Tackling Youth Unemployment: Evidence from a Labor Market Experiment in Uganda*

Livia Alfonsi **Oriana Bandiera** **Vittorio Bassi**
Robin Burgess **Imran Rasul** **Munshi Sulaiman** **Anna Vitali**†

August 2019

Abstract

We design a labor market experiment to compare demand- and supply-side policies to tackle youth unemployment, a key issue in low-income countries. The experiment tracks 1700 workers and 1500 firms over four years to compare the effect of offering workers either vocational training (VT) or firm-provided training (FT) for six months in a common setting where youth unemployment is above 60%. Relative to control workers we find that averaged over three post-intervention years, FT and VT workers: (i) enjoy large and similar upticks in sector-specific skills, (ii) significantly improve their employment rates, and, (iii) experience marked improvements in an index of labor market outcomes. These averages, however, mask differences in dynamics: FT gains materialize quickly but fade over time, while VT gains emerge slowly but are long-lasting leading VT worker employment and earning profiles to rise above those of FT workers. Estimating a job ladder model of worker search reveals the key reason for this: VT workers receive significantly higher rates of job offers when unemployed thus hastening their movement back into work. This likely stems from the fact that the skills of VT workers are certified and therefore can be demonstrated to potential employers. Tackling youth unemployment by skilling youth using vocational training pre-labor market entry, therefore appears to be more effective than incentivizing firms through wage subsidies to hire and train young labor market entrants. *JEL Classification: J2, M5.*

*We gratefully acknowledge financial support from the Mastercard Foundation, PEDL, the International Growth Centre and an anonymous donor. We thank Jerome Adda, Orazio Attanasio, Chris Blattman, Richard Blundell, Nick Cerkez, Thomas Dohmen, Dave Donaldson, Armin Falk, Marc Gurgand, John Hardman-Moore, Eliana La Ferrara, Joe Kaboski, Daniel Keniston, Thibaud Lamadon, Thomas Le Barbanchon, Ethan Ligon, Marco Manacorda, David McKenzie, Costas Meghir, Karen Macours, Dilip Mookherjee, Suresh Naidu, Derek Neal, Emily Nix, Jeffrey Nugent, Fabien Postel-Vinay, Barbara Petrongolo, Mark Rosenzweig, Johannes Schmieder, T.Paul Schultz, John Strauss, Miguel Urquiola, Gerard van der Berg, Eric Verhoogen, Alessandra Voena, Chris Woodruff, Fabrizio Zilibotti and numerous seminar participants for comments. All errors are our own.

†Alfonsi: BRAC, livia.alfonsi@berkeley.edu; Bandiera: LSE, o.bandiera@lse.ac.uk; Bassi: USC, vbassi@usc.edu; Burgess: LSE, r.burgess@lse.ac.uk; Rasul: UCL, i.rasul@ucl.ac.uk; Sulaiman: BRAC, Save the Children, munshi.slmn@gmail.com; Vitali: UCL, anna.vitali.16@ucl.ac.uk.

1 Introduction

Youth unemployment is a major challenge in the developing world. A growing mass of unskilled, young workers are failing to find work in manufacturing and service sectors consisting mainly of small-scale firms. This raises two questions about labor markets in these countries. First, on the supply side, why don't workers acquire the skills that can help them secure jobs? Second, on the demand side, what prevents firms from hiring these workers? Answering these questions is important – how development proceeds in the coming decades will be largely determined by whether or not these young workers can be employed in good jobs.

Nowhere is the youth unemployment challenge more keenly felt than in East Africa where the majority of the population is aged below 25, and youth represent 60% of the unemployed. We study interventions to tackle youth unemployment in urban labor markets in Uganda, the country with the second lowest median age in the world where 60% of the population is aged below 20, and where formal sector youth employment rates are below 30% [UN 2017].

To do this we design a two-sided experiment involving both workers and firms which allows us to compare supply- and demand-side interventions – vocational training and firm-provided training through apprenticeships – commonly used across the world to help workers transition into the labor market. As the vocational training and firm-provided training interventions are fielded in the same setting we can directly compare their impacts on workers. To investigate the mechanisms underpinning the effects we observe, we use the experimental results to estimate a job ladder model with treatment specific transition parameters. This is our core contribution.[1,2]

Both interventions are designed to improve skills, but they do so by relaxing different constraints. On the supply side, subsidized vocational training might help workers overcome credit market imperfections which prevent them from investing in skills or imperfect knowledge regarding the return to different skills [Jensen 2010]. Moreover, vocational training formally provided by vocational training institutes (VTIs) gives workers *certified* skills, showing which sector-specific

[1] Earlier studies have often evaluated a combination of in-class vocational and on-the-job training, e.g. JTPA in the US and the YTS in the UK. In low-income settings, Card *et al.* [2011] and Attanasio *et al.* [2011] both evaluate the impacts of combining three months of vocational training followed by three month apprenticeships, in the Dominican Republic and Colombia respectively. Lalonde [1995] and Heckman *et al.* [1999] survey the earlier literature on job training programs. On-the-job training, internships and wage subsidies are all common policy approaches that have been used to target disadvantaged groups in the labor market [Layard and Nickell 1980, Katz 1998]. The justification for such approaches are typically twofold [Ham and Lalonde 1996, Katz 1998, Bell *et al.* 1999, Blundell 2001]: (i) to reduce employer screening costs; (ii) to provide workers some labor market experience that can have persistent impacts. On the first channel, Autor [2001] and Hardy and McCasland [2017] present evidence on the use of apprenticeships as screening technologies. On the second channel, Pallais [2014] shows via an experiment on an online jobs platform that providing employment to an inexperienced worker helps improve their later employment outcomes, emphasizing that early labor market experiences convey information on the workers skills rather than raising productivity.

[2] Training programs like those evaluated here are popular interventions in both low- and high-income settings to assist workers transition into employment. The World Bank, for example, has invested $9bn in 93 such programs between 2002 and 2012 [Blattman and Ralston 2015]. Expanded training programs were the most common type of labor market policy implemented globally in response to the 2008 financial crisis [McKenzie and Robalino 2010].

skills they were trained in. This ameliorates adverse selection and can enhance the labor mobility of vocationally trained workers as long as there are firms willing and able to hire them.[3] If these do not exist, only a policy that relaxes firms' hiring constraints will increase employment rates.

On the demand side, subsidized apprenticeships might help firms overcome credit market imperfections which prevent them from incurring the costs of hiring and training workers [de Mel *et al.* 2019, Hardy and McCasland 2017] or of learning about the ability and match quality of inexperienced workers [Farber and Gibbons 1996, Altonji and Pierret 2001, Pallais 2014]. However, firms cannot avoid the time costs associated with training a new worker which may be large in a context where firms are small and where much of the training must be done by the firm owner.

Our research design provides evidence on these elements. Workers in our study are disadvantaged youth entering the labor market. On the demand side, we have small and medium size enterprises (SMEs) in both manufacturing and service sectors. These SMEs represent a core segment of firms in Uganda. We track 1700 workers and 1500 firms over four years, after randomly assigning workers to either control, vocational training (VT) for six months, or firm training (FT) for six months.

To evaluate treatment effects on skills we develop a sector-specific skills test together with skills assessors in Uganda. Our first finding is that, two to three years post-intervention, workers who have received training have accumulated sector specific skills (equivalent to a 30% or $.4sd$ increase over control workers). The magnitude of the improvement is almost identical across both VT and FT workers ($p = .904$). This finding is important because it helps to shut down one potential difference between the treatments.[4]

Our second finding is that there is substantial divergence in compliance: 68% of workers assigned to VT start this training, but only 24% of workers assigned to FT do. This gap is driven by firm, rather than worker, characteristics. In common with earlier studies, firm interest is a key limiting factor [Groh *et al.* 2016], here due to the fact that training was a time- and resource-costly requirement on firms particularly for firm owners as their involvement in the training of apprentices was monitored and enforced. This feature of the experiment is policy relevant when thinking through supply- and demand-side policies to tackle youth unemployment. It also shapes how we estimate treatment effects, the structural model, and how we conduct the IRR analysis. Throughout, we examine both ITT and ATE estimates for worker's labor market outcomes. ITT, because, by averaging over compliers and non-compliers, they reflect likely binding challenges to scaling-up VT and FT interventions in the same context, or of exporting them to other contexts.

[3] Evidence of the value of certification in labor markets is provided by Pallais [2014], MacLeod *et al.* [2015], Abebe *et al.* [2018], Bassi and Nansamba [2019], Abel *et al.* [2019], and Carranza *et al.* [2019].

[4] Our setting departs from the standard Beckerian framework in two ways: (i) we subsidize the apprenticeships through the wage subsidy, making firms more willing to provide skills that are not firm-specific; (ii) firms are contractually required to provide sector-specific skills to workers. This form of apprenticeship – where firms do not bear the full cost of training but are contractually obliged to train workers and monitored in doing so – is worthy of study because it is a training policy that can be replicated and scaled-up by government.

ATE, because by focusing on compliers, they map closely to theories of training and enable us to examine the channels via which VT and FT effects differ for trained workers.

Our third finding is that both treatments improve an index of worker labor market outcomes, that combines employment, total labor supply and earnings. Due to differences in compliance the ranking of the two treatments depends on whether we look at ITT or ATE but in both cases we fail to reject the null of equality. Indeed, the ATE estimates show increases by $.473sd$ for FT workers and by $.272sd$ for VT workers ($p = .202$) while ITT show increases by $.105sd$ for FT workers and by $.170sd$ for VT workers ($p = .169$).[5]

These similarities in averages mask differences in dynamic responses across treatments. On the dynamics of quarterly employment, FT workers find employment more quickly than VT workers, but over time their employment rate converges to the control group, while employment rates for VT workers increase over time.[6] This reversal of fortune between FT and VT workers is also found for earnings – FT workers do well initially but then over time, their earnings fall behind those of VT workers. VT workers steadily increase their earnings and diverge away from the control group. In other words, the similarity in ATE treatment effects between FT and VT workers when averaged over post-intervention survey waves is driven by the earlier quarters in which FT workers were hired by firms incentivized through wage subsidies. Subsequent to that, the patterns of employment and earnings differ across VT and FT workers, with FT workers having similar employment profiles as control workers.[7]

The second part of the analysis builds on this insight. Under the assumption that by endline (three years post-intervention), workers have reached their steady state wage trajectory, we structurally estimate a job ladder model of worker search.[8] This emphasizes three mechanisms driving labor market outcomes: (i) arrival rates of job offers when unemployed or when employed; (ii) job separation rates; (iii) skills. The model allows workers to be heterogeneous in two dimensions: their training (treatments), and their type-ε that determines their productivity on-the-job. We operationalize worker types by linking them to the measurable skills of workers. This form of observed worker heterogeneity provides an avenue for using the model for counterfactual analysis.

[5] In relation to earlier studies that have evaluated a combination of vocational and on-the-job training, Card *et al.* [2011] find no evidence of employment impacts; Attanasio *et al.* [2011] find a 7% increase in employment rates for women and a 20% earnings increase, with impacts being sustained in the long run [Attanasio *et al.* 2017]. Galasso *et al.* [2004], Levinsohn *et al.* [2014] and Groh *et al.* [2016] evaluate wage subsidy interventions. McKenzie [2017] reviews the evidence on training and wage subsidy programs in low-income settings. We later discuss our findings in relation to this literature.

[6] Employment spells are relatively short. For example, among those hired under the FT treatment, only 57% are employed for at least 6 months and so, by endline, almost none of the workers remain in the firm they were originally matched to. This implies that differences in VT and FT employment dynamics must stem in part from their ability to exit unemployment.

[7] Comparing two supply side interventions, Abebe *et al.* [2018] also find that the effect of subsidies – in their case to workers to fund transport – is short-lived while certification has lasting impacts.

[8] Comparing experimental and structural estimates of the returns to training remains rare in the literature. Notable exceptions are Card and Hyslop [2005] and Hoffman and Burks [2019], although the latter compares to quasi-experimental estimates, not those based on random assignment.

The job ladder model estimates reveal that: (i) VT workers have significantly higher steady state rates of unemployment-to-job (UJ) transitions than FT workers: if they fall off the job ladder into unemployment, they are more likely to get back on it; (ii) FT workers have very similar rates of UJ transition as the control group: their history of labor market attachment seems to count for little if they become unemployed. At the same time, accepted earnings conditional on employment are similar between FT and VT workers, consistent with them having similar skills, but VT workers can more easily certify their skills and climb back onto the job ladder if unemployed. In steady state, unconditional annual earnings of complier VT workers rise by 55% over controls, while the earnings of complier FT workers rise by just over half of that, 32%.

Combining these results gives us a precise interpretation to what drives the dynamic treatment effects: vocational trainees pull away from FT workers in their employment rates and earnings because they are more likely to get back onto the job ladder if they fall into unemployment. These dynamics are not due to greater job-to-job mobility, suggesting the returns to skills certifiability are higher when unemployed than when employed. Moreover, compliers across training routes move as far up the job ladder as each other – wages *conditional* on employment are similar for complier VT and FT workers because their skills are similar. The key distinction is that VT workers are more likely to get back onto the job ladder if they fall off it. Tackling youth unemployment by skilling youth using vocational training pre-labor market entry, therefore appears to be more effective than incentivizing firms through wage subsidies to hire and train young labor market entrants.

To understand how the results derived from the partial equilibrium model map to general equilibrium impact we exploit the two-sided experimental design which allows us to compare the two training routes from the dual perspectives of workers *and* firms. The firm-side of the experiment shows that there are no employment displacement (or crowding in) effects on other workers – in either the short run when wage subsidies are in place, or in the long run long after apprentices have left the firm. This is central to understanding general equilibrium effects of our training interventions.[9]

We use the model parameters to conduct two counterfactual simulations: (i) to assess the relative importance of the mechanisms at the heart of the model in explaining steady state outcomes; (ii) to simulate treatment effects if the training interventions were targeted to other workers in the economy; in particular, by drawing on data from other studies in the same context, we simulate what would have been the impacts of targeting our treatments to workers that are able to self-finance vocational training, and that are already employed in SMEs in manufacturing sectors similar to those in this current study.

We then combine program accounting costs with estimated steady state earnings benefits to derive the internal rate of return from each treatment. Assuming gains last 15 years, the IRR to vocational training is 24%, while the IRR for firm training is negative. It does *not* pay for

[9] Crepon and Premand [2019] design a labor market experiment in Cote d'Ivoire to quantify displacement effects of a subsidized apprenticeship program, and also find no evidence of displacement.

the social planner to replicate the kind of subsidized apprenticeship offered in the FT treatment. However, the reason for this negative IRR is the low compliance in the FT treatment: only 24% of workers assigned to this treatment end up being hired by firms they are matched to.[10] In these labor markets we do, however, observe workers paying firms for apprenticeships using precisely the kind of payment structure we set up in the FT treatment. To see why this is so, we redo the IRR calculations but based on the steady state earnings for compliers – namely those that acquire firm-provided or vocational training. Among this group the IRR for VT workers rises to 34%, and for FT workers the IRR is 26%. The rise in IRR for FT workers highlights the high social returns from being able to overcome firm's constraints in taking-on and training young workers. Under these IRRs, both training routes pay for themselves even if the steady state flow benefits last a decade or even less.

This of course raises the question of why workers are not availing themselves of these returns by paying for VT or FT themselves? A key reason may be credit constraints: the cost of vocational training, or of self-financing apprenticeships, are both orders of magnitude higher than young workers earnings at baseline.

Despite their popularity, the evidence base for training programs, or in contrasting alternative training programs in the same context, is thin. The meta-analyses of Card *et al.* [2015], Blattman and Ralston [2015] and McKenzie [2017] show relatively weak or short-lived impacts of training programs in low-income settings. We thus close our analysis by highlighting potential explanations for the impacts we document: the experimental design, the selection of workers, and the quality of the vocational training institutes worked with. On each dimension, we make suggestions for future research, and conclude by discussing the wider implications of our findings for the study of youth unemployment in low-income labor markets.[11]

The paper is organized as follows. Section 2 describes the setting, experimental design and data. Section 3 presents treatment effects on worker skills and labor market outcomes. Section 4 develops the job ladder model. Section 5 presents model estimates. Section 6 presents counterfactuals. Section 7 presents the IRR estimates and discusses external validity. Section 8 concludes. Robustness checks and further estimation details are in the Appendix.

[10] This low compliance is driven by a lack of firms taking-up the offer of the wage subsidy and the matched-to worker (workers are as likely to accept offers from firms as vocational trainees are to accept the offer of training from VTIs). Moreover, we document that less profitable firms are more likely to take-on workers through the FT treatment. This suggests that these firms are financially constrained in hiring young job seekers and that this is an important demand-side constraint.

[11] Card *et al.* [2015] discuss over 800 estimates from 200 studies documenting impacts of active labor market programs. They find in contrast to training or wage subsidies, job search assistance (matching) has comparatively large short and long run impacts that are more pronounced for disadvantaged workers. Search costs have been shown to be relevant for workers and firms in labor markets [Franklin 2016, Hardy and Macasland 2017, Pallais 2014, Abebe *et al.* 2018, Bassi and Nansamba 2019].

2 Setting, Design and Treatments

Our study is a collaboration with the NGO BRAC, who implemented all treatments, and five reputable vocational training institutes (VTIs). The VTI sector in Uganda is well established, with hundreds in operation. Each could offer a standard six-month training courses in eight sectors: welding, motor mechanics, electrical wiring, construction, plumbing, hairdressing, tailoring and catering. These sectors constitute a source of stable wage employment for young workers in Uganda: around 25% of employed workers aged 18-25 work in them.[12]

2.1 Setting

Workers Individuals were recruited into our evaluation as follows. Throughout Uganda, we advertized an offer of potentially receiving six months of sector-specific vocational training at one of the VTIs we collaborated with. As in other training interventions, the eligibility criteria targeted disadvantaged youth [Attanasio *et al.* 2011, Card *et al.* 2011]. We received eligible applications from 1714 individuals whose characteristics are described in Table A1: 44% are women, they are aged 20 on average, and the vast majority have never received vocational training.[13]

The first row of Table 1 shows baseline labor market outcomes for workers in our study: unemployment rates are over 60% for these youth (Columns 2 and 3) with insecure casual work comprising the most prevalent form of labor activity. Unconditionally, average monthly earnings are $6, corresponding to around 10% of the Ugandan per capita income at the time. Hence these are individuals that are unlikely to be able to self-finance investment into vocational training (that costs over $400), or to pay a firm for an apprenticeship. Table A2 compares our sample to those aged 18-25 in the Uganda National Household Survey from 2012/3. The intervention appears well targeted: our sample is worse off in terms of labor market outcomes at baseline, and that remains true when we compare to youth in the UNHS who report being labor market active.[14]

As we lacked funds to send all eligible applicants to vocational training, the experiment uses an oversubscription design. This is informative of the impact of marginally expanding such training.

[12] The VTIs we worked with were: (i) founded decades earlier; (ii) were mostly for-profit; (iii) trained hundreds of workers with an average student-teacher ratio of 10; (iv) in four VTIs, our worker sample shared classes with regular trainees. We derive the share of employed workers aged 18-25 working in these eight sectors using the 2012/3 Uganda National Household Survey.

[13] The program was advertized using standard channels, and there was no requirement to participate in other BRAC programs to be eligible. The eligibility criteria were based on: (i) being aged 18-25; (ii) having completed at least (most) a P7 (S4) level of education (corresponding to 7-11 years); (iii) not being in full-time schooling; (iv) a poverty score, based on family size, assets owned, type of building lived in, village location, fuel used at home, number of household members attending school, monthly wage, and education level of the household head. Applicants were ranked on a 1-5 score on each dimension and a total score computed. A relative threshold score (varying by geography) was used to select eligibles.

[14] Unemployment rates are often difficult to define in low-income contexts where workers might be working informally, receiving payments in kind, and it is hard to measure whether they are actively seeking work. It is thus perhaps more accurate to speak of rates of non-employment (as at least this includes those who are not looking for work), and for expositional ease, this is what we will implicitly have in mind when referring to unemployment.

Given Ugandan demographics, there is no shortage of the kind of disadvantaged youth that applied to our offer. At the same time, our study is targeted to disadvantaged youth, not representative youth. We return to this issue later, using our structural model to simulate counterfactual impacts of our training treatments on alternative groups of workers in the economy.

Firms To draw a sample of firms for the study, we conducted a firm census in 15 urban labor markets. From this census we selected firms: (i) operating in one of the eight manufacturing and service sectors in which we offered sector-specific vocational training; (ii) having between one and 15 employees (plus a firm owner). The first criteria limits skills mismatch in our study. The second restriction excludes micro-entrepreneurs and ensures we focus on SMEs that, as Figure A1 highlights, are central to employment generation in Uganda. We end up with a sample of 1538 SMEs, that employ 4551 workers in aggregate at baseline.

We asked SMEs about constraints to expansion. Prominent explanations related to credit and labor inputs: (i) 65% of firms reported the terms of available finance limiting their growth; (ii) 67% reported access to skilled labor as a constraint; (iii) 52% reported the inability to screen workers as a constraint. The offer of wage subsidies might help relax demand-side constraints on SMEs related to hiring young labor market entrants.

Returns to Vocational Training Table 2 provides evidence on the supply of, and returns to, vocational training in this setting. It shows: (i) the share of workers employed at baseline in these firms that self-report having ever received vocational training from a VTI; (ii) the coefficient on a dummy for this self-report in an otherwise standard Mincerian wage regression of log wages. The first row pools across all sectors and documents that at baseline, 31% of workers in our sample of SMEs have vocational training from some VTI. Vocational training is therefore a common route through which workers acquire skills in Uganda, and SME firm owners are familiar with recruiting workers with such training. The Mincerian wage returns to vocational training are over 50%, and this holds by sector: in each there is demand for these skills and there are potentially high returns to them. Of course the Mincerian returns are not causal, being upwards biased due to selection into employment. Our experimental results shed light on the causal impact of vocational training and quantify the selection bias in these Mincerian returns. This evidence just shows there is demand for vocational training in this setting, and potentially high returns to vocational training in the sectors that SMEs in our study operate in. This is in contrast to high-income settings where many training programs have had low returns or short-lived impacts on workers [Card *et al.* 2015].

Use of Apprenticeships Firm-sponsored training is a second important route through which workers accumulate human capital [Acemoglu and Pischke 1998, 1999, Autor 2001]. Apprenticeships are a common labor contract throughout Sub-Saharan Africa, including Uganda [Hardy and McCasland 2017]. Table 3 provides evidence on such contracts from our sample of SMEs. Panel A

shows that half the workers employed in these SMEs at baseline report having received on-the-job training in their current firm, with an average training duration of 10 months. Panel B shows a variety of payment structures for apprentices: the majority are unpaid, some are paid, and others pay for their training. For those paid during apprenticeships, they report an average monthly wage of $39. Firm owners were asked about the skills composition of apprentices. Using this to examine how these skills vary depending on the payment structure, we note that self-financed apprentices are more likely to be reported to have sector-specific rather than firm-specific skills.[15]

Panel C shows the main opportunity cost to SMEs taking on new hires is the firm owner's time: they are predominantly tasked to train apprentices. This is especially so for self-financed apprentices: 56% report being trained exclusively by the firm owner, and none report other employees training them. In contrast, among paid apprentices 40% are trained only by the firm owners, and the majority are trained by co-workers. The time that owners need to devote to training workers (and particularly those that pay for their apprenticeships) represents a key cost that firms must weigh up in deciding whether or not to take on workers.

Firm owners are likely to have the skills to train their employees: they have significantly more years of education than their workers, are significantly more likely to have received vocational training themselves, and have owned their firm for 6.6 years on average. As mentioned above, the majority of SMEs report an inability to screen workers as a constraint to expansion. Hence, if SMEs are credit constrained, it is these kinds of up-front screening costs, or firm owner's opportunity costs of training new hires, that are reduced in our apprenticeship treatment.[16]

2.2 Design

The left hand side of Figure 1 presents the design of the experiment from a worker's perspective. 1714 workers applied to the offer of vocational training. Focusing on the top branch of Figure 1A, those randomly assigned to vocational training were split into two treatments: the first group completed their six months of training and then transitioned into the labor market. This is the business-as-usual training model, where VTIs are paid to train workers, but not to find them jobs. The second group of vocationally trained workers, upon graduation, were matched to firms operating in the same sector as the worker had been trained in, and located in the same region as

[15] To get a sense of the cost of an apprenticeship we note that: (i) for 52% of all apprentices their main cost is the opportunity cost of labor market opportunities during the apprenticeship as well as fixed costs of work (e.g. travel, tools). (ii) For 29% of workers that pay for their apprenticeship, the average total payment is over $500. Whichever way we might calculate it, the expected cost of an apprenticeship is high and above the annual earnings of our sample of workers at baseline.

[16] The investment by firm owners in training workers is well recognized by employees. For example, from the firm-side surveys, we interviewed employees in our sample of SMEs pre-intervention. We asked them about the role of their firm owner in training workers. In the control group of firms, 79% of employees agreed with the statement, "Does the owner put special effort in training and retaining the best workers?", and when asked, "What do you feel makes it better to work at this firm relative to your competitors, if anything?", 43% of employees reported the better training/learning opportunities. This was the most frequent answer given, the next being higher wages, as reported by 22% of employees.

the worker's residence. Upon graduation, all vocational trainees leave their VTI with a certificate stating which VTI the individual attended, and the six-month training course taken.

Workers not offered vocational training, on the bottom branch of Figure 1A, were randomly assigned to three groups: (i) to be matched to firms; (ii) to be matched to firms and those firms offered a wage subsidy to hire and train them on-the-job for six months (i.e. as an apprentice);[17] (iii) held as a control group. Pairwise comparisons of treatment arms are informative of the returns to vocational training (VT), firm-provided training (FT), as well as labor market search frictions. This design allows us to thus compare and contrast supply- and demand-side interventions designed to raise skills and reduce youth unemployment, and to understand the nature of constraints on workers and firms that prevent such human capital investments being undertaken.[18]

Two other timing-related points are relevant. First, although workers were randomly assigned to treatment at the point of application, they were only informed about any match that might be offered once vocational trainees had completed their courses. This helps avoid lock-in or threat-effects on worker search [Black *et al.* 2003, Sianesi 2004]. Second, the design ensures vocational trainees and firm-trained workers both come into contact with firms at the same time: this is in line with the underlying motivation for our study, to understand labor market transitions of youth. However, inevitably this means that vocational trainees receive their training before firm-trained workers do. This six month divergence in training times is however unlikely to bias estimates based on the three years of follow-up data.[19]

We assign workers to treatment using a stratified randomization where strata are region of residence, gender and education. Table 1 shows the labor market characteristics of workers in each treatment, and Table A1 shows other background characteristics. In both cases, the samples are well balanced, and normalized differences in observables are small.

The right hand side of Figure 1 shows the design from a firm's perspective: firms were randomly assigned to either be matched with: (i) vocationally trained workers; (ii) untrained workers; (iii) untrained workers and given a wage subsidy to hire and train them; (iv) or held as a control.[20]

Timeline Figure 2 shows the study timeline: the baseline worker survey took place from June to September 2012 just after workers applied for vocational training. Eligible workers were tracked

[17] In terms of our discussion of apprenticeship this corresponds to where workers pay firms to be trained.

[18] As Figure 1A shows, the control group was purposively larger than the other groups in anticipation of higher attrition rates. At the point of application, workers provided a preference ranking over their top three sectors to be trained in. For those assigned to vocational training, 95% of them were trained in one of their top-3 sectors.

[19] We show evidence in support of this by exploiting a small second batch of vocational trainees that received their training between October 2013 and April 2014 when apprenticeships were implemented. The primary worker outcomes do not differ between the main and second batch of vocational trainees.

[20] In a companion paper, Bandiera *et al.* [2017], we provide a comprehensive analysis of the firm side impacts of all these treatments (and other treatments), and what light they shed on constraints to expansion that SMEs face. Of relevance for the current analysis is that: (i) firms are balanced on observables across treatments, including on monthly profits, employee numbers, the value of the capital stock, age and owner characteristics; (ii) we find that firms assigned to the wage subsidy treatment are more likely to attrit by the first follow up (but not by endline), and we account for this by weighting observations using inverse probability weights.

in surveys fielded 24, 36 and 48 months after baseline (12, 24 and 36 months after the end of vocational training/apprenticeships).[21] Attrition is low: 13% of workers attrit by the 48-month endline. The Appendix describes correlates of attrition in more detail, confirming attrition between baseline and endline is uncorrelated to treatment.[22]

The lower part of Figure 2 shows the timeline of firm surveys over four post-intervention waves. We use these waves of firm data to compare firm outcomes between those offered the apprenticeship and the control group. In particular we estimate short and long run employment displacement effects of the FT treatment. As we describe in more detail later, these results on (a lack of) displacement help us to extrapolate results from the structural model – that is in partial equilibrium – to understand the general equilibrium impacts of firm-provided training in these urban labor markets.

2.3 Treatments

Vocational Training Vocational training provides workers six months of sector-specific training in one of eight sectors. In treatment arms involving vocational training (T3, T4), BRAC entirely covered training costs, at $470 per trainee. Vocational training lessons were held Monday-Friday, for six hours per day; 30% of the training content was dedicated to theory, 70% to practical work covering sector-specific skills and managerial/business skills. As described above, these are all standard training courses provided by the VTIs we worked with. Upon course completion, trainees graduate with a certificate documenting what they have been skilled in.

VTIs signed contracts with BRAC to deliver training courses to workers. They were monitored by regular and unannounced visits by BRAC staff to ensure workers were present and being trained. For each worker, VTIs were paid half the training fee at the start of training, and half at the end, conditional on them having trained the worker (this staggered timing of payments ensured VT workers nearly always completed the full course of training conditional on enrolling).[23]

Firm Training In the firm-provided training treatment arm, we offered firms the chance to meet untrained workers and receive $50 a month for six months if they hired and trained one such worker on-the-job. This was an inflexible wage subsidy whereby $12/month was to be retained

[21] We surveyed those randomized out of vocational training just as vocational trainees were transitioning into the labor market. The tracker survey had a 23% attrition rate. The work status of respondents were as follows: 19% were currently involved in some work activity, 11% had been involved in a work activity in the last six months (but not on survey date), and 70% had not worked in the last six months.

[22] Our attrition rate compares favorably to other studies such as Attanasio *et al.* [2011] (18%), and Card *et al.* [2011] (38%). Indeed, in the meta-analysis of McKenzie [2017], all but one study have attrition rates above 18%.

[23] The cost per trainee breaks down as the cost to the VTI ($400), plus the worker's out-of-pocket costs during training, such as those related to travel and accommodation ($70). The staggered incentive contract solved drop out problems associated with training programs in low-income settings [Blattman and Ralston 2015]. There was no additional stipend paid to trainees during training, and no child care offered either (recall that around 10% of our worker sample have at least one child).

by the owner, and \$38/month was to be paid to the worker. This differs from the standard Beckerian apprenticeship model in that the firm does not bear the full cost of training. As such, the kinds of skills provided to workers might differ than in other apprenticeship structures. As Table 3 showed, this kind of apprenticeship structure exists in urban labor markets in Uganda (i.e. workers sometimes pay for their own training) and is associated with more intensive involvement of firm owners in the training of workers. Moreover, this represents a labor market policy intervention that could be implemented at scale by government.

