers is *positively* selected on soft skills, relative both to the eligible population of vocational training graduates and to typical young employees in SMEs, which is possible since participation in the experiment was voluntary and we worked with particularly reputable training institutions; and (*ii*) managers have

[1] On the importance of non-cognitive skills in the United States see Bowles *et al.* (2001), Heckman *et al.* (2006) and Deming (2017). Adhvaryu *et al.* (2018) showed that non-cognitive skills have high returns in India.

low expectations about workers' soft skills at baseline. These two facts together explain why the certificates create mostly *positive news* for managers. Further, exploiting our two-sided design, we study heterogeneity by manager characteristics, and document that the revision of beliefs is driven by managers of higher ability, who manage more profitable firms and value soft skills more in production.

We show that workers also react to the certificates by revising upwards their labour market expectations: in the two years post intervention, workers with a certificate report 7% higher expected earnings, 5% higher expected employment probability, a higher intention to bargain for wages and a larger size of their ideal employer. The certificates also lead workers to pull away from poorly paid casual work and to increase their investment in training, which points to the presence of complementarities between the certification and human capital accumulation. We interpret these findings as evidence that the certificates raise the perceived outside option of workers, who try to transition to better and higher paying jobs. We further show that the updating of workers is driven by a reduction in their perceived difficulty to signal skills in the market, rather than by workers learning about their skills or about the returns to skills. Again, these overall positive treatment effects for workers are in line with the positive selection into the experiment and with the low priors of managers about their skills.

Informed by these reduced-form findings on beliefs, we then develop a partial equilibrium search model with heterogeneous workers and firms, in order to: (*i*) guide the interpretation of the impacts on labour market outcomes and (*ii*) study how workers at the low end of the skill distribution, who selected out of the experiment, would be affected by the information revelation, which is important for cost-benefit analysis considerations. In the model, soft skills have higher returns when matched to higher ability managers, and so the efficient allocation exhibits positive assortative matching between workers and firms. However, search frictions and lack of information on workers' skills result in mismatch and loss of output. We model the intervention as an increase in the precision of the signal on workers' skills during job interviews, and we show that the certificates improve the allocation of labour. Workers with higher skills earn higher wages after the intervention, as the certificates allow them to send a precise signal about their type to those firms who are willing to pay them more. Conversely, the model highlights that workers with lower skills can be negatively affected by the certificates, as these reduce their chance of employment at higher ability firms.

In line with the model, we find that the certificates cause an increase in sorting, by raising the probability of employment between our experimental sample of workers—who are positively selected on soft skills—and higher ability managers. In addition, the certificates increase sorting between workers with a high value of specific soft skills and employers where those skills might be particularly needed. For instance, workers with high communication skills are more likely to be hired by firms with more employees and where interactions with customers are more frequent. While the certificates lead to a change in the allocation of labour, they do not result in an overall increase in the probability of employment. This is consistent with the model, which shows that total employment does not necessarily increase, as workers reallocate across different types of firms. However, in line with the allocation of labour being more efficient, workers with a certificate earn 11% more in the two years post intervention, conditional on being employed, and this effect is larger at the top of the skill distribution.

The estimated earnings benefits from the certificates outweigh the program costs. In line with the average experimental worker benefiting from the certificates, we show that workers in the control group would be willing to pay over 40% of their monthly earnings to obtain one. Interestingly, their willingness to pay is very close to the cost of the certificates, and so we discuss potential reasons why similar certificates are not already provided by the market.

Finally, we address the potential for government intervention in this area, and note that in a large-scale version of this program in which workers are not given the option to select out, some workers at the low end of the skill distribution might be negatively affected by the certificates, as indicated by the evidence on positive selection on soft skills into the experiment and by our model. This highlights how certification policies, while raising efficiency through the improved allocation of labour, can also have important distributional effects among workers. However, we discuss how a large-scale certification intervention might also provide incentives to invest in skills, which could weaken this type of distributional concern.

Related literature: Our paper extends a growing experimental literature on how labour market frictions in developing countries affect workers and firms. On the worker side, recent studies evaluate how helping job seekers signal their skills impacts job search, employment and earnings (Abel *et al.*, 2020; Carranza *et al.*, 2020; Abebe *et al.*, 2021). On the firm side, this literature studies whether screening and training costs at recruitment impede hiring and constrain firm growth (De Mel *et al.*, 2019; Hardy and McCasland, 2020).

Two papers closely related to ours are Carranza *et al.* (2020) and Abebe *et al.* (2021). Abebe *et al.* (2021) evaluated a job application workshop targeting a representative sample of unemployed youths in Addis Ababa. The workshop included an orientation on conducting effective job applications, as well as the provision of a certificate on cognitive skills, mathematical ability, language skills and performance on a generic work sample. As the certificates were provided only to job seekers, the study aimed to identify the gains for workers from a having a more precise signal that they can use during job search. They documented that the intervention leads to short-term improvements in formal employment, and to a significant increase in earnings after four years. Carranza *et al.* (2020) conducted a series of field experiments with economically disadvantaged young work seekers and firms in Johannesburg. Their main intervention provided workers with a certificate on cognitive ability, numeracy, English proficiency, grit, focus and planning. They also conducted a separate audit study with firms, where a skills certificate was attached to CVs in formal job applications. By focusing on workers and firms, the study aimed to identify the strength of information frictions on both sides of the market. They found that the certificates improve workers' job search direction, employment and earnings three–four months after treatment, and that they increase callbacks by firms.

We make the following contributions relative to these papers. Since our unique design includes both a matching *and* a signalling component, we are the first to identify the role of information frictions on skills, conditional on a real job interview between a worker and a firm taking place. That is, we can separately identify the impacts of certification arising from reduced uncertainty on the skills of workers at recruitment, from those resulting from changes in search behaviour of workers. This distinction is important to fully characterise how information frictions in the labour market operate.[2] Second, our two-sided experimental design and data collection allow us to study heterogeneity in responses on both sides of the market, and to directly estimate impacts on sorting between workers and firms. This is crucial to identify the channels through which

[2] While the firm-side audit study in Carranza *et al.* (2020) holds fixed job search behaviour, it does not speak directly to the impacts of revealing information during job interviews, where workers and firms interact and potentially generate information. Informal walk-ins leading to job interviews are a primary recruitment channel in SMEs in our context, while formal applications through mailing of CVs are rare.

information on skills can improve productivity, and to determine the efficiency of intervention. Third, while both Carranza *et al.* (2020) and Abebe *et al.* (2021) disclosed information on a mix of hard and soft skills, we focus on soft skills specifically. We show that firms in our context are already able to screen on hard/practical skills, which then allows us to highlight soft skills as the key source of information frictions. Finally, while both Carranza *et al.* (2020) and Abebe *et al.* (2021) focused on unemployed job seekers, our target sample are semi-skilled labour market entrants that have not experienced long unemployment spells and are not disadvantaged. Such vocational training graduates constitute an important source of employment for SMEs, and so our results speak directly to the role of information frictions in impeding matching between SMEs and the pool of applicants they typically face.

We also contribute to a classic literature on employer learning and information frictions about skills (Farber and Gibbons, 1996; Altonji and Pierret, 2001; Schönberg, 2007; Kahn and Lange, 2014; Pallais, 2014). A closely related paper is Pallais (2014), who showed that disclosing more information about workers' past performance increases total hiring and welfare in an online labour market. We highlight the role of mismatch between workers and jobs as one important channel through which information frictions on skills can reduce output and earnings.

Finally, we add to an established literature on the role of job referrals in hiring (Beaman and Magruder, 2012; Burks *et al.*, 2015; Dustmann *et al.*, 2015; Brown *et al.*, 2016; Pallais and Sands, 2016; Heath, 2018). While referrals can potentially mitigate information frictions and improve match quality, their effectiveness can be lower in developing countries, where social and economic networks are often overlapping, so that employees might face social incentives to refer network members who are not the best match for the job (Beaman and Magruder, 2012). In fact, while referrals are common in our context, we show that they do not fully solve the hiring problem, which then justifies the focus on certificates as an alternative policy tool to alleviate information frictions at recruitment.[3]

Structure of the paper: Section 1 presents the sample. Section 2 describes the experiment. Section 3 shows the impacts on beliefs. Section 4 introduces the model and then shows treatment effects on labour market outcomes. Section 5 discusses policy implications and concludes. Additional details are given in the Online Appendix, and other results and analysis not intended for publication can be found on the authors' websites in a document called Supplemental Material.

1. Setting, Sample Selection and Descriptives

The project was implemented in partnership with a large and reputable NGO, BRAC Uganda. This section describes the sample of both workers and firms, and presents descriptive evidence.

1.1. *Firm Census and Selection into the Experimental Sample*

As shown in Figure 1, we began the study by identifying the sample of firms and workers for the intervention. Firms were identified through a census of SMEs conducted in urban areas of Uganda, covering all four regions of the country. To be in the census, firms had to: (*i*) operate

[3] Our study is also related to a growing literature about how the lack of information on job opportunities and on the consequences of unemployment impedes effective job search (Altmann *et al.*, 2018; Belot *et al.*, 2019). More broadly, our study contributes to the empirical literature on the micro-foundations of the aggregate matching function (Petrongolo and Pissarides, 2001).

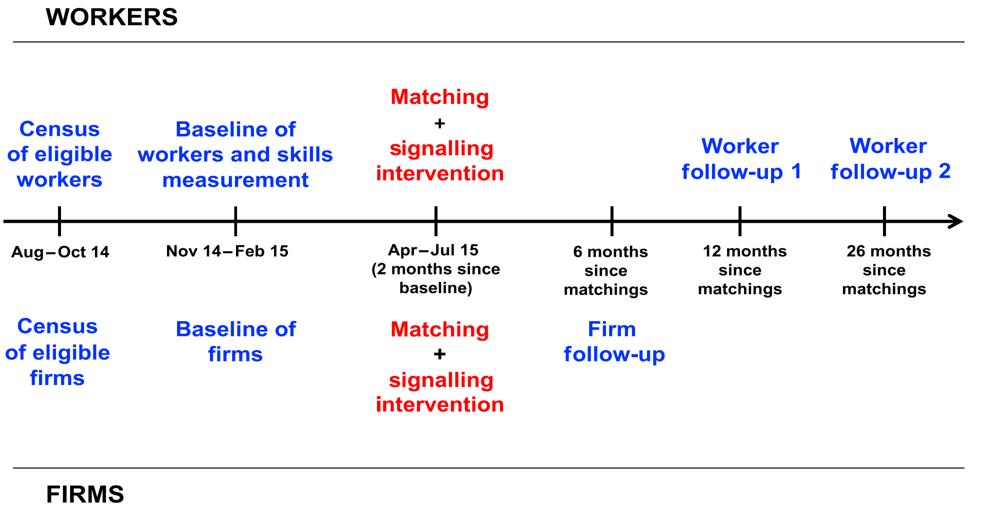

Fig. 1. *Timeline.*

in either carpentry, catering, hairdressing, motor mechanics, tailoring, welding; and (*ii*) employ at least two workers in addition to a firm owner. We identified 1,086 eligible SMEs through the census.[4] Table 1 reports summary statistics from the census. The median firm employs four workers and has been operating for five years, so the typical owner has experience recruiting and managing workers.[5]

We focus on SMEs in these sectors for two reasons. First, the majority of workers in Uganda and in developing countries more broadly are employed in small firms (Hsieh and Olken, 2014). SMEs in our six sectors are an important source of employment for young workers in Uganda, as shown by the fact that vocational training institutes (VTIs) typically offer courses in these sectors. Second, SMEs might be particularly affected by information frictions at recruitment given their limited access to screening technologies such as job platforms.

At the end of each interview in the census, firm owners were asked their interest in participating in the BRAC Job Placement Program: they were told that BRAC would facilitate job interviews with recent graduates from VTIs looking for employment in their sector and region.[6] All firms interested in participating in the program were administered a baseline survey, and the 422 firms that confirmed their interest and completed the survey form our experimental sample. Importantly, firm owners were only informed about the matching component of the intervention and were not told that certificates on soft skills would be disclosed during the job interviews. This limits the

[4] This is a separate and non-overlapping census to that discussed in Alfonsi *et al.* (2020). Our census was conducted in 17 urban areas in 11 of the 121 districts of Uganda. In each urban area, the census took place within a 4 km radius from the local BRAC branch. The census covered about 1% of the total area of the eleven districts we worked in. The 2010 Census of Business Establishments (UBOS, 2011) further reveals that in 2010 there were 23,366 firms operating in the same sectors and districts targeted by our census, so that we covered less than 5% of the firms. This limits concerns that our intervention could generate general equilibrium effects.

[5] Given the small average firm size, in the great majority of cases the firm owner is also the manager. Therefore, we use the terms 'firm owner' and 'manager' interchangeably.

[6] BRAC is one of the largest NGOs in Uganda, and is well known across the country for its programs targeting youths and firms. Therefore, concerns related to its credibility are not first order.

Table 1. *Firm Descriptives from Initial Census.*

	Mean (1)	SD (2)	Median (3)
Panel A: owner and firm characteristics			
Owner is female	0.397		
Number of employees	5.88	7.14	4
Business is registered	0.938		
Age of business (years)	7.09	5.90	5
Panel B: sector			
Carpentry	0.138		
Catering	0.157		
Hairdressing	0.302		
Motor mechanics	0.122		
Tailoring	0.082		
Welding	0.199		
Panel C: region			
Kampala	0.425		
North	0.123		
East	0.270		
West	0.181		

Notes: The table uses data from the initial census of 1,086 firms conducted for the job placement intervention. The census was conducted in seventeen urban areas of Uganda, and targeted all firms employing at least two employees and operating in six sectors: carpentry, catering, hairdressing, motor mechanics, tailoring and welding.

possibility that expected gains from the signalling intervention are a driver of selection into the sample.[7]

1.2. *Worker Census and Selection into the Experimental Sample*

1.2.1. *Worker census*

We defined as eligible for the intervention all trainees currently enrolled at fifteen large and reputable partner VTIs in one of the six business sectors covered by the project, and expected to graduate in time for the placement intervention.[8] We conducted an initial survey of all the 1,011 eligible trainees. The survey was administered before any information was given to trainees about the BRAC Job Placement Program. Table 2 shows that the median eligible trainee is 20 years old, has completed 11 years of education before enrolling at the VTI and is undertaking a two-year course. Training focuses on practical skills, which are also certified. On the other hand, VTIs do not provide formal training nor certificates on soft skills. Over 60% of workers plan to look for *wage* employment—as opposed to self-employment—as their first job, and the ideal firm size is

[7] The Supplemental Material reports an analysis of selection into the experiment on the firm side and shows that: (*i*) of the observable firm and manager characteristics collected in the census, which included the sector and region of operation, owner's gender, number of employees, whether the business is registered and the business's age, only the sector and region are significant predictors of selection; (*ii*) we do not find strong evidence that the sample of firms that self-selected into the experiment differs in terms of manager's skills from the representative sample of Ugandan SMEs in Bassi *et al.* (2021). As managers were informed about the matching component, but not about the certification component, self-selection into the experiment can indicate unmet demand for labour, and so these results are consistent with the tightness of the labour market varying across sectors and regions, but not across other firm or manager characteristics.

