ARE SMALL FIRMS LABOR CONSTRAINED? EXPERIMENTAL EVIDENCE FROM GHANA

Morgan Hardy Jamie McCasland

March 2022

Abstract

We report the results of a field experiment that randomly placed unemployed young people as apprentices with small firms in Ghana, and included no cash subsidy to firms (or workers) beyond in-kind recruitment services. Treated firms experienced increases in firm size of approximately half a worker and firm profits of approximately 10% for each apprentice placement offered, documenting frictions to novice hiring. We interpret the program as providing a novel worker screening technology to firms, as (voluntary) worker participation included non-monetary application costs, echoing the widespread use of an entrance fee mechanism for hiring apprentices in the existing labor market.

JEL Classifications: D22, D61, J23, J38, J46, M51, M53, O12, O14, O15

We are grateful to Michael Anderson, Lauren Bergquist, David Card, Fenella Carpena, Kenneth Chay, Garret Christensen, Ben Faber, Fred Finan, Andrew Foster, Willa Friedman, Francois Gerard, Paul Gertler, David Glancy, Tadeja Gracner, Bryan Graham, Seema Jayachandran, Jonas Hjort, Hedvig Horvath, Kelsey Jack, Pat Kline, Attila Linder, Jeremy Magruder, Isaac Mbiti, David McKenzie, Costas Meghir, Edward Miguel, Owen Ozier, Ana Rocca, Adina Rom, Michael Walker, Christopher Woodruff, three anonymous referees, and seminar participants at Berkeley, Brown, UBC, Kellogg, UC Davis, Mathematica, Simon Fraser, Columbia, Notre Dame, Michigan, Hitotsubashi, University of Tokyo, NBER Summer Institute, and the CSWEP CeMENT workshop for helpful comments and suggestions. We also thank Lois Aryee, Robert Obenya, Charles Sefenu, Yani Tyskerud, Thomas Zubevial, and Innovations for Poverty Action for excellent research assistance in the field. Ragini Jain, Sean Lothrop and Everett Stamm also provided excellent research assistance. This research was supported by funding from the National Science Foundation (NSF) Graduate Research Fellowship, the Ewing Marion Kauffman Foundation, the Institute for Business and Economics Research (IBER) at UC Berkeley, the Private Enterprise Development in Low-Income Countries (PEDL) program, the International Growth Centre (IGC), 3ie, USAID Development Innovation Ventures (DIV), and the Centre for Innovation Data in Economics Research (CIDER) at UBC. The fieldwork for this research was conducted under IRB approval from Innovations for Poverty Action (563.12August-004), Southern Methodist University (2012-060-MBII), the University of Virginia (2015-0007-00), and the University of British Columbia (H15- 02708). The University of California, Berkeley had an Inter-Institutional Agreement with Southern Methodist University during a portion of the data collection. This research was registered with the AEA RCT registry (AEARCTR-0000297). Data and replication files can be found at ICPSR (openicpsr-155121). All errors are our own. Jamie McCasland: Vancouver School of Economics, University of British Columbia (jamie.mccasland@ubc.ca). Morgan Hardy: Social Science Division, New York University Abu Dhabi (morgan.hardy@nyu.edu).

Two of the most ubiquitous features of economic activity in poor countries are an extremely right-skewed firm size distribution and high rates of youth unemployment.[1] These small firms are less productive than their larger counterparts and rarely grow; many of their would-be workers join growing cohorts of unemployed and underemployed young people with limited access to formal sector wage employment. What prevents small firms from hiring more workers? Identifying these hurdles and which types of labor market institutions could overcome them is critical to understanding the organization of production in low-income countries and to the well-being of the next generation of entrepreneurs and jobseekers.