To assess whether the level of the wage subsidy is reasonable, we relate to two anchors: (i) Table 3 showed that during apprenticeships in the sample of SMEs, if workers were paid their mean wage was \$39/month; (ii) using the wages of all unskilled workers employed in our SMEs at baseline, our wage subsidy treatment had a subsidy rate (wage subsidy/average wage) of 111% (Figure A2a shows the entire distribution of unskilled wages at baseline among those employed in our SMEs). This is high: for example, de Mel *et al.* [2010, 2019] evaluate a wage subsidy program with a 50% subsidy rate.

This FT treatment was designed as a formal training program, backed by an explicit contractual agreement between firm owners and our implementing partner, BRAC. This was intended to mirror the vocational training program as far as possible, but with training being conducted at a firm rather than at a VTI. The contract between BRAC and firm owners in the FT treatment is shown in Figure A3. The contract is succinct and clearly states: (i) firm owners agree to train the worker in a specific trade for six months; (ii) firm owners would pay back the entire subsidy if the trainee was found not to be receiving training, or was not showing up to the firm regularly. There was no explicit training curriculum: firm owners were free to train workers as they saw fit, as long as sector-specific training was provided (as VTIs were also tasked to do).

As with VTIs, monitoring checks were used to ensure these agreements were adhered to. Firm owners and trainee workers were monitored during the entire six months of training in two ways: (i) each firm was provided an attendance register, and every day both the firm owner and the worker had to sign it, providing the worker's time of arrival and departure; (ii) BRAC staff conducted monthly unannounced spot checks at firm premises to verify the worker was showing up and being trained. Payments were disbursed monthly at the local BRAC office, and both the firm owner and the worker had to be present at disbursement, where they were asked to sign an additional register to certify the worker was showing up regularly at the firm and receiving training. Figure A2b shows worker and firm reports on the wage subsidy being received by the worker, with a clear spike at \$38/month as intended.[24]

[24] In terms of credibility and enforcement, during our study period, BRAC Uganda was operating a large number of programs including a highly successful microfinance program where similar contracts were being written and enforced by BRAC. Hence BRAC staff had extensive experience with program monitoring and enforcement in the Ugandan context.

Matching In the matching treatments firms were presented lists of workers that were: (i) willing to work and vocationally trained (T4); (ii) willing to work but untrained (T2, T5). In case (i), the firms knew what sector the workers had been trained in, where they had been trained, but not that training had been paid for by BRAC. Hence, firms might expect workers presented to them to be similar to those able to self-finance such training. In all treatments involving matching, there were a maximum of two workers presented to firms on a list, and the randomly assigned matches took place with firms operating in the same sector as the worker had been trained in (or had expressed an initial desire to be trained in), and in the same region as the worker was located.

3 Treatment Effects

3.1 Compliance

Workers are observationally equivalent when assigned to treatment. Subsequently, there can be selective non-compliance by workers, and for the treatments involving worker-firm matches, there can also be non-compliance driven by firms because worker-firm matches only occur if both a worker and the firm express a willingness to meet. Table A4 shows worker and firm take-up rates by treatment.

Focusing first on treatments involving vocational training we see that: (i) over 95% of workers that initially apply for vocational training are later found and offered it (Column 1); (ii) 68% of workers take-up the offer of vocational training and complete the training (Column 2); conditional on enrolment, over 94% of them completed the training.[25]

For workers assigned to the FT treatment, 51% of them are actually offered a meeting with a firm (Column 3). In common with earlier studies, firm interest is a key limiting factor preventing worker-firm matches to occur [Groh *et al.* 2016]. The explanation for this low take-up rate in our context is that the FT treatment required firm owners to provide time and resource-costly training which monitored and enforced. We provide additional evidence on this point from process reports collected during the intervention roll-out. Firm owners who were not interested in taking on a worker in the FT treatment were asked why: 62% said the subsidy was not large enough to cover training costs, in line with demand-side credit constraints binding.[26]

However, conditional on the worker-firm match, 80% of meetings take place (Column 4) and 90% of interviewed workers are offered a job (Column 5), and two thirds of job offers are accepted (Column 6). This offer acceptance rate to firm-provided training is very close to the 68% compli-

[25] In the meta-analysis of McKenzie [2017], most studies have training completion rates between 70 and 85%. Among workers that did not take-up the offer: 13% reported not doing so because they had found a job, 8% were in education, 4% were in another form of training, 36% did not take-up for family reasons (they had a child, illness, or family emergency), 15% reported distance as being the main constraint, and 24% reported other reasons.

[26] Much of the earlier literature has emphasized stigma effects driving down take-up by firms, because firms perceive workers with attached subsidies as being of low quality [Bell *et al.* 1999]. This could also drive down firm compliance in our setting.

ance rate for vocational training. The difference in compliance rates between FT and VT is not driven by the share of workers taking-up training offers, but by the share receiving *offers from firms* to provide them training. As a result, only 24% of workers initially assigned to the FT treatment end up being employed and trained at the matched firm.

Firm's lack of interest in meeting workers is most severe in the treatments involving matching (T4, T5): only 13% and 19% of workers end up being offered even a meeting with a firm in these treatments (Column 3). This is not altogether surprising given the context: for example, we note from Table 1 that given youth unemployment rates of around 60% firms should have little difficulty in meeting untrained workers, and as Table 2 showed, around one third of employees in SMEs are vocationally trained and so SMEs might have no difficulty meeting trained workers. In short, there is not much evidence for search frictions related to meeting untrained workers or meeting skilled workers in these labor markets.[27]

Given the low worker-firm matching rate in the vocational training plus match treatment (T4), for the remainder of the analysis we combine these workers with those assigned to the vocational training treatment (T3). Moreover, given the low worker-firm matching rate in the pure matching treatment for untrained workers (T5), we drop this treatment arm for the bulk of the analysis.

This allows us to focus attention throughout on the comparison between vocationally trained (VT) workers (T3 and T4) and firm-trained (FT) workers.

3.1.1 Differential Compliance Between Vocational and Firm-provided Training

There is a divergence in compliance between the VT and FT treatments: 68% of workers assigned to vocational training actually start this training, and 24% of workers assigned to firm-provided training are hired by firms they are matched to. This feature of the experiment shapes how we estimate treatment effects and the structural estimation. It is also highly policy relevant, and will be an important consideration for when we conduct the IRR analysis for each training route.

To understand whether and how outcome differences between VT and FT can be due to lower compliance in FT, we first establish whether FT compliance relates to worker traits, that might in turn determine returns to training. The experimental design has two features that help rule out this possibility. First, eligibility requirements mechanically ensure that individuals in the sample are relatively homogeneous.[28] Second, by design, firms in the FT treatment could only train one of two workers assigned to them. Hence, their ability to choose on worker traits was limited among a relatively homogenous pair of workers presented to them.

[27] For treatments T4 and T5 we note from Table A4 that a high percentage of these matches end up leading to job interviews (Column 4), and fewer of them convert to job offers (Column 5) and acceptances (Column 6). This suggests worker preferences can also be a source of matches not translating into hires [Groh *et al.* 2016].

[28] Table A2 compares our sample to young workers in the nationally representative 2012/13 UNHS data. The distribution of education in our sample is more compressed than for youth more generally (the standard deviation of years of education in our sample is 2.06; it is 3.65 in the UNHS). The same reduced dispersion also applies to other continuous outcomes in Table A2 (age and total earnings from wage employment).

Table A5 provides regression evidence on the correlates of compliance within the FT treatment. Column 1 controls for worker characteristics, and shows that 10 out of 11 of these do not predict take-up (with the other being marginally significant). This is in line with there being limited variation in worker traits presented to firms. Column 2 shows this to be robust when we add program related characteristics (the p-value on the joint test of significance of worker characteristics is .156). In contrast, when firm characteristics are added as covariates, these are jointly significant in predicting compliance ($p = .002$). This is driven by firms with significantly lower profits per worker taking on FT workers. Hence, firms hiring workers when given a wage subsidy incentive appear to be negatively selected. Finally, Column 4 shows that when we control for firm fixed effects (that is possible given multiple workers were presented to the same firm), worker characteristics remain jointly insignificant predictors of compliance ($p = .976$). This reiterates that firms were matched with observationally similar workers in the FT treatment.[29]

This body of evidence suggests compliers are similar in FT and VT treatments because lower compliance in FT is explained by firm, not worker, characteristics. This allows us to compare ATE estimates of both training routes, narrowing down the interpretation of any differences between them as stemming from either the skills imparted, or the certifiability of those skills.

To explore this point further, we provide descriptive evidence on skill accumulation among workers taking up the VT and FT treatments. In the next Section we present formal evidence from estimating treatment effects on various dimensions of skills accumulation.

3.1.2 Descriptive Evidence on Skill Accumulation Among Compliers

Vocational Training The vocational training courses workers were assigned to are the business-as-usual six months sector-specific training that the VTIs we worked with typically offer, and were monitored to do so. We surveyed VT workers towards the end of the training and asked about their satisfaction with it: 76% were extremely happy/very happy with the experience; 86% were extremely happy/very happy with the skills gained; 96% reported skills acquisition as being better than or as expected, and 56% reported that six-months of training was enough time for them to learn the skills they had wanted to.

Firm-provided Training We provide two pieces of evidence to confirm workers hired under the FT treatment were trained. From the first firm follow-up survey, that was deliberately fielded around the end of the six-month apprenticeship period, firm owners were asked to indicate for each employee hired in the last six months: (i) their productivity (on a 1 to 5 scale) when hired; (ii) their productivity at the time of the survey (or when the worker left the firm, if they left earlier).

[29] Again, evidence taken from the process reports further bolsters this. In less than 4% of cases did firms report turning down a worker in the FT treatment because of worker characteristics. Firms interested in taking-on a trainee reported they were happy to take on any of the workers assigned to them, and were not searching for specific trainee characteristics. Moreover, most firms that did not take-on a worker did not even meet any worker matched to them, again consistent with worker characteristics not mattering in firm selection.

We use this to compare the growth in productivity of workers hired as part of the FT treatment to other workers hired in the last six months by firms in our sample.

Panel A of Figure A4 shows the productivity growth of hired FT workers alongside that for: (i) hired workers at control firms, (ii) hired workers at control firms who report having received on-the-job training, (iii) workers at control firms that did not receive on-the-job training; (iv) hired workers in FT firms excluding those hired through our FT treatment. We see that hired FT workers had an increase in productivity of 2.24 points during the wage subsidy period, higher than for all other comparison groups. This descriptive evidence points to workers hired under the FT treatment having received training over and above what these firms would normally provide, in line with their contractual requirements with BRAC, and reinforcing the notion that this was a resource and time-intensive treatment from firms' perspective.[30]

We also use data from the first worker follow-up survey, conducted around six months *after* the wage subsidy expired. For each job spell in the previous year workers were asked to report: (i) their ability to perform a typical sector-specific task at the start of the spell (e.g. if the job spell was in motor-mechanics, they were asked whether they were able to mend a tire tube); (ii) their ability to perform the same task at the end of the job spell (or on survey date if the spell was ongoing). The same questions were asked about an important firm-specific task at the firm where the worker was employed (the worker indicated this task). Firm owners were asked the same questions about each employee in the first firm follow-up. We compare the rate of learning on these tasks at the matched firm for workers hired under the FT treatment. This comparison is in Panel B of Figure A4. Reassuringly, worker and firm reports are well aligned in both cases, and both show the rate of learning was substantial for both sector-specific and firm-specific tasks.[31]

3.2 Estimation

We present ITT and ATE estimates for worker's labor market outcomes. The former are useful from a policy maker's perspective because, by averaging over compliers and non-compliers, they reflect likely binding challenges to scaling-up the training interventions in the same context, or to export them to other contexts. We present ATE estimates because these map closely to theories of training, and to show channels through which VT and FT differ for trained workers.

Our ITT estimates are based on the following ANCOVA specification for worker i in strata s

[30] We have also examined the distributional gains in worker productivity. We find that very few workers have less than or equal to a one-point increase in productivity, so the mean impact does not mask a mass of hired FT workers receiving no training.

[31] The similarity in worker and firm reports is reassuring because firms might have had an incentive to over-state worker learning (to avoid the subsidy being removed). Workers likely have weaker incentives to misreport as the worker follow-up survey was conducted six months after the end of the subsidy period, when the majority of FT workers had left the firm they were assigned to.

in survey wave $t = 1, 2, 3$,

$$y_{ist} = \sum_j \beta_j T_{ij} + \gamma y_{i0} + \delta \mathbf{x}_{i0} + \lambda_s + \vartheta_t + u_{ist}, \tag{1}$$

where y_{ist} is the outcome of interest, T_{ij} denotes worker i being randomly assigned to treatment j (vocational training or firm training), y_{i0} is the outcome at baseline, $\mathbf{x}_{i0}$ are the worker's baseline covariates. λ_s and ϑ_t are strata and survey wave fixed effects respectively. As randomization is at the worker level (i), we use robust standard errors, and we weight ITT estimates using inverse probability weights (IPWs) to account for worker attrition. In the Appendix we show the robustness of the main results to dropping all covariates except baseline outcomes, randomization strata, and survey wave fixed effects, and to not using IPWs.[32]

The ATE specification replaces treatment assignment with treatment take-up (with the same controls included), where take-up is defined as a dummy equal to one if the worker: (i) started firm training in FT; (ii) started vocational training in VT. We then use treatment assignment as an IV for treatment take-up and report 2SLS regression estimates, where we bootstrap standard errors using 1,000 replications. We also report unadjusted p-values alongside Romano-Wolf p-values that account for multiple hypothesis testing.[33]

The coefficient of interest is β_j: the treatment effect of T_{ij} as *averaged* over the three post-intervention survey waves. To transition from the treatment effects to the structural estimates, we then estimate dynamic treatment effects. To do so, we convert our data to a worker job spells data set, that is possible because in each survey workers were asked to provide their monthly labor market history since the previous survey. We use this format to estimate dynamic treatment effects by quarter. This sheds light on the evolution of treatment effects, as well as on whether workers are in steady state towards the end of our study period, underpinning the job ladder model we develop and estimate.

β_j measures the causal effect of treatment on outcomes under SUTVA. In this setting SUTVA will not hold if treatment displaces control workers because treated workers are relatively more attractive to firms. To assess whether this is likely we first need to establish the relevant labor market for these workers. We note that at baseline workers are geographically and sectorally

[32] The baseline worker characteristics $\mathbf{x}_{i0}$ controlled for are age, a dummy for whether the worker was married, a dummy for whether the worker had any children, a dummy for whether the worker was employed, and a dummy for whether the worker scored at the median or above on a cognitive test administered at baseline. We also control for the vocational training implementation round and month of interview. The weights for the IPW estimates are computed separately for attrition at first, second and third follow-up. The instruments for the IPW estimates are whether the worker was an orphan at baseline, a dummy if anyone in the household of the worker reported having a phone at baseline, a dummy for whether the worker reported being willing to work in more than one sector at the time of their original application to the VTIs and dummies for the survey team the worker's interview was assigned to in each of the three follow-up survey rounds.

[33] We implement the Romano-Wolf [2016] step-down procedure based on re-sampling bootstrap methods, using the publicly provided code (http://www.econ.uzh.ch/en/faculty/wolf/publications.html9). These p-value adjustments account for the fact that we are testing two treatments at the same time, and for the fact that we test across a family of labor market outcomes.

mobile: the majority are willing to travel to other labor markets to find work, and many are willing to consider working in different sectors.[34] Defining a labor market as a sector-region, our firm census shows that on average, there are 156 employed workers and 40 firms in each market. We match an average of 8 workers per market, corresponding to 5% of all workers. Hence we do not expect the control group to be contaminated by treated workers in the same labor market as they are unlikely to be competing for the same exact job.[35]

3.3 Skills

If training routes are to impact labor market outcomes, they first should impact worker's skills. We thus first present results on three different aspects of skills acquisition among VT and FT workers. When doing so, it is natural to focus on the ATEs, namely on those that attended vocational training, or were hired as apprentices.

The first dimension of skills relates to whether workers report having been trained by a firm in their first employment spell. We define two dummies: (i) whether the worker reports having received on-the-job training at her first employer; (ii) whether the worker reported being a 'trainee' in her first employment spell. Columns 1 and 2 of Table 4 show that for both outcomes, workers hired by firms in the FT treatment are between 54 and 66pp more likely than the control group to be firm trained.[36] More surprisingly: (i) vocationally trained workers are no more likely than the control group to report being trainees in their first employment spell; (ii) workers assigned to firm training are significantly more likely to report having received training or view themselves as trainees than vocationally trained workers ($p = .000$ in Columns 1 and 2).[37]

This suggests firms are less willing to train workers that have already been vocationally trained in sector-specific skills. We thus find no evidence of firms targeting on-the-job training to workers with higher levels of human capital to begin with, or a complementarity between firm-provided skills and skills provided by vocational training institutes. The finding is consistent with both training routes providing workers similar skills, but also with firms anticipating VT workers to be

[34] At baseline, 33% of workers reported that they had previously attempted to find a job in a different town than the one they come from. Of the ones that had attempted to find a job in another town, 27% had succeeded. Of the ones that did not try to find a job in a different town, 92% said they would like to find a job in a different town than the one of origin. On workers sectoral mobility, at baseline 96% of workers reported being willing to work in more than one sector. Moreover, only 15% of all main job spells of workers in the control group at first follow-up are in the same sector as the ideal sector mentioned at baseline.

[35] Crepon *et al.* [2013] provide experimental estimates of the equilibrium impacts of labor market policies in France using a design that randomizes the fraction of treated workers across labor markets, and individual treatment assignment within labor markets.

[36] This is over a baseline of 40% of workers in the control group reporting to have received training in their first employment spell (Column 1), a magnitude that matches up well with the descriptive evidence in Table 3 where 50% of workers employed in the SMEs at baseline reported having been apprentices in the firm.

[37] These results are for workers with at least one employment spell in the post-intervention period. If we expand the sample to include all workers by assigning a value of zero to the independent variable for those never employed post-intervention, the results are qualitatively similar, and the differences in treatment effects between FT and VT remain significant at the 1% level.

more mobile than others because their skills are certifiable.

The second dimension we consider is a sector-specific skills test we developed in conjunction with skills assessors and modulators of written and practical occupational tests in Uganda. This kind of skills test has not been conducted much in the training literature.[38] Each sector-specific test comprises seven questions (five multiple choice, one pairwise-matching questions, and one question requiring tasks to be correctly ordered): Figure A5 shows an example of the skills test for the motor mechanics sector. Workers had 20 minutes to complete the test, and we convert answers into a 0-100 score. If workers answer questions randomly, their expected score is 11. The test was conducted on all workers (including those assigned to the control group) at second and third follow-up, so this outcome measures the persistence of skills accumulation in the VT and FT treatments. There is no differential attrition by treatment into the test.[39]

Before administering the test, we asked a filtering question to workers on whether they had *any* skills relevant for sectors in our study. The dependent variable in Column 3 of Table 4 is a dummy equal to one if the worker reported having skills for a sector. The ATE estimates show that VT workers and FT workers all report being significantly more likely to have relevant skills than those in the control group. As reported at the foot of the Table, 60% of control group workers report having skills for some sector, and reassuringly this rises to 100% for FT workers that were hired by firms, and for VT workers that attended a VTI. This is entirely consistent with the earlier descriptive evidence in Figure A4 suggesting complier workers in VT and FT treatments received real training through VTIs and firms.

All workers that reported having sectoral skills took the test: others (mostly in the control group) were assigned a score of 11 assuming they would answer the test at random. Column 4 shows that VT and FT workers significantly increase their measurable sector-specific skills, as recorded two and three years after the training. Relative to the control group, VT workers increase sector-specific skills by 34% (or $.4sd$ of test scores). FT workers increase sector-specific skills by 32%. Strikingly, both training routes cause persistent skills accumulation, although there is no

[38] Berniell and de la Mata [2016] present evidence from a 12-month apprenticeship program in Argentina. In comparison to the control group, they find little evidence that the cognitive and non-cognitive skills of apprentices are impacted, but show that relative to the control group they are able to certify their skills to a great extent and this drives some of the higher employment probabilities for treated workers. They do not develop specific tests to measure sector- or firm-specific skills. Adhvaryu *et al.* [2017] present more detailed evidence on worker knowledge, preferences and task assignment of workers in the context of an evaluation of a soft-skills training program for garment workers in India.

[39] We developed the sector-specific skills tests over a two-day workshop with eight practicing skills assessors and modulators of written and practical occupational tests from the Directorate of Industrial Training (DIT), the Uganda Business and Technical Examinations Board (UBTEB) and the Worker's Practically Acquired Skills (PAS) Skills Testing Boards and Directorate. To ensure the test would not be biased towards merely capturing theoretical/attitudinal skills taught only in VTIs, workshop modulators were instructed to: (i) develop questions to assess psychomotor domain, e.g. trainees ability to perform a set of tasks on a sector-specific product/service; (ii) formulate questions to mimic real-life situations (e.g. "if a customer came to the firm with the following issue, what would you do?"); (iii) avoid using technical terms used in VTI training. We pre-tested the skills assessment tool both with trainees of VTIs, as well as workers employed in SMEs in the eight sectors we study (and neither group was taken from our worker evaluation sample).

significant difference in sector-specific skills accumulation between VT and FT workers ($p = .904$).

A contribution we make to the training literature is to quantify the causal productivity impacts on workers of on-the-job training. Much of the earlier evidence has been based on observational data and there has been a long-standing debate over whether there are substantive human capital impacts of such training [Blundell *et al.* 1999], especially once the endogenous selection of workers into training is corrected for [Leuven and Oosterbeek 2008].

The fact that FT workers are provided sector-specific skills goes against a standard Beckerian framework of firm-sponsored training. Our setting departs from this in two ways: (i) BRAC subsidizes the apprenticeships through the wage subsidy, making firms more willing to provide skills that are not firm-specific; (ii) firms are contractually required by BRAC to provide sector-specific skills to workers (Figure A3). As the subsidy likely remains below the full cost of training, this reinforces the notion that a key reason for low take-up by firms is that the FT treatment imposed a costly time and resource requirement on firm owners to provide training that was monitored and enforced.

The final dimension we consider probes whether VT and FT workers differ in the firm-specificity of their skill set at endline. It is hard to directly measure firm-specific skills for our study sectors. We thus approach this issue using data from the endline survey where we asked employed workers whether they considered their skills to be transferable across firms. As Column 5 shows, relative to control workers, VT workers are significantly more likely to report having transferable skills, although there is no statistical difference to FT workers ($p = .264$). This is again consistent with the earlier descriptive evidence in Figure A4, that showed substantial learning for both sector-specific and firm-specific tasks among FT workers.

The key implication from these results is the similarity in long run sector-specific skill accumulation in VT and FT treatments. This largely shuts down a channel through which differences in outcomes across treatments could have been generated, and reinforces the idea that such differences stem from the greater certifiability of skills obtained through vocational training.

We next examine how these persistent skill improvements translate into labor market outcomes.

3.4 Employment and Earnings

Table 5 presents ITT estimates for labor market outcomes, starting with the extensive margin of being in paid employment. Column 1 shows that, averaged over the three post-intervention survey waves, both forms of worker training raise employment probabilities: workers assigned to FT and VT treatments are 6pp and 9pp more likely to be employed, corresponding to 14% and 21% impacts over the control group, whose unemployment rate is 56%. Hence, these ITT impacts on youth unemployment rates are economically significant, and this is so for both training routes.

On the total effect margin, Column 2 shows VT and FT workers significantly increase the months worked in the year by .88 and .52, respectively, corresponding to 19% and 11% increases

over the control group. Hence the pattern of results indicates that through the offer of either training route, young workers increase their labor market attachment. For VT workers, Column 3 shows this is further enhanced by them significantly increasing their weekly hours worked.[40] This is evidence that the VT treatment has a stronger impact on employment.

Column 4 combines extensive margin and total effect margin effects to derive ITT impacts on total monthly earnings. Averaged over all post-intervention waves, the ITT earnings impact for VT workers is an increase of 25% over the control group. In contrast, there is no ITT earning impact for FT workers, and the difference in earnings between VT and FT is statistically significant ($p = .048$). Hence from a social planner's point of view, the use of wage subsidies attached to workers does not – on average – lead to earnings gains for those workers.

Column 5 combines these multiple labor market outcomes into one index. This follows the construction in Anderson [2008], so accounting for the covariance structure in components, and we normalize by the standard deviation of the index in the control group to ease interpretation. The labor market index rises significantly for both FT and VT workers, with the magnitude being slightly larger for VT workers, because such workers have improved outcomes along all four margins considered. However, all the findings point to both groups of worker increasing their labor market attachment: by being more likely to work, and by supplying more labor over time. This notion is reinforced by the fact that for both FT and VT workers, increases in employment are driven by increases in wage, not casual, employment.[41] Following up on this, we consider whether workers are employed in their sector of training (for VT workers), or the firm sector they were matched to (for FT workers), or their first or second preferred sectors of employment (for the control group). Column 6 shows: (i) for VT workers, this likelihood rises by 167%; (ii) for FT workers this rises by 67%; (iii) the difference between VT and FT workers is significant ($p = .000$).

Of course some of the differences in ITT impacts might be driven by differential compliance between VT and FT treatments. To account for this and also map to theories of training, Table 6 presents ATE estimates for the same outcomes. The treatment effects, averaged over all post-intervention survey waves, are statistically similar between VT workers that started vocational training, and FT workers that were hired and trained by firms incentivized by wage subsidies. The overall labor market index shows no difference in ATEs between FT and VT training ($p = .202$). Indeed, taking into account the differential compliance, the point estimate on the labor market index is actually higher for FT workers than VT workers. This is driven by FT workers having a higher likelihood of any paid work, and a higher number of months worked on average over the three-year follow-up period – reversing the ranking from the ITT estimates. In short, a comparison of compliers shows both training routes to lead to large impacts for worker's labor market outcomes

[40] Hours worked were not measured at first follow-up, hence the lower sample size in Column 3.

[41] The increases in employment are entirely driven by wage employment for FT workers. For VT workers, they are 5.5pp more likely to be in wage employment, a 20% increase over the control group, and 3.6pp more likely to be in self employment, a 23% increase over the control group.

along extensive and intensive margins, when averaging over the three-year follow-up period.

The ATE estimates on monthly earnings in Column 4 show high experimental returns to vocational training: 42%, averaged over the post-intervention period. This begs the question of why workers do not themselves invest in vocational training given such returns? One explanation is credit constraints: as documented earlier, worker monthly earnings at baseline are $6, while the vocational training costs over $400.

An alternative explanation is that workers have incorrect beliefs about the returns to vocational training [Jensen 2010]. We assess this using information collected from workers at baseline over their expected probability of finding work, and their expected earnings conditional on employment, if they received vocational training. This is shown in Table A6. Columns 1 and 2 focus on the extensive margin and show that: (i) at baseline, workers expect their employment probability to be 57% (that is optimistic given baseline employment rates of 40%); (ii) workers expect their likelihood of finding work to rise by 30pp or 53%, if they receive vocational training. This is also optimistic given the ATE impact on the extensive margin being closer to 31%.

In terms of earnings, Column 3 of Table A6 reports worker beliefs at baseline, over the average monthly earnings given their current skill set (assuming they were employed). These correspond to just under $60. We then asked workers what they expected their maximum and minimum monthly earnings to be if they received vocational training (and the likelihood they would be able to earn more than the midpoint of the two). Fitting a triangular distribution to their beliefs we derive an expected earnings from vocational training. This is shown in Column 4: on average, workers report earnings would more than double, so a greater than 100% return. This is double the Mincerian returns shown in Table 2, that are themselves upwards biased. Combining both margins we see that workers expect the returns to vocational training to be nearly 200%, many times more than the ATE estimate of returns, at 42%. In short, workers are overly optimistic with regards to the returns to vocational training, and such expectations do not explain their lack of investment in their own human capital.

The Appendix presents robustness checks on our baseline ITT findings, using the labor market index.

3.5 Dynamics

3.5.1 Retention

We now shed light on dynamic responses across treatments. This allows us to uncover the mechanisms driving the impacts and also helps us to bridge to the job ladder model that assumes workers are in steady state. To begin with we consider retention rates among apprentices. In each survey wave we asked workers hired under the FT treatment if they were still employed at the same firm they were originally matched to. Figure 3 plots the survival function for them: among those actually hired, 57% are employed for at least 6 months. Yet, their tenure does not

last much longer: the average duration of employment at the matched firm, conditional on being strictly higher than six months is 9 months. Crucially, by endline, almost none of these workers remain in the firm they were originally matched to. The fact that apprentices have relatively short employment spells at their matched-to firm suggests the FT treatment provided them skills that could be less firm-specific than if firms had not borne the cost of training or been contractually obliged by BRAC to provide training. The fact that FT workers transition away from firms they were trained by after the wage subsidy expires further limits any additional degree of firm-specific skills they accumulate relative to VT workers.[42]

As a point of comparison, Figure 3 also shows survival functions for the first employment spell among VT and control workers: both groups have considerably longer first employment spells than workers hired in the FT by firms incentivized through a wage subsidy.

3.5.2 Dynamic Treatment Effects

Figure 4 shows quarterly dynamics of: (i) number of months worked (Panel A); (ii) total earnings (Panel B); (iii) average hourly wage in wage employment (conditional on being employed). For each outcome, the left hand panel shows the descriptive evolution of the outcome, and the right hand panel provides dynamic treatment effect estimates for these quarterly outcomes. The following key insights are obtained.

On the dynamics of quarterly employment (Panel A) we see that FT workers find employment more quickly than VT workers, but over time their employment rate converges to the control group, while employment rates for VT workers increase over time. This reversal of fortune between FT and VT is shown clearly in the dynamic ITT effects in the right hand panel: FT workers find employment more quickly than VT workers in the first few quarters, but over time their employment rate converges to the control group, while employment rates for VT workers increase steady over each quarter and diverge away from the control group. The dynamic treatment effects on employment show that transition rates are stable by the fourth quarter of 2015. This helps underpin the assumption that workers are in steady state when estimating the structural model (that is by the end of 2015 so in the last two waves of data).

On the dynamics for quarterly earnings (Panel B), again FT workers do well initially but then over time, their earnings fall behind those of VT workers. In contrast, VT workers steadily increase their earnings (the gradient is near linear) and diverge away from the control group over time.

Hence, the earlier documented similarity in treatment effects between FT and VT when averaged over post-intervention survey waves is driven by the earlier quarters in which FT workers were hired by firms incentivized through wage subsidies. Subsequent to that, the patterns of employment and earnings differ across VT and FT workers, with FT workers having similar employment

[42] Apprentices retained after the duration of the wage subsidy do experience a slight rise in their earnings above the level of the wage subsidy of $38/month.

profiles as control workers.

Panel C shows quarterly hourly wages (*conditional* on employment). Here we see that conditional hourly earnings for both FT and VT workers rise relative to the control group, and are not different from each other. For FT workers, hourly wages rise by 12% relative to the control group, and for VT workers they rise by 11%. To understand the similar wage impacts (conditional on employment), we note the earlier results that sector-specific skills accumulation is similar among complier FT and VT workers. Hence, on-the-job productivity and wages should be similar for young workers entering the labor market through either training route.[43]

4 Job Ladder Model

We develop a job ladder model of worker search. This identifies and quantifies the mechanisms driving labor market impacts in steady state. We later use the estimated model to simulate counterfactual scenarios where we shut down specific mechanisms to understand their relative importance (such as labor mobility that can arise from skills certification), and to simulate treatment effects if the training interventions had been targeted to workers in the economy beyond those recruited into our evaluation sample through the oversubscription design. Finally we use the simulated steady state earnings impacts to calculate IRRs for each training intervention.