[8] More details on the selection and summary statistics of the fifteen VTIs are discussed in the Supplemental Material. In 2014–17 there were 204 VTIs accredited with the Directorate of Industrial Training and operating in the same districts as our partner VTIs. This further limits concerns that our intervention could generate general equilibrium effects.

Table 2. *Worker Descriptives from Initial Census.*

	Mean (1)	SD (2)	Median (3)
Panel A: worker characteristics			
Age (years)	20.2	2.50	20
Female	0.551		
Completed prior education (years)	10.3	2.05	11
Course duration (years)	1.41	0.934	2
Ever employed	0.260		
Has a job waiting at the end of training	0.085		
Plans to look for wage employment	0.629		
Ideal firm size is twenty employees or less	0.605		
Panel B: sector of training			
Carpentry	0.072		
Catering	0.129		
Hairdressing	0.266		
Motor mechanics	0.292		
Tailoring	0.179		
Welding	0.062		

Notes: The table uses data from the census of the 1,011 workers eligible to participate in the job placement intervention. The census took place at fifteen partner vocational training institutes throughout Uganda, and included all workers currently receiving training in either carpentry, catering, hairdressing, motor mechanics, tailoring or welding and expected to graduate in time for the matching intervention.

less than twenty employees for about the same fraction of trainees. Thus, the typical trainee will look for jobs in SMEs after training.[9]

We focus on young trainees for two reasons. First, the share of young workers is higher in developing countries, and this is particularly true for Uganda, which has one of the youngest populations in the world (UN, 2019). Second, young workers can be particularly affected by information frictions at recruitment given their lack of work experience (Pallais, 2014). It is important to note that our sample of workers is representative of the population of students who were able to self-finance training at our partner VTIs, and so are substantially more educated and wealthier than the average Ugandan youth.[10] In addition, since the census was conducted before graduation, we did not restrict the sample to individuals who were unemployed or were facing particular challenges in finding employment. Our sample includes entry-level workers in semi-skilled occupations, and so is different from those in related studies such as Abel *et al.* (2020), Alfonsi *et al.* (2020), Carranza *et al.* (2020) and Abebe *et al.* (2021), who instead specifically targeted disadvantaged youths. We can therefore expect a larger share of workers to find employment in the follow-up period relative to these other studies, which makes this an appropriate context to study the impact of certificates on the allocation of labour and wages, over and above the extensive margin of overall employment probability.

[9] Job placement activities by VTIs are very limited and informal.

[10] The median youth aged 18–25 in the 2012/13 Ugandan National Household Survey (UNHS) has six years of formal education, while the median level of formal education is 11 years in our sample (UBOS, 2014). In Online Appendix A we compare our experimental sample with youths in the UNHS, and show that the youths in our sample are also less likely to be married and come from households that are significantly better off on multiple measures of asset ownership.

1.2.2. *Selection into the experimental sample*

After completing the survey, all eligible trainees were informed about *both* the matching and signalling components of the intervention. Specifically, they were told that BRAC would schedule job interviews with potential employers among SMEs, but also that BRAC would conduct additional skill measurements—on both cognitive and soft skills—and that the information from the assessments might be disclosed to potential employers during the matching.[11] All trainees were then asked their interest to participate. The interested trainees were administered a baseline survey and were included in the skill assessments: the 787 trainees that confirmed their interest in the intervention and completed the main skills assessment form our final experimental sample.

There are two potential margins of selection into the experiment that might be relevant. First, the experimental sample might be selected relative to the eligible population of graduates from our fifteen partner VTIs, which is possible since participation was voluntary. Second, our sample might be selected relative to the broader population of typical entry-level employees in SMEs in our sectors, which is possible since we worked with particularly reputable training institutions. Exploring both margins is important to correctly interpret the skill distribution of the sample of workers that are presented to employers in the matching intervention, and how this compares with managers' priors.

On the first point, since workers were informed about the signalling component, if workers are aware of their skills then those with higher skills should have higher perceived returns from participation, and so a higher propensity to sign up. In addition to socio-demographics, the worker census included two measures of skills: (*i*) a cognitive test and (*ii*) a Big Five questionnaire.[12] Online Appendix Table A1.1 uses this information to study selection into the experiment, and shows that indeed the five Big Five variables are jointly significant in predicting selection. As expected, the selection is *positive*, so that workers with higher soft skills are more likely to be in the final sample.[13] The magnitude of this selection is sizeable. The score on each of the Big Five goes from 1 to 5, and so we summarise the extent of the selection by considering the share of workers who scored 3 or above on all Big Five: this proportion is 24% in the excluded workers, and increases to 33% in the group included in the experiment, a difference significant at the 5% level and corresponding to a 38% increase. As shown in column 3 of Table A1.1, a dummy for whether the worker scored 3 or above on all Big Five remains a sizeable predictor of selection when controlling for worker characteristics: workers with a high value on all the Big Five are 7 percentage points (pp) more likely to be in the final sample, a result significant at the 5% level. The Big Five are as important in driving selection as lack of past work experience, which the literature has highlighted as a key source of heterogeneity in explaining the impacts of signalling interventions (Pallais, 2014). Figure 2 plots the distributions of the three Big Five traits that are individually significant in column 2 of Table A1.1, i.e., agreeableness, conscientiousness and neuroticism, and confirms that the positive selection is all along the skill distribution.

[11] Also, on the worker side, we do not believe concerns related to the credibility of BRAC as an implementing agency to be first order given the scope of BRAC's work on youth programs in Uganda.

[12] The Big Five are five basic dimensions of personality: agreeableness, conscientiousness, extraversion, neuroticism and openness to experience. See John and Srivastava (1999) for a review of the main concepts and methods related to the definition and measurement of the Big Five traits. More details on the measurement and distribution of these skills in our sample are given in the Supplemental Material.

[13] Interestingly, we find no selection on cognitive ability. This suggests that cognitive skills may be harder to hide from employers, something that indeed we verify in Subsection 2.1.

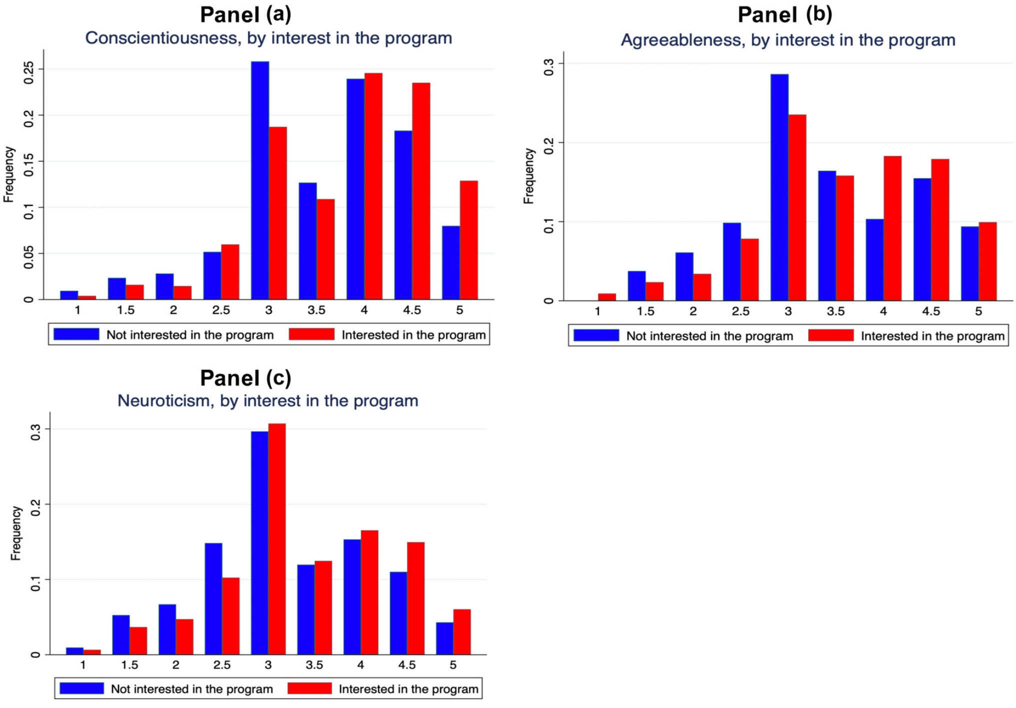

Fig. 2. *Distribution of Worker Soft Skills, by Participation in the Experiment.*
Notes: Agreeableness, conscientiousness and neuroticism are measured using a ten-item Big Five scale. The neuroticism variable is recoded so that a higher level of the variable corresponds to a lower level of neuroticism (i.e., to more self-control). The sample includes the 1,011 trainees eligible for the intervention.

On the second point, we compare our experimental sample to the representative sample of employees in small-scale Ugandan manufacturing in Bassi *et al.* (2021). The Bassi *et al.* (2021) survey covered SMEs in similar sectors to ours and collected the soft skills of every worker, using exactly the same Big Five scale that we used in our paper. In short, we find that our sample is positively selected on soft skills relative to comparable young employees in the Bassi *et al.* (2021) sample; for instance, while 21% of the young VTI graduates in Bassi *et al.* (2021) had a score of 3 or above on all Big Five, the same share is 33% in our sample, which is more than a 50% increase. This analysis, described in detail in Online Appendix A, indicates that our sample is positively selected not only within the workers eligible for the experiment, but also with respect to typical young VTI graduates employed in similar sectors.

This positive selection is important for interpreting the results of our intervention in two ways. First, if managers expect the matched workers to be a random sample of VTI graduates with similar characteristics to the workers that they typically hire (i.e., if they miss the positive selection), then the revelation of information might provide mostly positive news on the skills of workers. Second, this selection implies that we will not be able to estimate the impacts of the intervention for workers at the lower end of the distribution, who are selecting out. However, we will use evidence on the importance of soft skills in our sample of firms, along with a model, to discuss what the impacts of the information revelation would be for these workers in a counterfactual exercise.[14]

1.3. *Key Facts about SMEs at Baseline*

We now present four key facts from the baseline survey of SMEs, which inform our research design and empirical strategy. More details and the supporting evidence are provided in Online Appendix B.

(*i*) Soft skills are perceived by firm owners as having relatively high returns. This matches evidence that non-cognitive skills have high returns in labour markets in both developing countries and the United States (Heckman *et al.*, 2006; Deming, 2017; Adhvaryu *et al.*, 2018).
(*ii*) Firm owners report difficulties in observing workers' soft skills and theft by their own employees among their main perceived constraints.[15]
(*iii*) Firm owners have relatively low expectations on the distribution of soft skills among workers, and think that workers with good soft skills are hard to find. In addition, firm owners are not familiar with the VTIs in our sample and have low expectations on the soft skills of their graduates.
(*iv*) It is common for firms to recruit applicants who walk up to the firm and ask for a job, without any prior connection with the firm.

The first and second key facts justify our focus on soft skills. The second key fact in particular makes clear one reason why soft skills are important: workers with low soft skills can create a *loss* to the firm by, for example, stealing. Since soft skills are difficult to observe, as made precise later in the model, this can limit the propensity of managers to hire and can reduce wages. Given the relatively low expectations of firm owners on the distribution of soft skills, the third key fact suggests that the signalling intervention might be especially beneficial for those workers who can now credibly signal they are not among those low types that firm owners are afraid of hiring. The third key fact also suggests that firms might miss the positive selection of workers into the experiment, since they are not familiar with our partner VTIs and have low expectations on their graduates. Finally, the fourth key fact indicates that referrals do not fully solve the hiring problem, and so an intervention introducing workers to firms they have no prior contact with is likely to be informative of the regular hiring process in the labour market.

2. Intervention, Experimental Design and Data

2.1. *Intervention and Experimental Design*

The intervention we implemented has three components: (*i*) a screening component, whereby information was collected on the soft skills of workers while they were still enrolled at the VTIs;

[14] In the Supplemental Material we compare our sample to the US and Canadian samples in Srivastava *et al.* (2003), and show that our sample fares better on four of the five Big Five traits. With the caveat that the sampling strategies differ and that there might be cultural differences in how the questions are interpreted, these results are again in line with our sample being positively selected relative to the broader population.

[15] These results are in line with Bassi *et al.* (2021), who found that screening difficulties are the primary labour market constraint reported by managers in a representative survey of manufacturing firms in Uganda.

(*ii*) a matching component, whereby job interviews were scheduled between workers and firms; (*iii*) a signalling component, by introducing *experimental variation* in whether information from the screening assessments was disclosed to *both* workers and firms during the job interview process, through the provision of skill certificates.

2.1.1. *Worker screening*

Our screening activities targeted seven soft skills identified as relevant in initial focus groups with firm owners: creativity, communication skills, willingness to help others/pro-sociality, pro-activity, trustworthiness, discipline and attendance/time keeping.[16] Creativity, to be intended as the ability to come up with creative solutions to problems, is relevant for *all* sectors in our study, as workers are often asked to use in a creative way the limited tools available. Because employees are often asked to work in teams, and to take care of customers, skills such as communication, willingness to help others and pro-activity were mentioned as relevant in the focus groups. We discussed in the previous section how trustworthiness is important to firms. Existing research further suggests that firms value discipline, and that absenteeism is widespread in developing countries.[17]

We used teacher surveys to measure those skills that are easier to assess for an external examiner, namely attendance, discipline, communication, pro-sociality and pro-activity. To measure creativity and trustworthiness, we developed our own assessments: for creativity, we used a battery of questions; for trustworthiness, we made trainees play incentivised trust games.[18] We limited the revelation of information to five skills, to address concerns related to attention constraints. We selected the five soft skills based on the stated preferences of managers in the baseline survey. Specifically, we asked firm owners how much they would value additional information on the seven soft skills in our assessments, if they were to interview workers fresh out of VTIs. We then selected the top four skills—creativity, trustworthiness, communication and willingness to help others—plus attendance.[19] To facilitate the reporting, we followed the Ugandan education system, and graded each skill on a A–E scale. Both trainees and firm owners are used to this grading scale in this context, and grades of C and above are considered pass grades in Uganda. Grades were given using an absolute scale, and so were *not* curved within our experimental sample.[20]

To provide further evidence on selection into the experimental sample, in Online Appendix Table A2.2 we verify how our five soft skills are correlated with the Big Five. We document a positive correlation for some of the skills, in particular, conscientiousness and agreeableness,

[16] For the information revelation, we chose to focus on soft skills such as creativity, rather than on the Big Five, because the former were easier to explain to firm owners and workers, as revealed by our piloting exercises. We discuss the correlation between our chosen soft skills and the Big Five later in this section.

[17] Bowles and Gintis (2011) and Sackett and Walmsley (2014) showed that dependability and conscientiousness are among the skills most valued by US employers. Adhvaryu *et al.* (2020) and Krishnaswamy (2019) found absence rates of 8%–14% in Indian manufacturing, and Chaudhury *et al.* (2006) documented absence rates of 19% and 35% in the health and education sectors, respectively, across many developing countries, including Uganda.

[18] Specifically, as our measure of trustworthiness we use the amount sent back by the ‘recipient’ in the trust game, which is standard in the literature (see, for example, Glaeser *et al.*, 2000). The Supplemental Material reports more details on the skill assessments, including extracts from the actual scripts/questionnaires used.