4.1 Set-up

The labour market features a continuum of measure 1 of risk neutral workers. Time is continuous and we assume workers have reached their steady state labor market trajectories by the end of our study period, as suggested by the dynamics in Figure 4.[44] Workers are heterogeneous in two dimensions: their training (treatment) status, T (where $T = VT$, FT or C), and the amount of effective labor ε they supply each period given their training status T. We later relate a worker's effective labor to their sector-specific skills. Given the treatment effects on skills in Table 4, we assume that relative to control workers, ε is increasing in training T, and is fully transferable across firms. Conditional on T, ε is assumed fixed so there is no human capital accumulation over and above that provided by training. We can thus think of ε as characterizing a worker's fixed productivity or type. The cross sectional distribution of ε conditional on training is denoted $H(\varepsilon|T)$, with density $h(\varepsilon|T)$.

Following Barlevy [2008] we assume firms post wage contracts indexed to worker type-ε, namely they post a piece rate r paying a constant price per unit of effective labour. This fits our context,

[43] Given the majority of workers are paid piece rates in SMEs in the sectors we study, these earnings impacts can reflect real productivity impacts of the training routes.

[44] There is an established literature on job ladder search models, the defining characteristic of which is always that workers agree on the ranking of available jobs, hence the notion of a job ladder [Bontemps *et al.* 2000, Moscarini and Postel-Vinay 2018].

where the majority of workers are paid piece rates, and this is so within manufacturing and service sectors. A worker of type-ε employed at a firm posting a piece rate r earns a wage $w = r\varepsilon$. We assume the offered piece rate comes from a distribution $F(r)$ with density $f(r)$, and denote the lower (upper) bound of the support of $F(r)$ as $\underline{r}$ ($\overline{r}$). All workers sample piece rates from this same distribution. The firm commits to pay w each period until the worker is laid off or quits.[45]

Two points are of note. First, the distribution from which piece rates are drawn $F(\cdot)$ does not depend on treatment T. Once hired, worker productivity is realized and so higher type-ε workers are paid a higher wage (at the same piece rate r) because $w = r\varepsilon$. We relax this assumption in a later extension, so allowing for $F(r|T)$: this enables us to investigate, in a very reduced form way, whether, for example, workers in different treatments search differently across firms, who might then offer piece rates from different underlying distributions. Second, firms play no role in the model. Hence to understand how the results derived from this partial equilibrium framework map to general equilibrium impacts, we later present evidence from the firm side of the experiment. This identifies how firms react to treatments along one key margin: the displacement (or crowding in) of other workers in the economy, that is central to understanding general equilibrium effects of our training interventions.

4.1.1 Value Functions

Workers can be unemployed or employed each period. Unemployed workers earn zero income each period of unemployment. $\lambda_0(T)$ is the probability of receiving a job offer for an unemployed worker with training status T. The worker takes up this job offer if the expected value of the job is higher than the value of remaining unemployed. With discount rate ρ, the value of unemployment for a type-ε worker with training status T is:

$$\rho U(\varepsilon, T) = \lambda_0(T) \int_{R(\varepsilon,T)}^{\overline{r}} [V(x, \varepsilon, T) - U(\varepsilon, T)]\, dF(x). \tag{2}$$

Employed workers of type-ε earn $w = r\varepsilon$ at their firm in each period. They face an exogenous job destruction rate $\delta(T)$, that depends on their training status. This captures both the quality of jobs and the expected duration of the employment relation. On-the-job search is allowed. $\lambda_1(T)$ is the probability of receiving a job offer by an employed worker with training status T. She takes up this opportunity if the expected value of the job offer exceeds the current job value. A type-ε worker with training status T has the following valuation of a job paying piece rate r:

$$\rho V(r, \varepsilon, T) = r\varepsilon + \delta(T)\,[U(\varepsilon, T) - V(r, \varepsilon, T)] + \lambda_1(T) \int_{r}^{\overline{r}} [V(x, \varepsilon, T) - V(r, \varepsilon, T)]\, dF(x). \tag{3}$$

[45] There is no wage bargaining in this set-up. Our empirical setting is not well suited to such a version of the model: unionization rates are less than 1% in Uganda, and the demographic structure ensures there is no shortage of potential labor hires available. Both factors dampen worker bargaining power [Rud and Trapeznikova 2018].

Combining (2) and (3), the key endogenous choice for workers – their reservation wage – solves the following for a type-ε worker with training status T:

$$R(\varepsilon, T) = [\lambda_0(T) - \lambda_1(T)] \int_{R(\varepsilon,T)}^{\bar{r}} \frac{\bar{F}(x)}{\rho + \delta(T) + \lambda_1(T)\bar{F}(x)} dF(x), \tag{4}$$

where $\bar{F}(x) = 1 - F(x)$. Unemployed workers accept piece rates above $R(\varepsilon, T)$. We assume $\underline{r} \geq R(\varepsilon, T)$, so unemployed workers accept any job offer. Employed workers accept piece rates drawn from $F(r)$ that are higher than their current one irrespective of ε.

It is important to be clear on the distinct roles that two sources of worker heterogeneity play in the model: worker types-ε (as later proxied by their skills) determine wages conditional on employment, but do not play a role for labor market transitions. Treatment status instead is allowed to impact both worker types and the labor market transition parameters, $\lambda_0(T)$, $\lambda_1(T)$ and $\delta(T)$. In particular, worker types are affected by skills accumulated during training, and that makes workers more productive when employed. Transitions are impacted by treatment if this makes it easier for workers to receive job offers, for instance if treatment facilitates job search by allowing workers to better demonstrate their skills to employers. Finally, note that since transition parameters do not depend on ε, then the reservation wage for a type-ε worker does not depend on ε, so $R(\varepsilon, T) = R(T)$.

We do not explicitly model search effort. Rather this is encompassed within the labor mobility parameters (λ_0, λ_1) that capture both job search when unemployed, on-the-job search effort, and underlying factors that drive search effort in these states and differ across treatments such as the certifiability of skills.

4.1.2 Steady State

We close the model by deriving steady state conditions where for expositional ease, we omit conditioning on training status. The following steady state relationship characterizes when outflows and inflows for unemployment are equal for workers of type-ε:

$$\lambda_0 u(\varepsilon) = \delta[h(\varepsilon) - u(\varepsilon)], \tag{5}$$

$$u(\varepsilon) = \frac{\delta}{\delta + \lambda_0} h(\varepsilon). \tag{6}$$

where $u(\varepsilon)$ is the unemployment rate for type-ε workers. The population unemployment rate, $u(\varepsilon)/h(\varepsilon)$, is independent of type ε, which is unsurprising as worker labor market mobility is independent of ε, and depends only on T (through $\delta(T)$ and $\lambda_0(T)$).

For employed workers, we first note that because they can search on-the-job, the cross-sectional distribution of *observed* piece rates for type-ε workers $G(r|\varepsilon)$ differs from the offer sampling distribution $F(r)$, because observed piece rates are those accepted by workers. For type-ε employed

workers with piece rate $\leq r$, the steady state relationship for employment is:

$$[\delta + \lambda_1 \bar{F}(r)][h(\varepsilon) - u(\varepsilon)]G(r|\varepsilon) = \lambda_0 F(r)u(\varepsilon). \tag{7}$$

The LHS of (7) is the outflow from the stock of type-ε workers employed at a piece rate less than r. To see this note that of the $(h(\varepsilon) - u(\varepsilon))G(r|\varepsilon)$ workers employed at a piece rate $\leq r$, a fraction δ have the job exogenously terminated, while a fraction $\lambda_1 \bar{F}(r)$ receive and accept an offer greater than r. The RHS of (7) is the inflow into employment from unemployment.

Using these steady state relationships, we can derive the link between $G(\cdot)$ and $F(\cdot)$:

$$F(r) = \frac{(\delta + \lambda_1)G(r|\varepsilon)}{\delta + \lambda_1 G(r|\varepsilon)}, \tag{8}$$

$$G(r|\varepsilon) = \frac{\delta F(r)}{\delta + \lambda_1 \bar{F}(r)}, \tag{9}$$

where $G(r|\varepsilon) = G(r)$ is independent of ε (which is expected given that worker mobility is independent of ε). We see that $G(r)$ FOSD $F(r)$ unless there are no job-to-job transitions ($\lambda_1 = 0$), i.e. because on-the-job search leads to outside offers, there exists a wedge between offered and accepted piece rates.

4.2 Estimation

4.2.1 Data

In each worker survey, we asked respondents to provide their monthly labor market history since the last survey. We use this information to convert our panel data into a job spells format data set: for each worker i we construct a complete monthly history of their employment status $e_i \in \{0, 1\}$ from August 2014, one year after the end of vocational training, to our endline in November 2016. We assume workers have reached their steady state trajectories by November 2015. Consistent with the model, we set one wage per employment spell, w_i, and then estimate transition probabilities between job and unemployment states (τ_{JU_i}, τ_{JJ_i}, τ_{UJ_i}) using a maximum of two spells since the steady state has been reached. Hence the model is estimated off the last two survey waves.[46]

[46] Figure A6 shows some possible cases how the worker spells data is constructed for worker i, where the spell duration is d_i and transition indicators between spells are τ_{JU_i}, τ_{JJ_i}, τ_{UJ_i}. In the top panel, we show a scenario in which worker i is unemployed in November 2015, the unemployment spell is not left censored, and the worker transitions into employment after d_i months of unemployment. The bottom panel considers a case where worker i is employed in November 2015, the employment spell is left censored, and the worker transitions into a new state after d_i months of employment. In our data, initial unemployment spells (those being experienced in November 2015) are more likely to be right censored than initial employment spells (60% vs. 38%). Initial unemployment spells are also more likely to be left censored than initial employment spells (79% vs. 34%).

4.2.2 Identification

Our worker surveys record worker i's wage in each employment spell j, w_{ij}. However, the distributions of worker productivity $H(\varepsilon)$, piece rates $G(r)$ and piece rate offers $F(r)$ are not observed. We tackle these identification issues as follows. First, the data includes a proxy for worker productivity: their sector-specific skills test score, s. We assume measurable skills relate to true worker productivity as follows: $\varepsilon = s^{\alpha}e^{\epsilon}$, where ϵ is idiosyncratic measurement error. Taking the natural logarithm of wages we obtain the following expression for worker i in spell j: $ln(w_{ij}) = ln(r_j) + \alpha ln(s_i) + \epsilon_{ij}$, where r_j denotes the piece rate paid to the worker in spell j.

As r_j is unobserved, to identify α using OLS we need r_j to be independent of s_i. However, in our model training simultaneously impacts worker type-ε (skills) and observed wages w_{ij} (because $G(r|\varepsilon)$ depends on $\delta(T)$ and $\lambda_1(T)$). To correct for this omitted variable bias, we run the following regression of wages on skills, controlling for treatment status:

$$ln(w_{ij}) = \gamma_0 + \alpha ln(s_i) + \sum_k \gamma_k T_{ik} + u_{ij}, \tag{10}$$

where T_{ik} is a dummy equal to one if worker i is assigned to treatment k (i.e. vocational training or firm training), and $u_{ij} = ln(r_j) + \epsilon_{ij}$. We estimate this for workers transitioning from unemployment into employment.

Table A8 shows estimates of α, using the same survey waves we estimate the model parameters from. Namely, we assume workers are in steady state and so include all job spells workers have been involved in starting from November 2015. Our baseline estimate is $\widehat{\alpha} = .289$ from Column 1. This falls slightly to $\widehat{\alpha} = .235$ when we condition on worker characteristics and strata dummies.[47] Given $\widehat{\alpha}$, i's measured sector-specific skills (s_i), and i's wages in employment spell j (w_{ij}), we recover the estimated piece rate for each worker-spell as $\hat{r_j} = w_{ij}/s_i^{\hat{\alpha}}$. We thus recover $G(r)$ and use the steady state conditions to estimate the distribution of piece rate offers, $F(r)$. These functions are recovered for each group of workers (controls, non-compliers, compliers), so we do not impose a common $F(\cdot)$ across treatments.

4.2.3 Parameterization

We assume the model parameters have the following parametric form:

$$\lambda_0 = \lambda_{00} + \sum_k \lambda_{0k} T_k, \tag{11}$$

[47] Column 3 adds interactions of the logarithm of the skills test score with treatment dummies: we see both point estimates are close to zero, suggesting there is no need to include a $T_{ik} \times s_i$ interaction in (10). Columns 4 and 5 impute earnings in the first month of employment for individuals who have this information missing, using two alternative approaches. We see that in this larger sample, $\widehat{\alpha}$ rises slightly but is more precisely estimated.

$$\lambda_1 = \lambda_{10} + \sum_k \lambda_{1k} T_k, \tag{12}$$

$$\delta = \delta_0 + \sum_k \delta_k T_k, \tag{13}$$

where, as above, T_k denotes worker's treatment status. In line with SUTVA we assume workers across treatments do not interact with each other in the labor market. Given differences in compliance between FT and VT treatments, we treat compliers and non-compliers as separate groups and simultaneously estimate different parameters (λ_0, λ_1, δ, ε) for each of the resulting five groups (control, non-compliers in FT and VT, compliers in FT and VT). From the treatment effect estimates on skills accumulation, we expect complier and non-compliers to have different $h(\varepsilon)$ distributions: indeed, all non-compliers might be largely comparable to control workers. In the Appendix we detail the construction of the likelihood function, and estimate the parameters using maximum likelihood following the two-step procedure in Bontemps *et al.* [2000], recovering asymptotic standard errors for the 20 parameters simultaneously estimated (λ_0, λ_1, δ, ε in the five groups of workers).

4.2.4 Supportive Evidence

We provide three classes of evidence supporting and motivating the model structure. First, a key assumption of the model is that labor mobility does not depend on worker type ε. We test this in the sample used to estimate the model parameters, so including job spells workers have been involved in since November 2015. Table A9 shows estimates from an OLS regression of a dummy equal to 1 if the individual is employed in November 2015, on the worker's sector-specific skills test score, controlling for treatment status and where the unit of observation is the job spell. Across worker samples in Columns 1 to 3, there is zero correlation between skills (our proxy for ε) and the likelihood the worker's initial steady state spell is in employment. Column 4 confirms this to be the case when we allow skills to differentially impact this employment probability by treatment.

Second, we test a key prediction of the model: that wage growth occurs between, not within, job spells because a worker's wage only increases when making transitions. Within a spell, the wage is fixed at w. To examine this prediction we decompose workers' wage growth into that occurring within and between job spells. We find the average wage growth of job movers is at least twice as high as that of job stayers, irrespective of the exact reference period used.[48]

[48] To decompose worker's wage growth, we first exploit the fact that for each job spell we have information on the wage in the first month and the last month of the spell. We then choose some reference date and linearly interpolate wages from the first and last month of the spell ongoing at the reference date. We then calculate the wage growth between two reference dates (e.g. between April 2015 and April 2016) for: (i) workers employed in the same job throughout the reference period (job stayers); (ii) workers who change job at least once in the reference period (job movers). To avoid sensitivity to outliers, the top 1% of wages are excluded. Self-employed workers and workers with at least one unemployment spell in the reference period are excluded. We then take the ratio of the average wage growth of the job movers to job stayers. Using the reference period of April 2015 to April 2016 this ratio is 2.06 (the ratio of medians is 2.31).

Third, we take the sample used to estimate the model parameters and estimate ITT regressions on outcomes closely linked to labor mobility, a key mechanism in the model. The results in Table 7 show that VT workers experience more job spells (Column 1), and this leads them to gain over two months of extra employment relative to the control group (Column 2). On the other hand, and consistent with the dynamic evidence, in steady state, FT workers do not experience more job spells and are not employed for longer than the control group. These differences between VT and FT are significant at the 1% level.

We further see that while VT workers are more likely to be employed in the first spell in steady state (Column 3), VT workers who are unemployed in the first spell are significantly more likely than control workers to transition to employment (Column 4). The estimate in Column 4 relates to λ_0, and shows the positive treatment effect on the number of work spells for VT workers in Column 1 is not only driven by them being more likely to be initially employed in the first spell. Again, we find no impact of FT on transitions away from unemployment, and the difference between FT and VT workers is significant at the 5% level.

Finally, in Columns 5 and 6 we estimate treatment effects relating to λ_1 and δ respectively. Column 5 shows no significant impact of treatments on the number of transitions conditional on being employed in the first spell. Column 6 shows that conditional on being employed in the first spell, FT workers have significantly fewer transitions into unemployment, but we are unable to reject this estimate is different for VT workers ($p = .403$).

While suggestive of the mechanisms that might be at play – that vocational training leads to higher labor market mobility primarily because of higher rates of unemployment-to-job transitions (λ_0), rather than JJ transitions or job destruction rates (λ_1, δ) – the findings in Columns 4 to 6 condition on employment status in the first spell and so cannot be easily interpreted, despite the experimental variation in treatment assignment. We therefore move to presenting the full set of model parameter estimates.

5 Model Estimates

5.1 Parameters

Table 8 presents the baseline results. Panel A shows the mean worker type from $\overline{\varepsilon}$ and parameter estimates ($\widehat{\delta}$, $\widehat{\lambda}_0$, $\widehat{\lambda}_1$) for controls (Column 1), non-compliers in each training arm (Columns 2 and 3), and compliers in each training arm (Columns 4 and 5). This split is informative given the treatment effects on skills accumulation: we expect compliers and non-compliers in each treatment to differ in their skills and hence their estimated type ε. This is confirmed in the first row of Panel A: the distribution of worker types has higher means for compliers from either training route (2.92, 2.84), with both being at least 11% higher than the mean type-ε among controls and non-compliers (2.51, 2.48, 2.56).

We first note that steady state job destruction rates δ are identical for FT and VT workers (.023). To the extent that δ captures job quality, this is consistent with the similar skills acquired through both training routes leading to similar job qualities.[49]

On labor mobility, for the arrival rate of job offers when unemployed (λ_0): (i) controls and non-compliers have similar estimates (.018 to .019); (ii) remarkably, the arrival rate of job offers when unemployed is almost identical for complier firm-trained workers and control workers ($\widehat{\lambda}_0 = .019$, .020): the additional skills and labor market experience gained by firm-trained workers count for little if they fall off the job ladder into unemployment; (iii) complier VT workers have transition rates that are 40% higher than complier FT workers. These λ_0 estimates are statistically different from each other ($p = .082$), $\widehat{\lambda}_0(VT)$ is also statistically higher than for the control group ($p = .004$).

On job offer arrival rates when employed (λ_1) among compliers, we see that VT workers have arrival rates 22% higher than FT workers, although this is not statistically different ($p = .371$).[50]

Among trained workers, we see the difference in labor mobility of vocationally trained workers relative to those transitioning into the labor market through firm-provided training is through a higher rate of UJ transitions: when unemployed, VT workers get back onto the job ladder more quickly. This pattern of mobility is in line with VT workers having more certifiable skills than FT workers, and this certifiability having especially high returns when workers are unemployed. In contrast, when employed, the certifiability of skills matters less because potential employers have a signal of a worker's skills by virtue of them being employed in the first place. Hence JJ transitions play less of a role in explaining differences between training routes.

5.2 Unemployment, Wages and Earnings

Panel B shows implied impacts for key labor market statistics for workers, and beneath each we report the percentage impact relative to the control group. To being with, the first row in Panel B measures the intensity of interfirm competition for workers (labor market tightness): this is the number of outside offers a worker receives before being laid off: $\frac{\lambda_1}{\delta}$. Inter-firm competition for complier VT workers is slightly higher than for FT workers. This is driven by VT workers receiving slightly more outside job offers when employed; recall both groups of workers have identical job destruction rates. Relative to controls, interfirm competition for complier VT workers rises by 15%. This is in line with the earlier ATE estimates on skills accumulation, that showed VT workers reporting having more transferable skills across firms. In contrast, the steady state interfirm competition for FT workers is very similar to control workers (it is just 3% lower).

On steady state unemployment rates, both training routes substantially reduce youth unem-

[49] The estimated destruction rates match closely other literature estimates. In particular, Rud and Trapeznikova [2018] estimate a different structural model of job search using UNHS 2010-11 data for Uganda and find a very similar implied annual job destruction rate of 32% as we find for the control group.

[50] We note that for nearly all groups of worker, the monthly rate of JJ transitions is higher than observed for US workers, that is usually below 3% across the business cycle [Moscarini and Postel-Vinay 2018].

ployment rates for compliers relative to controls. FT compliers have 10% lower unemployment rates in steady state; for VT compliers the reduction is 23%. Both impacts are of economic significance given the high levels of youth unemployment in Uganda. Unemployment durations fall for both sets of compliers, but this fall is far larger for VT workers (32% vs. 5.5%): this is because of the significantly greater unemployment-to-job mobility of VT workers, so they more quickly get back onto the job ladder if they fall into unemployed (recall job destruction rates are identical across compliers).

Ultimately, what we are concerned with is whether workers transition to better paying jobs. These results are shown in Panel C. To derive these earnings impacts, we take the appropriately weighted mean of the kernel density estimate of each group specific $F(r)$, to impose a common $F(r)$. We invert this using the steady-state relationships (8) and (9) to obtain $G(r|\cdot)$ and $G(w|\cdot)$ distributions for each group of workers.

Panel C shows that when unemployed, the mean offered wage for control and non-complier workers are similar. This is in line with the assumption unemployed workers accept any job offer, and with the estimated mean type ε being similar across controls and non-compliers. In contrast, steady state wages offered to complier FT and VT workers are substantially higher: this is because these workers are more skilled, so even if firms draw from the same piece rate distribution $F(r)$, this translates into higher wages for workers because of the complementarity between piece rates and skills ($w = r\varepsilon$). Of course, workers only accept job offers if the value of the offered job is greater than their current one. The mean accepted wage for complier VT and FT workers is thus much higher than offered wages, but not much different between the two (73.8 versus 74.0). This is in line with the skills (hence type-ε) of both trained workers being similar, so earnings conditional on employment are similar. The impact of training on skills also explains why accepted wages for compliers are higher than for controls/non-compliers.

In steady state, the final row in Panel C shows that the annual earnings (unconditional) of complier VT workers rise by 55% over the control group, while the earnings of complier FT workers rise by just over half of that, 32%.

Combining these results provides a precise interpretation to the driving force behind the dynamic treatment effects: vocational trainees pull away from FT workers in their employment rates and earnings (Panels A and B in Figure 4) because they are more likely to get back onto the job ladder if they fall into unemployment. These dynamics are not so much due to any great job-to-job mobility, suggesting the returns to skills certifiability are higher when unemployed than when employed. Moreover, compliers across training routes move as far up the job ladder as each other – wages conditional on employment are similar for complier VT and FT workers because their skills are similar (Panel C in Figure 4). The key distinction is that VT workers are more likely to get back onto the job ladder if they fall off it.

Three further points are of note. First, it is useful to contrast the experimental and model estimates of the returns to training. The estimated returns to vocational training in steady

state (55%) are higher than the ATE estimate (42%) that is averaged over the post-intervention period. In contrast, the steady state returns to firm-provided training (32%) are lower than the experimental returns (48%). This contrast arises because the steady state calculations account for the lower UJ transition rates of FT workers. In steady state they get back on the job ladder at the same rate as controls, slowly closing the gap between them in terms of employment rates.[51]

Second, the earnings impacts are larger than the percentage impacts on unemployment rates in Table 8. This reinforces the fact that each training route not only reduces unemployment risk, but also leads to higher wages when employed, consistent with the reduced form evidence on the human capital impacts and ATE results of each training route.

Finally, we note that controls and non-compliers have the highest job arrival rates when employed (Panel A, Columns 1 to 3). These selected groups of less skilled workers that find employment churn between jobs at a high rate, but they do not progress up the job ladder. To see that note that, within treatment, the accepted wage for non-compliers is on average lower than for compliers: hence they do not move far up the job ladder despite higher JJ mobility and interfirm competition (especially for VT non-compliers). This is because unskilled workers do not improve their type over time, say through skills accumulation. Hence, as shown in the final row of Table 8, ultimately the steady state earnings impact is almost identical between controls and non-compliers (remaining well below earnings impacts for either complier groups).

In the Appendix we present robustness checks where we: (i) allow piece rates to be drawn from treatment specific functions $F(r|T)$; (ii) show how our results vary when using the highest and lowest $\widehat{\alpha}$'s across specifications in Table A8.

6 Extensions

6.1 Employment Displacement

To understand how the results derived from the partial equilibrium model map to general equilibrium impacts, we present evidence from the firm side of the experiment. This identifies how firms react to treatments along one key margin: the displacement (or crowding in) of other workers, that is central to understanding general equilibrium effects of our training interventions. The right hand side of Figure 1 summarizes the experimental design from firms' perspective. Recall that we drew a sample of 1538 SMEs, operating in one of the eight sectors of interest, and having between one and 15 employees (plus a firm owner). We focus on the comparison between firms assigned to the

[51] An alternative hypothesis for these dynamics is that the training routes differ in how they enable workers to learn-how-to-learn, rather than enhancing their productive capacity *per se* [Neal 2017]. Dynamic impacts are then driven by intertemporal complementarity in worker's capacity to learn. Although it is difficult to find skills that impact learning capacity but not productivity, we partially explore this hypothesis by estimating whether workers cognitive abilities, and other preference parameters, change over time and differentially by training route. We do not find any evidence of such mechanisms in our setting.

wage subsidy offer and the control group of firms. For the other treatments involving matching, workers are retained within firms too infrequently to say anything on employment displacement.[52]

We estimate the effect on firms of being offered to meet an untrained worker and a wage subsidy to hire and train that worker, using the following ITT specification for firm f in randomization strata s:

$$y_{fst} = \beta Firm\text{-}Trained_f + \gamma y_{f0} + \delta \mathbf{x}_{f0} + \lambda_s + \vartheta_t + u_{fst}. \tag{14}$$

y_{fst} is the firm outcome of interest in post-intervention survey wave t, $Firm\text{-}Trained_f$ is a dummy equal to one if the firm is in the FT treatment. We also estimate ATE impacts, where we instrument hiring a worker the firm is matched to with treatment assignment. y_{f0} is the firm outcome at baseline, $\mathbf{x}_{f0}$ are the firm's baseline covariates and λ_s and ϑ_t are strata and survey wave fixed effects respectively. We cluster standard errors by sector-BRAC branch, and to account for attrition we weight observations using IPWs, and present Lee Bounds. Finally, we examine dynamic impacts by estimating effects: (i) in the short run, only using the first firm follow-up survey conducted towards the end of the six-month FT training intervention; (ii) in the long run, averaging treatment effects over survey waves two to four, that run to years after wage subsidies have expired, and long after any initially hired workers have left FT firms (Figures 2 and 3).[53]

The results are in Table 9. Column 1 shows no evidence of other hires being crowded out by FT workers in the short run when the wage subsidy is in place. The change in the number of employees hired almost equals the number of post-intervention hires (Columns 1 and 2), and there is no evidence of more workers being fired post-intervention among FT firms (Column 3). There are also no long-run impacts on employment, hires or fires for FT firms relative to control firms (Columns 5 to 7). The (scaled-up) similarity in the pattern of results between the ITT and ATE estimates suggest impacts are driven by firms that hire a worker they are matched to.

In short, there is no employment displacement of other workers not in our evaluation sample in the short run, and there are no net employment effects of wage subsidies in the long run. This lack of employment crowd-in or crowd-out in the long run implies the partial and general equilibrium employment effects of the model coincide.[54]

[52] Measuring employment displacement effects of hiring VT workers was part of our original design with the VT+match treatment, T4. However, as described earlier, take-up rates in this treatment are too low to say anything about outcomes (even at first follow up) for firms that VT+match workers were hired by.

[53] $\mathbf{x}_{f0}$ controls include owner's gender and years of education, and firm size. The strata are BRAC branch and sector fixed effects. The instruments for the IPW estimates are a dummies for whether the respondent provided a phone number at baseline, and for whether he/she was an employee of the firm (rather than the firm owner or the manager), the number of network firms and dummies for interviewers at baseline.

[54] Our two-sided experimental design adds to a nascent literature examining impacts of wage subsidy programmes on firms (above and beyond the impacts on workers). De Mel *et al.* [2019] conduct a field experiment in Sri Lanka that provides wage subsidies to SMEs. They find firms increased employment during the subsidy period, but there was no lasting impact on employment, profitability, or sales. McKenzie *et al.* [2016] also find positive (short run) employment impacts of a wage subsidy and matching intervention with firms in Yemen. Hardy and McCasland [2017] evaluate an apprenticeship program with small firms in Ghana, and find firms retain this extra labor for at least six months, and earn higher profits in doing so.

One final point arises from these firm-side results: Column 8 shows positive long run profit impacts of having been earlier offered a worker with a wage subsidy. The impact corresponds to an ITT of 11% (or .113 x 183 = \$20.7 monthly increase in long run profits), with impacts increasing in magnitude over time. This cannot be attributed to the wage subsidy because that expires before this period. With the obvious caveat that the Lee Bounds on these profit estimate are wide, Table A11 provides suggestive evidence on what links short run exposure to workers with wage subsides to long run profits. This focuses on long run outcomes related to worker recruitment and shows: (i) FT firms report being significantly less constrained in being able to find skilled workers (Column 2); (ii) new firm hires are significantly more likely to be able to perform sector-specific tasks when hired (Column 3). This suggests long run changes in recruiting effort might link the original assignment of firms to the FT treatment to longer run profit increases.

6.2 Counterfactuals

We use our parameter estimates for controls and compliers to simulate counterfactuals (from Panel A of Table 8). We simulate two classes of counterfactual.[55]

6.2.1 Understanding the Relative Importance of Mechanisms

We assess the relative importance of the mechanisms at the heart of the model to explain steady state unemployment rates, earnings conditional on employment, and unconditional earnings. These mechanisms relate to differences in: (i) arrival rates of job offers (λ_0, λ_1); (ii) separation rates (δ); (iii) skills (s) that drive worker type ε. In the baseline model these parameters vary across treatments. For these counterfactuals, we hold two sets constant and allow only the third to vary with treatment. We thus assess the qualitative importance of each mechanism (due to the parameter interactions and non-linearities in the model, these are not exact decompositions).

The results are in Table 10. Panel A shows the baseline level of each outcome across controls, FT and VT compliers. Panel B compares FT and VT compliers each to control workers. Hence when we equate parameters we set them all to the value in the control group. Panel C compares compliers in FT and VT to each other, so when we equate parameters we set them all to the value

[55] To set the initial conditions for the simulations, we use the parameter estimates in Panel A of Table 8 and kernel density estimates of $F(r)$ and $h(\varepsilon)$ for each group of workers. We construct a common $F(r)$ function by taking an appropriately weighted average of the group-specific $F(r)$ functions. We then simulate a panel of $50,000$ workers observed over 48 months in steady-state, where workers are randomly assigned to treatment in the same proportions as in our experiment. As we showed earlier that compliance is uncorrelated with worker observables, workers are also randomly assigned to take-up their treatment in the same proportion as in the experiment. The simulation allows ($h(\varepsilon)$, δ, λ_0, λ_1) to vary across treatments, and in line with the baseline model, we assume all piece rate offers are drawn from the same $F(r)$ distribution across treatments. In each simulation, the average $G(r)$ is calculated as the mean piece-rate in the population of employed workers across the 48 months of the simulation. The average of $F(r)$ is calculated as the mean piece-rate received by workers transitioning from unemployment to employment. Final statistics are computed as the average results across 10 simulations.

for FT compliers. We do not conduct this comparison for earnings conditional on employment given Panel A shows minimal differences along this margin between training routes.