[19] This design allows us to check whether attendance is then given a lower weight during recruitment, relative to the other skills. In practice, however, as shown in Online Appendix Table A2.1, the five soft skills are correlated, which limits our ability to separately identify the effect of revealing information on each skill. In Figure A5 in the Online Appendix we report the average importance given to each skill by managers.

[20] More details on the grading procedure are given in the Supplemental Material.

which predict selection in the final sample, are positively correlated with trustworthiness and creativity, respectively. While these correlations are not particularly strong, they do suggest that our sample of workers is positively selected not just on the Big Five, as documented in Section 1, but also on at least some of the specific soft skills revealed during the intervention. Online Appendix Figure A6 reports the distribution of grades for the five soft skills: indeed, only 2% of workers have a grade of D or below (which would be considered fail grades in Uganda) on all soft skills, and 88% of workers have a grade or A or B on at least one skill, so that the skill distribution is left skewed.

2.1.2. *Validating the soft skill measures*

To validate our skill measures, we check whether they predict employment and earnings in our sample. We focus on the control group, where information on skills was never revealed to workers, and report employment and earnings regressions estimated by pooling the two post-intervention worker follow-up surveys. Even if there are sizeable information frictions, we would expect soft skills to matter for earnings as employers learn them over time (Altonji and Pierret, 2001). On the other hand, soft skills might not predict the extensive margin of employment if they are hard to observe at recruitment. To account for the correlation among soft skills, we aggregate them by creating a dummy equal to one if the worker scored C or above (that is, had a pass grade) on all five skills.[21]

Online Appendix Table A3 shows that soft skills are a significant predictor of earnings in the two-year study period, even within our positively selected sample. This holds conditional on observables such as gender, age and education, and also conditional on cognitive skills. For instance, column 6 shows that workers with a pass grade on all skills earn $9 more per month, corresponding to an increase of 19% over the group with at least one fail mark. Such an increase in earnings is equivalent to having 2.5 additional years of education. While not causal, these results show that our measurements capture skills that are associated with labour market success.[22]

Columns 1 and 2 show that soft skills do not predict whether the worker is engaged in any paid work (including from wage employment, self employment and casual work), and nor does education, another strong predictor of earnings. This can in part be explained by the fact that engagement in paid work in the post-intervention period is already relatively high at 75%, which is expected given that by design we focus on youths that are more educated than the Ugandan average and are not disadvantaged, so that they do not face particular challenges in finding paid work. In columns 3 and 4 we focus on wage employment specifically, and find that soft skills are again not a significant predictor. On the other hand, column 4 shows that cognitive skills, years of education, vocational training duration and past employment are jointly significant at the 5% level in predicting wage employment. These four variables are all positively correlated and can be interpreted as proxies for hard/practical skills. This result is consistent with managers already being able to screen on hard/practical skills but not on soft skills. In line with this, Online Appendix Figure A2 shows that managers perceive practical skills as easier to screen than soft skills.

[21] Forty-one percent of workers had a pass grade on all skills.

[22] While the focus of this paper is on screening and signalling soft skills, these results also point to the potential value of *teaching* soft skills. In line with this, Adhvaryu *et al.* (2018) documented high returns to soft skills training for garment workers in India.

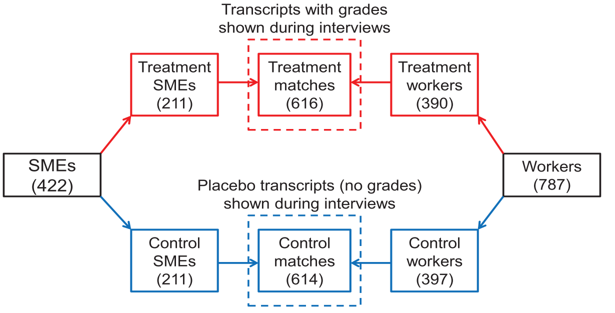

Fig. 3. *Summary of Our Experimental Design.*

2.1.3. *Treatment assignment and matching*

The second component of the intervention involved scheduling job interviews. This was done in two steps. First, workers and firms were randomly assigned to a treatment and a control group. The randomisation was done at the individual worker and firm levels, and stratified by submarkets, where a submarket is a sector-BRAC branch combination. Workers and firms were then randomly matched within each strata and treatment group. Figure 3 shows a summary of our experimental design. The random assignment to treatment and control groups produced a balanced sample. Panel A of Online Appendix Table A4.1 reports balance checks on the firm side, showing a balanced sample on nine of the ten variables considered.[23] Panel A of Table A4.2 shows that the randomisation produced a balanced sample also for workers. Importantly, the sample is well balanced on all soft skills.

A total of 1,230 job interviews were scheduled: 616 in the treatment group and 614 in the control group. The median firm (worker) was matched with three workers (one firm). There were no cross-treatment matches: treatment firms (workers) only met treatment workers (firms); control firms (workers) only met control workers (firms). Both groups got the matching component, but only the treatment group got the signalling component, as discussed below.[24]

2.1.4. *Skills signalling*

We created certificates reporting the grades of treatment workers on the five soft skills. Panel A of Online Appendix Figure A7 shows an example. The order of the skills on the certificates was randomised. On the back page, the certificates reported a brief description of the skills assessment procedure, as well as guidelines to interpret the grades.[25] To stress the credibility of the certificate, the front page reported the signatures of two high BRAC officials. To control for any potential effects of simply releasing any new document, a placebo certificate was produced for workers in the control group. An example is shown in panel B of Figure A7: the document simply states that the trainee was willing to be put in contact with potential employers, which is something that both the worker and the matched firms already knew, while it does not report

[23] To limit concerns related to potential lack of balance, we note that the normalised difference for the other variable (i.e., whether the owner has received training from a VTI) is small (0.12), and we further control for this variable and other firm characteristics in our main specifications. Also, we are not able to reject the null hypothesis that the ten variables considered for the balance checks are all jointly insignificant in predicting treatment assignment (p-value = 0.393).

[24] The Supplemental Material reports more details on treatment assignment and matching procedures. It also reports balance checks at the match level, and shows that the sample remains well balanced.

[25] The description made clear that trainees had not received any soft skills training as part of this project.

information on skills. The certificate is otherwise identical to the treatment one. Any treatment effects will thus be due to the *content* of the certificate, rather than to just having an additional document in their application files.[26]

The timing of the revelation of the certificates was the same in the treatment and control groups. On the interview day, the worker was first met by the BRAC staff, who showed the certificate and explained its content to the worker. The worker was informed that the firm owner would be shown the same certificate at the start of the interview. The worker then had the option to pull out, in case she had changed her mind about wanting to meet the firm owner. In practice, only two workers met the BRAC staff but decided not to proceed to the job interviews.[27] The worker was then introduced to the matched firm by the BRAC staff, and the firm owner was also shown the transcript by the BRAC staff, who made sure the firm owner understood how to interpret the grades. The transcript was then left to the worker to keep. After the initial introduction, the firm owner and the worker were left to interact as they pleased, and the BRAC staff played no further role in the interview.[28]

2.2. *Data Collection, Compliance and Attrition*

As shown in Figure 1, we collected four post-intervention surveys. First, a 'matching survey' was conducted during the matching intervention. For each scheduled job interview, we have information on: (*i*) whether the interview took place (and the reasons why it did not take place); (*ii*) basic interview descriptives such as whether the trainee brought additional documents to the interview and (*iii*) the beliefs of managers on the skills of the matched workers. The firm follow-up survey was conducted six months after the matching intervention, and contains information on firm-level outcomes. Finally, two worker follow-up surveys were conducted 12 and 26 months after the matchings, and contain information on worker-level labour market outcomes, expectations and search behaviour.

Figure 4 shows a summary of compliance and attrition. Starting from compliance, of the 1,230 scheduled job interviews, 515 (or 42%) took place. Lack of compliance is mainly due to workers having lost interest in being matched (32% of cases) or to the firm having lost interest (30% of cases) by the time they were called for the interviews.[29] Panel A of Online Appendix Table A5 explores the determinants of compliance, and shows little evidence of selection on observables.[30] Importantly, treatment assignment does not predict whether job interviews took place. This is not surprising, as the certificates were shown to firms and workers only *conditional* on the job interview taking place. Consistently with this, the sample of job interviews that took place

[26] It is possible that managers interpret the placebo certificate as a form of 'endorsement' of the worker by BRAC. This would bias our analysis towards finding no positive effect of the treatment certificates.

[27] This is consistent with the sample of workers being positively selected on soft skills, as discussed above.

[28] Our design allows us to estimate the impact of the soft skill certificates over and above any application documents already presented by the worker to the firm. Data collected during the matching intervention shows that 37% of workers brought at least one additional education/training certificate and 22% brought at least one reference letter to the interview (with 52% bringing either). This is in contrast to the results in Abel *et al.* (2020), who showed very low rates of usage of reference letters at 2% among their sample of disadvantaged job seekers in South Africa. This again confirms that our sample of youths is not disadvantaged.

[29] As shown in Figure 1, for logistical reasons, there was a lag of about three months from the graduation of trainees to the roll out of the matching intervention. This can explain the loss of interest. Another common reason why job interviews did not take place was failure to contact the worker/firm (18% of cases).

[30] The only significant predictor is the duration of the vocational course, so that workers in longer programs are less likely to participate. This is likely due to misunderstandings about when the intervention would take place, so that some workers had in fact not graduated yet by the time of roll out (and so were unable to participate).

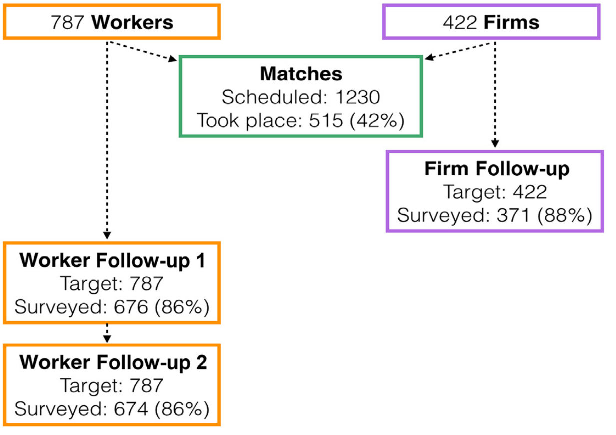

Fig. 4. *Compliance and Attrition.*

remains balanced on the main observable worker and firm characteristics.[31] All the treatment workers who showed up to the job interviews were given the certificates (corresponding to 49% of treatment workers). The remaining certificates were disbursed to the workers shortly after the first follow-up survey, so that 81% of treatment workers received the certificate in total.[32]

Moving on to attrition, the follow-up surveys targeted all firms and workers in the experimental sample, irrespective of whether the scheduled job interviews took place. We have very moderate attrition rates: these are 12% in the firm follow-up and 14% in both worker follow-ups. Panel B of Online Appendix Table A5 shows that attrition is not related to treatment assignment in either sample, and there is also very little evidence of observable characteristics determining attrition. Therefore, we do not correct for attrition in our main regression specifications.[33]

2.3. *Empirical Strategy and Regression Specifications*

2.3.1. *Empirical strategy*

Going into this experiment, we had two main priors about the labour market impacts of the certificates intervention: (*i*) the overall benefits of the intervention in terms of employment and earnings should be larger for workers with higher soft skills, and could be negative for workers at the bottom of the skill distribution, although the voluntary nature of the intervention should limit the possibility of any such negative impacts; (*ii*) more productive firms with higher returns to soft skills should respond more to the certificates, leading to an increase in sorting and efficiency.

[31] See the Supplemental Material for more details (Tables S3.1 and S3.2).

[32] The remaining 19% did not receive it as they could not be contacted during the disbursement.

[33] Our attrition rates are in line with other studies in similar settings. For example, the attrition rate is 17% in Abel *et al.* (2020) and 15% in Abebe *et al.* (2021). Panel B of Online Appendix Table A4.1 and panels B and C of Online Appendix Table A4.2 confirm that the samples of workers and firms overall remain balanced on baseline characteristics at follow-up. The firm follow-up shows some lack of balance on owner's age and past VTI training. However, we cannot reject that all observable firm characteristics are jointly insignificant in predicting treatment assignment at follow-up (p-value = 0.259). To further limit concerns related to attrition, we control for owner's age and VTI training in our main specifications. In the Supplemental Material (Table S7) we further show robustness of our main results to correcting for attrition using inverse probability weights (Wooldridge, 2010).

We planned our matched worker-firm design and our two-sided data collection to explore both dimensions of heterogeneity on the worker and firm sides.

Ultimately, however, the impacts on employment and earnings depend on the priors of managers and workers, and how these are revised with the certificates. The extent to which managers will increase or decrease hiring depends on whether they adjust their beliefs upwards or downwards for different workers. Similarly, the wage adjustment accepted by workers will depend on the extent to which workers revise their labour market expectations. As we have shown, the sample of workers is positively selected. Thus, it is crucial to document impacts on beliefs, as we cannot merely assume that the revision will be positive for workers with above-average soft skills and negative for workers with below-average soft skills.

Therefore, we start by presenting treatment effects on the beliefs of both workers and firms, allowing for heterogeneity on both sides of the market. Then, we discuss the implications of the documented impacts on beliefs for the treatment effects that we should expect on employment and earnings. We do so by writing down a simple job search model with two-sided heterogeneity and beliefs updating, and use the model to guide the interpretation of the treatment effects on employment and earnings along the skill distribution. We return to how our results match our initial priors at the end of Section 4, after presenting our main results.

2.3.2. *Regression specifications*

Our data and experimental design allow us to run regressions both at the match level and the worker/firm level. For the match-level analysis, we start by estimating the OLS regression equation

$$y_{ij} = \beta_0 + \beta_1 T_i + \beta_2 S_i + \beta_3 T_i \times S_i + \gamma \mathbf{X}_i + \delta \mathbf{X}_j + \theta \mathbf{Strata}_{ij} + \alpha Int_{ij} + \varepsilon_{ij}, \tag{1}$$

where y_{ij} is the outcome of the worker i-firm j match, such as the beliefs of firm j about worker i; T_i is a treatment group indicator; S_i captures the soft skills of worker i—in our baseline specification, we use the average grade on the five soft skills disclosed on the certificates, standardised to have mean 0 and standard deviation 1. As the soft skill measure is standardised, β_1 recovers the treatment effect for the average worker, and β_3 is then the heterogeneous effect by soft skills. We also estimate alternative specifications where S_i is a dummy equal to one if the worker had a pass grade (i.e., C or above) on all soft skills. In this case, β_1 is the treatment effect for workers who had a fail grade on at least one skill, and β_3 is the differential effect for workers who passed on all skills. The $\mathbf{X}_i$ are baseline worker controls.[34] The $\mathbf{X}_j$ are baseline firm controls.[35] The $\mathbf{Strata}_{ij}$ are dummies for the stratification variables (sector and BRAC branch). The Int_{ij} are dummies for the month of interview. We cluster SEs both at the level of the firm and the worker, to reflect our research design whereby workers were potentially matched to more than one firm, and firms were potentially matched to more than one worker (Cameron *et al.*, 2011; Abebe *et al.*, 2020).

The estimation sample for the match-level analysis includes the 515 matches that took place. This is our preferred sample because, as discussed above, (*i*) treatment assignment does not predict which job interviews take place and (*ii*) the sample of workers and firms remain balanced conditional on meeting. As the certificates were disclosed in all but two of the job interviews that

[34] These include a dummy for female, age and age squared, a dummy for any work experience, VTI course duration (in years) and completed years of formal education.

[35] These include the number of employees and the following owner characteristics: dummy for female, age and age squared, dummy for having attended a VTI.

took place, we interpret β_1 as the average treatment effect (ATE) on the population that meets. Panel A of Online Appendix Table A5 shows very little evidence that observable worker and firm characteristics drive which job interviews take place. However, if unobservable characteristics drive the decision to participate then the sample of realised matches might not be representative of the initially scheduled matches. We address this concern in the Supplemental Material, where we show that our results are robust to estimating a two-sided econometric model of sample selection.[36]

The main advantage of our two-sided experimental design is that it allows us to study impacts on the *allocation* of labour. To do so, we also need a measure of heterogeneity on the firm side. What we are after is a measure of manager productivity and returns to soft skills, as this would allow us to test whether the intervention increased positive assortative matching between workers and jobs. In the baseline survey we collected three measures of manager skills: (*i*) cognitive skills (through a Raven matrices test); (*ii*) years of education and (*iii*) the Big Five. Panels B and C of Online Appendix Table A6.1 show that firm owners with higher cognitive ability (*i*) manage more skilled employees; (*ii*) have higher profits; (*iii*) value soft skills relatively more and (*iv*) are better able to delegate tasks This suggests that indeed they are more productive managers.[37] In particular, since they employ more skilled workers and value soft skills more, higher ability managers might be particularly responsive to the certificates, as hiring skilled workers is especially important for this group of managers. Indeed, we note that our trainees have similar education to the employees of high-ability owners, and are instead substantially more educated than those of low-ability firms, having over one more year of education than them on average. This justifies using a measure of manager ability in the heterogeneous analysis. As shown in panel B of Table A6.1, cognitive ability is positively correlated with both the Big Five and years of schooling. Therefore, we start by aggregating all these skill measures (i.e., cognitive skills, the Big Five and years of education), and create a dummy equal to one if the manager has a value of the first principal component of these skills above the median.[38] This allows us to run heterogeneous effects by manager ability, and also to estimate a sorting regression similar to (1) but where a full set of interactions is included between worker skills, manager ability and a treatment indicator (so a triple difference specification).

As another approach to study sorting, we also estimate heterogeneous specifications that exploit the score on the individual soft skills of workers and their match with the needs of different employers. For example, we study whether workers with high communication skills are more likely to be hired by firms where interactions with customers are more frequent.

For the worker-level analysis, we estimate the following ANCOVA specification by OLS, on the two follow-up survey waves, $t = 1, 2$:

$$y_{it} = \beta_0 + \beta_1 T_i + \beta_2 S_i + \beta_3 T_i \times S_i + \beta_4 y_{i0} + \gamma \mathbf{X}_i + \theta \mathbf{Strata}_i + \alpha Int_{it} + \vartheta_t + \varepsilon_{it}. \quad (2)$$

[36] In the Supplemental Material we also show robustness to estimating equations like (1) on the full sample of 1,230 *scheduled* matches by assigning a value of zero to the outcome of those job interviews that did not take place (see Table S5.2). Here β_1 recovers the intention to treat (ITT) parameter in this alternative specification.

[37] The fact that we find significant dispersion in skills and productivity within our firm sample is in line with Bassi *et al.* (2021), who found large differences in labour productivity and managerial quality across firms within three manufacturing sectors in Uganda. More generally, the literature on productivity dispersion tends to find substantial heterogeneity in productivity even within narrowly defined industries (Syverson, 2011).

[38] We also check robustness to using a continuous version of the first principal component and to using only cognitive skills as a measure of manager ability.

Here y_{it} is the outcome of worker i in follow-up t, for instance their total earnings, y_{i0} is the baseline value of the outcome (when available), ϑ_t is an indicator for the second follow-up, and the other variables are as previously defined. In our preferred specification we pool observations from the two follow-up surveys and cluster SEs at the worker level, to maximise power.[39] Some of the outcomes are only available at second follow-up, and we indicate when that is the case in the notes to the relevant tables. In this specification, β_1 recovers the ITT parameter since the sample includes all experimental workers, regardless of whether they met any firms for the job interviews. We also show heterogeneous effects by whether the worker met any firms as part of the matching intervention. For the firm-level analysis, we estimate equations like (2) but at the firm level, and so using only one round of follow-up firm surveys.

3. Impacts on Beliefs

We begin the empirical analysis by studying whether the certificates led firm owners to revise their beliefs on the skills of the matched workers, and whether workers reacted by updating their labour market expectations. We can expect managers to react to the certificates because, as documented in Subsection 1.3, they report soft skills as important but difficult to observe. At the same time, the labour market expectations of workers might change if through the certificates workers are better able to signal their skills to potential employers, or if they obtain more precise information on their own skills. Therefore, it is possible that *both* sides of the market updated their beliefs as a result of the new information.

3.1. *Impacts on Managers' Beliefs*

After each job interview, we elicited the beliefs of managers on the skills of the worker. Managers were asked whether: (*i*) there was anything they particularly liked about the worker; (*ii*) there was anything they particularly did not like about the worker; (*iii*) they thought the worker was more skilled than usual applicants; (*iv*) they thought the worker was less skilled than usual applicants. We create four dummies using the answers to these questions, and combine these into a standardised index, with a more positive value indicating a more positive assessment. Importantly, this index allows us to document both positive and negative updating. We run OLS regressions analogous to (1) with this index as the dependent variable. As discussed above, the sample includes the 515 job interviews that took place.

Table 3 reports the results, both with and without controls. Column (1) shows that: (*i*) the estimated treatment effect for the average worker ($\hat{\beta_1}$) is positive and relatively large at around 0.1 standard deviations, although it is not significant at conventional levels (p-value $= 0.165$), and (*ii*) there is significant heterogeneity by soft skills, with the revision being more positive for workers with higher skills: the treatment effect is 0.26 standard deviations larger for those with soft skills one standard deviation above the average, a result significant at the 1% level. This conclusion is unaffected by whether we omit control variables, as shown in column (2).

In columns (3) and (4) we consider our alternative measure of skill heterogeneity, and find that the revision of beliefs is positive and significant for workers with a pass grade on all soft skills: as shown by the interactions in row (*v*), the treatment effect for these workers is 0.36–0.41

[39] We show dynamic treatment effects for our main outcomes in the Supplemental Material (Table S7).

Table 3. *Impacts on Managers' Beliefs about Matched Workers.*

OLS coefficients; SEs adjusted for two-way clustering in parentheses

Dependent variable:	Assessment of worker skills by owner (standardised index)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(*i*) Treatment	0.117	0.130	−0.041	−0.007	−0.047	−0.062	−0.197	−0.193
	(0.084)	(0.089)	(0.121)	(0.120)	(0.104)	(0.105)	(0.152)	(0.148)
(*ii*) Average worker soft skills grade (standardised)	−0.091	−0.084						
	(0.063)	(0.063)						
(*iii*) Average worker soft skills grade (standardised) × Treatment	0.264***	0.250***						
	(0.097)	(0.095)						
(*iv*) Worker has pass grades on all soft skills			−0.086	−0.081	0.133	0.118	0.040	0.057
			(0.134)	(0.134)	(0.098)	(0.098)	(0.168)	(0.164)
(*v*) Worker has pass grades on all soft skills × Treatment			0.414**	0.360**			0.390*	0.340
			(0.189)	(0.183)			(0.233)	(0.224)
(*vi*) Owner is high ability					−0.298**	−0.308**	−0.156	−0.148
					(0.146)	(0.138)	(0.174)	(0.171)
(*vii*) Owner is high ability × Treatment					0.570***	0.621***	0.522**	0.567***
					(0.174)	(0.167)	(0.217)	(0.216)
(*viii*) Worker has pass grades on all soft skills × Owner is high ability × Treatment							0.083	0.092
							(0.339)	(0.344)
Mean of dep. var. in control group	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Worker and firm controls	Yes	No	Yes	No	Yes	No	Yes	No
p-value from test: $(i) + (v) = 0$			[0.005]	[0.009]				
p-value from test: $(i) + (vii) = 0$					[0.000]	[0.000]		
Number of observations (matches)	515	515	515	515	515	515	515	515

Notes: ***, ** and * denote significance at the 1%, 5% and 10% levels, respectively. Results from the matching surveys are reported. SEs are adjusted for two-way clustering (at the level of both the firm and the worker), following Cameron *et al.* (2011). The dependent variable is the standardised index of the following variables: dummy equal to one if the worker was reported by the owner as more skilled than usual applicants, and equal to zero otherwise; dummy equal to zero if the worker was reported as less skilled than usual applicants, and equal to one otherwise; dummy equal to one if there was anything that the owner particularly liked about the worker, and equal to zero otherwise; dummy equal to zero if there was anything that the owner particularly did not like about the worker, and equal to one otherwise. The index is constructed by converting each component into a z-score, averaging these and taking the z-score of the average. The z-scores are computed using means and SDs from the control group. All regressions control for stratification variables (dummies for BRAC branch and sector) and for dummies for the month of interview. The regressions in the odd columns also control for the following baseline worker characteristics: age and age squared; female dummy; years of education; duration (in years) of the vocational training program the worker was attending at baseline; dummy for any past work experience. The odd columns also control for the following baseline firm characteristics: female owner dummy; age and age squared of the owner; dummy for whether owner attended a VTI in the past; number of employees. We use two measures of worker skills. The first, used in columns (1)–(2), is the average grade (standardised) on the five soft skills measured in the baseline assessments. The second, used in the rest of the table, is a dummy for whether the worker had a pass grade (C or above) on all five soft skills. To proxy for owner ability, we use the first principal component of the following skills measured at baseline: cognitive skills, years of education and the Big Five. Firm owners with a value of the first principal component on or above the median are assigned to the high-ability group. Each regression also controls for the variable by which heterogeneous effects are considered, uninteracted with Treatment. All regressions further control for dummies for missing values in each of the independent variables.

standard deviations larger than for workers with at least one fail grade, a result significant at the 5% level. However, we find no strong evidence of negative updating for workers with lower skills, as indicated by the small and insignificant coefficients in row (*i*). That is, the certificates strengthen the correlation between workers' skills and managers' assessments, but this is driven by workers with higher soft skills. We confirm this in panel A of Online Appendix Figure A8, which reports a non-parametric regression of the managers' assessments on the average soft skill grade of the worker, by treatment group. The figure shows a relatively flat relationship in the control group, which is in line with managers being unable to screen on soft skills without the certificates. The correlation instead becomes positive in the treatment group, and the figure confirms that the revision of beliefs is positive along most of the skill distribution. There is some evidence of negative revision at the very low end of the distribution, but this is at most limited to the 10% of workers who have an average soft skill grade of 2.4 or below (the median soft skill grade is 3.4 and the distribution is left skewed).[40] The positive revision for workers with higher skills, and the limited revision for workers with lower skills are consistent with the sample being positively selected on soft skills, and with the low priors of managers on the skills of workers documented in Subsection 1.3, so that the revelation of information provides mostly *positive news* to managers. The fact that, as discussed in Subsection 1.3, managers are not familiar with the VTIs in our sample can explain why they miss the positive selection of workers into the experiment and why the reaction to the certificates is mostly positive.

In columns (5) and (6) we consider heterogeneity by manager ability, and find that the positive response is concentrated among higher ability managers, who revise upwards their beliefs significantly even for the average worker. This is in line with the evidence documented earlier that higher ability managers are more productive and value soft skills more, so that they are more responsive to the information.[41] Note that, as shown in Online Appendix Table A6.2, there are no differences by manager ability in how familiar managers are with our partner VTIs, and in their expectations about the quality of workers graduating from them. This can also explain why higher ability managers miss the positive selection of workers into the experiment.[42] Finally, in columns (7) and (8) we include a full set of interactions between worker skills, manager ability and the treatment dummy. The coefficients in rows (*i*) and (*v*) show that, while low-ability managers do also revise their beliefs, this revision is weaker and imprecisely estimated. The coefficients in rows (*vii*) and (*viii*) of columns (7) and (8) decompose the difference between high- and low-ability managers documented in row (*vii*) of columns (5) and (6) into that part coming from matches with lower skilled workers (row (*vii*) of columns (7) and (8)), and the *additional* effect from meeting a worker with higher skills (row (*viii*) of columns (7) and (8)). The results confirm that the positive revision of higher ability managers is stronger for all types of workers. The positive triple-difference estimates in row (*viii*) suggest that the difference in

[40] In Online Appendix Table A7.1 we confirm that these results are not sensitive to alternative ways of aggregating the five soft skills. In the Supplemental Material, we also estimate more flexible specifications that allow for heterogeneous effects by terciles or quartiles of the average soft skills grade (see Table S4). These reinforce the main conclusions from Online Appendix Figure A8.

[41] Online Appendix Table A7.1 shows that these results are robust to alternative ways of defining manager ability.

[42] Row (*vi*) of Table 3 shows that in the control group, higher ability managers have a lower assessment of workers' skills. On this point, we note that soft skill grades do not predict managers' assessments in the control group (for either type of manager), so this result is not driven by higher ability managers being more or less able to screen the soft skills of control workers. As higher ability managers are more skilled themselves and employ workers with higher skills, this more negative assessment in the control group might simply reflect the fact that high-ability managers compare the matched workers against a higher benchmark.

treatment effects between high- and low-ability managers is stronger for workers with higher skills, but not significantly so.[43]

One potential concern with this heterogeneous analysis is that manager ability might proxy for other firm characteristics. To address this, we run a regression where the treatment dummy is interacted with a number of firm and manager characteristics, all at the same time. Online Appendix Table A7.2 reports the coefficients on such interactions, and confirms that manager ability is the key source of heterogeneity, as the coefficient on the interaction with the high-ability dummy remains large, stable and significant at the 1% level as other interactions are progressively added.[44] As discussed later in Section 4, we will show that these heterogeneous impacts on beliefs then map into heterogeneous impacts on hiring. As hiring is a high-stakes outcome, this alleviates the additional potential concern that the impacts on beliefs might be driven by experimenter demand effects related to high-ability owners being better able to understand what the experimenter is after and responding to this.[45]

In summary, the overall positive revision of beliefs is consistent with workers being positively selected and with managers having low priors on the distribution of soft skills. The fact that the revision is stronger for managers of higher ability is in line with these managers being more productive and valuing soft skills more, and so paying more attention to the certificates. The results on manager updating are in line with the findings of the literature on recruitment in firms, which shows that the introduction of job tests helps managers select workers with *higher* skills (Autor and Scarborough, 2008), and that better managers are more likely to follow the recommendations of job tests (Hoffman *et al.*, 2018).

3.2. *Impacts on Job-Seekers' Beliefs and Outside Options*

We now turn to the impacts of the certificates on the labour market expectations of workers. We can expect such an impact because the certificates were left to the workers to keep after the job interview. So if the certificates help workers signal their skills in the labour market, or if workers learn about their soft skills through the certificates, then this could affect their expectations and in particular proxies for their outside options.