On unemployment rates: (i) Panel B shows the impact of firm-training over controls is mostly driven by lower separation rates; (ii) Panel C shows that the differential impact on unemployment of VT over FT is nearly all due to differences in job offer rates (explaining 112% of the gap). This confirms the central importance of skills certification for youth unemployment in this setting.

On earnings conditional on employment, Panel B shows: (i) the impact of FT over controls is mostly due to skill differentials, with separation rates also being relevant but qualitatively around one third as important; (ii) the impact of VT over controls is mostly due to skill differentials; separation rates are of less but still non-negligible importance in explaining this difference.

Combining both outcomes, on unconditional earnings we see that: (i) skill and separation rate differences are equally important in explaining the gap between FT and controls; (ii) all three mechanisms – skills, separation rates and job offer arrival rates – explain the gap between VT and controls; (iii) the gap between VT and FT workers is overwhelmingly due to differences in job offer arrival rates.

This exercise shows that: (i) all three mechanisms play an important role in explaining why training impacts outcomes over the control group; (ii) the key driver of differential outcomes between training routes along these important margins is the higher job offer arrival rates that VT workers experience when unemployed, in line with the greater certifiability of their skills, and the returns to such certifiability being higher when unemployed.

6.2.2 Extending Training to Other Workers in the Economy

We use the model to construct counterfactual treatment effects if the training interventions were targeted to other workers in the economy. The issue is relevant because workers were recruited into our evaluation sample based on the potential offer of vocational training, and eligibility criteria targeting disadvantaged youth. As Table A2 shows, relative to labor market active youth in Uganda, our sample is worse off in terms of labor market outcomes at baseline.

Workers in our sample might be especially acutely selected on two traits relative to other young labor market entrants: ability and patience. The first is relevant because our sample workers are unemployed at baseline with worse labor market histories. The second trait is relevant because only those prepared to forgo the opportunity cost of labor market offers during the six-months of training would have been willing to apply for our treatment offer. To build counterfactual scenarios targeting youth with different ability or patience than in our sample, we exploit the fact that the job ladder model has built in observed heterogeneity of workers type-ε, depending on their skill. To begin with, we examine heterogeneous skills accumulation in FT and VT treatments by ability and patience. In our sample, worker's ability is measured by Raven matrices tests, and worker's patience is measured using answers to a series of questions about their willingness to wait

to receive (hypothetical) monetary rewards. For each trait we then classify a worker of being high or low type depending on whether they are above/below the sample median.[56]

Table A12 shows the results on heterogenous skills accumulation. Column 1 highlights there are different levels of skill accumulation by high/low ability. This difference is statistically significant in the FT treatment ($p = .076$) and marginally so within the VT treatment ($p = .102$). Hence there is a complementarity between firm-provided training and underlying ability in the accumulation of sector-specific skills. As expected, Column 2 shows such complementarities also exist between firm-provided training and high/low patience workers, but these are less precisely estimated.

We use these estimates to construct counterfactuals assuming: (i) the distribution of worker types, $H(\varepsilon)$, varies by workers with above/below the median trait in controls, complier VT workers, and complier FT workers (so among six groups overall); (ii) there is a share θ of high trait (ability or patience) workers in the economy; (ii) the parameters (δ, λ_0, λ_1) are the same as in the baseline model for controls and compliers in each treatment (*irrespective* of their ability/patience, so traits only impact outcomes through skill accumulation). We then simulate counterfactual impacts in the economy varying the share of high-trait workers θ.

Recall that in the model, worker types-ε determine wages conditional on employment, but play no role for labor market transitions, so that unemployment rates do not vary with traits (or hence θ) by assumption. Therefore, we focus on simulated impacts on earnings conditional on employment, and unconditional earnings. These are in Figure 5 (along with 95% confidence intervals). Panels A and B vary the share of high ability workers treated. Panels C and D vary the share of high patience workers treated.

Panel A shows that the overall impact on earnings when employed for FT workers rises steeply in the share of high ability workers treated. At the extreme, if firm-training were only taken-up by low ability workers, treatment effects on earnings conditional on employment would be 5%, and if only high ability workers were targeted they would be closer to 20%. At this extreme, the point estimate treatment effect of FT would actually be higher than for VT. Panel B shows that for unconditional earnings, the impact varies from 20% to 38% as θ varies from zero to one.

Panels C and D show less pronounced impacts for patience, but treatment effect impacts for both training routes are increasing in θ. This is especially the case for earnings conditional on employment for FT: if only high patience workers were targeted, then the treatment effects of both training routes converge to be almost identical at 17%.

These counterfactuals offer an explanation for why studies in the literature might differ in their estimated returns to firm-provided apprenticeships: as θ varies across samples in the literature, the absolute impacts of training vary, and their relative ranking can also reverse.

We use this range of simulated estimates to discuss implied treatment impacts on two alternative groups of worker from our same context. To do so, we draw on worker samples from

[56] Patience is measured at baseline. Cognitive ability is measured at first follow-up. We verify that there are no direct treatment effect impacts on these traits.

related studies with comparable information on the cognitive abilities of workers. First, Bassi and Nansamba [2019] survey 1000 young workers currently receiving training in similar sectors and at similar VTIs in Uganda. These trainees have self-financed their vocational training. Exactly the same 10-question Raven matrices test was used to measure their cognitive ability. In their sample, the share of high ability workers is 60% (higher than in our sample). Second, Bassi *et al.* [2019] survey over 2000 employees in a representative sample of firms in welding, furniture making and grain milling sectors, operating in urban areas in Eastern, Central and Western Uganda. Cognitive ability was measured using a subset of the 10-question Raven matrices we used. Using the overlapping matrices, we find the share of high ability workers in this employee sample to be 49% (lower than in our sample).[57]

Panels A and B in Figure 5 superimpose the simulated impact our FT and VT treatments would have had they been targeted to workers self-financing their vocational training, or to workers currently employed in similar manufacturing sectors, under the assumption that the difference between samples is the share of high cognitive ability workers in each. Two important comparisons are as follows. First, for self-financed VTI attendees, the impacts of VT are only marginally higher than for our workers. This suggests factors associated with not being credit constrained (and so allowing workers to self-finance vocational training as in Bassi and Nansamba [2019]), are not much correlated to the returns to vocational training. Second, for currently employed workers, the impacts of firm-provided training are slightly lower than for apprenticeships in our sample. This suggests firms might not be hiring workers optimally. This is in line with evidence presented earlier that: (i) 52% of SME firm owners reported the inability to screen workers as a constraint; (ii) the firm side results showed potential long run impacts of SME's recruitment effort once they were exposed to workers in our FT treatment.

7 IRR and External Validity

7.1 IRR

The supply- and demand-side training interventions we evaluate are costly big-push style policies. Hence, it is important to establish whether the returns are sufficiently high to warrant a social planner implementing either policy. Table 11 presents IRR calculations for each treatment arm, where our benchmark case assumes a social discount rate of 5%, and that the steady state earnings gains to workers last for 15 years. In this exercise, we assume no employment displacement effects of either training route. For FT workers this is justified given the earlier firm-side results. For VT workers this is an implicit assumption we have to make. We present two sets of IRR calculations

[57] The share of workers scoring at the median or above in our sample when using only the four overlapping questions is 61%. The corresponding share in the Bassi *et al.* [2019] sample is 55%. We map this back to the percentage of high types in the original 10 Raven matrices by solving the following proportion: 61%:54% = 55%:X. Solving for X gives the 49% figure reported.

based on steady state earnings impacts on: (i) all workers (Columns 1 and 2); (ii) only workers that comply with their treatment (Columns 3 and 4). As before, the former is most appropriate from a social planner's point of view. The latter provides a sense of the private returns to workers if they could overcome constraints to make these kinds of human capital investment themselves.[58]

Panel A in Table 11 shows the per intended beneficiary cost breakdown of each treatment. Total costs comprise: (i) training costs: the cost per individual of vocational training was \$470, while the wage subsidy amounted to \$302 per trainee (\$50.3/month for six months); (ii) program overhead costs: these vary by treatment depending on whether worker-firm matches needed to be organized, the firm monitored etc.; (iii) the opportunity cost to workers of attending the vocational training: these turn out to be relatively small (comprising less than 10% of the total cost) because levels of youth unemployment and underemployment are so high.[59]

Panel B shows the NPV of the lifetime earning gains, as derived from the job ladder model, assuming these gains last 15 years. Focusing on the impacts for all workers we see that the gains to those assigned to FT are around 12% of those assigned to VT. However, the benefit cost-ratio is below one for FT (at .43), and the IRR is negative. It does *not* pay for the social planner to replicate the kind of subsidized apprenticeship offered in the FT treatment. However, the reason for this negative IRR is the low compliance in the FT treatment. This low compliance is driven by a lack of firms taking-up the offer of the wage subsidy and the matched to worker (workers are as likely to accept offers from firms as vocational trainees are to accept the offer of training from VTIs). Moreover, we documented that less profitable firms are more likely to take-on workers through the FT treatment, suggesting these select firms are financially constrained in hiring young job seekers and that this is an important demand-side constraint.

However, in these labor markets we do observe workers paying firms for an apprenticeship – exactly the kind of payment structure we set up in the FT treatment. To see why this is so, we redo the IRR calculations but based on the steady state earnings for compliers – namely those that acquire firm-provided or vocational training. For those workers that are hired and trained by firms under the FT treatment, Column 3 shows the benefit-cost ratio is well above one (2.80) and the IRR is 26%. The rise in IRR for FT workers highlights the high social returns from being able to overcome firm constraints in taking-on and hiring young workers. A core problem remains to design such interventions that raise take-up by firms. This might mean offering higher

[58] Table A13 shows the model estimates when we pool non-compliers and compliers in each treatment. Panel A shows that this narrows the gap in skill differences between control and FT workers, so they have similar $\overline{\varepsilon}$ values, while because of the high compliance in VT, it remains the case that VT workers have higher mean ε than controls ($p = .000$). The arrival rate of job offers when unemployed is still higher for VT than for FT ($p = .073$), and FT and controls have the same $\widehat{\lambda}_0$ estimate. Panel B shows qualitatively similar patterns on unemployment rate impacts and durations as the baseline model, and Panel C shows overall steady state earnings to be 5%, and 40% higher for all workers assigned to FT and VT with respect to those in Control.

[59] These cost structures are per intended beneficiary and do not change across the two sets of IRR calculations based on all workers or only those that comply with their treatment. This is because we take the view that such costs are incurred by the social planner *ex ante*, prior to compliance being observed.

subsidy rates, or some other incentive – such as perhaps making salient the potential long run profit impacts of being exposed to such workers (Table 9).

For the VT treatment, Column 2 shows that based on steady state gains for all workers assigned to this treatment, the cost benefit-cost ratio is well above one (2.58) and the IRR is 24%. Even with take-up rates of 68%, vocational training generates high returns under the assumptions made, and certainly compares favorably to a menu of other anti-poverty policies focusing on human capital accumulation. When using the benefits for those workers that actually take-up the vocational training treatment, Column 4 shows the benefit-cost ratio rises above 3 and the IRR rises to almost 34%.[60]

Panel C shows the sensitivity of these IRR estimates to alternative assumptions on: (i) the remaining productive life of beneficiaries; (ii) varying the foregone earnings from attending vocational training. We see that the IRR for FT drops off more quickly with shorter productive lives, while for VT it remains at 7.7% or above under the alternative scenarios. This is as expected given the different wage profiles to the interventions. The VT intervention always pays for itself within a decade. On foregone earnings, only under very extreme assumptions does the IRR for VT ever fall below 10%.

These calculations are based on the cost structure of the NGO BRAC that we collaborated with. This is an established NGO in Uganda. Hence their program overhead costs represent the *marginal* costs to them of extending their activities to the kind of training program evaluated. To get a more accurate sense of the social return of starting such programs from scratch, Panel D shows what the *total* cost per individual would have to be in order for the IRR to equal the social discount rate, 5%, focusing on the scenarios in Columns 2 to 4 where the baseline IRR is positive to begin with.

For vocational training, in Column 2 we see total costs per beneficiary would have to increase almost threefold for the intervention to break even. The final row performs the same calculation assuming a 10% social discount rate. In this case the costs for vocational training would still need to nearly double for the social planner not to intervene.

7.2 External Validity

In meta-analyses of training interventions in low-income settings, Blattman and Ralston [2015] and McKenzie [2017] document most interventions have a very low IRR. Figure 6 compares our ITT treatment impacts relative to the experimental studies discussed in McKenzie [2017], on employment and earnings outcomes. Our effect sizes are large relative to earlier studies, although

[60] Two further points are of note. First, these IRR figures match up quite well with the IRRs for the combined in-class vocational and on-the-job training intervention evaluated in Colombia by Attanasio *et al.* [2011, 2017]. Second, they likely underestimate the utility gains from each intervention as we measure benefits only through earnings, and take no account of reduced earnings risk, or how such human capital investments can reduce worker vulnerability to macroeconomic and other shocks over the longer term.

the ranking across treatment types is in line with earlier work. We speculate over four reasons why our returns are high relative to other studies, each of which opens up avenues for future work.

First, there are design issues: our experiment has a precise sectoral focus limited to eight sectors. All workers receive vocational training in one of these sectors, and all sampled firms operate in one of these sectors (and have at least one employee plus a owner at baseline). This reduces the possibility of worker-firm mismatch. Moreover, our treatments are intensive. Specifically, we separate out in-class vocational training from a wage subsidy program, both treatments last six months, the wage subsidy treatment had a subsidy rate higher than some other studies, and in the wage subsidy treatment firms were contractually obliged to train hired workers.

Second, only 13% of workers attrit over our four year evaluation. This attrition rate compares favorably to other studies such as Attanasio *et al.* [2011] (18%), and Card *et al.* [2011] (38%). Indeed, in the meta-analysis of McKenzie [2017], all but one study has attrition rates above 18%. As Figure 6 shows, other studies have similar or larger point estimates, but more imprecise treatment effects, that might in part arise from attrition. Moreover, our payment structures to VTIs ensured that the vast majority of workers completed training conditional on starting it, mitigating drop-out problems that earlier studies have faced.

Third, workers selected into our sample given the oversubscription design might differ from other young workers. We exploited this fact earlier to build counterfactual treatment effect estimates of targeting our interventions to others workers in the economy. The potential selection of unemployed youth into our evaluation has important implications for how we think about training interventions. Given youth unemployment rates of 60%, there might be an improvement in the allocation of talent in the economy if we think of the large pool of unemployed workers as heterogeneous, and those attracted to the sample through the offer of vocational training as being positively selected relative to the average unemployed youth in Uganda. It is exactly these kinds of motivated young jobseeker that the economic gains from matching to jobs might be highest for.

Fourth, we worked with a limited set of VTIs in Uganda, pre-selected to be of high quality based on their reputation. There is no shortage of VTIs in Uganda, and as in other low-income contexts, there are concerns over a long tail of low-quality training providers existing in equilibrium. Hence, although our treatments essentially relax credit constraints for workers, it is not obvious the results would be replicated through an unconditional cash transfer: this would rely on workers having knowledge over training providers. Rather a conditional cash transfer (conditioned on having to attend one of these VTIs) is likely to have higher returns, all else equal. This might explain why similar programs providing vouchers to workers redeemable at any training provider within the VTI market have met with more limited success [Galasso *et al.* 2004, Groh *et al.* 2016].

8 Conclusion

The development path of low income countries in the coming decades will largely depend on whether or not young workers can be matched to good jobs. High levels of youth unemployment are a symptom of the mismatch between supply and demand for labor in these countries – a growing mass of young, mainly unskilled workers are failing to find work in manufacturing and service sectors consisting mainly of small-scale firms.

This paper contributes to the classic literature on the value of human capital [Becker 1964, Schultz, 1981, 1993] by looking at whether and how different forms of worker training can ease the transition into manufacturing and service sector jobs. Transitions into the labor market mark a key stage in the life cycle, and a body of evidence documents how initial experiences and first job opportunities during this transition have persistent impacts on later employment trajectories and welfare [Becker 1994, Pissarides 1994]. This paper provides experimental and structural evidence on this transition from a novel two-sided experiment in the context of urban labor markets in a low-income country: Uganda.

Training of young workers whether through vocational training institutes or apprenticeships has a particular salience in low income economies for three main reasons: (i) very young populations imply that transitioning new workers into the labor market is the dominant challenge, (ii) the quality and duration of schooling is low and therefore young people are ill-equipped to access jobs in the manufacturing and service sectors of the economy and (iii) there are limited opportunities to use colleges, universities or other forms of tertiary education as a means of transitioning young people into good jobs.

What the paper reveals is that both types of training when provided over an extended period can have highly positive effects on employment and earnings within a disadvantaged set of young people transitioning into the labor market in Uganda. This is in sharp contrast with workers who receive neither type of training and who remain largely unemployed or employed in casual work, as is common among unskilled workers across the developing world.

What is even more revealing is that the steady state effects on employment and earnings for VT workers are almost twice as large as those for FT workers. This result speaks directly to value of the certifiability of skills, that is a key difference between skills gained through vocational training and those gained via firm-provided apprenticeships. Structural estimation of a job ladder model of worker search reveals labor market mobility as being the main mechanism for the divergence in employment and earnings profiles between VT and FT workers. In particular, VT workers receive significantly higher rates of job offers when unemployed. In other words, they are more likely to get back onto the job ladder if they fall off it and into unemployment. The lack of skills certification leads FT workers to have flatter earnings profiles as well as more limited labor market experience over the their working lives.

Two final points are of note. First, as exploited in the second counterfactual, the strength of

complementarity between cognitive ability and skills differs between FT and VT treatments. This hints at training being imparted differently at firms and VTIs: at firms higher ability workers learn more skills, while at VTIs the teaching approach seems to ensure a more workers gain skills. This is in line with the objectives of VTIs relative to firms. However, it is important for policy design to account for the differential targeting by worker ability by alternative skills providers in the economy.

Second, the costs of having uncertifiable skills is greater in lower income settings also because the firm size distribution is highly skewed in such economies: young workers are reliant on hiring by SMEs that have limited potential for promotion *within* firms, that would be a natural alternative to labor mobility between firms as the way to climb the job ladder.

These findings and implications open up a rich set of research possibilities for analyzing how vocational education might be best organized in these countries, how government might intervene to incentivize firms to provide apprenticeships, and how skills certification fits into this.[61]

A Appendix

A.1 Attrition

Table A3 presents evidence on the correlates of worker attrition. Attrition is low, with only 13% of workers attriting by the 48-month endline. Focusing on attrition between baseline and endline, Column 1 shows that: (i) attrition is uncorrelated to treatment assignment; (ii) worker characteristics do not predict attrition in general but workers that score higher on a cognitive ability test at baseline are more likely to attrit. Column 2 shows there to be little evidence of heterogeneous attrition across treatments by baseline cognitive scores at baseline. Any bias that might arise from selective attrition on unobservables cannot be signed *a priori*. Tracked workers would be negatively selected if attriters are more likely to find employment themselves, or they would be positively selected if attriters are least motivated to find work and remain attached to the labor market. To account for attrition, we weight our ITT estimates using inverse probability weights (IPWs). We also show the robustness of the main treatment impacts when using conditional Lee bounds [Lee 2009].

On the IPWs, we proceed as follows. At each survey wave t we define a dummy s_{it} such that we observe (y_{it}, x_{it}) for observations for which $s_{it} = 1$. We then first estimate a probit of s_{it} on z_{it} for each post-intervention survey wave separately, where z_{it} includes: (i) $\mathbf{x}_{i0}$: the vector of baseline covariates used as controls throughout in (1); (ii) strata and implementation round dummies; (iii) $\mathbf{z}_{i0}$, baseline measures excluded from regression analysis: dummies for orphan, anyone in household

[61] For example, if as economies develop, workers acquire credible means by which to certify their skills (both vocational and those acquired on-the-job) and to signal their labor market histories, then this might explain why in many high-income settings, training programs more commonly provide a combination of vocational training and apprenticeships, such as JTPA in the US and the YTS in the UK.

has a phone, willing to work in multiple sectors, and; (iv) the survey team the respondent was assigned to in each survey round ($Team_{it}$). The underlying assumption is that conditional on z_{it}, y_{it} is independent of s_{it}. $\hat{p}_{it}$ are fitted probabilities from this regression using survey wave t, and so at a second stage, we weight our OLS ITT estimates with weights $1/\hat{p}_{i1}$, $1/\hat{p}_{i2}$, $1/\hat{p}_{i3}$.

A.2 Robustness Checks

To conduct robustness checks we first combine multiple labor market outcomes into the same index shown in Columns 5 of Table 5. Column 1 of Table A7 repeats the baseline ITT estimates as a point of comparison. In addition to the ITT estimates, we also report conditional Lee bounds on the treatment effects (where we use the convention that the bound is underlined if it is statistically different from zero).[62]

Columns 2 and 3 split the labor market index by gender. Women have been found to benefit more from some training interventions [Friedlander 1997, Attanasio *et al.* 2011], although this finding is far from universal [McKenzie 2017]. We generally find larger ITT impacts on men. Columns 4 and 5 split treatment effects by sector: we generally find larger labor market impacts in manufacturing. Given the correlation between gender and sector (manufacturing sectors tend to be male dominated), it is hard to definitively separate out whether the impacts are driven by gender or sector. Fourth, we consider impacts in labor markets outside of Kampala, where 81% of workers reside: the result in Column 6 largely replicates the main findings.

Finally, we examine the sensitivity of the treatment effects to the timing of labor market entry. To do so, we exploit the fact that we have two batches of vocationally trained workers: the majority of trainees from the first round of applicants started training in January 2013. For logistical reasons, a second round of randomized-in applicants received vocational training between October 2013 and April 2014 (and so receive their training at the same time as when the apprenticeships are being implemented). In Column 7 we allow the impacts of vocational training to differ by the first and second batch of trainees: we see no evidence that workers in the second batch have different outcomes as measured by the labor market index.[63]

[62] We bound the treatment effect estimates using the trimming procedure proposed by Lee [2009]. The procedure trims observations from above (below) in the group with lower attrition, to equalize the number of observations in treatment and control groups. It then re-estimates the program impact in the trimmed sample to deliver the lower (upper) bounds for the true treatment effect. The bounding procedure relies on the assumptions that treatment is assigned randomly and that treatment affects attrition in only one direction so there are no heterogeneous effects of the treatment on attrition/selection, in line with the evidence in Table A3. As Lee [2009] discusses, using covariates to trim the samples yields tighter bounds. The covariates we use are the strata dummies.

[63] To further examine this concern, we also estimated employment rates in August 2013 (when VT workers were graduating from the VTIs and the FT treatment was being rolled out): we find no significant differences in employment rates between workers assigned to the FT, VT and control groups at that point. Moreover, recall that in terms of compliance with the FT treatment, the results in Table A5 already showed that being employed in August 2013 does not predict compliance (so workers that might have found jobs earlier are no less likely to still take up the FT treatment). This is robust to alternative specifications for compliance (Columns 1 to 4 of Table A5). Finally, descriptive evidence from the process reports collected just prior to the FT intervention shows that in

In Columns 2 to 7, in most cases the Lee bounds remain significantly different from zero.

Finally, the final two Columns show the robustness of the main results to dropping all covariates except baseline outcomes, randomization strata, and survey wave fixed effects, and to additionally not using IPWs.

A.3 Likelihood

We assume all random events $(\lambda_0, \lambda_1, \delta)$ are realizations of Poisson processes, so the residual durations are exponentially distributed. As unemployed workers are always assumed to be made job offers they accept, the unemployment spell hazard is λ_0. There are two competing causes of job spell termination: workers can be laid off (at rate δ), or workers can make a JJ transition (at rate $\lambda_1 \bar{F}(r)$). Hence the hazard rate of job spells with piece rate r is $(\delta + \lambda_1 \bar{F}(r))$. Thus, conditional on initial employment status ($e_i = 0$ or 1) and on an initial piece rate r_{i1}, the individual likelihood contributions are the following.

For type-ε employed workers in treatment group k:

$$\begin{aligned} l(\mathbf{x}_i|e_i = 1, \varepsilon_i, T_k) &= g(r_{1i}|T_k) \times (\delta + \lambda_1 \bar{F}(r_{1i}|T_k))^{(1-c_i)} e^{-(\delta+\lambda_1 \bar{F}(r_{1i}|T_k))d_i} \\ &\times \left(\frac{\delta}{\delta + \lambda_1 \bar{F}(r_{1i}|T_k)} \right)^{\tau_{JU_i}} \times \left(\frac{\lambda_1 \bar{F}(r_{1i}|T_k)}{\delta + \lambda_1 \bar{F}(r_{1i}|T_k)} \right)^{\tau_{JJ_i}}, \end{aligned} \tag{15}$$

where λ_0, λ_1 and δ are parametrized as in (11) to (13), and are therefore functions of the treatments, $g(\cdot)$ is the density of $G(\cdot)$, c_i is an indicator for right censoring, d_i is the duration (in months) of the spell, τ_{JU_i} is an indicator for job to unemployment transition, and τ_{JJ_i} is an indicator for job to job transition.

For unemployed workers:

$$l(\mathbf{x}_i|e_i = 0, \varepsilon_i, T_k) = \lambda_0^{1-c_i} e^{-\lambda_0 d_i} \times f(r_{0i}|T_k)^{1-c_i}, \tag{16}$$

where $f(\cdot)$ is the density of $F(\cdot)$.

Given there is no selection into employment conditional on training status T, the generic likelihood contribution of an observation $\mathbf{x_i}$ given its type ε and treatment group T_k is given by:

$$l(\mathbf{x}_i|\varepsilon_i, T_k) = \left(\frac{\lambda_0}{\delta + \lambda_0} (\mathbf{x}_i|e_i = 1, \varepsilon_i, T_k) \right)^{e_i} \times \left(\frac{\delta}{\delta + \lambda_0} (\mathbf{x}_i|e_i = 0, \varepsilon_i, T_k) \right)^{1-e_i}. \tag{17}$$

The likelihood is an explicit function of the transition parameters δ, λ_0, λ_1, and of both distributions $F(\cdot)$ and $G(\cdot)$. The empirical cross-sectional cdf of piece rates among employed workers at

the great majority of cases, workers were interested and willing to start training at the FT firms, so that selection is mostly on the firm side. Only a handful of workers reported not being interested in meeting a firm because they already had a job.

the initial sampling date provides a nonparametric estimator of $G(\cdot)$:

$$\hat{G}(r|T_k) = \frac{1}{\sum_i T_{ik}} \sum_i 1(r_{1i} \leq r) T_{ik}. \tag{18}$$

Under the steady state assumptions, the relationship between $F(\cdot)$ and $G(\cdot)$ provides a nonparametric estimator of the piece rate sampling distribution F, for any given value of λ_1 and δ:

$$\hat{F}(r|T_k) = \frac{(\delta + \lambda_1)\hat{G}(r|T_k)}{\delta + \lambda_1 \hat{G}(r|T_k)} \tag{19}$$

We use maximum likelihood to estimate the parameters δ, λ_0 and λ_1, and their asymptotic standard errors.

A.4 Robustness of the Model Estimates

In the baseline model, the distribution from which piece rate offers are drawn $F(\cdot)$ does not depend on treatment T: rather, all workers draw from this distribution, but once hired, workers are realized to be of higher type-ε, and paid a higher wage (at the same piece rate r). We now allow $F(\cdot)$ to also depend on compliance and treatment T. This enables us to investigate, in a very reduced form way, whether across treatments, workers search differently across firms in the economy who might then draw from different piece rate distributions.

Table A10 shows these results: only Panel C changes from the baseline model because we no longer impose a common $F(\cdot)$ across groups. By allowing for treatment specific piece rate offer distributions, we see differences in terms of offered wages, especially for complier FT workers. The mean offered wage is \$42 allowing for $F(r|T)$, while it was \$50 in our baseline model that assumed $F(r)$. For VT workers the means are far more similar (47 vs. 49). To understand what might drive this, recall the earlier results on FT compliance showed that firm characteristics predict whether a worker is taken on and trained by a firm offered a wage subsidy. Moreover, negatively selected firms (those with lower profits per worker) are more likely to hire the worker when offered a wage subsidy. These results suggest this *initial* match with a low productivity firm as part of the FT treatment might have *persistent* impacts on the wage offers these FT workers receive in steady state. This hysteresis shows up in the annual earnings impacts: these are 16% for FT workers, far lower than the baseline estimate of 32% (for VT workers the estimate of 49% is more similar to the baseline estimate of 55%). Indeed, the gap in earnings impacts of FT compliers and FT non-compliers narrows considerably (16% vs. 11%) while the earnings gap is stable between VT compliers and VT non-compliers. This kind of persistence might be suggestive of directed search of workers, and is something we study in greater detail on ongoing work [Bandiera *et al.* 2019].

We also conducted robustness checks examining how the estimates and simulated steady state impacts change with alternative $\widehat{\alpha}$ estimates: recall this parameter relates to how worker skills map

to worker productivity or type. The baseline results set $\widehat{\alpha} = .289$ from Column 1 of Table A8. We can also take the lowest and highest values of $\widehat{\alpha}$ from this table. Doing so reveals a qualitatively similar pattern of results. In particular, for both low and high $\widehat{\alpha}$: VT workers have significantly higher job offer arrival rates than FT workers when unemployed. The bottom line is that for low $\widehat{\alpha}$ the steady state earnings impacts are 28% for FT and 53% for VT; for high $\widehat{\alpha}$ these are 34% and 57% respectively. As we would expect, a higher $\widehat{\alpha}$ translates into larger earnings impacts because skills translate into higher productivity and wages.

References

[1] ABEBE.G, S.CARIA, M.FAFCHAMPS, P.FALCO, S.FRANKLIN AND S.QUINN (2018) Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City, mimeo, University of Bristol.

[2] ABEL.M, R.BURGER AND P.PIRAINO (2019) "The Value of Reference Letters: Experimental Evidence from South Africa," *American Economic Journal: Applied Economics*, forthcoming.

[3] ACEMOGLU.D AND J.S.PISCHKE (1998) "Why Do Firms Train? Theory and Evidence," *Quarterly Journal of Economics* 113: 79-119.

[4] ACEMOGLU.D AND J.S.PISCHKE (1999) "The Structure of Wages and Investment in General Training," *Journal of Political Economy* 107: 539-72.

[5] ADHVARYU.A, N.KALA AND A.NYSHADHAM (2017) The Skills to Pay the Bills: Returns to On-the-job Soft Skills Training, mimeo Michigan.

[6] ALTONJI.J.G AND C.R.PIERRET (2001) "Employer Learning and Statistical Discrimination," *Quarterly Journal of Economics* 116: 313-50.

[7] ANGRIST.J, E.BETTINGER AND M.KREMER (2006) "Long-term Educational Consequences of Secondary School Vouchers: Evidence from Administrative Records in Colombia," *American Economic Review* 96: 847-62.

[8] ATTANASIO.O, A.KUGLER AND C.MEGHIR (2011) "Subsidizing Vocational Training for Disadvantaged Youth in Colombia: Evidence from a Randomized Trial," *American Economic Journal: Applied Economics* 3: 188-220.

[9] ATTANASIO.O, A.GUARIN, C.MEDINA AND C.MEGHIR (2017) "Vocational Training for Disadvantaged Youth in Colombia: A Long Term Follow up," *American Economic Journal: Applied Economics* 9: 131-143.

[10] AUTOR.D.H (2001) "Why Do Temporary Help Firms Provide Free General Skills Training," *Quarterly Journal of Economics* 116: 1409-48.

[11] BANDIERA.O, V.BASSI, R.BURGESS, I.RASUL, M.SULAIMAN AND A.VITALI (2017) The Missing Middle? Experiment Evidence on Constraints to SME Expansion in Uganda, mimeo UCL.

[12] BANDIERA.O, V.BASSI, R.BURGESS, I.RASUL, M.SULAIMAN AND A.VITALI (2019) Skills, Signals and Job Search in Low-Income Labor Markets: Evidence from a Two-Sided Six Year Field Experiment, mimeo UCL.

[13] BARLEVY.G (2008) "Identification of Search Models Using Record Statistics," *Review of Economic Studies* 75: 29-64.

[14] BASSI.V AND A.NANSAMBA (2019) Screening and Signaling Non-Cognitive Skills: Experimental Evidence from Uganda, mimeo USC.

[15] BASSI.V, R.MUOIO, T.PORZIO, R.SEN AND E.TUGUME (2019) Achieving Scale Collectively: Evidence from Manufacturing Firms in Urban Uganda, mimeo USC.

[16] BEAM.E (2016) "Do Job Fairs Matter? Experimental Evidence on the Impact of Job Fair Attendance," *Journal of Development Economics* 120: 32-40.

[17] BECKER.G.S (1964) *Human Capital*, Chicago: University of Chicago Press.

[18] BECKER.G.S (1994) "Human Capital Revisited," in *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education*, Chicago: University of Chicago Press.

[19] BELL.B, R.BLUNDELL, AND J.VAN REENEN (1999) "Getting the Unemployed Back to Work: The Role of Targeted Wage Subsidies," *International Tax and Public Finance* 6: 339-60.

[20] BELL.D.N AND D.G.BLANCHFLOWER (2011) "Youth Unemployment in Europe and the United States," *Nordic Economic Policy Review* 1: 11-37.

[21] BERNIELL.L AND D.DE LA MATA (2016) Starting on the Right Track: Experimental Evidence from a Large-scale Apprenticeship Program, mimeo CAF.

[22] BLACK.D.A, J.A.SMITH, M.C.BERGER AND B.J.NOEL (2003) "Is the Threat of Reemployment Services More Effective than the Services Themselves? Evidence from Random Assignment in the UI System," *American Economic Review* 93: 1313-27.

[23] BLATTMAN.C AND L.RALSTON (2015) Generating Employment in Poor and Fragile States: Evidence from Labor Market and Entrepreneurship Programs, mimeo Chicago.

[24] BLUNDELL.R (2001) "Welfare Reform for Low Income Workers," *Oxford Economic Papers* 53: 189-214.

[25] BLUNDELL.R, L.DEARDEN, C.MEGHIR AND B.SIANESI (1999) "Human Capital Investment: The Returns from Education and Training to the Individual, the Firm and the Economy," *Fiscal Studies* 20: 1-23.

[26] BONTEMPS.C, J.M.ROBIN AND G.J.VAN DEN BERG (2000) "Equilibrium Search with Continuous Productivity Dispersion: Theory and Nonparametric Estimation," *International Economic Review* 41: 305-58.

[27] BURDETT.K AND D.T.MORTENSEN (1998) "Wage Differentials, Employer Size, and Unemployment," *International Economic Review* 39: 257-73.

[28] CARD.D, AND D.R.HYSLOP (2005) "Estimating the Effects of a Time-Limited Earnings Subsidy for Welfare-Leavers," *Econometrica* 73: 1723-70.

[29] CARD.D, P.IBARRAN, F.REGALIA, D.ROSAS-SHADY AND Y.SOARES (2011) "The Labor Market Impacts of Youth Training in the Dominican Republic," *Journal of Labor Economics* 29: 267-300.

[30] CARD.D, J.KLUVE AND A.WEBER (2015) What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations, NBER WP21431.

[31] CARRANZA.E, R.GARLICK, K.ORKIN AND N.RANKIN (2019) Job Search, Hiring, and Matching with Two-sided Limited Information about Workseekers' Skills, mimeo, Duke University.

[32] CONTI.G (2005) "Training, Productivity and Wages in Italy," *Labour Economics* 12: 557-76.

[33] CREPON.B, E.DUFLO, M.GURGAND, R.RATHELOT AND P.ZAMORA (2013) "Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment," *Quarterly Journal of Economics* 128: 531-80.

[34] CREPON.B AND P.PREMAND (2018) Creating New Positions? Direct and Indirect Effects of an Apprenticeship Program, mimeo, CREST.

[35] DASGUPTA.U, L.GANGADHARAN, P.MAITRA, S.MANI AND S.SUBRAMANIAN (2012) Choosing to Be Trained: Evidence from a Field Experiment, mimeo Wagner College.

[36] DE MEL.S, D.MCKENZIE AND C.WOODRUFF (2010) "Wage Subsidies for Microenterprises," *American Economic Review Papers and Proceedings* 100: 614-8.

[37] DE MEL.S, D.MCKENZIE AND C.WOODRUFF (2019) Labor Drops: Experimental Evidence on the Return to Additional Labor in Microenterprises, *American Economic Journal: Applied* 11: 202-35

[38] FARBER.H.S AND R.GIBBONS (1996) "Learning and Wage Dynamics," *Quarterly Journal of Economics* 111: 1007-47.

[39] FRANKLIN.S (2018) "Location, Search Costs and Youth Unemployment: Experimental Evidence from Transport Subsidies," *Economic Journal* 128: 2353-79.

[40] FRIEDLANDER.D, D.H.GREENBERG AND P.K.ROBINS (1997) "Evaluating Government Training Programs for the Economically Disadvantaged," *Journal of Economic Literature* 35: 1809-55.

[41] GALASSO.E, M.RAVALLION AND A.SALVIA (2004) "Assisting the Transition from Work-fare to Work: A Randomized Experiment," *Industrial and Labor Relations Review* 58: 128-42.

[42] GOLDSTONE.J.A (2002) "Population and Security: How Demographic Change can lead to Violent Conflict," *Journal of International Affairs* 56: 3-21.

[43] GROH.M, N.KRISHNAN, D.MCKENZIE AND T.VISHWANATH (2016) "Do Wage Subsidies Provide a Stepping Stone to Employment for Recent College Graduates? Evidence from a Randomized Experiment in Jordan," *Review of Economics and Statistics* 98: 488-502.

[44] GROH.M, D.MCKENZIE, N.SHAMMOUT AND T.VISHWANATH (2015) "Testing the Importance of Search Frictions and Matching through a Randomized Experiment in Jordan," *IZA Journal of Labor Economics* 4:7.

[45] HAM.J.C AND R.J.LALONDE (1996) "The Effect of Sample Selection and Initial Conditions in Duration Models: Evidence from Experimental Data on Training," *Econometrica* 64: 175-205.

[46] HARDY.M AND J.MCCASLAND (2017) Are Small Firms Labor Constrained? Experimental Evidence from Ghana, mimeo UBC.

[47] HECKMAN.J.J, R.J.LALONDE AND J.A.SMITH (1999) "The Economics and Econometrics of Active Labor Market Programs," *Handbook of Labor Economics* 3: 1865-2097.

[48] HOFFMAN.M AND S.V.BURKS (2019) "Worker Overconfidence: Field Evidence and Implications for Employee Turnover and Returns from Training," *Quantitative Economics*, forthcoming.

[49] HSIEH.C AND B.OLKEN (2014) "The "Missing" Missing Middle," *Journal of Economic Perspectives* 28: 89-108.

[50] JENSEN.R (2010) "The (Perceived) Returns to Education and the Demand for Schooling," *Quarterly Journal of Economics* 125: 515-48.

[51] KATZ.L.F (1998) "Wage Subsidies for the Disadvantaged," in R.Freeman and P.Gottschalk (eds.) *Generating Jobs,* Russell Sage Foundation, New York.

[52] KAUFMANN.K.M (2014) "Understanding the Income Gradient in College Attendance in Mexico: The Role of Heterogeneity in Expected Returns," *Quantitative Economics* 5: 583-630.

[53] LALONDE.R.J (1995) "The Promise of Public Sector-sponsored Training Programs," *Journal of Economic Perspectives* 9: 149-68.

[54] LAYARD.P.R.G. AND S.J.NICKELL (1980) "The Case for Subsidising Extra Jobs," *Economic Journal* 90: 51-73.

[55] LEE.D.S (2009) "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects," *Review of Economic Studies* 76: 1071-102.

[56] LEUVEN.E AND H.OOSTERBEEK (2008) "An Alternative Approach to Estimate the Wage Returns to Private-sector Training," *Journal of Applied Econometrics* 23: 423-34.

[57] LEVINSOHN.J., N.RANKIN, G.ROBERTS AND V.SCHOER (2014) Wage Subsidies and Youth Employment in South Africa: Evidence from a Randomised Control Trial, mimeo Yale.

[58] MACLEOD.W.B., E.RIEHL, J.E.SAAVEDRA AND M.URQUIOLA (2015) "The Big Sort: College Reputation and Labor Market Outcomes," *American Economic Journal: Applied Economics* 9: 223-61.

[59] MAITRA.P AND S.MANI (2017) "Learning and Earning: Evidence from a Randomized Evaluation in India," *Labour Economics* 45: 116-30.

[60] MCKENZIE.D (2017) "How Effective are Active Labor Market Policies in Developing Countries? A Critical Review of Recent Evidence," *World Bank Research Observer* 32: 127-54.

[61] MCKENZIE.D AND D.ROBALINO (2010) Jobs and the Crisis: What Has Been Done and Where to go From Here?, Viewpoint Note 325, World Bank Finance and Private Sector Development Vice Presidency.

[62] MCKENZIE.D, N.ASSAF AND A.P.CUSOLITO (2016) "The Demand for, and Impact of, Youth Internships: Evidence from a Randomized Experiment in Yemen," *IZA Journal of Labor and Development* 5: 1-15.

[63] MOSCARINI.G AND F.POSTEL-VINAY (2018) "The Cyclical Job Ladder," *Annual Review of Economics* 10: 165-88.

[64] PALLAIS.A (2014) "Inefficient Hiring in Entry-Level Labor Markets," *American Economic Review* 104: 3565-99.

[65] PISSARIDES.C.A (1994) "Search Unemployment with On-the-job Search," *Review of Economic Studies* 61: 457-75.

[66] ROMANO.J.P AND M.WOLF (2016) "Efficient Computation of Adjusted P-values for Resampling-based Stepdown Multiple Testing," *Statistics and Probability Letters* 113: 38-40.

[67] RUD.J.P. AND I.TRAPEZNIKOVA (2018) Wage Dispersion, Job Creation and Development: Evidence from Sub-Saharan Africa, mimeo RHUL.

[68] SCHULTZ.T.W (1981) *Investing in People*, University of California Press.

[69] SCHULTZ.T.W (1993) *The Economics of Being Poor*, Cambridge, Mass.: Blackwell.

[70] SIANESI.B (2004) "An Evaluation of the Swedish System of Active Labor Market Programs in the 1990s," *Review of Economics and Statistics* 86: 133-55.

[71] UN DESA POPULATION DIVISION (2017) *World Population Prospects: The 2017 Revision Population Database.*

[72] WORLD BANK (2009) *Africa Development Indicators 2008/9: Youth and Employment in Africa*, Washington DC: The World Bank.

Table 1: Baseline Balance on Worker Labor Market Outcomes

Means, robust standard errors from OLS regressions in parentheses
P-value on t-test of equality of means with control group in brackets
P-value on F-tests in braces

	Number of workers	Currently working	Has worked in the last month	Has done any wage employment in the last month	Any self employment in the last month	Has done any casual work in the last month	Total earnings in the last month [USD]	F-test of joint significance
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
All Workers	1714	.360	.383	.130	.046	.257	5.93	
		(.045)	(.044)	(.023)	(.013)	(.508)	(1.11)	
T1: Control	451	.381	.401	.120	.038	.296	5.11	
		(.049)	(.048)	(.025)	(.015)	(.047)	(1.27)	
T2: Firm Trained	283	.369	.387	.103	.064*	.266	6.44	{.968}
		(.035)	(.035)	(.023)	(.017)	(.032)	(1.35)	
		[.979]	[.985]	[.520]	[.096]	[.592]	[.239]	
T3: Vocationally Trained	390	.358	.389	.149	.034	.253	7.29*	{.881}
		(.032)	(.032)	(.023)	(.013)	(.029)	(1.26)	
		[.763]	[.990]	[.188]	[.802]	[.265]	[.063]	
T4: Vocationally Trained + Matched	307	.320	.360	.149	.050	.205*	5.25	{.758}
		(.033)	(.034)	(.026)	(.015)	(.030)	(1.20)	
		[.316]	[.747]	[.229]	[.266]	[.070]	[.758]	
T5: Untrained, Matched	283	.364	.367	.127	.057	.251	5.58	{.996}
		(.033)	(.034)	(.025)	(.016)	(.031)	(1.25)	
		[.707]	[.386]	[.821]	[.207]	[.210]	[.713]	
F-test of joint significance		{.882}	{.908}	{.301}	{.214}	{.433}	{.379}	

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. All data is from the baseline survey to workers. Column 1 reports the number of workers assigned to each treatment. Columns 2 to 7 report the mean value of each worker characteristic, derived from an OLS regression of the characteristic of interest on a series of dummy variables for each treatment group. All regressions include strata dummies and a dummy for the implementation round. The excluded (comparison) group in these regressions is the Control group. Robust standard errors are reported throughout. Column 8 reports the p-values from F-Tests of joint significance of all the regressors from an OLS regression where the dependent variable is a dummy variable taking value 0 if the worker is assigned to the Control group, and it takes value 1 for workers assigned to treatment group j (with j going from 2 to 5) and the independent variables are the variables in Columns 2 to 7. Robust standard errors are also calculated in these regressions. The p-values reported in the last row are from the F-test of joint significance of the treatment dummies in each Column regression where the sample includes all workers. In Column 6 casual work includes any work conducted in the following tasks where workers are hired on a daily basis: loading and unloading trucks, transporting goods on bicycles, fetching water, land fencing and slashing the compound. Casual work also include any type of agricultural labor such as farming, animal rearing, fishing and agricultural day labor. In Column 7 workers who report doing no work in the past month (or only did unpaid work in the last month) have a value of zero for total earnings. The top 1% of earnings values are excluded. All monetary variables are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD.

Table 2: The Mincerian Returns to Vocational Training, by Sector

		Worker is skilled: self-reported VTI attendance		
	Share of firms in sector	% workers skilled in sector	Coefficient and SE from worker wage regressions [USD]	Coefficient and SE from worker log(wage) regressions [USD]
	(1)	(2)	(3)	(4)
All Sectors		**31.0%**	**26.2*** **	**.515*** **
			(3.15)	(.045)
Manufacturing				
Welding	14.57%	24.9%	34.5***	.381***
			(6.40)	(.084)
Motor-mechanics	9.80%	23.5%	16.1*	.294*
			(9.41)	(.153)
Electrical wiring	6.37%	41.9%	27.3***	.486**
			(7.60)	(.189)
Construction	4.38%	28.8%	11.5	.289*
			(9.39)	(.170)
Plumbing	3.08%	49.1%	60.9***	.719**
			(19.0)	(.281)
Services				
Hairdressing	39.64%	29.2%	22.9***	.444***
			(5.97)	(.069)
Tailoring	14.96%	41.6%	15.9	.898***
			(9.76)	(.182)
Catering	7.20%	40.2%	26.8**	.330***
			(11.6)	(.109)

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. The data used is from the Census of firms, which includes 2309 firms and 6306 workers. A worker is defined as skilled if he/she was reported as having attended formal vocational training at any point in the past. Coefficients and standard errors in Columns 3 and 4 are from a regression of workers' total earnings in the last month (or the logarithm of workers' total earnings in the last month) on a dummy for being a skilled worker (as defined above). Control variables in these regressions include: employee's age and age squared, gender, tenure and tenure squared, firm size, BRAC branch dummies and firm sector dummies. Robust standard errors are reported. All monetary variables are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD. The top 1% wages and capital stock values are excluded.

Table 3: Characteristics of Apprenticeships

A. Availability	
Worker received on-the-job training at the current firm	.498
Duration of on-the-job training [months]	10
B. Payments	
In the first month of training, the worker:	
Was paid	.198
Was unpaid	.515
Was paying the firm owner	.288
Earnings (conditional on > 0) [US$] (median)	39.2 (40.1)
Amount worker was paying to owner (conditional on > 0) [US$] (median)	51.9 (33.3)
C. Trainers	
Who was mainly involved in training the worker:	
Firm owner only	.457
Other employees only	.091
Firm owner as well as other employees	.452

Notes: The data is from the first firm follow-up, and the sample is restricted to those workers employed in Control firms. The sample includes 955 workers employed in 332 firms. All monetary variables are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD. The top 1% monetary values are excluded.

Table 4: ATE Estimates, Training and Skills

2SLS regression coefficients, bootstrapped standard errors in parentheses

Bootstrap p-values in braces: unadjusted p-values (left) and Romano and Wolf [2016] adjusted p-values (right)

	Being Trained by Firms		Sector-Specific Skills		Skills Transferability
Treatment effects on:	**Received On the Job Training**	**Position is "Trainee"**	**Any Skills (0/1)**	**Test Score (0-100)**	
Measured at:	**First Job**	**First Job**	**Two-Three Years after Training**	**Two-Three Years after Training**	**Three Years after Training, Conditional on Employment**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Firm Trained	.544***	.658***	.422***	9.67*	-.072
	(.115)	(.104)	(.100)	(5.33)	(.341)
	{.003 ; .032}	{.001 ; .004}	{.001 ; .009}	{.087 ; .288}	{.821 ; .821}
Vocationally Trained	-.014	-.011	.407***	10.3***	.253**
	(.053)	(.033)	(.033)	(1.69)	(.104)
	{.421 ; .803}	{.482 ; .803}	{.001 ; .001}	{.001 ; .002}	{.048 ; .136}
Mean (SD) Outcome in Control Group	.404	.092	.596	30.1 (22.9)	-
Control for Baseline Value	No	No	No	No	No
P-values on tests of equality:					
Firm Trained = Vocationally Trained	[.000]	[.000]	[.864]	[.904]	[.264]
N. of observations	792	794	1,818	1,818	650

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. The data used is from the baseline and three worker follow-up surveys. We report 2SLS regression estimates, where treatment assignment is used as IV for treatment take-up. Treatment take-up is defined as a dummy equal to one if the worker (i) started firm training in FT or (ii) started vocational training in VT. Bootstrap standard errors are calculated using 1,000 replications. We also report p-values adjusted for multiple testing in braces. These are computed using the step-down procedure discussed in Romano and Wolf [2016], with 1,000 bootstrap replications. The number of hypotheses being tested simultaneously in the five columns is ten – for each outcome we test the impact of two treatments, FT and VT, and we do so for five outcomes. Therefore, the p-values in columns 1 to 5 are adjusted for testing on ten hypotheses. All regressions include strata dummies, survey wave dummies, a dummy for the implementation round and dummies for the month of interview. We also control for the following baseline characteristics of workers: age at baseline, a dummy for whether the worker was married at baseline, a dummy for whether the worker had any children at baseline, a dummy for whether the worker was employed at baseline, and a dummy for whether the worker scored at the median or above on the cognitive test administered at baseline. In Columns 1 and 2 we use information on the first employment spell reported by a worker in the post-intervention period (so the sample only includes workers that had at least one job in the post-intervention period). In Column 1 the dependent variable is a dummy=1 if the worker reported having received on the job training at her first employer. In Column 2 the dependent variable is a dummy=1 if the worker reported being a "Trainee" when asked about her position at her first employer. In Column 3 the dependent variable is a dummy for whether the respondent reports having any sector specific skills or not at second and third follow-up. In Column 4 the dependent variable is the skills test score, from the test administered to workers in the second and third worker follow-up. In Column 5 the dependent variable is based on a question on the perceived transferability of the skills learned at the current firm. This question is asked only to individuals who are working and is only available at third followup. The variable is standardized using the mean and standard deviation in Control. A higher value of the variable corresponds to more transferable skills. For the regressions in Columns 3 and 4, workers that reported not having any sector specific skills are assigned a test score equal to what they would have got had they answered the test at random. Workers that refused to take the skills test are excluded from the regressions in Columns 3-4. At the foot of each Column we report p-values on the null that the impact of the vocational training is equal to the impact of firm training.

Table 5: ITT Estimates, Labor Market Outcomes

OLS IPW regression coefficients and robust standard errors in parentheses
Bootstrap p-values in braces: unadjusted p-values (left) and Romano and Wolf [2016] adjusted p-values (right)

	Any paid work in the last month	Number of months worked in the last year	Hours worked in the last week	Total earnings in the last month [USD]	Labor market index	Worked in sector of training/matching in the last month
	(1)	(2)	(3)	(4)	(5)	(6)
Firm Trained	.063**	.518**	-.196	1.89	.105**	.045***
	(.025)	(.259)	(2.27)	(2.20)	(.051)	(.015)
	{.013 ; .042}	{.044 ; .128}	{.921 ; .921}	{.391 ; .581}	{.036 ; .036}	{.035 ; .035}
Vocationally Trained	.090***	.879***	3.76**	6.10***	.170***	.112***
	(.020)	(.207)	(1.84)	(1.80)	(.041)	(.013)
	{.001 ; .001}	{.001 ; .002}	{.040 ; .128}	{.001 ; .005}	{.001 ; .001}	{.001 ; .001}
Mean Outcome in Control Group	.438	4.52	28.2	24.7	.003	.067
Control for Baseline Value	Yes	No	Yes	Yes	Yes	Yes
P-values on tests of equality:						
Firm Trained = Vocationally Trained	[.255]	[.134]	[.059]	[.048]	[.169]	[.000]
N. of observations	3,256	3,256	2,057	3,115	3,256	3,256

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. The data used is from the baseline and three worker follow-up surveys in all columns apart from Column 3 which uses data from baseline and second and third follow-up, since hours worked in the last month are not available at first follow-up. We report OLS IPW regression estimates in all Columns, together with robust standard errors. We also report p-values adjusted for multiple testing in braces. These are computed using the step-down procedure discussed in Romano and Wolf [2016], with 1,000 bootstrap replications. The number of hypotheses being tested simultaneously in the first four columns is eight – for each outcome we test the impact of two treatments, FT and VT, and we do so for four outcomes. Therefore, the p-values in columns 1 to 4 are adjusted for testing on eight hypotheses. In addition, we adjust the p-values in column 5 for two hypotheses, that is, for testing two treatments at the same time. We again adjust the p-values for testing on two hypotheses in column 6. All regressions control for the value of the outcome at baseline (except in Column 2), as well as strata dummies, survey wave dummies, a dummy for the implementation round and dummies for the month of interview. We also control for the following baseline characteristics of workers: age at baseline, a dummy for whether the worker was married at baseline, a dummy for whether the worker had any children at baseline, a dummy for whether the worker was employed at baseline, and a dummy for whether the worker scored at the median or above on the cognitive test administered at baseline. In Column 4 the dependent variable is total earnings from any wage or self-employment in the last month. Individuals reporting no wage employment earnings and no self-employment earnings are assigned a value of zero. The top 1% of earnings values are excluded. All monetary variables are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD. The Labor Market Index in column 5 is a standardized index of the variables in Columns 1-4, where we use the procedure discussed in Anderson [2008] to construct the index. In Column 6 the dependent variable is a dummy equal to one if the individual conducted any work in the last month in: the sector of training (for VT); the sector of matching (for FT); either the first or second most preferred sectors of employment, as indicated in the baseline survey (for Control). The weights for the Inverse Probability Weights (IPW) are computed separately for attrition at first, second and third follow-up. The instruments for the IPW estimates are whether the worker was an orphan at baseline, a dummy=1 if anyone in the household of the worker reported having a phone at baseline, a dummy for whether the worker reported being willing to work in more than one sector at the time of their original application to the VTIs and dummies for the survey team the worker's interview was assigned to in each of the three follow-up survey rounds. At the foot of each Column we report p-values on the null that the impact of the vocational training is equal to the impact of firm training.

Table 6: ATE Estimates, Labor Market Outcomes

2SLS regression coefficients, bootstrapped standard errors in parentheses
Bootstrap p-values in braces: unadjusted p-values (left) and Romano and Wolf [2016] adjusted p-values (right)

Dependent variable:	Any paid work in the last month	Number of months worked in the last year	Hours worked in the last week	Total earnings in the last month [USD]	Labor market index	Worked in sector of training/matching in the last month
	(1)	(2)	(3)	(4)	(5)	(6)
Firm Trained	.246***	2.31**	4.13	11.9	.473**	.245***
	(.085)	(.917)	(7.56)	(8.08)	(.176)	(.062)
	{.004 ; .023}	{.013 ; .029}	{.662 ; .662}	{.145 ; .241}	{.010 ; .010}	{.001 ; .001}
Vocationally Trained	.135***	1.38***	7.12**	10.3***	.272***	.190***
	(.028)	(.302)	(2.61)	(2.65)	(.059)	(.019)
	{.001 ; .001}	{.001 ; .001}	{.013 ; .026}	{.001 ; .001}	{.001 ; .001}	{.001 ; .001}
Mean Outcome in Control Group	.438	4.52	28.2	24.7	.003	.067
Control for Baseline Value	Yes	No	Yes	Yes	Yes	Yes
P-values on tests of equality:						
Firm Trained = Vocationally Trained	[.141]	[.255]	[.661]	[.830]	[.202]	[.343]
N. of observations	3,256	3,256	2,057	3,115	3,256	3,256

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. The data used is from the baseline and three worker follow-up surveys in all columns apart from column 3 which uses data from baseline and second and third follow-up, since hours worked in the last month are not available at first follow-up. We report 2SLS regression estimates in columns 1-6, where treatment assignment is used as IV for treatment take-up. Treatment take-up is defined as a dummy equal to one if the worker (i) started firm training in FT or (ii) started vocational training in VT. Bootstrap standard errors are calculated using 1,000 replications. We also report p-values adjusted for multiple testing in braces. These are computed using the step-down procedure discussed in Romano and Wolf [2016], with 1,000 bootstrap replications. The number of hypotheses being tested simultaneously in the first four columns is eight – for each outcome we test the impact of two treatments, FT and VT, and we do so for four outcomes. Therefore, the p-values in columns 1 to 4 are adjusted for testing on eight hypotheses. In addition, we adjust the p-values in column 5 for two hypotheses, that is, for testing two treatments at the same time. We again adjust the p-values for testing on two hypotheses in column 6. All regressions control for the value of the outcome at baseline (except in Column 2), as well as strata dummies, survey wave dummies, a dummy for the implementation round and dummies for the month of interview. We also control for the following baseline characteristics of workers: age at baseline, a dummy for whether the worker was married at baseline, a dummy for whether the worker had any children at baseline, a dummy for whether the worker was employed at baseline, and a dummy for whether the worker scored at the median or above on the cognitive test administered at baseline. In Column 4 the dependent variable is total earnings from any wage or self-employment in the last month. Individuals reporting no wage employment earnings and no self-employment earnings are assigned a value of zero. The top 1% of earnings values are excluded. All monetary variables are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD. The Labor Market Index in column 5 is a standardized index of the variables in columns 1-4, where we use the procedure discussed in Anderson [2008] to construct the index. In column 6 the dependent variable is a dummy equal to one if the individual conducted any work in the last month in: the sector of training (for VT); the sector of matching (for FT); either the first or second most preferred sectors of employment, as indicated in the baseline survey (for Control). At the foot of each Column we report p-values on the null that the impact of the vocational training is equal to the impact of firm training.

Table 7: ITT Estimates, Labor Mobility

OLS regression coefficients, robust standard errors in parentheses

	Number of non-casual work spells	Total months of non-casual work	Non-casual work in first spell	Number of UJ transitions	Number of JJ transitions	Number of JU transitions
Sample of workers:	**All**	**All**	**All**	**Unemployed in First Spell**	**Employed in First Spell**	**Employed in First Spell**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Firm Trained	-.036	-.016	-.001	-.013	.106	-.238**
	(.063)	(.744)	(.032)	(.055)	(.103)	(.119)
Vocationally Trained	.184***	2.24***	.069**	.122***	.063	-.147
	(.054)	(.610)	(.027)	(.046)	(.079)	(.093)
Mean Outcome in Control Group	.947	9.73	.168	.734	.3	.6
P-value FT=VT	.000	.001	.023	.011	.655	.403
Observations	1,122	1,122	1,122	894	228	228

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. We report OLS regression coefficients, and robust standard errors in parentheses. The dataset is a cross-section of workers built using data from the second and third worker follow-up surveys. For each worker it contains information on: spell type (employment, unemployment), spell duration (in months), earnings in employment spells (in USD), dates of transitions between spells and type of transition: (i) job to unemployment, (ii) unemployment to job, or (iii) job to job. The initial spell is identified as the (employment or unemployment) spell that was ongoing in November 2015. Spells are right censored at the date of the third follow-up interview (which ended in December 2016). Spells are left censored at 1 August 2014. Casual and agricultural occupations are coded as unemployment. Self-employment is coded as employment (but self-employment spells are assigned a separate spell). The sample is restricted to individuals who were unemployed in their first spell in Column 4, and individuals who were employed in their first spell in Columns 5 and 6. The outcomes in Columns 4, 5 and 6 are respectively the number of unemployment-to-job, job-to-job and job-to-unemployment transitions the worker experienced between November 2015 and the date of the interview. All regressions include strata dummies and a dummy for the implementation round. We also control for the following baseline characteristics of workers: age at baseline, a dummy for whether the worker was married at baseline, a dummy for whether the worker had any children at baseline, a dummy for whether the worker was employed at baseline, and a dummy for whether the worker scored at the median or above on the cognitive test administered at baseline. At the foot of each Column we report p-values on the null that the impact of the vocational training is equal to the impact of firm training.