We pool data from the two worker follow-up surveys, and estimate regressions analogous to (2). Table 4 reports the results. Panel A focuses on *beliefs*. Column (1) shows that the average treated worker reports expected monthly earnings that are \$7.9 higher than the control group, a result significant at the 5% level and corresponding to a 7% increase. The rest of panel A shows that treatment workers report 5% higher expected probability of employment (column (2)), a higher intention to bargain for wages, a result at the margin of significance (column (3)), and are also 7pp more likely to report that their ideal job is in a firm with ten or more employees (i.e.,

[43] Overall, Table 3 shows that the inclusion of control variables does not alter the results significantly. Therefore, we relegate remaining results without controls to the Supplemental Material.

[44] In particular, Table A7.2 controls for interactions between the treatment dummy and (*i*) the manager's risk aversion, (*ii*) the manager's English proficiency and (*iii*) the baseline firm performance indicators, namely the number of employees and profits. This helps rule out that the stronger reaction of higher ability managers to the certificates is capturing differences in risk preferences, in whether the content of the certificates was understood properly, or in labour demand due to differential firm growth. The results in Table A7.2 additionally rule out that manager ability is capturing sectoral heterogeneity, as shown by the *p*-value on the test of joint significance of the interactions of the sector dummies with the treatment dummy, reported at the bottom of the table.

[45] The consistency of the impacts on beliefs with those on hiring, and the stability of the coefficients on the interactions between the high-ability dummy and treatment in Table A7.2 as other interactions are added reassure us that these heterogeneous results on beliefs are not simply driven by noise in small samples.

Table 4. *Impacts on Worker Labour Market Expectations and Outside Options.*

OLS coefficients; SEs in parentheses are robust in columns (3), (4) and (7), and clustered at the worker level in the other columns

	Panel A: beliefs				Panel B: behaviours			
	Monthly expected earnings (USD) (1)	Expected probability of employment in the next six months (0 to 10 scale) (2)	Expected bargaining over wages (standardised index) (3)	Ideal job is in a large firm (4)	Any casual work in the last week (5)	Attended further education or training in the last year (6)	Looked for a job in the public/NGO sector in the last year (7)	Looked for a job in the last year (8)
Treatment	7.90**	0.280**	0.104	0.068**	−0.048*	0.038*	0.105***	−0.018
	(3.13)	(0.116)	(0.063)	(0.033)	(0.026)	(0.020)	(0.036)	(0.025)
Average soft skills grade (standardised)	−0.442	0.126	−0.028	0.024	0.018	0.019	0.018	0.008
	(1.98)	(0.083)	(0.045)	(0.023)	(0.019)	(0.012)	(0.025)	(0.019)
Average soft skills grade (standardised)	2.85	−0.039	0.064	−0.010	−0.027	−0.016	−0.038	−0.012
× Treatment	(3.11)	(0.117)	(0.063)	(0.033)	(0.026)	(0.019)	(0.035)	(0.026)
Mean of dep. var. in control	114.7	5.53	−0.001	0.624	0.323	0.118	0.268	0.749
Controls for baseline value of outcome	Yes	Yes	Yes	No	Yes	No	No	Yes
Uses data from first follow-up	Yes	Yes	No	No	Yes	Yes	No	Yes
Uses data from second follow-up	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of observations	1,330	1,349	666	668	1,350	1,348	674	1,350

Notes: ***, ** and * denote significance at the 1%, 5% and 10% levels, respectively. Results from the worker follow-ups are reported. The outcome variables in columns (3), (4) and (7) are available only at second follow-up, which explains why the number of observations is lower in those columns. All regressions control for stratification variables (dummies for BRAC branch and sector) and dummies for the month of interview. The regressions in columns (1), (2), (5), (6) and (8) further control for a dummy for the second follow-up. In addition, all regressions control for the following baseline worker characteristics: age and age squared; female dummy; years of education; duration (in years) of the vocational training program the worker was attending at baseline; dummy for any past work experience. In column (1) the dependent variable is constructed as follows: respondents were asked to report (*i*) their minimum and maximum expected earnings and (*ii*) the probability that they could earn at least the midpoint. We use this information to fit a triangular probability distribution of expected earnings. Monetary amounts are deflated and expressed in terms of the price level in January 2015, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted in January 2015 USD. The top 1% values of expected earnings are excluded. The dependent variable in column (3) is constructed using two variables from the second follow-up: a dummy for whether the worker would not accept a job without negotiating on the wage; a variable reporting how much the worker would expect wages to be influenced by negotiation (0–10 scale). The index is constructed by converting each component into a z-score, averaging these and taking the z-score of the average. The z-scores are computed using means and SDs from the control group. In column (4) the dependent variable is a dummy equal to one if the worker reported an ideal firm size on or above the median (i.e., at least ten workers). Since at baseline all workers were enrolled in vocational training and only 1% of them were currently doing any paid work, for the employment outcomes, we consider as baseline value of the outcome the expected probability of employment in the 6sixmonths after graduation, as reported at baseline. So we control for this variable in column (5). In column (6) the dependent variable is a dummy equal to one if the worker attended further formal education or vocational training in the year prior to the survey. All regressions further control for dummies for missing values in each of the independent variables.

a ‘large’ firm) (column (4)). These results suggest that the certificates raise the perceptions of workers about what they can achieve in the labour market.

While impacts for the average worker are positive, there is no significant evidence of heterogeneity by worker skills, as shown by the estimates of β_3. This is consistent with the positive selection of workers into the experiment, which limits the extent of skill heterogeneity in our sample, and with the fact that higher ability managers revise their beliefs upwards even for the average worker. Another factor that reduces the scope for heterogeneity in Table 4 is that information comes from the follow-up surveys, when workers are searching in the wider labour market: as at that point workers can strategically decide whether to show the certificates during job search, we do not expect a negative effect on labour market expectations even at the bottom of the skill distribution, which then further limits the potential for heterogeneous effects. In line with this, panel C of Online Appendix Figure A8 reports non-parametric treatment effects on expected earnings along the skill distribution, and shows that the impact is positive along most of the distribution, and there is no evidence of significant negative revision at the bottom of the distribution. This is in contrast to panel A, which shows some evidence that in the matching intervention treated managers revise downwards their beliefs for workers with very low skills.

Panel B of Table 4 shows that treated workers change their labour market *behaviours* in a way consistent with their revision of expectations: the average treatment workers are 15% less likely to engage in casual work (column (5)), which is a poorly paid and insecure form of employment; in addition, they are 3.8pp more likely to have attended further education or training in the post-intervention period, from a control mean of 12% (column (6)). This effect on human capital accumulation is particularly revealing, as it points to the presence of complementarities between certification and investment in human capital (Spence, 1973). We investigate this further in Online Appendix Table A9, where we show that the treatment effects on education/training are entirely driven by individuals that we predict to have high ability for schooling, based on their baseline characteristics.[46] These results are then consistent with talented workers limiting their human capital investments because of information frictions, and with the certificates improving sorting also in terms of the allocation of labour between work and education. Moving to job search, treated workers are also 39% more likely to have looked for a job in the public or NGO sector in the last year (column (7)); however, they do not increase their overall search intensity (column (8)), so that the impact is primarily on search direction. These changes in behaviour are all consistent with treated workers trying to transition to better and more productive jobs. Again, we do not find any significant heterogeneity by worker soft skills.

Table 5 explores *why* workers reacted to the certificates. We distinguish between three potential explanations: (*i*) workers might have learned about their own soft skills; (*ii*) the certificates might have affected the perceptions of workers about the labour market returns to soft skills; (*iii*) the certificates might have improved the ability of workers to signal their skills. In the second follow-up, we asked workers in both treatment and control groups to assess themselves on the five soft skills reported on the certificates, using the same A–E scale used for grading. Column (1) shows that there is a positive and significant correlation between the worker self-evaluations and our measurements. This shows that workers are already aware of their soft skills, and is in line with the positive selection into the experiment documented above. Column (1) further shows

[46] In particular, we proxy schooling ability by a standardised index of four variables measured at baseline: (*i*) completed years of formal education prior to enrolling at the VTI; (*ii*) cognitive skills (measured through Raven matrices); (*iii*) a dummy for whether the worker was planning to gain further formal training or education in the future; (*iv*) the highest level of formal education that the worker would like to achieve in the future.

Table 5. *Impacts on Workers' Beliefs about own Skills, Returns to Skills and Signalling.*

OLS coefficients; robust SEs in parentheses

	Average self-assessed soft skills grade (1 to 5 scale) (1)	Average perceived returns to soft skills (0 to 10 scale) (2)	Average perceived constraints in signalling skills (1 to 5 scale) (3)
Treatment	0.019	−0.004	−0.185**
	(0.033)	(0.106)	(0.087)
Average soft skills grade (standardised)	0.059**	0.036	0.061
	(0.024)	(0.075)	(0.065)
Average soft skills grade (standardised)	−0.040	−0.006	0.098
× Treatment	(0.032)	(0.106)	(0.088)
Mean of dep. var. in control group	4.35	8.06	2.73
Controls for baseline value of outcome	No	No	Yes
Uses data from first follow-up	No	No	No
Uses data from second follow-up	Yes	Yes	Yes
Number of observations	673	673	673

Notes: ***, ** and * denote significance at the 1%, 5% and 10% levels, respectively. Results from the worker follow-ups are reported. The outcome variables in this table are available only at second follow-up. All regressions control for stratification variables (dummies for BRAC branch and sector) and dummies for the month of interview. In addition, all regressions control for the following baseline worker characteristics: age and age squared; female dummy; years of education; duration (in years) of the vocational training program the worker was attending at baseline; dummy for any past work experience. The dependent variable in column (1) is constructed as follows: workers were asked to self-assess themselves on the five soft skills measured in the baseline assessments and disclosed on the certificates, using the same A–E scale. We compute the average self-assessed grade on the five skills. The dependent variable in column (2) is constructed as follows: workers were asked how important each of the five soft skills disclosed on the certificates were for making someone a productive worker, using a 0–10 scale. We compute the average perceived return to the five skills. The dependent variable in column (3) is constructed as follows: workers were asked to report their perceived importance of constraints related to signalling skills in the labour market, using a 1–5 scale. Workers were asked this question separately for practical and soft skills. Note that this question was asked about soft skills in general. We compute the average importance of constraints related to signalling the two types of skills (practical and soft) and use this as the dependent variable in column (3). All regressions further control for dummies for missing values in each of the independent variables.

that the certificates do not alter the correlation between the self-reports and our measurements. So workers are *not* learning about their skills. Column (2) shows that there is no impact on perceived returns to soft skills, thus confirming that workers are also not changing their beliefs about the importance of soft skills. Finally, column (3) shows that treated workers believe they face fewer challenges in signalling their skills to employers. The effect corresponds to a 7% reduction over the control mean (significant at the 5% level). This evidence indicates that the impacts on the perceived outside options of workers are driven by the signalling value of the certificates. Consistently, we further note that among treatment workers (*i*) 92% still had the certificate after two years and (*ii*) 74% reported using it in job search.[47]

In Online Appendix Table A8.1 we report heterogeneous effects on the outcomes in Tables 4 and 5 by whether the worker met any firms in the matching intervention. This is a valid comparison because, as discussed in Subsection 2.2, treatment assignment does not predict whether the job interviews took place. The table shows that the treatment effects are driven by the matched sample. In particular, treatment effects for the matched sample are largely significant and in almost all cases stronger than in the whole sample, as indicated by the coefficient on the Treatment dummy. On the other hand, impacts for the unmatched sample are largely not significant, as shown by the

[47] This is consistent with employers outside our experiment also valuing the BRAC certificate, which is not surprising given that BRAC is very well known throughout Uganda for its youth and firm programs.

p-values from the test of significance of the sum of the coefficients on the Treatment dummy and the interaction with the dummy for whether the worker did not meet any firms, reported at the bottom of the table. The larger impacts for the matched sample suggest that workers learned about the importance/value of the certificates from the reaction they got from the matched employers. As managers mostly responded positively, this can then explain why the matched workers revise their labour expectations upwards on average. This result again reinforces our interpretation that both sides of the market face information frictions and react to the certificates, and that job interviews are a crucial stage when information is revealed.

4. Impacts on Sorting, Employment and Earnings

In the previous section, we showed that, upon receiving the certificates, (*i*) managers revise *upwards* their beliefs on the skills of workers, with a stronger effect among higher ability managers, and (*ii*) workers increase their labour market expectations. We now present the reduced-form impacts on sorting, overall employment and earnings. To guide the discussion, we interpret these results in the context of a search model with two-sided heterogeneity, which makes precise the impacts on labour market outcomes that we should expect given the revision of beliefs documented above. Altogether, the empirical evidence in this section shows that the treatment certificates lead to an increase in sorting, as well as to an increase in earnings conditional on employment, which is particularly pronounced at the top of the skill distribution.

4.1. *Conceptual Framework and Mapping to Data*

We develop a partial equilibrium search model with asymmetric information on workers' skills and two-sided updating, building on our evidence on two-sided belief updating documented above. We provide full details of the model in Online Appendix C. Here, we sketch the main argument and provide the economic intuition behind the main results.

There are two types of workers, who differ in their soft skills, and two types of firms, who differ in their production technology. High-skill workers generate higher output only when matched to high-type firms. Therefore, the output-maximising allocation exhibits positive assortative matching. Search frictions and asymmetric information on workers' skills result in worker-firm mismatches and output losses. We model the intervention as an increase in the precision of the signal on workers' soft skills during job interviews.[48] The model generates the following comparative statics implications about the introduction of the certificates.

IMPLICATION 1. *There is an increase in positive assortative matching.*

IMPLICATION 2. *The change in employment probability is ambiguous for both types of workers.*

IMPLICATION 3. *Average wages conditional on employment increase for high-skill workers and decrease for low-skill workers.*

The effect on assortative matching follows from the reduction in information asymmetries caused by the certificates, which improve the precision of workers' signals. On the one hand, high-skill workers are more likely to receive and accept an offer from high-type firms, who now

[48] For simplicity, the model abstracts from the positive workers' selection and low managers' priors documented above.

offer a higher wage due to higher expected productivity. On the other hand, low-skill workers are more likely to accept offers from low-type firms, due to lower employment opportunities at high-type firms. As the treatment leads workers to *reallocate* across firms, effects on overall employment probability are ambiguous. Earnings conditional on employment increase for high-skill workers, since they are more likely to get employed at high-type firms, who offer a higher wage. On the other hand, low-skill workers earn less, as their probability of employment at high-type firms is reduced.

The fact that the certificate is observed by both sides of the market is crucial for the intervention to have an impact on wages: in the treatment group, high-skill workers can truthfully reveal their skills to all high-type employers. This increases their outside option, and so each high-type employer has to increase the offered wage for high-skill workers to accept their offer.

The model also highlights that, while high-skill workers always benefit from the certificates, the latter can worsen the outcomes of workers at the low end of the distribution by reducing their employment probability at high-type firms. So while the certificates improve the efficiency of the allocation of workers to jobs, they can increase wage dispersion and inequality among workers. Since low-skill workers represent a small share of the sample due to the positive selection on soft skills that we documented, we are unable to use the experiment to study impacts on these workers. In Section 5 we use the insights from this model together with the reduced-form evidence from the experiment to discuss the implications of scaling up the certification intervention to workers of all skill levels.