Table 8: Baseline Estimates of the Job Ladder Search Model

Two-step estimation procedure in Bontemps, Robin and van den Berg [2000]
Asymptotic standard errors in parentheses
Steady State: November 2015 (Data from Second and Third Follow Up)

			Non-Compliers		Compliers	
Panel A: Parameter Estimates (Monthly)		**Control**	**Firm Trained**	**Vocationally Trained**	**Firm Trained**	**Vocationally Trained**
		(1)	**(2)**	**(3)**	**(4)**	**(5)**
Average units of effective labor [USD]		2.51	2.48	2.56	2.92	2.84
Job destruction rate, δ		.027	.028	.026	.023	.023
		(.003)	(.006)	(.005)	(.007)	(.004)
Arrival rate of job offers if UNEMPLOYED, λ_0		.019	.019	.018	.020	.028
		(.002)	(.003)	(.003)	(.005)	(.003)
Arrival rate of job offers if EMPLOYED, λ_1		.040	.039	.054	.032	.039
		(.010)	(.019)	(.022)	(.022)	(.013)
Panel B: Competition for Workers and Unemployment						
Interfirm competition for workers		1.46	1.43	2.08	1.41	1.67
	% Impact:		**-2.24%**	**42.1%**	**-3.33%**	**14.5%**
Unemployment rate		.589	.592	.593	.531	.455
	% Impact:		**.618%**	**.823%**	**-9.80%**	**-22.7%**
Unemployment duration (months)		52.9	52.7	56.2	50.0	35.9
	% Impact:		**-.359%**	**6.15%**	**-5.48%**	**-32.2%**
Employment duration (months)		37.0	36.3	38.5	44.2	42.9
	% Impact:		**-1.85%**	**4.04%**	**19.47%**	**16.1%**
Panel C: Wages and Earnings						
Average monthly OFFERED wages [USD]		43.0	42.5	43.9	50.0	48.6
Average monthly ACCEPTED wages [USD]		63.2	62.1	70.4	72.9	73.8
Impact on annual earnings [USD]			**-7.64**	**32.36**	**98.6**	**171.1**
	% Impact:		**-2.5%**	**10%**	**32%**	**55%**

Notes: The dataset is a cross-section of workers, and for each worker it contains information on: spell type (employment, unemployment), spell duration (in months), earnings in employment spells (in USD), dates of transitions between spells and type of transition: (i) job to unemployment, (ii) unemployment to job, or (iii) job to job. Wages are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD. The dataset contains at most two spells (and one transition) per individual. The data comes from the second and third follow-up survey of workers, and the initial spell is identified as the (employment or unemployment) spell that was ongoing in November 2015. Spells are right censored at the date of the third follow-up interview (which ended in December 2016). Spells are left censored at 1 August 2014. Casual and agricultural occupations are coded as unemployment. Self-employment is coded as employment (but self-employment spells are assigned a separate spell). The estimation protocol follows the two-step procedure in Bontemps, Robin and van den Berg [2000]: in the first step the G function is estimated non-parametrically from the data (so this is just the empirical CDF of observed wages for those workers that are employed in their first spell), and is then substituted into the likelihood function. In the second step, maximum likelihood is then conducted using information from both the first and second spells for each individual to recover the parameter estimates. As shown in Panel A, we estimate separate parameters for Control and Treatment groups, and, within treatments, for compliers and non-compliers. Outputs in Panel B are derived from the model and computed as functions of the estimated parameters: (i) interfirm competition for workers=λ_1/δ; (ii) unemployment rate=$\delta/(\delta+\lambda_0)$; (iii) unemployment duration=$1/\lambda_0$; employment duration=$1/\delta$. In Panel C average monthly offered and accepted wages are computed as the product of average offered and accepted piece-rates, and average units of effective labor. We assume workers draw piece-rates from the same offer distribution F(r). F(r) is the kernel density estimate of a weighted average of the distributions of offered piece-rates across treatments - F(r|T) - where such distributions are obtained from their steady-state relationship with non-parametrically estimated G(r|T). Weights are equal the share of individuals in each treatment. For each treatment we then re-invert F(r) using estimated parameters and steady-state relationships to obtain G(r|T) under the assumption that workers draw piece-rates from the same offer distribution.

Table 9: Employment and Other Firm Outcomes

IPW regression coefficients, standard errors clustered by sector-branch in parenthesis, Lee Bounds in brackets

	Short Run (first follow-up)				Long Run (second to fourth follow-ups)			
	Number of Employees	Number of Post-intervention Hires	Number of Post-intervention Fires	Log (Average Monthly Profits)	Number of Employees	Number of Post-intervention Hires	Number of Post-intervention Fires	Log (Average Monthly Profits)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
PANEL A: ITT Estimates								
Firm Trained	.350*	.370***	-.118	.011	-.116	-.054	-.093	.113**
	(.205)	(.137)	(.160)	(.114)	(.154)	(.077)	(.150)	(.050)
	[.553 ; 1.16]	[.430 ; .668]	[-.272 ; .111]	[-.089 ; .204]	[-.133 ; .237]	[-.087 ; .176]	[-.007 ; .435]	[-.069 ; .188]
PANEL B: ATE Estimates								
Firm Trained	1.343*	1.417***	-.453	.036	-.358	-.127	-.182	.313*
	(.770)	(.441)	(.623)	(.375)	(.431)	(.229)	(.451)	(.169)
Mean outcome in Control firms	2.41	.647	.647	209	2.29	.889	.889	183
Number of observations	569	569	569	444	1,611	1,606	1,611	1,178

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. The data used is from the firm follow-up data surveys. Panel A reports OLS IPW regression estimates together with standard errors adjusted for heteroskedasticity and clustered at the branch-trade level in parenthesis. We report Lee [2009] bounds in brackets, where we implement a conditional Lee Bounds procedure that is able to condition on dummies for the interview round and baseline trade. Underlined bounds are significantly different from zero at the 95% condfidence level. Panel B reports 2SLS IPW regression estimates, where treatment assignment is used as IV for treatment take-up. Treatment take-up is defined as a dummy equal to one if the firm hired one of the workers it was matched with. All regressions control for the value of outcome at baseline (when available), and include branch and trade fixed-effects, survey wave dummies and dummies for the month of interview. Baseline controls also include the owner's sex, business age (measured as number of years since the business was established) and business age squared, firm size and owner's years of education. The weights for the Inverse Probability Weights (IPW) are computed separately for attrition at second, third and fourth follow-up. The instruments for the IPW estimates are dummies for whether the respondent provided a phone number at baseline, and for whether he/she was an employee of the firm (rather than the firm owner or the manager), the number of network firms and dummies for interviewers at baseline. All monetary amounts are deflated and expressed in terms of the price level in January 2013 using the monthly Producer Price Index for the manufacturing sector (local market), published by the Uganda Bureau of Statistics. The monetary amounts are then converted in January 2013 USD (1USD=2385UGX). Monthly profits and revenues are truncated at 99th percentile.

Table 10: Counterfactual Analysis on Relative Importance of Mechanisms

	Unemployment			Earnings Conditional on Employment			Unconditional Earnings		
	Different Arrival Rates	Different Separation Rates	Different Skills	Different Arrival Rates	Different Separation Rates	Different Skills	Different Arrival Rates	Different Separation Rates	Different Skills
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Baseline Levels									
Control		.589			64.3			26.4	
Firm Trained		.531			74.0			34.7	
Vocationally Trained		.455			75.1			40.9	
Panel B: FT=VT=Control									
Control									
Firm Trained	***25%***	***74%***	***0%***	***-64%***	***44%***	***145%***	***-14%***	***64%***	***68%***
Vocationally Trained	***73%***	***28%***	***0%***	***-2%***	***27%***	***80%***	***51%***	***30%***	***32%***
Panel C: FT=VT									
Vocationally Trained	***112%***	***-10%***	***0%***	***-***	***-***	***-***	***138%***	***-12%***	***-17%***

Notes: The table reports OLS estimates from simulated data generated from the model. We run 10 simulations of the behavior of 50,000 workers followed over a period of 48 months. In each simulation, we randomly assign individuals to treatment in the same proportions as in our experiment. Workers are also randomly assigned to take-up their treatment in the same proportion as in the experiment. In each simulation we calculate treatment effects as the average monthly impact of FT and VT on employment and earnings across the 48 months from OLS regressions. We then aggregate estimates across the different simulations. In Columns 1 to 3 we allow arrival rates λ_0 and λ_1, separation rates δ and the distribution of effective units of labor $h(\varepsilon)$ to vary across treatments. In all the remaining Columns, Panel A shows percentage changes in treatment effects beween the baseline and the counterfactual simulations when we set the parameters indicated at the top of the table to be the same for individuals in the FT, VT and Control group. In Panel B we set the parameters of FT workers to be equal to those of VT workers. So, in Panel B the parameters of individuals in VT and Control remain the same as in the baseline simulation (Columns 1 to 3). In Columns 1 to 3 we set arrival rates λ_0 and λ_1 to be equal across treatments. In Columns 4 to 6 we set separation rates δ to be equal across treatments. In Columns 7 to 9 we set the distribution of effective units of labor $h(\varepsilon)$ to be equal across treatments. The percentages in Panel A are calculated as the percentage change in FT and VT coefficients between baseline and counterfactual simulation. The percentages in Panel B are instead calculated as the percentage change in the difference between the VT and FT coefficients in the baseline and counterfactual simulations.

Table 11: Internal Rate of Return

	All Workers		Compliers	
	Firm Trained	Vocationally Trained	Firm Trained	Vocationally Trained
	(1)	(2)	(3)	(4)
Social discount rate = 5%				
Remaining expected productive life of beneficiaries	15 years	15 years	15 years	15 years
Panel A. External parameters				
Total cost per individual at year 0 [USD]:	**368**	**510**	**368**	**510**
(i) Training costs (for 6 months)	302	470	302	470
(ii) Program overheads costs	31	4	31	4
(iii) Foregone earnings (for 6 months) - average at baseline	36	36	36	36
Panel B. Estimated total earnings benefits				
1 NPV change in steady state earnings (from model estimates)	159	1315	1032	1810
2 Benefits/cost ratio	**.431**	**2.58**	**2.80**	**3.55**
3 Internal Rate of Return (IRR)	**-.054**	**.239**	**.262**	**.338**
Panel C. Sensitivity				
Sensitivity to different expected remaining productive life of beneficiaries				
Remaining expected productive life = 10 years	**-.134**	**.212**	**.238**	**.321**
Remaining expected productive life = 5 years	**-.367**	**.077**	**.109**	**.211**
Sensitivity to different earnings				
Foregone earnings = 90th percentile at baseline (120USD)	**-.075**	**.199**	**.206**	**.287**
Foregone earnings = double 90th percentile at baseline (241USD)	**-.097**	**.157**	**.153**	**.233**
Foregone earnings = max earnings at baseline (794USD)	**-.154**	**.055**	**.037**	**.108**
Foregone earnings = double max earnings at baseline (1588USD)	**-.194**	**-.010**	**-.030**	**.031**
Panel D. Program Costs for IRR to equate social discount rate				
5 Total cost per individual at year 0 [USD]	**-**	**1315**	**1032**	**1810**
Sensitivity to different discount rates/time horizons				
Social discount rate = 10%	-	963	756	1327

Notes: The Vocationally Trained group combines both T3 and T4. Forgone earnings are calculated as the average monthly earnings at baseline (6 USD) multiplied by six (as the duration of both types of training was six months). The computation of the IRR uses as input for the benefit the ITT and ATE impacts of Firm and Vocational Training on annual income from the structural model. All monetary variables are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD.

Figure 1: Experimental Design

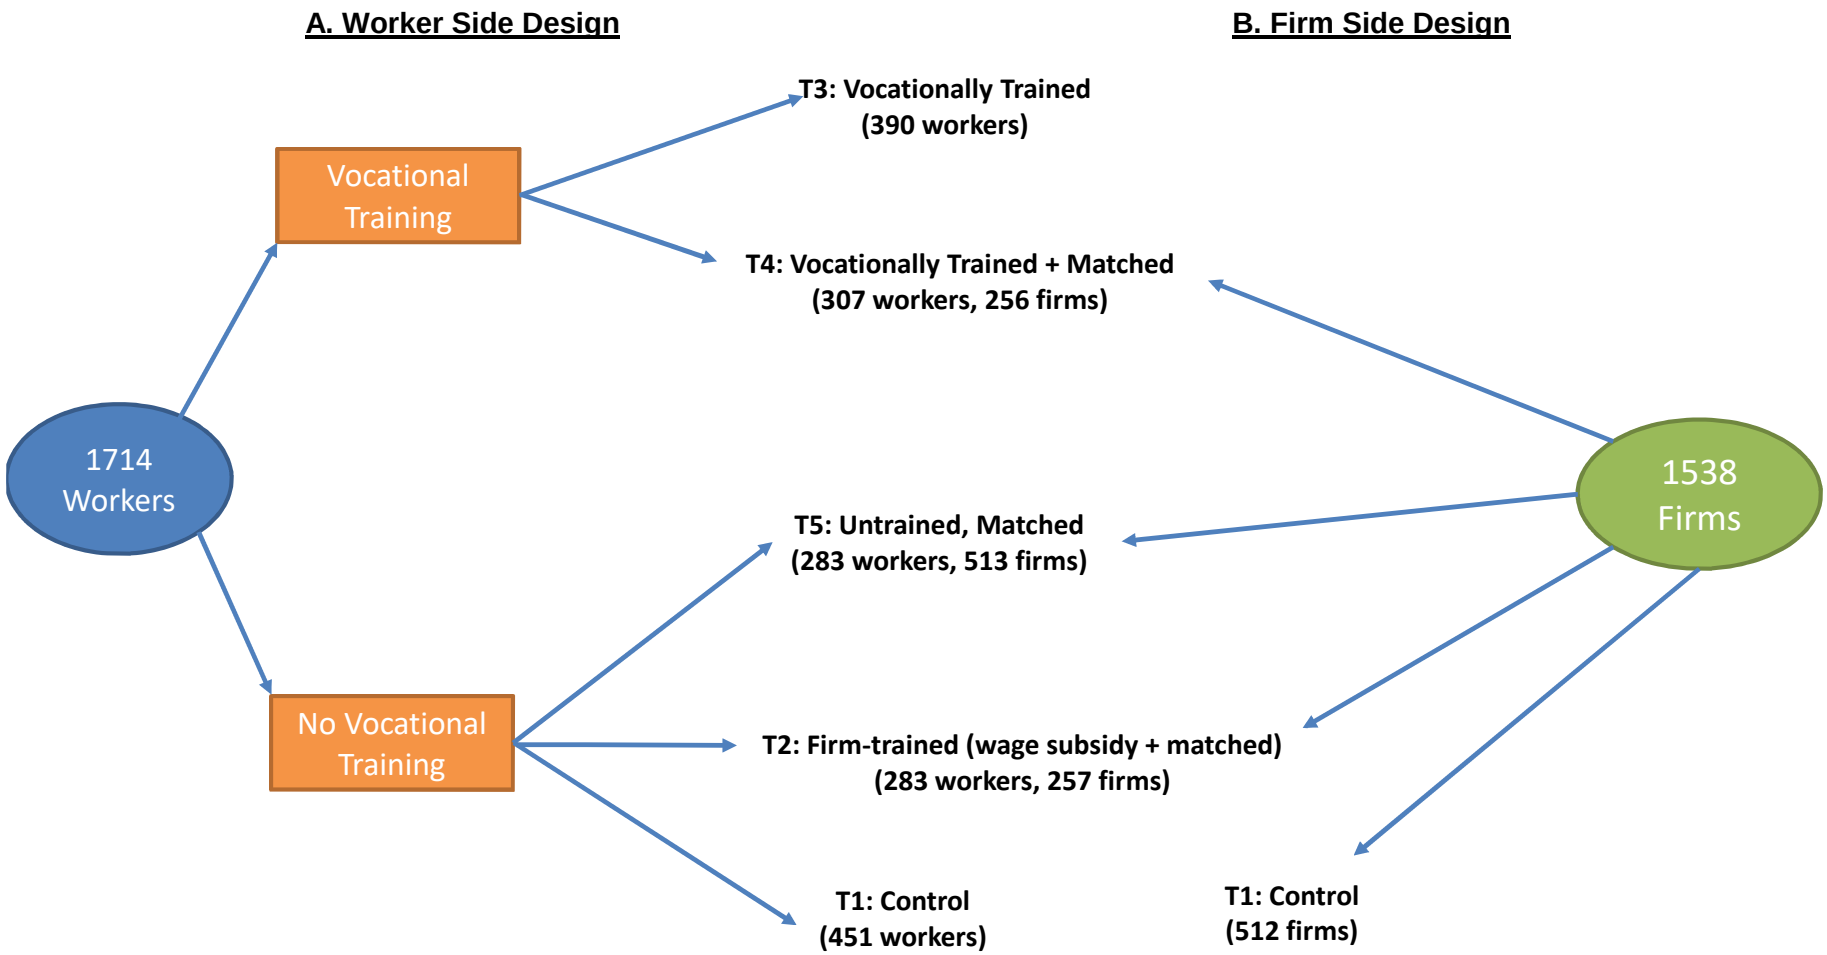

Note: Numbers in parentheses refer to the number of applicants originally assigned to each treatment, and the number of firms assigned to each treatment.

Figure 2: Timeline

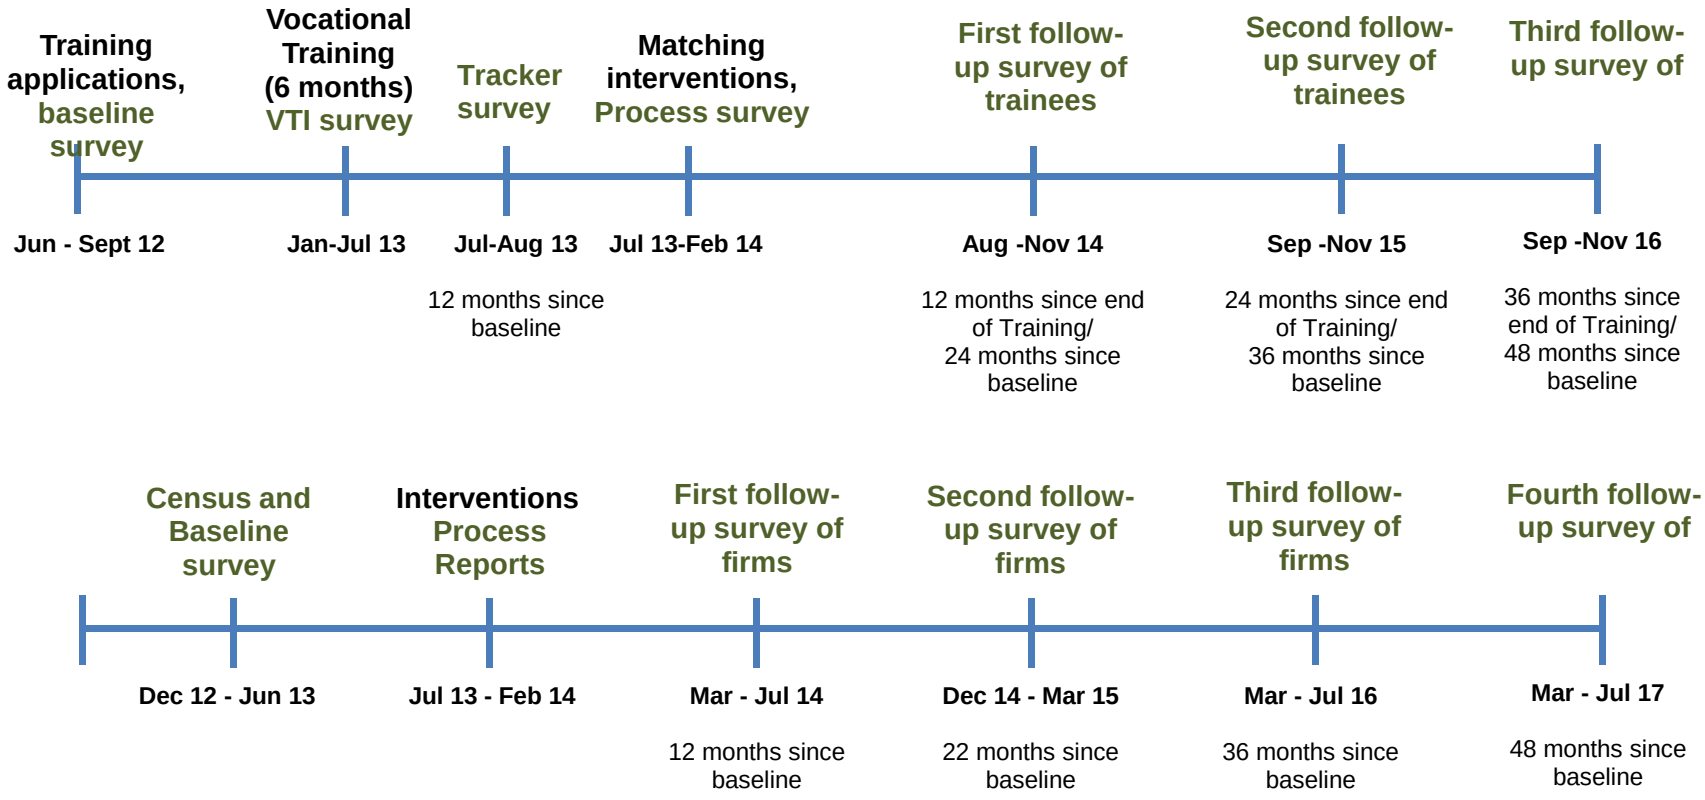

Notes: The timeline highlights the dates relevant for the main batch of worker applications and baseline surveys. A second smaller round of applications and baseline surveys were conducted in May and June 2013. The majority of trainees from the first round of applicants started training in January 2013, as shown in the timeline. For logistical reasons, a smaller group received training between April and October 2013. The trainees from the second round of applications received vocational training between October 2013 and April 2014. VTI surveys were collected towards the end of the training period while trainees were still enrolled at the VTIs. Workers from the second round of applicants were not included in the Tracker Survey. The remaining interventions (the matching treatments and firm training placements) and all follow-up surveys were conducted at the same time for workers from the first and second round of applicants. On the firms' timeline, the firm level interventions include: Matching, Vocational Training + Matching, and Firm Training. There were two rounds of Matching and Vocational Training + Matching interventions, in line with the two batches of trainees from the vocational training institutes. The first round of the Vocational training + Matching interventions took place in August-September 2013. The second round took place in December 2013-February 2014. The Firm Training intervention took place in September-November 2013.

Figure 3: Survival Analysis for Employment

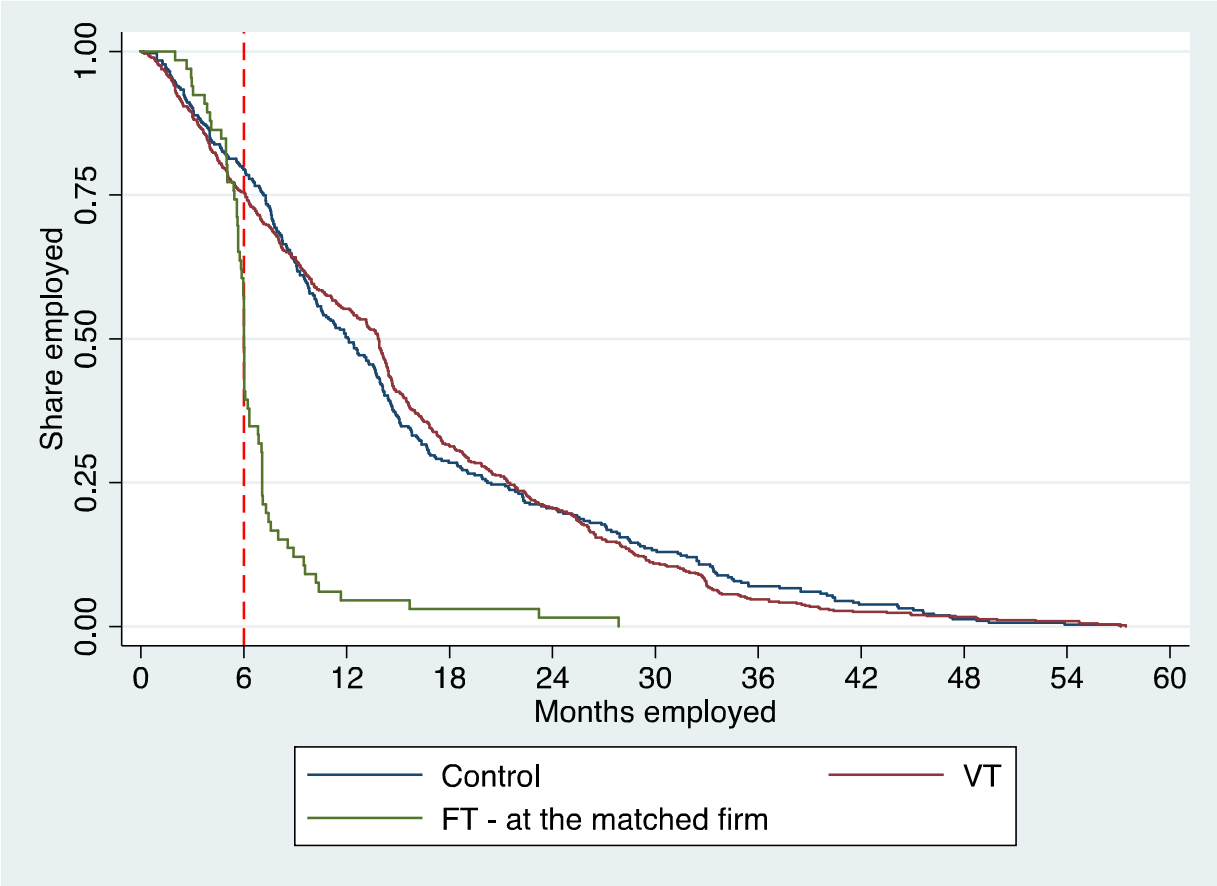

Notes: The Figure plots survival functions for the first employment spell. For Firm Trained workers, we plot the survival function for workers who started training at the matched firm. For Control and Vocationally Trained workers, we plot survival functions in the first non-casual and non-agricultural employment spell in the post-training period (since August 2013).

Figure 4: Dynamics of Employment, Earnings and Wages

Panel A: Number of Months Worked per Quarter

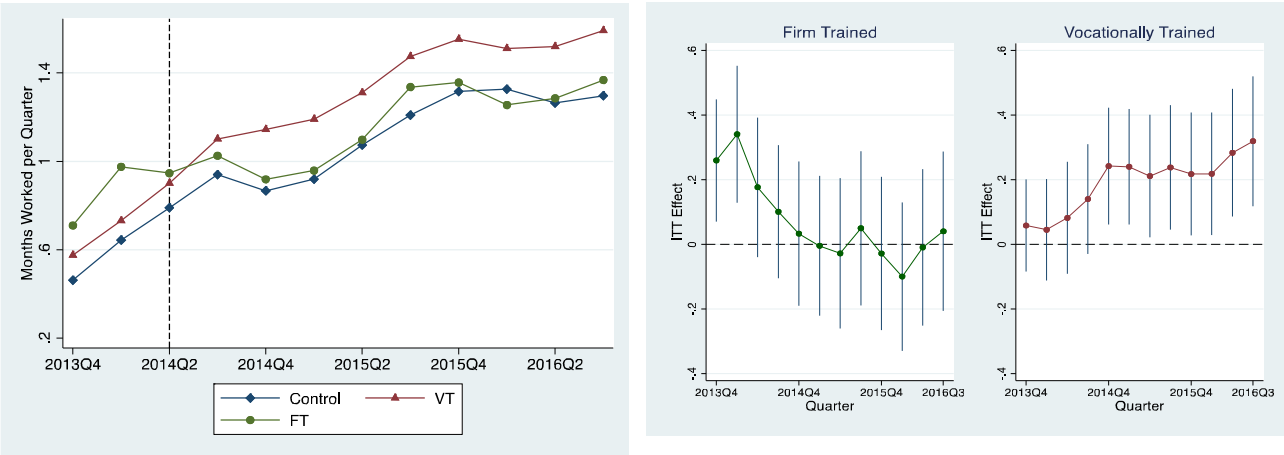

Panel B: Total Quarterly Earnings [USD]

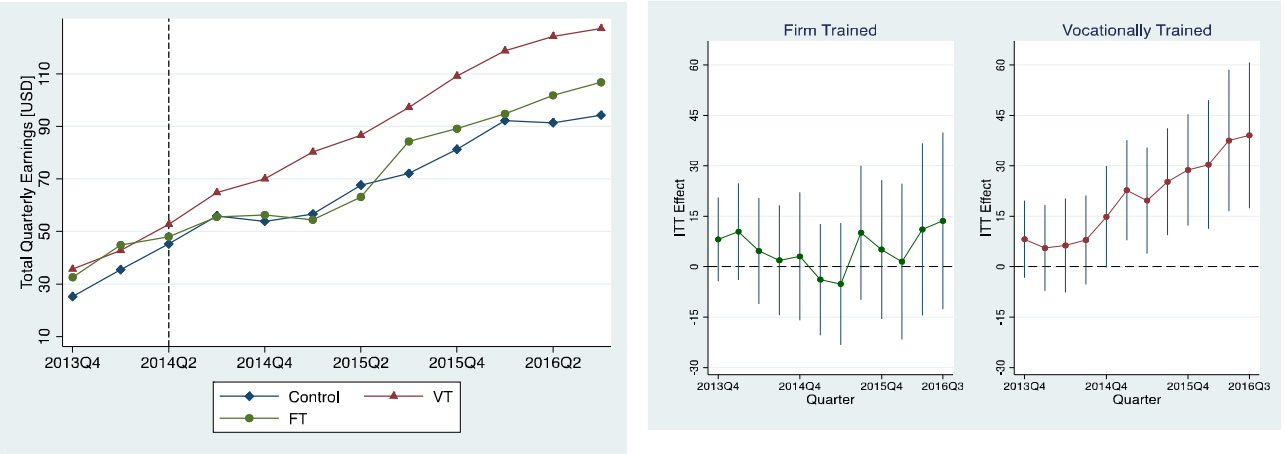

Panel C: Average Hourly Wage per Quarter in Wage Employment [USD]

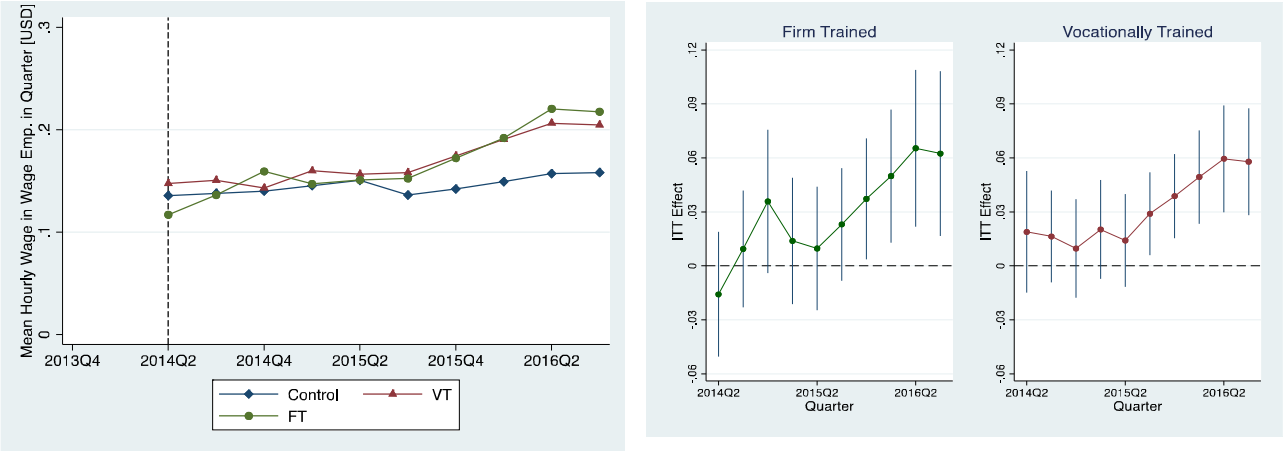

Notes: Data is from the first, second and third follow-up surveys. We use information on all employment and self-employment job spells reported by the workers in the twelve months prior to each survey. So the period considered goes from the fourth quarter of 2013, which is the first quarter covered in the first follow-up survey, until the third quarter of 2016, which is the last quarter covered in the third follow-up survey. Figures on the left of each panel report average number of months, earnings, and hourly wages for each quarter. The vertical dashed line corresponds to the end of the Firm Training intervention. Figures on the right of each panel report quarterly ITT treatment effects of Firm Training and Vocational Training on various outcomes, with 95% confidence intervals. All coefficients reported in each panel are estimated from the same dynamic treatment effects regression, where the FT and VT treatment indicators are interacted with dummies for each quarter considered, with robust standard errors. All regressions further include strata dummies, dummies for quarters, and a dummy for the implementation round. We also control for the following baseline characteristics of workers: age at baseline, a dummy for whether the worker was married at baseline, a dummy for whether the worker had any children at baseline, a dummy for whether the worker was employed at baseline, and a dummy for whether the worker scored at the median or above on the cognitive test administered at baseline. Casual and agricultural occupations are coded as unemployment. Wages and earnings are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD. The top 1% values of earnings and wages are excluded.