4.2. *Impacts on Sorting*

We present two sets of results on sorting impacts from the matching intervention: first, we study whether the certificates increase the sorting of workers with high soft skills to high-ability managers, who we have shown are more productive and value soft skills more in general; second, we leverage the scores on specific soft skills and their match with the needs of different employers, to examine whether firms in need of certain skills are more likely to hire a worker with a higher grade on those skills.

4.2.1. *Sorting between soft skills and manager ability*

Table 6 reports the results of match-level regressions analogous to (1), but where the outcome is a dummy for whether the worker was hired by the matched firm. The data are from the first worker follow-up, and so the observations are 412 (instead of 515), due to the attrition discussed in Subsection 2.2.[49] Columns (1) and (2) show that: (*i*) there is no significant treatment effect for the average worker and (*ii*) there is no evidence of significant treatment heterogeneity by worker skills.[50] This is consistent with model Implication 2, which states that the impact on overall employment is ambiguous if workers reallocate between different types of firms. That is, given

[49] Information on hiring outcomes in the matching intervention was collected in both the firm and worker follow-ups. We prefer to use information from the worker follow-ups as measurement error is likely lower there for at least two reasons: (*i*) while the median firm was matched to three workers, the median worker was matched to one firm, so possible recall errors related to the respondent getting confused about the different job interviews are lower on the worker side; (*ii*) in 13% of the cases, the person answering the firm follow-up survey is different from the owner that conducted the job interviews. However, we note that the correlation between hiring outcomes in the two surveys is 0.82, and results using firm reports (not shown) are qualitatively similar.

[50] Online Appendix Table A7.1 shows that this result does not depend on the way skills are aggregated. In line with this, panel B of Online Appendix Figure A8 also shows a lack of significant heterogeneity in non-parametric regressions of hiring probability on worker skills.

Table 6. *Impacts on Sorting Between Worker Soft Skills and Manager Ability.*

OLS coefficients; SEs adjusted for two-way clustering in parentheses

Dependent variable:	Worker was hired by the matched firm			
	(1)	(2)	(3)	(4)
(*i*) Treatment	−0.001	0.009	−0.047	−0.047
	(0.034)	(0.041)	(0.041)	(0.054)
(*ii*) Average worker soft skills grade (standardised)	−0.004			
	(0.025)			
(*iii*) Average worker soft skills grade (standardised) × Treatment	−0.025			
	(0.035)			
(*iv*) Worker has pass grades on all soft skills		−0.018	−0.036	−0.050
		(0.042)	(0.032)	(0.048)
(*v*) Worker has pass grades on all soft skills × Treatment		−0.025		−0.001
		(0.064)		(0.073)
(*vi*) Owner is high ability			−0.066	−0.102*
			(0.048)	(0.059)
(*vii*) Owner is high ability × Treatment			0.133**	0.166**
			(0.065)	(0.081)
(*viii*) Worker has pass grades on all soft skills × Owner is high ability × Treatment				−0.076
				(0.135)
Mean of dep. var. in control group	0.116	0.116	0.116	0.116
Worker and firm controls	Yes	Yes	Yes	Yes
p-value from test: $(i) + (v) = 0$		[0.766]		
p-value from test: $(i) + (vii) = 0$			[0.102]	
Number of observations (matches)	412	412	412	412

Notes: ***, ** and * denote significance at the 1%, 5% and 10% levels, respectively. Results from the worker first follow-up are reported. The number of observations is lower than in Table 3 due to attrition. SEs are adjusted for two-way clustering (at the level of both the firm and the worker), following Cameron *et al.* (2011). All regressions control for stratification variables (dummies for BRAC branch and sector) and dummies for the month of interview. In addition, all regressions control for the following baseline worker characteristics: age and age squared; female dummy; years of education; duration (in years) of the vocational training program the worker was attending at baseline; dummy for any past work experience. All regressions also control for the following baseline firm characteristics: female owner dummy; age and age squared of the owner; dummy for whether owner attended a VTI in the past; number of employees. We use two measures of worker skills. The first, used in column (1), is the average grade (standardised) on the five soft skills measured in the baseline assessments. The second, used in the rest of the table, is a dummy for whether the worker had a pass grade (C or above) on all five soft skills measured in the baseline assessments. To proxy for owner ability, we use the first principal component of the following skills measured at baseline: cognitive skills, years of education and the Big Five. Firm owners with a value of the first principal component on or above the median are assigned to the high-ability group. Each regression also controls for the variable by which heterogeneous effects are considered, uninteracted with Treatment. All regressions further control for dummies for missing values in each of the independent variables.

that managers revise upwards their beliefs on the skills of workers (columns (1)–(4) of Table 3), we would expect this to increase the employment probability of treated workers at firms that value those skills, and to reduce it at firms that do not value them, so that the net effect on overall employment is ambiguous.

In column (3) we consider heterogeneity by manager ability: the coefficient in row (*i*) is the treatment effect for low-ability managers. This is negative, even though not significant. On the other hand, higher ability managers in the treatment group are 13pp more likely than low-ability managers to hire a worker, as shown by the positive and significant estimate in row (*vii*). The sum of the coefficients in rows (*i*) and (*vii*) is the treatment effect for high-ability managers, and the *p*-value at the bottom of column (3) shows that this is positive and at the margin of significance. As these effects are imprecisely estimated, we check robustness in Online Appendix Table A7.1, where we use only cognitive ability (rather than the first principal component) as a measure of ability. We find that the treatment effect for low-ability managers is negative and significant at the 10% level, and that for high-ability managers is positive and significant at the 5% level (Table A7.1, column 8). As discussed, we interpret our positively selected sample as mostly comprising high-skill workers. These results are then in line with Implication 1, which states that the treatment should result in an increase in the probability of employment between high-skill workers and high-type firms, and to a *decrease* in employment between high-skill workers and low-type firms. Finally, in column (4) we include the full set of interactions between worker skills, manager ability and the treatment dummy. This does not add further insights compared to column (3), again confirming that heterogeneity by worker skills is less relevant, while manager ability is the key source of heterogeneity.[51]

Taken together, these results show that the certificates lead to an increase in positive assortative matching between higher ability managers and our positively selected sample of experimental workers. As discussed in Subsection 2.3, our experimental workers have similar education levels to the employees of higher ability managers, but are more educated than those working for low-ability ones. The results in Table 6 therefore show that the certificates also lead to an increase in positive assortative matching on education.

4.2.2. *Sorting between soft skill components and specific managers' needs*

While the workers in our sample are positively selected on soft skills on average, we can exploit variation in the grades on specific soft skills to study whether the treatment leads to an increase in positive assortative matching between workers with high values of specific skills and firms particularly in need of those skills. We proxy firms' needs with various firm characteristics and examine the following dimensions of sorting between worker skills and firm characteristics: (*i*) communication skills and the number of employees; (*ii*) willingness to help others and the number of employees; (*iii*) communication skills and the number of customers per worker. Larger firms may be particularly in need of workers who can communicate effectively and are pro-social, as employees in this context typically work in the same small space on similar tasks, and often in teams. In addition, firms with more customers may particularly value communication skills

[51] To further rule out that manager ability is proxying for other characteristics, Online Appendix Table A7.2 progressively adds interactions between the treatment dummy and other firm and manager characteristics. The stability of the interaction between the high-ability dummy and treatment reassures us that manager ability is the key source of heterogeneity. In particular, these results rule out that risk aversion plays a significant role in explaining the results. This indicates that the stronger reaction of high-ability managers to the certificates is more in line with the sample of workers being positively selected and high-ability managers valuing soft skills more, rather than high-ability managers valuing a reduction in uncertainty per se due to risk aversion.

because in this context there is no clear separation between production and retail space. Therefore, interactions with customers happen directly at the firm premises and can involve employees, as the firm owner might be away, for instance.[52]

The results are provided in Table 7. We regress a dummy for whether the worker was hired by the matched firm on our triple-interaction specification between workers' skills, our proxy for managers' needs for that skill and treatment. Workers are divided into having a high or low value of the specific skill considered (e.g., communication skills), by whether they scored above or below the median. For both communication skills and willingness to help others, scoring above the median corresponds to having a pass grade on that skill (i.e., C or above). Similarly, managers are divided into low and high needs for a given skill by whether the specific firm characteristic related to that skill (e.g., firm size) is above or below the median.

Columns (1)–(3) show that we find significant evidence that the treatment leads workers with higher communication skills and willingness to help others to match with larger firms, and to reallocate away from smaller firms where those skills may be less needed. This is indicated by the positive triple interaction in row (*vi*) and the negative interactions between the high-skill dummy and treatment in row (*iii*). As shown by the coefficients in rows (*i*) and (*v*), there is also some evidence that treated workers with low communication/willingness to help others are more likely to be hired in smaller firms, and less likely to be hired in larger firms. However, the effect on lower ability workers is weaker. This is again consistent with our sample being positively selected, so that treatment effects for workers at the bottom of the distribution are more muted. Finally, column (4) also shows increased sorting when looking at the match between the number of customers per worker and communication skills, as the coefficient on the triple interaction is positive and significant.[53]

The results in Table 7 indicate that heterogeneity within our sample, while limited, is still relevant, and leads to sorting on specific dimensions such as communication skills and firm size. To reconcile these results with the positive treatment effects for the average worker documented in most of the other tables, note again that most workers have a high value of at least one skill: only 2 workers out of the 787 in our sample have a fail grade on every skill, and 88% of workers have a grade or A or B on at least one skill. Therefore, we expect the average worker to have high scores on at least some dimensions valued by at least some firms. This can then explain why we see positive effects on outside options for the average worker. This is also consistent with high-ability managers reacting positively for the average worker—because high-ability managers value soft skills more on average, and the average worker has high grades on at least some dimensions of

[52] Bassi *et al.* (2021) provided direct evidence that in this context workers' specialisation across tasks is limited, teamwork is frequent and interactions with customers happen directly at the firm premises. In particular, their analysis for the carpentry sector in Uganda showed that: (*i*) 80% of orders are placed at the firm premises through walk-ins by customers; (*ii*) the typical employee works on more than half of the production steps and (*iii*) production steps feature direct teamwork between employees in about 20% of cases.

[53] We also examined sorting on trustworthiness and the following proxies for firms' need of trustworthiness: (*i*) presence of expensive machines; (*ii*) (reported) importance of stealing as a constraint and (*iii*) how often the owner is away from the firm premises, as a proxy for the amount of supervision. We do not find any significant evidence of sorting on these dimensions (results not reported). We note that the interpretation of these results is less straightforward however, as lack of trust related to stealing is something that firms may be able to hedge against through protective investments; for instance, if firms with expensive machines have taken actions to secure them, such as locking them, then this could weaken their response to information on trustworthiness, making any predictions on this margin of sorting ambiguous. Indeed, we find a *negative* correlation between the presence of expensive machines and reported episodes of stealing, which is consistent with such protective investments taking place. These results suggest that providing information on skills associated with behaviours that firms can respond to through protective investments might not change hiring decisions.

Table 7. *Impacts on Sorting between Specific Soft Skills and Managers' Needs.*

OLS coefficients; SEs adjusted for two-way clustering in parentheses				
Dependent variable:	Worker was hired by the matched firm			
Worker skill:	Communication skills	Willingness to help others	Communication skills and willingness to help others	Communication skills
Firm characteristic:	Number of employees	Number of employees	Number of employees	Number of customers per worker
	(1)	(2)	(3)	(4)
(*i*) Treatment	0.109*	0.105	0.106*	−0.040
	(0.066)	(0.068)	(0.062)	(0.079)
(*ii*) Worker has pass grade on skill	0.019	0.027	0.055	0.020
	(0.056)	(0.050)	(0.050)	(0.067)
(*iii*) Worker has pass grade on skill × Treatment	−0.152**	−0.151*	−0.162**	−0.067
	(0.076)	(0.079)	(0.075)	(0.093)
(*iv*) Firm characteristic above median	0.124	0.075	0.113*	0.056
	(0.084)	(0.066)	(0.068)	(0.100)
(*v*) Firm characteristic above median × Treatment	−0.214**	−0.064	−0.145	−0.005
	(0.102)	(0.092)	(0.090)	(0.123)
(*vi*) Worker has pass grade on skill ×	0.332***	0.107	0.248**	0.273**
Firm characteristic above median × Treatment	(0.121)	(0.115)	(0.116)	(0.136)
Mean of dep. var. in control group	0.116	0.116	0.116	0.116
Worker and firm controls	Yes	Yes	Yes	Yes
Number of observations (matches)	412	412	412	412

Notes: ***, ** and * denote significance at the 1%, 5% and 10% levels, respectively. Results from the worker first follow-up are reported. The number of observations is lower than in Table 3 due to attrition. SEs are adjusted for two-way clustering (at the level of both the firm and the worker), following Cameron *et al.* (2011). A pass grade is C or above. The median grade is C for all skills considered in this table. The firm characteristics used in these regressions are defined as follows: the number of workers is the number of employees at baseline; the number of customers per worker is the number of customers that placed orders at the business in the seven days prior to the baseline, divided by number of employees. All regressions control for stratification variables (dummies for BRAC branch and sector) and dummies for the month of interview. In addition, all regressions control for the following baseline worker characteristics: age and age squared; female dummy; years of education; duration (in years) of the vocational training program the worker was attending at baseline; dummy for any past work experience. In addition, all regressions include dummies for whether the worker scored on the median or above on each of the five soft skills measured in the baseline assessments. All regressions also control for the following baseline firm characteristics: female owner dummy; age and age squared of the owner; dummy for whether owner attended a VTI in the past; number of employees. Each regression also controls for the variable by which heterogeneous effects are considered, uninteracted with Treatment. All regressions further control for dummies for missing values in each of the independent variables.

soft skills. So, overall, our results are consistent with the information improving sorting—of our positively selected sample of workers in general to higher ability managers (who value soft skills more in general), and of workers with a high value of specific skills to firms with a need for those skills.

4.3. *Impacts on Employment*

Implication 2 states that the impact of the certificates on overall employment of high-skill workers is ambiguous: the certificates increase their probability of employment at high-type firms; however, since these workers also revise upwards their reservation wages, this reduces their employment at low-type firms, who value soft skills less and so are not willing to pay them more. As discussed above, column (1) of Table 6 shows that indeed the certificates do not lead to a change in the overall probability of employment in the matching intervention. We further study treatment effects on employment in the two years post intervention. Table 8 reports the results of worker-level regressions analogous to (2) where we pool observations from the two follow-up surveys. The table shows that treated workers are not significantly more likely to be in wage or self-employment at follow-up (columns (1) and (2)), although the effect on wage employment is positive. Also, there is no significant effect on hours worked in the last job (columns (4)–(7)).[54] The interactions between soft skills and treatment throughout the table show that there is no significant heterogeneity. These results confirm the evidence from Tables 6 and 7 that the certificates change the allocation of workers to firms, but do not significantly increase overall employment, in line with model Implication 2.