Figure 5: Counterfactuals

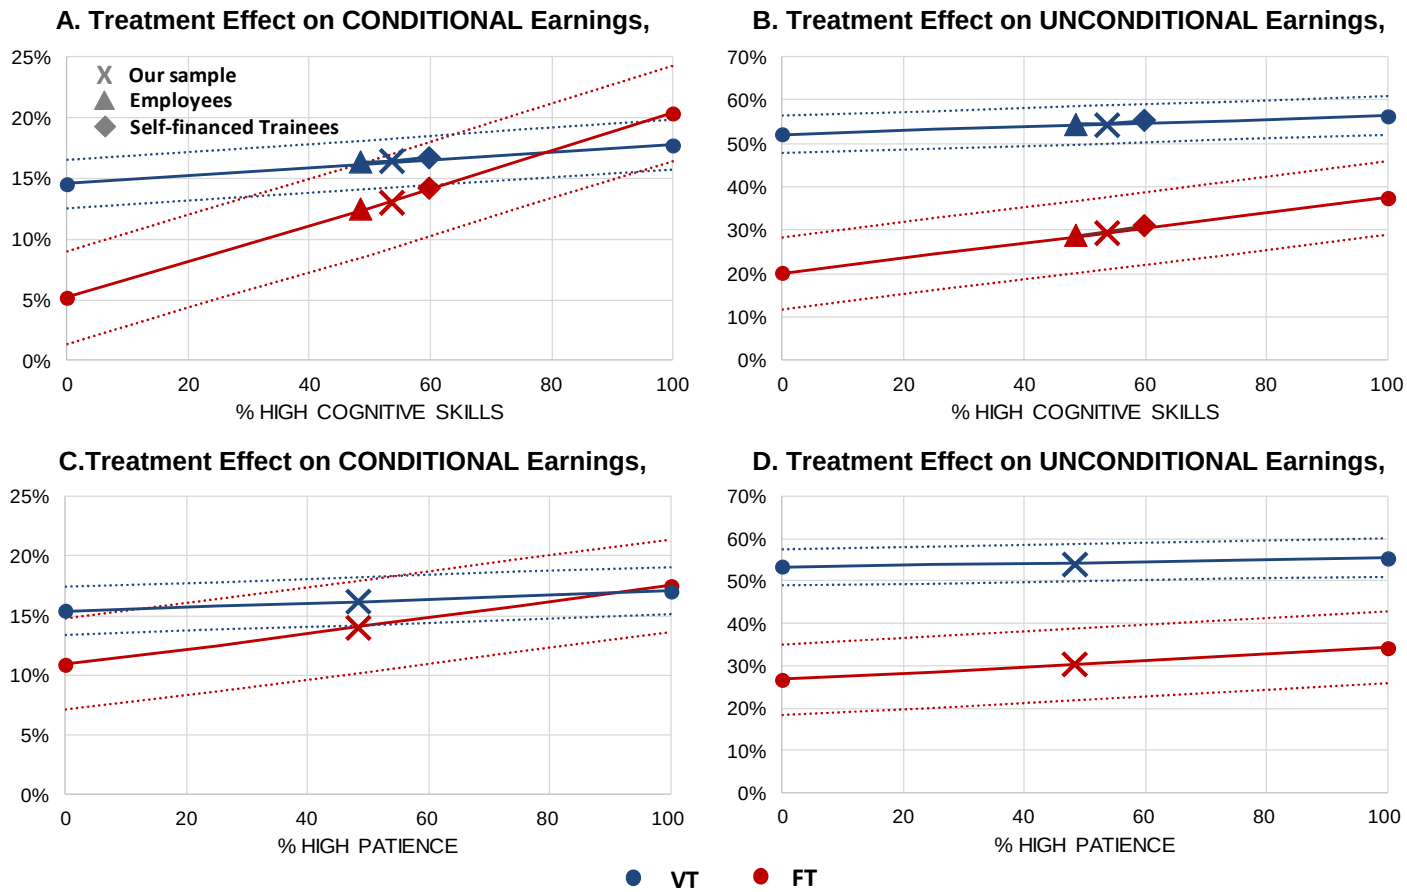

Notes: Figure 5 shows the percentage impact of FT and VT on employment, conditional and unconditional earnings from OLS estimates run on simulated data generated from the model. The dashed lines show 95% confidence intervals. We run 10 simulations of the behavior of 50,000 workers followed over a period of 48 months. In each simulation, we randomly assign individuals to treatment in the same proportions as in our experiment. Workers are also randomly assigned to take-up treatment in the same proportion as in the experiment. In each simulation we calculate treatment effects as the average monthly impact of FT and VT on employment and earnings across the 48 months from OLS regressions. We then aggregate estimates across the different simulations. We only show estimated impacts on compliers. Each Panel shows the treatment effects when we vary the share of individuals with high cognitive skills and high patience in the population (from 0 to 100%). To do that, we first divide workers in our data into high/low Raven matrices using their score on the Raven Matrices test implemented at first follow-up. Workers are assigned to the High Raven group if they scored on or above the median of the Raven Matrices test. Similarly, workers are divided into high/low Patience using their answers to a series of questions about their willingness to wait to receive (hypothetical) monetary rewards at baseline. Workers are assigned to the High Patience group if they had a value of Patience on or above the median. We then obtain kernel density estimates of the distribution of effective units of labor h(ε) for each of these groups. In the simulations individuals are randomly assigned to be High/Low Raven and High/Low Patience in the proportion indicated on the Figures, and draw their effective units of labor from the corresponding distribution. The crosses indicate the exact percentages of High Cognitive Skills and High Patience individuals in our sample. These are respectively 53.5% and 48.3%. The shaded triangles indicate the percentage of High Cognitive Skills individuals in two other Ugandan samples from related studies. The first sample - "self-financed trainees" - are youth analysed in Bassi and Nasamba (2019) and includes trainees that have self-financed their own training. The second sample - "employees" - are employees in a representative sample of firms in welding, furniture making and grain milling from Bassi et al (2019). The share of High Cognitive Skills individuals is 59.8% in the "self-financed trainees" sample, and 48.5% in the "employees" samples.

Figure 6: Comparison of Treatment Impacts to Meta-analysis by McKenzie [2017]

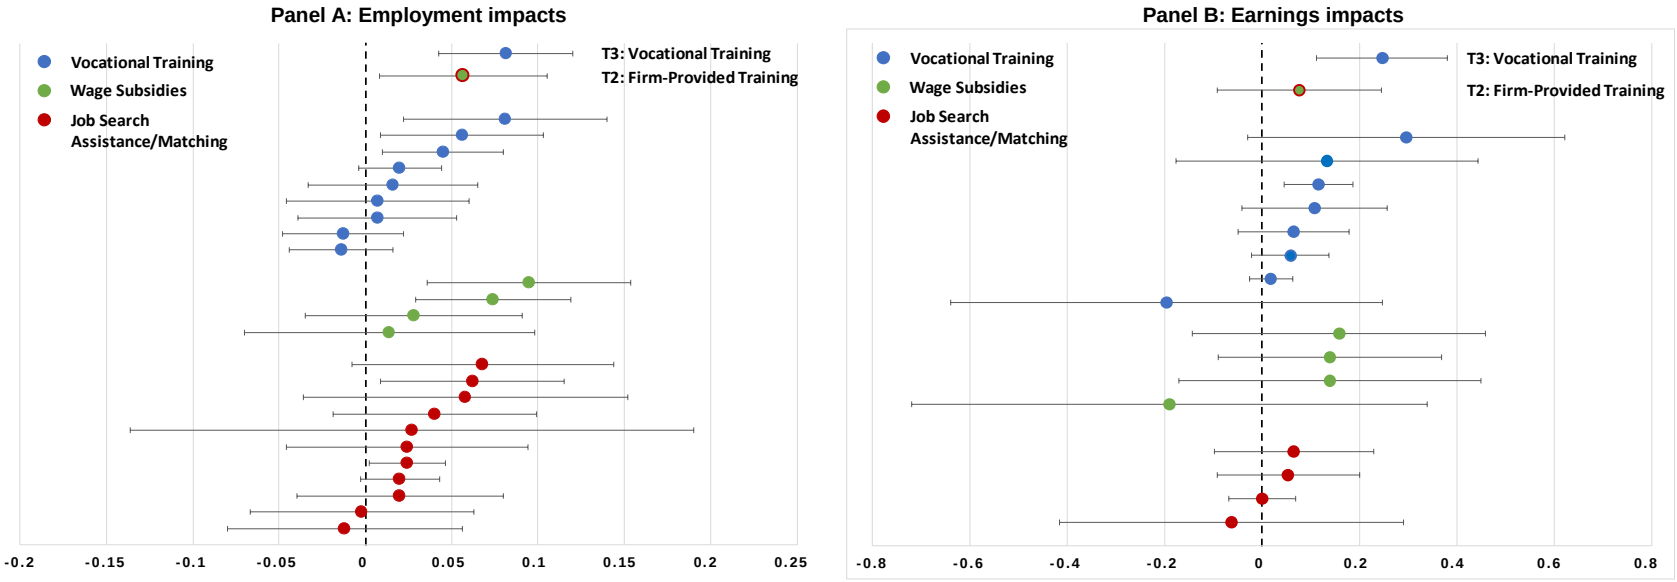

Notes: The Figures compare the treatment impacts from this study to the treatment impacts reported in the meta-analysis by McKenzie [2017]. The green estimates correspond to wage subsidy programs, the blue estimates to vocational training programs, and the red estimates to job search and matching assistance programs. Panel A reports treatment impacts (ITT) on the probability of paid employment, together with 95% confidence intervals. The estimates from our study are taken from Column 2 of Table 4, where we use as outcome variable "Any wage employment in the last month". Alongside our estimates, Panel A further reports 22 estimates of treatment impacts taken from Table 1, 3 and 4 of McKenzie [2017]. These correspond to all the available program estimates for this outcome reported in McKenzie [2017], a part from the estimate from Galasso et al. [2004], which is omitted as no standard error is provided, and the estimate from Groh et al. [2016] with time frame 6 months, as that is estimated while the wage subsidy was still ongoing (while our estimates for T2: FT and all the other estimates for wage subsidy programs reported in the Figure refer to the period after the wage subsidy ended). Panel B reports treatment impacts (ITT) on earnings, in terms of percentage increase relative to the earnings level of the Control group, together with 95% confidence intervals. The estimates from our study are taken from Column 4 of Table 5, where we use as outcome variable "Total earnings in the last month". Alongside our estimates, Panel B further reports 15 estimates of treatment impacts taken from Table 1, 3 and 4 of McKenzie [2017]. These correspond to all the available program estimates for this outcome reported in McKenzie [2017], apart from the estimate from Groh et al. [2016] with time frame six months, as that is estimated while the wage subsidy was still ongoing (while our estimates for T2: FT and all the other estimates for wage subsidy programs reported in the Figure refer to the period after the wage subsidy ended), and the estimate from Maitra and Mani [2017], which is excluded as that is very large relative to all the other estimates: Maitra and Mani [2017] estimate a treatment impact on earnings of .957, with confidence interval [.056 ; 1.86]. However, this corresponds to only a $2.40 monthly increase in earnings in absolute terms, and so the large treatment impact is due to the women in their sample having extremely low earnings to begin with.

Table A1: Baseline Balance on Worker Characteristics

Means, robust standard errors from OLS regressions in parentheses
P-value on t-test of equality of means with control group in brackets
P-value on F-tests in braces

	Number of workers	Age [Years]	Married	Has child(ren)	Currently in school	Ever attended vocational training	Cognitive Test Score	F-test of joint significance
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
All Workers	1714	20.0	.040	.118	.016	.036	.561	
		(.198)	(.016)	(.024)	(.008)	(.018)	(.496)	
T1: Control	451	20.1	.027	.102	.011	.042	.560	
		(.211)	(.016)	(.025)	(.009)	(.020)	(.047)	
T2: Firm Trained	283	20.1	.040	.121	.018	.038	.554	{.999}
		(.139)	(.014)	(.024)	(.009)	(.015)	(.037)	
		[.970]	[.271]	[.260]	[.576]	[.897]	[.640]	
T3: Vocationally Trained	390	20.0	.056*	.127	.018	.032	.529	{.849}
		(.134)	(.014)	(.022)	(.008)	(.013)	(.033)	
		[.781]	[.056]	[.339]	[.553]	[.461]	[.573]	
T4: Vocationally Trained + Matched	307	20.0	.030	.123*	.029	.038	.603	{.878}
		(.146)	(.012)	(.023)	(.011)	(.015)	(.037)	
		[.975]	[.128]	[.075]	[.248]	[.792]	[.772]	
T5: Untrained, Matched	283	20.0	.047*	.122	.007	.027	.568	{.937}
		(.148)	(.015)	(.024)	(.007)	(.014)	(.037)	
		[.429]	[.084]	[.201]	[.468]	[.359]	[1.00]	
F-test of joint significance		{.933}	{.243}	{.449}	{.445}	{.752}	{.974}	

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. All data is from the baseline survey to workers. Column 1 reports the number of workers assigned to each treatment. Columns 2 to 7 report the mean value of each worker characteristic, derived from an OLS regression of the characteristic of interest on a series of dummy variables for each treatment group. All regressions include strata dummies and a dummy for the implementation round. The excluded (comparison) group in these regressions is the Control group. Robust standard errors are reported throughout. The variable in Column 7 is a dummy equal to 1 if the applicant scored at the median or above on a cognitive test administered with the baseline survey. The test consisted in six literacy and six numeracy questions. Column 8 reports the p-values from F-Tests of joint significance of all the regressors from an OLS regression where the dependent variable is a dummy variable taking value 0 if the worker is assigned to the Control group, and it takes value 1 for workers assigned to treatment group j (with j going from 2 to 5) and the independent variables are the variables in Columns 2 to 7. Robust standard errors are also calculated in these regressions. The p-values reported in the last row are from the F-test of joint significance of the treatment dummies in each Column regression where the sample includes all workers.

Table A2: External Validity

Means, standard deviations in parentheses

	Number of individuals	Age [Years]	Gender [Male=1]	Married	Currently in school	Years of Education	Ever attended vocational training	Has worked in the last week [Yes=1]	Has had any wage employment in the last week	Has done any casual work in the last week	Total earnings from wage employment in the last month [USD]
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**	**(9)**	**(10)**	**(11)**
A. Baseline, aged 18-25	1,608	20.1	.567	.038	.014	9.77	.037	.362	.142	.156	2.60
		(1.86)	(.496)	(.190)	(.116)	(2.06)	(.189)	(.481)	(.350)	(.363)	(9.74)
Uganda National Household Survey 2012/13:											
B. All, aged 18-25	4,696	21.1	.465	.395	.309	7.42	.062	.681	.293	.512	9.13
		(2.32)	(.499)	(.489)	(.462)	(3.65)	(.241)	(.466)	(.455)	(.500)	(28.2)
C. Labor Market Active, aged 18-25	3,456	21.4	.475	.448	.207	6.95	.064	.902	.389	.679	12.2
		(2.33)	(.499)	(.497)	(.405)	(3.50)	(.245)	(.297)	(.489)	(.467)	(32.0)

Notes: We present characteristics of individuals from three samples: (i) those individuals in our baseline sample aged 18-25; (ii) individuals aged 18-25 and interviewed in the Uganda National Household Survey 2012/13 (UNHS) conducted by the Ugandan Bureau of Statistics; (iii) individuals aged 18-25 and interviewed in the UNHS who self-report being active in the labor market (either because they are employed or actively seeking employment). The UNHS was fielded between June 2012 and June 2013. Our baseline survey was fielded between June and September 2012. In the UNHS respondents are considered to have attended vocational training if the highest grade completed is post-primary specialized training/diploma/certificate or post-secondary specialized training/diploma/certificate. In the baseline survey questions on employment status did not refer to work activities performed in the last week, but to work activities performed at the time of the survey. Therefore, for the baseline survey the variable "Has worked in the last week" corresponds to the worker being "Currently employed or involved in a work activity". Similarly, Columns 8-10 for the baseline survey are based on the most recent activity performed by the individual, conditional on him/her saying to be currently employed or involved in a work activity. For UNHS, the outcomes in Columns 8-10 are based on the main activity performed in the week before the survey. In Column 9 casual work includes occupations that are casual in nature, as well as agricultural occupations. In Column 10 workers who report doing no wage employment in the past month (or only did unpaid work in the last month) have a value of zero for total earnings.

Table A3: Attrition

OLS regression coefficients, robust standard errors in parentheses

	Worker attrited by endline	
	With covariates	Heterogeneous
	(1)	(2)
T2: Firm Trained	-.000	.002
	(.026)	(.035)
T3: Vocationally Trained	-.018	.022
	(.024)	(.034)
T4: Vocationally Trained + Matched	-.011	-.012
	(.027)	(.036)
T5: Untrained, Matched	.013	.014
	(.027)	(.035)
High Score on Cognitive Test at Baseline [Yes=1]	.045**	.061*
	(.018)	(.032)
T2: Firm Trained X High Cognitive Score		-.005
		(.051)
T3: Vocationally Trained X High Cognitive Score		-.071
		(.047)
T4: Vocationally Trained + Matched X High Cognitive Score		.001
		(.051)
T5: Untrained, Matched X High Cognitive Score		-.002
		(.053)
Mean of outcome in T1 Control group	.134	.134
Strata and Implementation round dummies	Yes	Yes
Other baseline characteristics	Yes	Yes
Test of joint significance of baseline characteristics		
F-statistic	2.35	1.57
P-value	.071	.196
Test of joint significance of Treatment X High Score interactions		
F-statistic		.79
P-value		.529
Number of observations (workers)	1,561	1,561

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. Data is from baseline, first, second and third follow-up of applicants to the vocational scholarships. Standard errors are adjusted for heteroskedasticity in all regressions. Other baseline characteristics include: age at baseline, a dummy for whether the worker was married at baseline, a dummy for whether the worker had any children at baseline, and a dummy for whether the worker was employed at baseline. The variable High Score on Cognitive Test at Baseline is a dummy=1 if the applicant scored at the median or above on the cognitive test administered with the baseline survey.

Table A4: Take-Up of Treatments

	Vocational Training		Matching and Firm Training			
Sample of Workers:	**All Workers**	**Offered Training**	**All Workers**	**Invited to interview**	**Met at least one Firm**	**Worker received a Job Offer**
Outcome:	**% Workers Offered Training**	**% Workers Trained**	**% Workers Invited to Interview**	**% Workers That Met at Least One Firm**	**% Workers Who Received a Job Offer**	**% Workers Hired**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
T3: Vocationally Trained	**97.9**	73.8	-	-	-	-
T4: Vocationally Trained + Matched	**95.4**	63.1	**12.7**	74.4	58.6	23.5
T2: Firm Trained	-	-	**50.5**	80.4	90.4	66.4
T5: Untrained, Matched	-	-	**19.1**	85.2	34.8	18.8

Notes: The data used is from the tracker survey and process reports. The tracker survey was collected in July-August 2013, at the end of the main round of vocational training. Process reports were collected during the implementation of the firm-level interventions (September 2013-February 2014). In Columns 1 and 3 the sample includes all workers assigned to the respective treatment groups. In Column 1 only workers that were traced and successfully informed about the treatment offer are considered as having been offered treatment. In Columns 2 the sample includes those workers who could be traced and were offered the treatment by BRAC staff, and the percentage of workers who took up training includes the workers who completed the 6 months vocational training. For Matching and Firm Training (Column 3) the treatment offer is defined as firms having invited the worker for an interview (so those workers matched to firms that were not interested in the program are not included, as they were not offered treatment). In Column 4 the sample includes workers who were invited for an interview, in Column 5 it includes those workers who met with at least one firm, in Column 6 the sample includes workers who received an offer to start at the firm. In Column 6 the percentage of workers who took up treatment is calculated as the percentage of workers who accepted the offer received by the firm, and so started work/training at the firm.

Table A5: Compliance with the Firm Training Treatment

OLS regression coefficients, robust standard errors in parentheses in all Columns except column 4 where standard errors are clustered at the firm level

Dependent variable: worker started training at the firm assigned to in the FT treatment

	Worker Characteristics	Worker and Program Characteristics	Worker, Program and Firm Characteristics	Firm Fixed Effects
	(1)	(2)	(3)	(4)
Female	.011	-.132	-.108	.019
	(.065)	(.092)	(.094)	(.138)
Age	.013	.006	.002	.004
	(.017)	(.017)	(.017)	(.023)
Any child	.017	.039	.073	.058
	(.092)	(.089)	(.084)	(.120)
High education	-.070	-.043	-.030	.013
	(.058)	(.058)	(.059)	(.086)
High cognitive test score	-.081	-.067	-.064	.047
	(.057)	(.056)	(.054)	(.089)
Employed	-.063	-.068	-.035	-.079
	(.060)	(.065)	(.066)	(.158)
Ideal job is wage employment	-.103*	-.070	-.079	-.032
	(.060)	(.061)	(.060)	(.100)
High risk attitude	-.054	-.066	-.080*	-.040
	(.053)	(.050)	(.048)	(.070)
High patience	.086	.107**	.100*	.099
	(.055)	(.054)	(.052)	(.089)
Employed in August 2013	.075	.071	.060	.066
	(.071)	(.069)	(.066)	(.117)
Second round		.278***	.251***	.147
		(.085)	(.086)	(.132)
Matched to more than one firm		-.040	.002	-.288
		(.075)	(.077)	(.187)
Average firm size of matched firms			.000	
			(.020)	
Average log profit per worker of matched firms			-.119**	
			(.052)	
Average log capital per worker of matched firms			-.023	
			(.057)	
Mean of dep. var. in control	.244	.244	.244	.244
P-value: worker covariates	.065	.156	.194	.976
P-value: firm covariates			.002	
Region of application dummies	Yes	Yes	Yes	Yes
Sector of match dummies	No	Yes	Yes	No
BRAC branch of match dummies	No	Yes	Yes	No
Firm fixed effects	No	No	No	Yes
Adjusted R-squared	.083	.177	.213	.143
Observations	259	259	259	417

Notes: OLS regression coefficients, robust standard errors in parentheses in all columns except in column 4 where standard errors are clustered at the firm level. Data is from the first follow-up worker survey and from the matching surveys, which are used to construct compliance measures. Compliance is defined as having started training at the firm. The sample includes workers assigned to Firm Training. The regression in Column 4 is run on a dataset at the match level. So the dataset includes all the scheduled assignments between workers and firms in FT. The p-values reported at the bottom of each column are from join F-tests of significance of the worker and firm covariates, as indicated in the table. Risk attitudes and patience are measured with hypothetical survey questions. All variables termed as "High" correspond to dummies equal to one if the worker had a value of the underlying variable on or above the sample median at baseline.

Table A6: Worker Expectations

Means, standard deviations in parenthesis

All amounts in 2012 USD

	Expected probability of finding a job in the next 12 months		Average expected monthly earnings (triangular distribution)	
	With Current Skill Set	**If Received VT**	**With Current Skill Set**	**If Received VT**
	(1)	**(2)**	**(3)**	**(4)**
All Workers (Baseline Interview)	.567	.867	57.8	118
	(.288)	(.144)	(46.9)	(71.5)
N. of observations	1,611	1,589	1,243	1,411

Notes: Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. The data used is from the baseline and first three follow-up worker surveys. Columns 1 to 4 report the mean and standard deviation (in parentheses) of the average expected probability of finding a job and the average monthly earnings (assuming a triangular distribution of expected earnings) with the current skill set (columns 1 and 3), or if the worker were to receive vocational training (columns 2 and 4). This is based on all workers interviewed at baseline (across all treatments). All monetary variables are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD. The top 1% values of each variable are excluded from the analysis.

Table A7: Robustness Checks

Dependent Variable: Labor Market Index

OLS regression coefficients, IPW estimates in Columns 1 to 7, robust standard errors in parentheses

Lee [2009] Bounds in brackets

	(1) All	(2) Women	(3) Men	(4) Services	(5) Manufacturing	(6) Non-Kampala	(7) Batches	(8) No Covariates	(9) No IPW, No Covariates
Firm Trained	.105**	.070	.135**	.027	.167**	.189***	.106**	.106**	.115**
	(.051)	(.077)	(.067)	(.080)	(.067)	(.056)	(.051)	(.051)	(.050)
	[.029 ; .122]	[.067 ; .112]	[.009 ; .129]	[.022 ; .095]	[.038 ; .147]	[.108 ; .180]			
Vocationally Trained	.170***	.134**	.196***	.117*	.214***	.198***	.179***	.170***	.179***
	(.041)	(.061)	(.055)	(.065)	(.053)	(.044)	(.045)	(.041)	(.040)
	[.112 ; .204]	[.138 ; .198]	[.094 ; .208]	[.096 ; .198]	[.109 ; .202]	[.143 ; .221]			
Vocationally Trained x Second Batch of Trainees							-.050		
							(.088)		
Mean Outcome in Control Group	.003	-.115	.092	-.109	.081	-.067	.003	.003	.003
Control for Baseline Value	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
P-values on tests of equality:									
Firm Trained = Vocationally Trained	[.169]	[.362]	[.338]	[.222]	[.460]	[.869]	[.146]	[.175]	[.180]
N. of observations	3,256	1,424	1,832	1,320	1,925	2,578	3,256	3,256	3,256

**Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. The data used is from the baseline and first three follow-up worker surveys. We report OLS regressions, where we use inverse probability weighting (Columns 1 to 8) and robust standard errors are reported in parentheses. We report Lee [2009] bounds in brackets, where we implement a conditional Lee Bounds procedure that is able to condition on strata dummies in Columns 1-3 and 6-9, and to condition on region dummies and a dummy for having a level of education at the median or above at baseline in Columns 4-5. The dependent variable is the Labor Market Index that is computed using the following variables: any paid work in the last month (dummy), months worked in the last year, hours worked in the last week and total earnings in the last month. Total earnings are set to zero for workers with no earnings. The index is constructed following Anderson's [2008] approach. Manufacturing sectors are: motor-mechanics, plumbing, construction, electrical wiring and welding. Service sectors are: hairdressing, catering and tailoring. Workers are assigned to Manufacturing or Service sectors according to stated preferences over their ideal job, reported at baseline. In Column 6 we restrict the sample to labor markets outside of Kampala. All regressions include strata dummies, survey wave dummies, a dummy for the implementation round and dummies for the month of interview. In columns 1 to 7 we also control for the following baseline characteristics of workers: age at baseline, a dummy for whether the worker was married at baseline, a dummy for whether the worker had any children at baseline, a dummy for whether the worker was employed at baseline, and a dummy for whether the worker scored at the median or above on the cognitive test administered at baseline. Columns 1 and 4-9 further control for a complete set of strata dummies. Columns 2 and 3 further control for region dummies, and a dummy for having a level of education at the median or above at baseline. The weights for the IPW estimates are computed separately for attrition at first, second and third follow-up. The instruments for the IPW estimates are whether the worker was an orphan at baseline, a dummy if anyone in the household of the worker reported having a phone at baseline, a dummy for whether the worker reported being willing to work in more than one sector at the time of their original application to the VTIs and dummies for the survey team the worker's interview was assigned to in each of the three follow-up survey rounds. At the foot of each Column we report p-values on the null that the impact of the vocational training is equal to the impact of firm training (T2=T3). All monetary variables are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD.

Table A8: Alpha

OLS Regression Estimates, Robust Standard Errors

	Ln (Earnings in First Month of Employment)				
Sample:	**All Treatments, U2J**	**All Treatments, U2J**	**All Treatments, U2J**	**All Treatments, U2J**	**All Treatments, U2J**
Outcome Variable:	**Actual Earnings**	**Actual Earnings**	**Actual Earnings**	**Imputed Earnings**	**Imputed Earnings**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Ln (Skills Test Score)	.289	.235	.328	.312**	.239*
	(.192)	(.194)	(.407)	(.147)	(.144)
Transitioned from Unemployment					
Vocationally Trained X Ln(Skills Test Score)			-.077		
			(.486)		
Firm Trained X Ln(Skills Test Score)			.006		
			(.563)		
Baseline Controls	No	Yes	No	No	No
N. Observations	109	109	109	169	169

Notes: The data is from the second and third follow-up survey of workers and includes information on all job spells workers have been involved in starting from November 2015. The unit of observation for the analysis is the job spell. The table shows coefficients and standard errors from an OLS regression of the logarithm of earnings in the first month of employment on the logarithm of the score obtained by the worker in a sector-specific skills test. The sample includes workers who transitiond from unemployment into employment. All regressions control for treatment dummies. In Column 2 we also control for age, gender and education at baseline, as well as strata dummies. In Column 3 we add interactions of the logarithm of the skills test score with treatment dummies. In Columns 4 and 5, we impute earnings in the first month of employment for individuals for whom we do not have this information. In Column 4 the imputation is done in two steps: first, we run a probit regression of a dummy for whether an individual was paid in the first month of employment on the logarithm of earnings in the last month of employment, spell duration, a dummy equal to 1 if the worker transitioned from unemployment to employment, and treatment and sector dummies. In the second step, we run the same regression using the logarithm of earnings in the first month of employment as outcome variable. We then predict earnings in the first month of employment from the second regression, and set such earnings equal to zero for individuals with the lowest predicted probability of being paid, so the share of unpaid individuals in the predicted data corresponds to the percentage of unpaid individuals in the actual data. In Column 5 we impute earnings in the first month of employment by regressing the logarithm earnings in the last month of employment on spell duration, a dummy equal to 1 if the worker transitioned from unemployment to employment, and treatment and sector dummies. We then impute the logarithm of earnings in the first month of employment as the difference between the logarithm of earnings in the last month and the OLS coefficient on spell duration multiplied by the spell duration.

Table A9: Effect of Skills on Employment

Dependent variable: =1 if worker is employed in November 2015
Robust standard errors in parentheses
Unit of observation: worker spells

Worker sample:	All Treatments	All Treatments	Control Group	All Treatments
	(1)	(2)	(3)	(4)
Skills Test Score	.001	.001	-.000	-.000
	(.001)	(.001)	(.001)	(.001)
Vocationally Trained X Skills Test Score				.002
				(.001)
Firm Trained X Skills Test Score				.001
				(.002)
Baseline Controls	No	Yes	No	No
N. Observations	1,373	1,358	419	1,373

Notes: The data is from the second and third follow-up survey of workers and includes information on all job spells workers have been involved in starting from November 2015. The unit of observation for the analysis is the job spell. The table shows coefficients and standard errors from an OLS regression of a dummy equal to 1 if the individual was employed in November 2015 on the score obtained by the worker in a sector-specific skills test. The sample in Columns 1 and 2 includes individuals from all treatment groups, while the sample in Column 3 is restricted to workers in the Control group. The regressions in Columns 1, 2 and 4 control for treatment dummies. In Column 2 we also control for age, gender and education at baseline, as well as strata dummies.

Table A10: Estimates in the Job Ladder Search Model, with F(r|T)

Two-step estimation procedure in Bontemps, Robin and van den Berg [2000]

Asymptotic standard errors in parentheses

Steady State: November 2015 (Data from Second and Third Follow Up)

		Non-Compliers		Compliers	
	Control	Firm Trained	Vocationally Trained	Firm Trained	Vocationally Trained
Panel C: Wages and Earnings	(1)	(2)	(3)	(4)	(5)
Average monthly OFFERED wages [USD]	43.9	48.5	41.5	41.9	47.0
Average monthly ACCEPTED wages [USD]	63.4	70.8	67.1	64.4	71.5
Impact on annual earnings [USD]		**33.7**	**14.7**	**49.8**	**154.7**
% Impact:		**10.8%**	**4.7%**	**15.9%**	**49.4%**

Notes: The dataset is a cross-section of workers, and for each worker it contains information on: spell type (employment, unemployment), spell duration (in months), earnings in employment spells (in USD), dates of transitions between spells and type of transition: (i) job to unemployment, (ii) unemployment to job, or (iii) job to job. Wages are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD. The dataset contains at most two spells (and one transition) per individual. The data comes from the second and third follow-up survey of workers, and the initial spell is identified as the (employment or unemployment) spell that was ongoing in November 2015. Spells are right censored at the date of the third follow-up interview (which ended in December 2016). Spells are left censored at 1 August 2014. Casual and agricultural occupations are coded as unemployment. Self-employment is coded as employment (but self-employment spells are assigned a separate spell). The estimation protocol follows the two-step procedure in Bontemps, Robin and van den Berg [2000]: in the first step the G function is estimated non-parametrically from the data (so this is just the empirical CDF of observed wages for those workers that are employed in their first spell), and is then substituted into the likelihood function. In the second step, maximum likelihood is then conducted using information from both the first and second spells for each individual to recover the parameter estimates. In Panel C average monthly offered and accepted wages are computed as the product of average offered and accepted piece-rates, and average units of effective labor. We assume workers draw piece-rates from the same offer distribution F(r). F(r) is the kernel density estimate of a weighted average of the distributions of offered piece-rates across treatments - F(r|T) - where such distributions are obtained from their steady-state relationship with non-parametrically estimated G(r|T). Weights are equal the share of individuals in each treatment.

Table A11: Firm Recruiting Constraints, Hires and Channels

IPW regression coefficients, standard errors clustered by sector-branch in parenthesis, Lee Bounds in brackets

	Ability to find UNSKILLED workers is a constraint	Ability to find SKILLED workers is a constraint	Able to perform sector-specific task
	(1)	(2)	(3)
PANEL A: ITT Estimates			
Firm Trained	.004	-.054**	.086**
	(.028)	(.025)	(.043)
	[-.013 ; .078]	[-.089 ; .002]	[.073 ; .108]
PANEL B: ATE Estimates			
Firm Trained	-.032	-.147*	.189**
	(.084)	(.076)	(.089)
Mean outcome in Control firms	.386	.678	.301
Number of observations	1,579	1,584	651

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. The data used is from the second, third and fourth firm follow-up data surveys. Panel A reports OLS IPW regression estimates, together with standard errors adjusted for heteroskedasticity and clustered at the branch-trade level in parenthesis. We report Lee [2009] bounds in brackets, where we implement a conditional Lee Bounds procedure that is able to condition on dummies for the interview round and baseline trade. Underlined bounds are significantly different from zero at the 95% condfidence level. Panel B reports 2SLS IPW regression estimates, where treatment assignment is used as IV for treatment take-up. Treatment take-up is defined as a dummy equal to one if the firm hired one of the workers it was matched with. All regressions control for the value of outcome at baseline (when available), and include branch and trade fixed-effects, survey wave dummies and dummies for the month of interview. Baseline controls also include the owner's sex, business age (measured as number of years since the business was established) and business age squared, firm size and owner's years of education. The weights for the Inverse Probability Weights (IPW) are computed separately for attrition at second, third and fourth follow-up. The instruments for the IPW estimates are dummies for whether the respondent provided a phone number at baseline, and for whether he/she was an employee of the firm (rather than the firm owner or the manager), the number of network firms and dummies for interviewers at baseline. In Columns 1 and 2 the outcomes are dummies equal to one if the would hire more workers if he/she received more applications from unskilled and skilled workers respectively. In Columns 3 we identified a specific task for each of the study sectors and asked the owner whether the worker was able to perform that task when he joined the firm. ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level.

Table A12: Heterogeneous Impacts on Skills

2SLS regression coefficients, bootstrapped standard errors in parentheses

Dependent Variable: Sector-Specific Test Score (0-100)

Heterogeneous effects by:	Raven Matrices	Patience
	(1)	(2)
Firm Trained X Below Median Trait	2.63	3.50
	(8.81)	(7.54)
Firm Trained X Above Median Trait	21.1***	12.9*
	(5.72)	(6.88)
Vocational Training X Below Median Trait	7.79***	7.96***
	(2.67)	(2.40)
Vocational Training X Above Median Trait	13.2***	12.2***
	(2.19)	(2.30)
Mean outcome in Control group	30.1	30.1
p-value FT X Low = FT X High	.076	.358
p-value VT X Low = VT X High	.102	.204
Observations	1,485	1,799

Notes: ***denotes significance at the 1% level, ** at the 5% level, * at the 10% level. The data used is from the baseline, second and third follow-up worker surveys in all columns. We report 2SLS regression estimates, where treatment assignment is used as IV for treatment take-up. Treatment take-up is defined as a dummy equal to one if the worker (i) started firm training in FT or (ii) started vocational training in VT. Bootstrap standard errors are calculated using 1,000 replications. All regressions control for strata dummies, survey wave dummies, a dummy for the implementation round and dummies for the month of interview. We also control for the following baseline characteristics of workers: age at baseline, a dummy for whether the worker was married at baseline, a dummy for whether the worker had any children at baseline, a dummy for whether the worker was employed at baseline, and a dummy for whether the worker scored at the median or above on the cognitive test administered at baseline. At the foot of each Column we report p-values on the null that the impact of the vocational training is equal to the impact of firm training, by the various variables considered in each of the columns. Workers are divided into high/low Raven matrices using their score on the Raven Matrices test implemented at first follow-up. Workers are assigned to the High Raven group if they scored on or above the median of the Raven Matrices test. Workers are divided into high/low Patience using their answers to a series of questions about their willingness to wait to receive (hypothetical) monetary rewards at baseline. Workers are assigned to the High Patience group if they had a value of Patience on or above the median.

Table A13: Parameter Estimates of the Job Ladder Search Model

Two-step estimation procedure in Bontemps, Robin and van den Berg [2000]
Asymptotic standard errors in parentheses
Steady State: November 2015 (Data from Second and Third Follow Up)

Panel A: Parameter Estimates (Monthly)	Control	Firm Trained	Vocationally Trained
	(1)	(2)	(3)
Average units of effective labor [USD]	2.51	2.59	2.74
Job destruction rate, δ	.027	.026	.024
	(.003)	(.005)	(.004)
Arrival rate of job offers if UNEMPLOYED, λ_0	.019	.019	.024
	(.002)	(.003)	(.003)
Arrival rate of job offers if EMPLOYED, λ_1	.040	.043	.037
	(.010)	(.016)	(.013)

Panel B: Competition for Workers and Unemployment				
Interfirm competition for workers		1.46	1.64	1.55
	% Impact:		**12.0%**	**5.88%**
Unemployment rate		.587	.578	.500
	% Impact:		**-2.08%**	**-14.7%**
Unemployment duration (months)		52.9	52.2	42.2
	% Impact:		**-1.45%**	**-20.2%**
Employment duration (months)		37.0	38.3	41.9
	% Impact:		**3.65%**	**13.1%**

Panel C: Wages and Earnings				
Average monthly OFFERED wages [USD]		42.8	44.2	46.8
Average monthly ACCEPTED wages [USD]		62.9	64.4	72.2
Impact on annual earnings [USD]			**15.5**	**124**
	% Impact:		**5.02%**	**40.1%**

Notes: The dataset is a cross-section of workers, and for each worker it contains information on: spell type (employment, unemployment), spell duration (in months), earnings in employment spells (in USD), dates of transitions between spells and type of transition: (i) job to unemployment, (ii) unemployment to job, or (iii) job to job. Wages are deflated and expressed in terms of August 2012 prices, using the monthly consumer price index published by the Uganda Bureau of Statistics. Deflated monetary amounts are then converted into August 2012 USD. The dataset contains at most two spells (and one transition) per individual. The data comes from the second and third follow-up survey of workers, and the initial spell is identified as the (employment or unemployment) spell that was ongoing in November 2015. Spells are right censored at the date of the third follow-up interview (which ended in December 2016). Spells are left censored at 1 August 2014. Casual and agricultural occupations are coded as unemployment. Self-employment is coded as employment (but self-employment spells are assigned a separate spell). The estimation protocol follows the two-step procedure in Bontemps, Robin and van den Berg [2000]: in the first step the G function is estimated non-parametrically from the data (so this is just the empirical CDF of observed wages for those workers that are employed in their first spell), and is then substituted into the likelihood function. In the second step, maximum likelihood is then conducted using information from both the first and second spells for each individual to recover the parameter estimates. As shown in Panel A, we estimate separate parameters for Control and Treatment groups, but we pool together compliers and non-compliers. Outputs in Panel B are derived from the model and computed as functions of the estimated parameters: (i) interfirm competition for workers=λ_1/δ; (ii) unemployment rate=$\delta/(\delta+\lambda_0)$; (iii) unemployment duration=$1/\lambda_0$; employment duration=$1/\delta$. In Panel C average monthly offered and accepted wages are computed as the product of average offered and accepted piece-rates, and average units of effective labor. We assume workers draw piece-rates from the same offer distribution F(r). F(r) is the kernel density estimate of a weighted average of the distributions of offered piece-rates across treatments - F(r|T) - where such distributions are obtained from their steady-state relationship with non-parametrically estimated G(r|T). Weights are equal the share of individuals in each treatment. For each treatment we then re-invert F(r) using estimated parameters and steady-state relationships to obtain G(r|T) under the assumption that workers draw piece-rates from the same offer distribution.

Figure A1: Firm Size and Employment Distributions, by Country

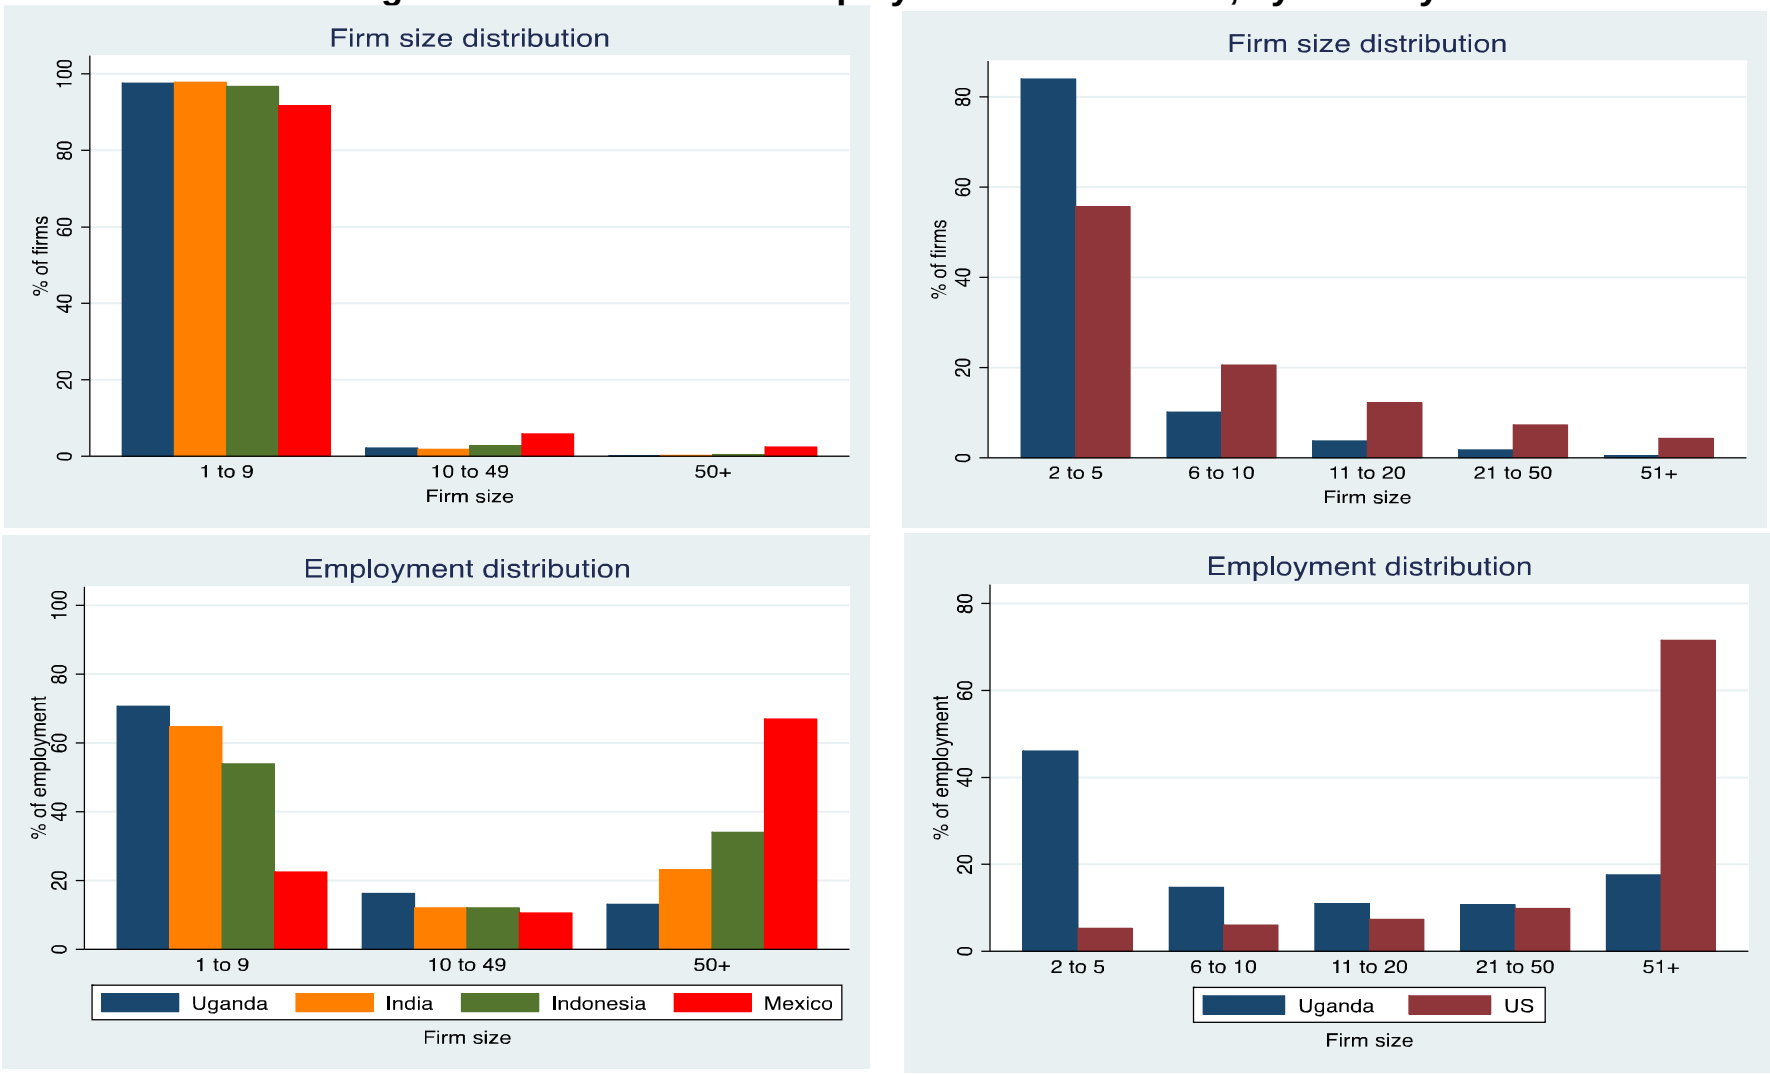

Notes: The data used to produce these figures is from the following sources: Uganda – 2010 Census of Business Establishments collected by the Uganda Bureau of Statistics; US - 2010 Business Dynamic Statistics collected by the US Census Bureau; India, Indonesia and Mexico - data are from Hsieh and Olken (2014). Data are from 2011 for India, 2006 for Indonesia and 2008 for Mexico. The firm size distribution reflects the % of firms that employ 0 to 9, 10 to 49, and 50+ employees. The employment distribution reflects the percentage of workers employed in firms that employ 0 to 9, 10 to 49, and 50+ employees.

Figure A2a: Wage Distribution of Unskilled Workers at Baseline

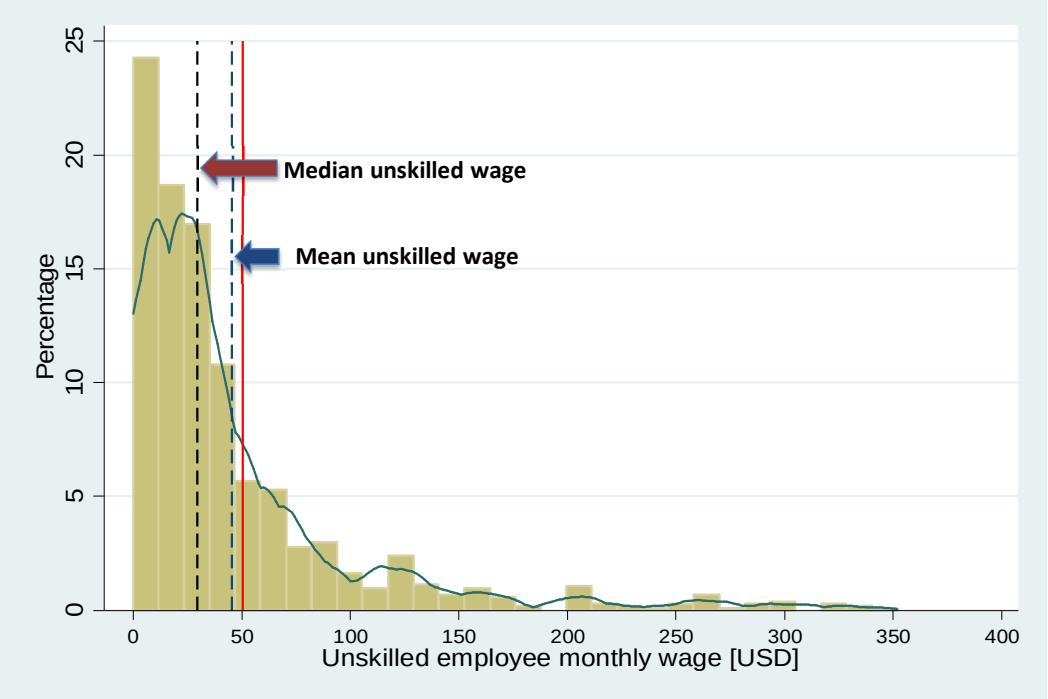

Figure A2b: Worker-Firm Wage Subsidy Splits

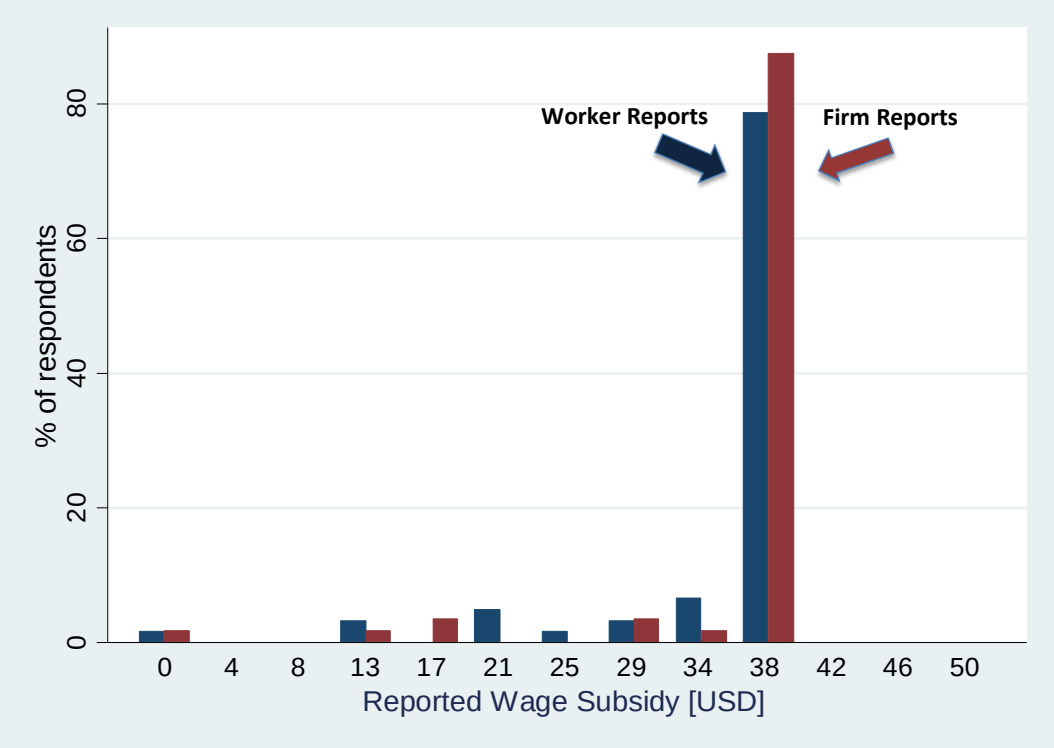

Notes: The top graph shows the distribution of unskilled workers' wages at baseline. The solid line is drawn in correspondence to the total amount of wage subsidy under the Firm Training treatment, and the dashed line indicates the median (unskilled) wage at baseline. A Kernel density estimate of the distribution of wages is also shown. The lower histogram shows the reported monthly earnings of workers hired through the Firm Training treatment, where the first bar is always the worker's self-reported wage, and the second bar is what the firm reports paying the worker.

Figure A3: Firm-provided Training Contract

Append Firm Owner's Photo Here

brac

Branch:

CONTRACT

Small Firm Expansion and Job Creation Program: Mentorship

I _____**(Firm Owner)** owner of_____**(Business Name)**

In _____**(Village/Area)** in _____**(Sub-County)**

Of _____**(County)** and _____**(District)**

hereby promise that I will conduct training for the undersigned trainee regularly as per the following terms:

1. I am bound to provide a training course on _____ (trade type) for the duration of six month, starting ___/ ___ / 2013
2. I will refund the full allocation given by BRAC for training cost should I fail to provide the training.
3. BRAC Uganda will provide 30,000 UGSH monthly to supplement training costs. I expect no other payments from BRAC, Uganda.
4. I will attend monthly meetings and training sessions as per the schedule provided the Job Placement Officer (JPO) or ELA Staff.
5. I will keep track of the trainee's attendance in the attendance register provided by the JPO/ELA.
6. I will abide by the decisions of BRAC Uganda regarding changes to the schedule, content of training, or assignment of trainee.
7. I agree provide honest information to BRAC Uganda. I understand that falsifying documents, deceiving BRAC staff or bearing false witness will result in the immediate termination of the contract.
8. I commit to try my best to ensure to good quality skill development for the trainee through this training.
9. I confirm that I am physically able to conduct and compete the training.

I here by sign the promissory note with full conciseness after reading, fully understanding and accepting the conditions, without any influence from any one.

_____
Firm Owner's Full Name

_____
Trainee's Full Name

_____
Firm Owner's Signature

_____
Trainee's Signature

_____
Date

_____
Date

_____
JPO / ELA Staff Signature

_____
Date

Worker ID

Firm ID

Figure A4: Change in Worker Productivity Between Recruitment and First Follow-up

Means and 95% Confidence Intervals

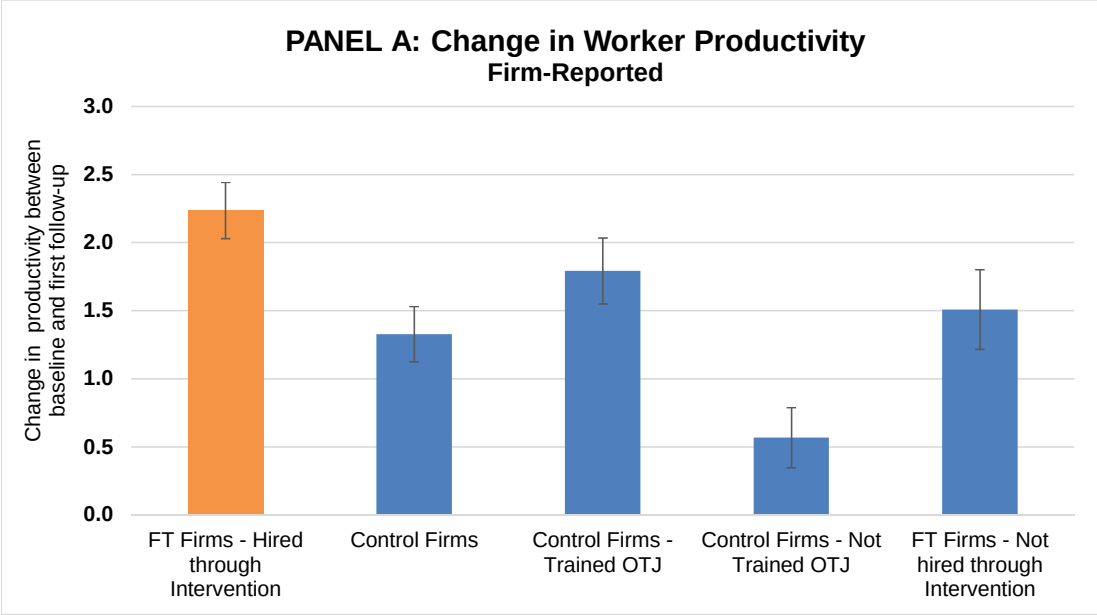

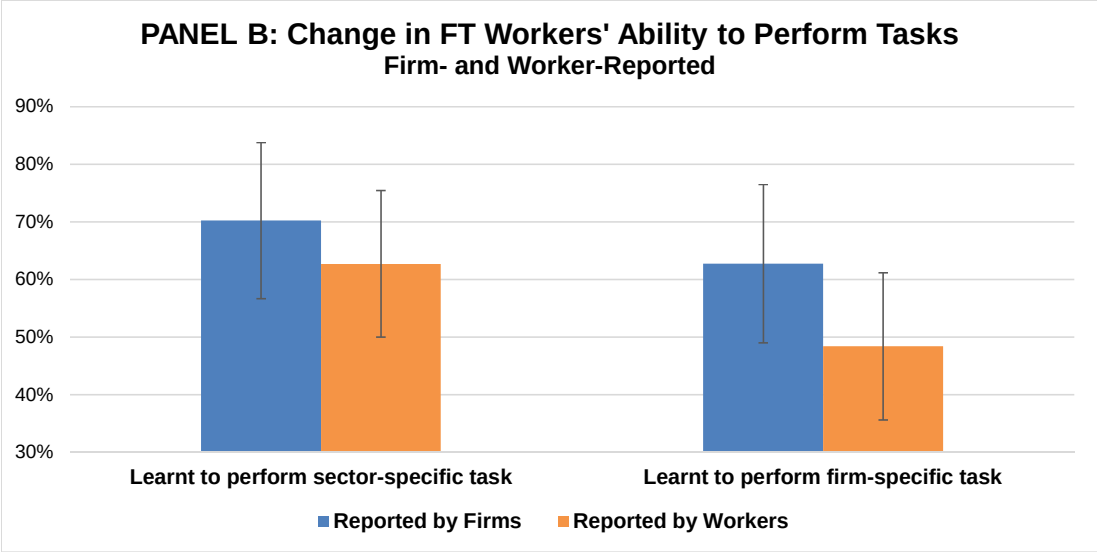

Notes: The data used is from the first follow-up survey of firms and workers. In the firm data, the unit of observation is the employee, and the sample only includes workers hired between 4 and 6 months prior the survey. In Panel A, the sample Control Firms - trained OTJ includes workers in Control firms who received on-the-job training; the sample FT Firms - Not hired through Intervention includes all workers hired in firms assigned to the Firm Training treatment, but not directly through the intervention; the sample FT Firms - Hired through Intervention includes workers hired through the Firm Training intervention only. For each worker, the respondent (i.e. the firm owner in most cases) was asked to rate the employee's productivity at recruitment and at the time of the interview (or at the time when the worker left the firm) on a scale from 1 to 5. The average productivity growth of workers across the different samples is shown in Panel A. In Panel B, we identified a specific task for each of the study sectors and asked the respondent (i.e. the firm owner or the worker) whether the worker was able to perform that task when he joined the firm and at follow-up (or at the time when the worker left the firm). For firm-specific tasks respondents were asked to identify a task considered particularly important for the firm, and were then asked whether the worker was able to perform that task when he joined the firm and at follow-up (or at the time when the worker left the firm). Panel B shows the percentage of workers who learnt how to perform the task between baseline and follow-up (or between baseline and the time when the worker left the firm) for workers in the firm training intervention who took-up the treatment. The blue bars report the learning rate as reported by firms; the orange bars report the learning rate as reported by workers.

Figure A5: Sector Skills Test for Motor Mechanics

1. MOTOR-MECHANICS		
1	*multiple-choice* What are you advised to do when servicing the engine by changing oil?	A. Top up lubricating oil B. Replace oil filter C. Over hand engine D. Over hand cylinder head **Correct Answer: B**
2	*multiple-choice* What immediate remedy can you give to a vehicle with a problem of excessive tyre wear in the center more than other parts?	A. Increase tyre pressure B. Reduce tyre pressure C. Inflate pressure D. Remove the vehicle tire **Correct Answer: B**
3	*multiple-choice* If a customer reports to you that his/her vehicle charging system works at lower rate, how can you help him?	A. Replacing the charging system B. Adjusting the alternator tension C. Replacing alternator housing D. Renewing wire insulator **Correct Answer: B**
4	*multiple-choice* Which of the following set of systems or component call for mechanical adjustment during general vehicle service?	A. Tyres, cooling system, master cylinder B. Break shoes, alternator, and valve clearance C. Distributor, radiator, propeller shaft D. Tank, crank shaft, Turbo charger **Correct Answer: B**
5	*multiple-choice* What solution would you give a customer with a vehicle engine producing blue smoke?	A. Top up lubricant B. Time the engine C. Replace piston rings D. Remove carbon deposits **Correct Answer: C**
6	*matching* What should you do to stop the following vehicle troubles?	1 Battery over charging — A Leaking fuel tank 2 Engine over heating — B Renew regulator 3 Lubricant leakage — C Reduce oil to the correct level 4 Smoke in exhaust — D Renew piston rings 5 Engine fails to start — E Charge the battery **Correct Answer: 1B, 2A, 3C, 4D, 5E**
7	*order* When changing engine oil, in which order should you perform the following steps?	A. Drain oil through drain plug B. Remove oil filter cup C. Run engine to check leaks D. Fill new oil through filler cup to level E. Remove oil filter F. Warm up the engine **Correct Answer: B, E, A, D, F, C**

Figure A6: Worker Spells Data