Finally, in column (3) of Table 8 we study impacts on the probability of being involved in non-casual employment or education/training, that is, in 'productive' activities. We find a positive and significant treatment effect on this margin for the average worker, corresponding to an 8% increase. This result is driven by a reallocation away from casual work (Table 4, column (5)) and onto education/training (Table 4, column (6)) as well as to a lesser extent wage employment (Table 8, column (1)). In a broad sense, this highlights how the certificates improve not only the sorting of workers to jobs, but also the allocation of labour across sectors.[55]

4.4. *Impacts on Earnings*

The overall number of hires in the matching intervention was low: fewer than fifty workers were hired across the two experimental groups.[56] Such low take-up makes it difficult to study wages at the matched firm. Nevertheless, when using information on weekly earnings at the matched firm, we find that treatment workers earned \$4.05 on average in the first week of employment, compared to \$2.46 in the control group. While this difference in means is not statistically

[54] Conditional on employment, workers in the control group work 50–60 hours a week, so there is not much scope for the treatment to increase the intensive margin of hours worked.

[55] We have very limited information on the characteristics of the firms where workers found employment outside the matching intervention. Specifically, the information on manager ability and number of customers used to study sorting in Tables 6 and 7 is not available for these firms. This limits our ability to study sorting in the end-line surveys. We do however have data on the size of the worker's last employer at follow-up (including the current job). We do not find significant impacts on this outcome, but note that we cannot reject a positive impact on firm size similar to the effect documented for ideal firm size in Table 4 (results not reported).

[56] This result is in line with recent evaluations of matching interventions, such as those by Groh *et al.* (2015), Abebe *et al.* (2020), Alfonsi *et al.* (2020) and Bandiera *et al.* (2021), who all found very limited take-up. Taken together, these results suggest that small firms in developing countries do not face particular challenges in *meeting* workers.

Table 8. *Impacts on Employment and Engagement in Productive Activities.*

OLS coefficients; SEs in parentheses are clustered at the worker level

Dependent variable:	Any work as wage employee in the last week	Any work as self-employed in the last week	Main activity in last week is non-casual employment or education/training	Weekly hours worked in last job as employee		Weekly hours worked in last job as self-employed	
Sample of workers:	All (1)	All (2)	All (3)	All (4)	Employed (5)	All (6)	Self-employed (7)
Treatment	0.029	−0.002	0.048*	1.91	2.18	−1.08	−1.17
	(0.029)	(0.024)	(0.028)	(1.95)	(1.54)	(1.61)	(2.99)
Average soft skills grade (standardised)	−0.020	−0.012	−0.015	−0.183	−0.155	−1.23	1.03
	(0.021)	(0.016)	(0.020)	(1.49)	(1.16)	(1.11)	(2.36)
Average soft skills grade (standardised) × Treatment	0.036	−0.010	0.032	2.72	1.35	−1.29	−3.89
	(0.030)	(0.023)	(0.029)	(2.04)	(1.59)	(1.57)	(3.12)
Mean of dep. var. in control	0.428	0.211	0.600	37.0	61.1	16.0	51.9
Controls for baseline value of outcome	Yes	Yes	Yes	No	No	No	No
Uses data from first follow-up	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Uses data from second follow-up	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of observations	1,350	1,350	1,350	1,347	816	1,349	411

Notes: ***, ** and * denote significance at the 1%, 5% and 10% levels, respectively. Results from the worker follow-ups are reported. All regressions control for stratification variables (dummies for BRAC branch and sector), a dummy for the second follow-up and dummies for the month of interview. In addition, all regressions control for the following baseline worker characteristics: age and age squared; female dummy; years of education; duration (in years) of the vocational training program the worker was attending at baseline; dummy for any past work experience. Since at baseline all workers were enrolled in vocational training and only 1% of them were currently doing any paid work, for the employment outcomes, we consider as baseline value of the outcome the expected probability of employment in the six months after graduation, as reported at baseline. So we control for this variable in columns (1) and (2). All regressions further control for dummies for missing values in each of the independent variables.

significant given the small sample, these results are in line with Implication 3, which states that the certificates should increase wages for high-skill workers, since high-type firms are willing to pay these workers more, and their outside option has increased.

We further probe this finding by running earnings regressions on all workers in the two years post intervention, regardless of the outcomes of the matching intervention. Since very few workers are still employed at the originally matched firms after one year, we interpret any results on earnings as arising mostly from other firms that workers match to later on. Table 9 shows the results. Column (1) confirms that there is no impact on paid employment. In columns (2)–(7) the dependent variable is total earnings in the month before the survey. Column (2) shows that, when all workers are included in the regression, so that those with no earnings are assigned a value of zero, we find an 8% increase in earnings for the average worker, but this is not significant. Columns (3)–(6) report quantile regressions on the full sample, and show that there is a positive and significant impact on earnings at the 75th quantile and above. Figure 5 reports quantile regressions along the entire distribution of earnings, and confirms a positive and significant treatment effect in the upper quartile of the earnings distribution.[57] Finally, since we showed that treatment assignment does not affect selection into paid employment, in column (7) of Table 9 we report impacts conditional on employment at follow-up: we find that the certificates lead to an increase of about \$7 per month for the average worker, a result significant at the 5% level, and corresponding to an 11% increase in monthly earnings, relative to the control mean.[58]

On heterogeneity by soft skills, we note that all interactions between skills and treatment in Table 9 are positive, with the one in column (6) being significant. In panel D of Online Appendix Figure A8 we probe this further by running a non-parametric regression of earnings (conditional on employment) on the average soft skills grade, by treatment group. The figure shows that the earnings gains are significantly larger for workers at the top of the distribution. These results are consistent with Figure 5, which shows that the positive effects are concentrated at the upper quantiles of earnings. Taken together, these results suggest that highly skilled workers gained more from the intervention, which is again in line with the information strengthening positive assortative matching between workers and jobs. The fact that the earnings gains are larger at the top of the distribution also explains why we find significant impacts on earnings conditional on employment, but not on unconditional earnings: since there is no impact on involvement in paid employment and no effect at the lower quantiles of earnings, including unemployed workers in the earnings regression just dilutes the positive effect found further up the distribution.[59]

Taken together, the results from Tables 4, 6, 7 and 9 give support to the claim that the increase in earnings is due to workers transitioning to more productive employment: Table 4 shows that workers reallocate away from poorly paid casual work, they increase their investment in skills and they search more ambitiously in the two years post intervention, which suggests that they are transitioning to more productive jobs. Tables 6 and 7 further substantiate this claim by showing

[57] The x axis in the figure starts at the 20th quantile since about 25% of the observations have zero earnings.

[58] Column (1) of Table 9 shows that the treatment does not change the share of workers employed. In the Supplemental Material, we show that the sample of workers employed at follow-up remains well balanced on baseline characteristics, across treatment and control groups (Table S6). This limits potential concerns about the interpretation of results conditional on employment (Lee, 2009; Attanasio *et al.*, 2011), since there is little evidence that the certificates affect selection into employment in our case.

[59] The Supplemental Material shows that the earnings impacts are robust to (*i*) applying a log transformation; (*ii*) excluding control variables; (*iii*) using inverse probability weights to correct for attrition (Wooldridge, 2010), and that we do not find significant evidence of the impacts on earnings varying across follow-ups (Table S7).

Table 9. *Impacts on Labour Market Earnings.*

OLS coefficients in columns (1), (2) and (7); quantile regression coefficients in columns (3)–(6)
SEs clustered at the worker level in parentheses in columns (1), (2) and (7)
Bootstrapped SEs in parentheses in columns (3)–(6)

Dependent variable:	Any paid work in the last month	Total earnings in the last month (USD)					
Specification:	OLS (1)	OLS (2)	Q(50) (3)	Q(25) (4)	Q(75) (5)	Q(90) (6)	OLS, conditional on any paid work (7)
Treatment	−0.014	3.72	1.27	0.398	7.71*	12.9**	7.10**
	(0.025)	(3.20)	(3.26)	(1.27)	(4.10)	(6.25)	(3.56)
Average soft skills grade (standardised)	−0.013	−0.866	−1.48	−0.085	−1.55	−5.40	−0.654
	(0.018)	(1.95)	(2.06)	(1.06)	(2.67)	(4.01)	(2.08)
Average soft skills grade (standardised)	0.010	3.03	1.48	0.097	3.46	11.0*	3.78
× Treatment	(0.025)	(3.10)	(3.06)	(1.34)	(3.71)	(6.62)	(3.45)
Mean of dep. var. in control group	0.750	47.2	47.2	47.2	47.2	47.2	63.1
Controls for baseline value of outcome	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Uses data from first follow-up	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Uses data from second follow-up	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of observations	1,338	1,329	1,329	1,329	1,329	1,329	988

Notes: ***, ** and * denote significance at the 1%, 5% and 10% levels, respectively. Results from the worker follow-ups are reported. In the quantile regressions SEs are bootstrapped (with six hundred replications). All regressions control for stratification variables (dummies for BRAC branch and sector), a dummy for the second follow-up and dummies for the month of interview. In addition, all regressions control for the following baseline worker characteristics: age and age squared; female dummy; years of education; duration (in years) of the vocational training program the worker was attending at baseline; dummy for any past work experience. In column (1) the dependent variable is a dummy for whether the worker conducted any paid work in the month prior to survey. In columns (2)–(7) the dependent variable is total labour earnings in the month prior to the survey. This variable is set to zero for workers with no labour earnings. Since at baseline all workers were enrolled in vocational training and only 1% of them were currently doing any paid work, for the employment outcomes, we consider as baseline value of the outcome the expected probability of employment in the six months after graduation, as reported at baseline. So we control for this variable in column (1). Similarly, in columns (2)–(7) we consider as baseline value of the outcome expected earnings at baseline. Monetary amounts are deflated and expressed in terms of the price level in January 2015, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted in January 2015 USD. The top 1% values of total earnings in the last month are excluded in columns (2)–(7). All regressions further control for dummies for missing values in each of the independent variables.

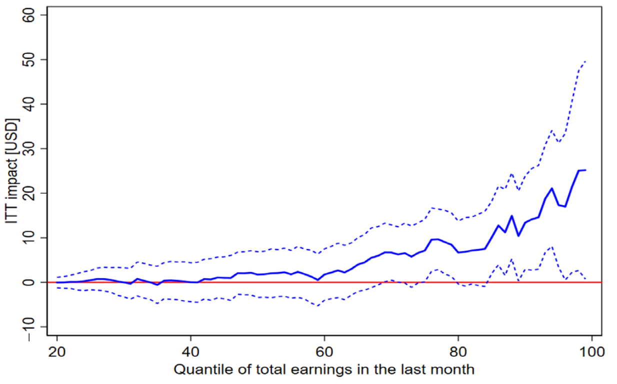

Fig. 5. *Quantile Treatment Effects on Labour Market Earnings.*
Notes: The figure reports quantile regression estimates of treatment effects on total labour earnings in the month prior to the survey, with 90% confidence intervals. SEs are bootstrapped (with six hundred replications). The sample includes all workers from first and second follow-ups. The regressions control for stratification variables (dummies for region and sector), a dummy for the second follow-up and dummies for the month of interview. In addition, all regressions control for the following baseline worker characteristics: a dummy for whether the worker had a pass grade (C or above) on all five soft skills measured in the baseline assessments and disclosed on the certificates; age and age squared; female dummy; years of education; duration (in years) of the vocational training program the worker was attending at baseline; dummy for any past work experience; expected earnings at baseline. Monetary amounts are deflated and expressed in terms of the price level in January 2015, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted in January 2015 USD. The top 1% values of total earnings are excluded.

that in the matching intervention, more productive managers and those in particular need of specific soft skills are those that respond to the certificates and increase their hiring rates. Since we have shown that our sample of workers is positively selected on skills, these results are then consistent with the provision of information resulting in an increase in sorting between workers and jobs, and with higher wages as a result.[60]

4.5. *Discussion*

While the results on sorting, employment and earnings can be explained by our model, we highlight three points on how they fit with our initial priors. First, the positive effects on earnings for the average worker were not necessarily expected: sensible priors would have been that only workers with relatively high skills gain from the certificates. This result can be explained by the positive selection of workers into the experiment and the low managers' priors that we have documented. Thinking about external validity, the fact that we document impacts even for the average worker likely generalises to other settings where recruiters would plausibly find it difficult to understand the selection of candidates in the applicant pool, such as in developing-country

[60] The earnings results cannot be explained solely by the impact on human capital accumulation documented in Table 4, and so must also reflect a change in the allocation of labour: assuming a return to education of 7.5% per year (taken from the Mincerian regression in Online Appendix Table A3) and assuming that the 3.8pp increase in the probability of further enrolment in education/training documented in column (6) of Table 4 corresponds to two full extra years of education, the treatment effect on skills accumulation can then explain an earnings impact of around 1% (i.e., $0.038 \times (7.5\% \times 2) \approx 1\%$), while the impact in column (7) of Table 9 is 11%.

labour markets dominated by small firms with limited experience hiring workers. On the other hand, we would not necessarily expect the certificates to have an effect for the average applicant in settings with larger and more sophisticated recruiters, who may have less biased beliefs about the underlying ability distribution of applicants (Autor and Scarborough, 2008).

Second, the increase in sorting is in line with sensible priors based on standard sorting models. Third, the lack of positive effects on overall employment was also not necessarily expected. This result highlights that the primary impact of the certificates is on worker *reallocation* across jobs, rather than on the extensive margin of overall employment. This can be explained by the relatively high employment rates in the control group, so that any scope for impacts on overall employment is limited in this context. A larger impact could be expected in settings with more disadvantaged job seekers and lower employment rates, such as those in the related studies by Abebe *et al.* (2021) and Carranza *et al.* (2020).

On this last point, we further note that column 1 of Online Appendix Table A8.2 shows that we do find a positive and significant treatment effect on wage employment for the sample of workers who met at least one firm in the matching intervention, corresponding to a 17% increase over the control mean. This result confirms that the impacts are stronger for the matched sample, which is in line with Online Appendix Table A8.1, and further helps to reconcile the findings on employment with our initial priors.

5. Policy Implications and Conclusion

5.1. *Policy Implications*

As this was a relatively inexpensive intervention, with a cost per worker of \$19, we find that the certificates were cost effective at increasing earnings for those workers who found employment, even if we assume that the gains only lasted for the two-year study period.[61] This then raises the question of why soft skill certificates are not already provided by the market. We can rule out lack of demand by workers. At first follow-up we asked workers in control—who never saw the results of their skill assessments—for their willingness to pay for certificates similar to the treatment ones. We find that workers would be willing to pay \$18 on average for the certificates, corresponding to 44% of their monthly earnings. So their willingness to pay is substantial, and interestingly, very similar to the cost of the certificates (\$19).[62]

It is possible however that credibility concerns explain why there is no firm assessing workers on soft skills and selling this information to the market, as the profitability of this activity clearly relies on building a reputation for providing truthful information. We were able to overcome credibility concerns because BRAC is the largest NGO in Uganda and has a strong reputation,

[61] Details of the cost-benefit analysis are given in Online Appendix D. We only consider the costs of the certification intervention. So we do not include here the costs of matching workers to firms. We assume zero impacts on profits, as we do not have reliable profit data. If the intervention produced productivity gains through the improved allocation of labour, then setting these to zero creates a lower bound to the benefits on the firm side. While we find no impact on firm size at follow-up, Online Appendix Table A10 shows some evidence that higher ability owners in the treatment group report lower challenges in screening workers (column 2), and also revise upwards their ideal size (column 8), suggesting possible long-lasting benefits to firms. Providing more direct evidence on the firm-side impacts of certification interventions remains an important area for future research.

[62] Willingness to pay was elicited using a 'take-it-or-leave-it' (TIOLI) approach (Berry *et al.*, 2020): workers were presented with options to purchase the soft skills certificate at a series of prices. We use the highest price indicated by workers as a measure of the lower bound of their willingness to pay. These results are in contrast to Abel *et al.* (2020), who showed that lack of demand by workers limits usage of reference letters by disadvantaged job seekers in South Africa.

but a new market entrant might take years to establish credibility, increasing the entrepreneurial risk of this activity. Finally, we note that providing soft skill certificates might also not be profit maximising for VTIs: as discussed, about 20% of workers decided not to participate in the intervention, and this is consistent with them realising they would not have benefited from it. Therefore, if vocational institutes started advertising new certificates on soft skills, this might affect enrolment decisions in the first place, potentially reducing VTI profits.

In summary, it is unclear that any private agent has enough incentives to create and sell certificates on soft skills, even though these are valuable to at least some firms and workers. When thinking about the incentives of the government to produce this kind of information, our reduced-form results and our model show that, while more information on skills can improve the allocation of labour and productivity, it is possible that some workers with low skills might lose out from this, thus also increasing wage inequality among workers.[63] The willingness of the government to step in and provide this information would then depend on political economy considerations regarding the trade-off between equity and efficiency in the labour market.

Finally, we highlight four important considerations about scaling up this type of intervention. First, in terms of feasibility, while there might be concerns related to the scalability of trust games or psychometric assessments (as workers might learn to 'game' these tests over time), teacher surveys seem a scalable alternative. Second, making participation voluntary in a large-scale certification program might produce equivalent results to introducing a mandatory policy due to unravelling effects (Jin and Leslie, 2003), as not presenting a soft skills certificate at a job interview would then be perceived by firms as a negative signal. Third, the total effects of this type of intervention also depend on any additional changes in the hiring behaviour of firms. We provide tentative evidence on this in Online Appendix Table A10. There is some evidence that treatment assignment reduces the perceived constraints in screening soft skills for higher ability managers at follow-up (column 2), and this is associated with an increase in the duration of the typical interview (column 4). While these results are imprecise due to the small sample of firms, they suggest that the total impacts of signalling interventions could extend beyond their direct effect through additional improvements in the recruitment behaviour of firms coming, for example, from longer interviews as firms learn about the importance of screening on soft skills.[64] Finally, perhaps the most important consideration about scaling up is that such policies would provide an incentive for new generations to increase their investment in skills, thus potentially weakening the trade-off between efficiency and equity in the long run. The positive treatment effect on human capital accumulation documented in Table 4 is particularly revealing in this respect, and so this is an area that deserves more attention in future research.[65]

[63] We formalise this in Online Appendix Figure A9, where we perform a bounding exercise, and show how the cost-benefit calculations for the average *eligible* worker change as those 22% of workers that initially selected out are allowed to experience a (conjectured) reduction in earnings from participation in the program. We find that if those workers that selected out experienced a loss of less than 20% of their earnings, the intervention would still be cost effective for the average *eligible* worker. However, if their earnings losses were larger than 20% then the benefits/cost ratio would fall below one. In the limit case in which all these workers became unemployed as a result of the intervention, the benefits/cost ratio would show very negative results.

[64] Abebe *et al.* (2020) showed that firms in Ethiopia invested more in recruitment after participating in a job fair that produced very few hires.

[65] If improved information on skills boosts aggregate demand for labour then this would provide an additional channel weakening the efficiency-equity trade-off of providing information at scale.

5.2. *Conclusion*

This paper studies how the lack of credible information on the skills of workers affects job matching in a developing country. We do so using a field experiment that aims to understand how employers and job seekers react to certificates on the non-cognitive skills of workers at recruitment. Our main finding is that both sides of the market respond to the certificates in terms of beliefs, and that this improves the allocation of labour, as firms are better able to screen productive workers, and workers are better able to signal their skills in the market. Consistent with better matching, earnings increase in our sample of employed workers. Taken together, our findings highlight that: (*i*) the lack of information on the skills of workers at recruitment creates a significant friction that contributes to keeping wages and productivity low; (*ii*) a credible signal from a trusted institution can significantly reduce the information friction. We believe that these results have important implications for the design of labour market policies in the developing world.

In this paper we have taken a first step towards understanding how both sides of the labour market are affected by information frictions at recruitment. Looking ahead, a promising extension would be to understand the general equilibrium effects of scaling up this type of information intervention. This would require a randomisation at the regional level, introducing certificates in some local labour markets but not others. While the challenges of implementing such a design would be substantial due to the high spatial mobility of workers in developing countries, this type of study would generate important directives for labour market policy, and so is something worth attempting in future research.

University of Southern California, USA
BRAC NGO, Uganda

Additional Supporting Information may be found in the online version of this article:

Online Appendix
Replication Package

References

Abebe, G., Caria, S., Fafchamps, M., Falco, P., Franklin, S. and Quinn, S. (2020). ‘Matching frictions and distorted beliefs: Evidence from a job fairs experiment’, Working Paper, Stanford University.

Abebe, G., Caria, S., Fafchamps, M., Falco, P., Franklin, S. and Quinn, S. (2021). ‘Anonymity or distance? Job search and labor market exclusion in a growing African city’, *Review of Economic Studies*, vol. 88(3), pp. 1279–310.

Abel, M., Burger, R. and Piraino, P. (2020). ‘The value of reference letters—experimental evidence from South Africa’, *American Economic Journal: Applied Economics*, vol. 12(3), pp. 40–71.

Adhvaryu, A., Bassi, V., Nyshadham, A. and Tamayo, J. (2020). ‘No line left behind: Assortative matching inside the firm’, Technical Report w27006, National Bureau of Economic Research.

Adhvaryu, A., Kala, N. and Nyshadham, A. (2018). ‘The skills to pay the bills: Returns to on-the-job soft skills training’, Technical Report w24313, National Bureau of Economic Research.

Alfonsi, L., Bandiera, O., Bassi, V., Burgess, R., Rasul, I., Sulaiman, M. and Vitali, A. (2020). ‘Tackling youth unemployment: Evidence from a labor market experiment in Uganda’, *Econometrica*, vol. 84(6), pp. 2369–414.

Altmann, S., Falk, A., Jäger, S. and Zimmermann, F. (2018). ‘Learning about job search: A field experiment with job seekers in Germany’, *Journal of Public Economics*, vol. 164, pp. 33–49.

Altonji, J.G. and Pierret, C.R. (2001). ‘Employer learning and statistical discrimination’, *The Quarterly Journal of Economics*, vol. 116(1), pp. 313–50.

Attanasio, O., Kugler, A. and Meghir, C. (2011). ‘Subsidizing vocational training for disadvantaged youth in Colombia: Evidence from a randomized trial’, *American Economic Journal: Applied Economics*, vol. 3(3), pp. 188–220.

Autor, D.H. and Scarborough, D. (2008). ‘Does job testing harm minority workers? Evidence from retail establishments’, *The Quarterly Journal of Economics*, vol. 123(1), pp. 219–77.

Bandiera, O., Bassi, V., Burgess, R., Rasul, I., Sulaiman, M. and Vitali, A. (2021). 'The search for good jobs: Evidence from a six-year field experiment in Uganda', Working Paper, University College London.

Bassi, V., Muoio, R., Porzio, T., Sen, R. and Tugume, E. (2021). 'Achieving scale collectively', Technical Report w28928, National Bureau of Economic Research.

Beaman, L. and Magruder, J. (2012). 'Who gets the job referral? Evidence from a social networks experiment', *The American Economic Review*, vol. 102(7), pp. 3574–93.

Belot, M., Kircher, P. and Muller, P. (2019). 'Providing advice to job seekers at low cost: An experimental study on on-line advice', *Review of Economic Studies*, vol. 86(4), pp. 1411–47.

Berry, J., Fisher, G. and Guiteras, R.P. (2020). 'Eliciting and utilizing willingness-to-pay: Evidence from field trials in Northern Ghana', *Journal of Political Economy*, vol. 128(4), pp. 1436–73.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D. and Roberts, J. (2013). 'Does management matter? Evidence from India', *The Quarterly Journal of Economics*, vol. 128(1), pp. 1–51.

Bloom, N., Mahajan, A., McKenzie, D. and Roberts, J. (2010). 'Why do firms in developing countries have low productivity?', *American Economic Review*, vol. 100(2), pp. 619–23.

Bloom, N., Sadun, R. and Reenen, J.V. (2012). 'The organization of firms across countries', *The Quarterly Journal of Economics*, vol. 127(4), pp. 1663–705.

Bowles, S. and Gintis, H. (2011). *Schooling in Capitalist America: Educational Reform and the Contradictions of Economic Life*, Chicago, IL: Haymarket Books.

Bowles, S., Gintis, H. and Osborne, M. (2001). 'Incentive-enhancing preferences: Personality, behavior, and earnings', *The American Economic Review*, vol. 91(2), pp. 155–8.

Brown, M., Setren, E. and Topa, G. (2016). 'Do informal referrals lead to better matches? Evidence from a firm's employee referral system', *Journal of Labor Economics*, vol. 34(1), pp. 161–209.

Burks, S.V., Cowgill, B., Hoffman, M. and Housman, M. (2015). 'The value of hiring through employee referrals', *The Quarterly Journal of Economics*, vol. 130(2), pp. 805–39.

Cameron, A.C., Gelbach, J.B. and Miller, D.L. (2011). 'Robust inference with multiway clustering', *Journal of Business and Economic Statistics*, vol. 29(2), pp. 238–49.

Carranza, E., Garlick, R., Orkin, K. and Rankin, N. (2020). 'Job search and hiring with two-sided limited information about workseeker's skills', Working Paper, Duke.

Caselli, F. (2005). 'Accounting for cross-country income differences', *Handbook of Economic Growth*, vol. 1, pp. 679–741.

Chade, H. and Eeckhout, J. (2017). 'Stochastic sorting', Technical Report, Society for Economic Dynamics.

Chaudhury, N., Hammer, J., Kremer, M., Muralidharan, K. and Rogers, F.H. (2006). 'Missing in action: Teacher and health worker absence in developing countries', *The Journal of Economic Perspectives*, vol. 20(1), pp. 91–116.

De Mel, S., McKenzie, D. and Woodruff, C. (2019). 'Labor drops: Experimental evidence on the return to additional labor in microenterprises', *American Economic Journal: Applied Economics*, vol. 11(1), pp. 202–35.

Deming, D.J. (2017). 'The growing importance of social skills in the labor market', *The Quarterly Journal of Economics*, vol. 132(4), pp. 1593–640.

Dustmann, C., Glitz, A., Schönberg, U. and Brücker, H. (2015). 'Referral-based job search networks', *The Review of Economic Studies*, vol. 83(2), pp. 514–46.

Farber, H.S. and Gibbons, R. (1996). 'Learning and wage dynamics', *The Quarterly Journal of Economics*, vol. 111(4), pp. 1007–47.

Glaeser, E.L., Laibson, D.I., Scheinkman, J.A. and Soutter, C.L. (2000). 'Measuring trust', *The Quarterly Journal of Economics*, vol. 115(3), pp. 811–46.

Groh, M., McKenzie, D., Shammout, N. and Vishwanath, T. (2015). 'Testing the importance of search frictions and matching through a randomized experiment in Jordan', *IZA Journal of Labor Economics*, vol. 4(1), p. 1.

Hall, R.E. and Jones, C.I. (1999). 'Why do some countries produce so much more output per worker than others?', *The Quarterly Journal of Economics*, vol. 114(1), pp. 83–116.

Hardy, M. and McCasland, J. (2020). 'Are small firms labor constrained? Experimental evidence from Ghana', Working Paper, University of British Columbia.

Heath, R. (2018). 'Why do firms hire using referrals? Evidence from Bangladeshi garment factories', *The Journal of Political Economy*, vol. 126(4), pp. 1691–746.

Heckman, J.J., Stixrud, J. and Urzua, S. (2006). 'The effects of cognitive and noncognitive abilities on labor market outcomes and social behavior', *Journal of Labor Economics*, vol. 24(3), pp. 411–82.

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

Hsieh, C.-T. and Olken, B.A. (2014). 'The missing 'missing middle'', *The Journal of Economic Perspectives*, vol. 28(3), pp. 89–108.

ILO. (2020). *World Employment and Social Outlook: Trends 2020*, International Labour Organization.

Jin, G.Z. and Leslie, P. (2003). 'The effect of information on product quality: Evidence from restaurant hygiene grade cards', *The Quarterly Journal of Economics*, vol. 118(2), pp. 409–51.

John, O.P. and Srivastava, S. (1999). 'The big five trait taxonomy: History, measurement, and theoretical perspectives', *Handbook of Personality: Theory and Research*, vol. 2(1999), pp. 102–38.

Jovanovic, B. (1979). 'Job matching and the theory of turnover', *The Journal of Political Economy*, vol. 85(5), pp. 972–90.

Kahn, L.B. and Lange, F. (2014). 'Employer learning, productivity, and the earnings distribution: Evidence from performance measures', *The Review of Economic Studies*, vol. 81(4), pp. 1575–613.

Krishnaswamy, N. (2019). 'Worker absence and firm outcomes', Working Paper, University of Southern California.

Lee, D.S. (2009). 'Training, wages, and sample selection: Estimating sharp bounds on treatment effects', *The Review of Economic Studies*, vol. 76(3), pp. 1071–102.

McKenzie, D.J. (2017). 'How effective are active labor market policies in developing countries? A critical review of recent evidence', Working Paper 8011, World Bank Policy Research.

Pallais, A. (2014). 'Inefficient hiring in entry-level labor markets', *The American Economic Review*, vol. 104(11), pp. 3565–99.

Pallais, A. and Sands, E.G. (2016). 'Why the referential treatment: Evidence from field experiments on referrals', *The Journal of Political Economy*, vol. 124(6), pp. 1793–828.

Petrongolo, B. and Pissarides, C. (2001). 'Looking into the black box: A survey of the matching function', *Journal of Economic Literature*, vol. 39(2), pp. 390–431.

Sackett, P.R. and Walmsley, P.T. (2014). 'Which personality attributes are most important in the workplace?', *Perspectives on Psychological Science*, vol. 9(5), pp. 538–51.

Schönberg, U. (2007). 'Testing for asymmetric employer learning', *Journal of Labor Economics*, vol. 25(4), pp. 651–91.

Spence, M. (1973). 'Job market signaling', *The Quarterly Journal of Economics*, vol. 87(3), pp. 355–74.

Srivastava, S., John, O.P. and Gosling, S.D. (2003). 'Development of personality in early and middle adulthood: Set like plaster or persistent change?', *Journal of Personality and Social Psychology*, vol. 84(5), p. 1041.

Syverson, C. (2011). 'What determines productivity?', *Journal of Economic Literature*, vol. 49(2), pp. 326–65.

UBOS. (2011). *Census of Business Establishments, 2010/11*, Uganda Bureau of Statistics.

UBOS. (2014). *Uganda National Household Survey 2012/13*, Uganda Bureau of Statistics.

UN. (2019). *World Population Prospects: The 2019 Revision Population Database*, UN Desa Population Division.

Wooldridge, J.M. (2010). *Econometric Analysis of Cross Section and Panel Data*, Cambridge, MA: MIT press.