We take this question to a particularly intriguing setting, apprenticeships in West Africa, where inexperienced workers pay an entrance fee to firms to gain employment. In this setting, we study a national-scale government-initiated and -implemented worker placement program that was designed to overcome the liquidity problem imposed by the entrance fee norm. Our object of interest in the paper is the firm, and we use the embedded placement experiment to determine whether labor market frictions hinder firm hiring. We then explore how the existing apprenticeship labor market institution can inform our interpretation of these frictions.

The experimental program recruited unemployed young people interested in apprenticeships and placed them with small firms in Ghana, with a particular emphasis on identifying would-be apprentices too poor to pay the entrance fee. Program roll-out ultimately eliminated the entrance fee entirely, and included no cash subsidy to firms (or workers) beyond in-kind recruitment services. In order to qualify for the program, would-be apprentices were required to submit an application, attend a series of meetings and interviews, and continue to show interest in the apprenticeship despite an unexpected delay in program roll-out. Firm recruitment followed the recruitment of would-be apprentices, and was centered around occupational trades preferred by program applicants

[1] The *World Bank Enterprise Surveys*, firm-level data from 135 countries which include primarily formal firms and only those with five or more employees, nonetheless show a strikingly higher density of small firms in poorer countries and poorer regions. In Ghana, the National Industrial Census (NIC) attempts to capture at least some proportion of informal manufacturing firms and shows 94% of manufacturing firms have fewer than twenty workers and these account for 48% of manufacturing employment (in 2000). Both the Enterprise Surveys and the NIC have been used to argue that firms in Sub-Saharan Africa start small and do not grow over time, in contrast to surviving firms in other regions (Iacovone, Ramachandran and Schmidt, 2014; Sandefur, 2010). Hsieh and Olken (2014) present more comprehensive data of both formal and informal firms of all sizes (which is generally unavailable for countries in Sub-Saharan Africa) from India, Indonesia, and Mexico, where 98%, 97%, and 92% of firms have fewer than 10 employees, and 65%, 54%, and 22% of the labor force work in firms with fewer than 10 employees, respectively.

International Labor Organization measures put youth (age 15-24) unemployment at 11.8% in Sub-Saharan Africa and 12.6% in Ghana in 2012 (ILO, 2013). The unemployment rate may also understate the difficulties young people face in the labor market, as many are classified as employed but working only a few hours in agriculture or petty trade. Inactivity rates are also quite high, reaching 50% in some countries, and at least 20% in a majority of Sub-Saharan African countries with data, even among young men (Garcia and Fares, 2008).

and geographic areas with high concentrations of program applicants.

Firm owners interested in hiring apprentices through the program and would-be program apprentices were required to attend a matching meeting, conducted at the district by trade level. At these meetings, firm owners introduced themselves and apprentices were given the opportunity to list the firms with which they would be willing and able to work, based on geographic feasibility and general interest. In practice this design meant that some apprentices listed more firms than others. These listed preferences generated apprentice-specific firm sets and within these apprentice-specific firm sets, each apprentice was randomly assigned to one of their listed firms. Each randomization was independent and apprentices had equal probability of being assigned to each of their listed firms. For example, if a would-be apprentice listed two firms, she had a 50% chance of being assigned to each of them. If a would-be apprentice listed three firms, she had a 33% chance of being assigned to each of them.

The apprentice-level randomization generates random variation in the number of apprentices assigned to each firm, conditional on non-random apprentice preferences, which we observe completely. Two key differences arise between a simple unconditional firm-level randomization over a binary treatment status and the design implemented via the matching meetings in this study. First, our treatment is multi-valued rather than binary. Dropping firms that received no interest in the matching meetings, the firms in our sample face non-zero probabilities of being randomly assigned between zero and seven apprentices, depending on the number of would-be apprentices who expressed interest in the firm. Second, treatment assignment at the firm level is conditional not only on the number of would-be apprentices interested in the firm, but also on the probability that each apprentice is assigned by the randomization to that firm (i.e. the number of firms listed by that apprentice). Functionally, treatment assignment is conditional on a set of predetermined characteristics that must be controlled for in the estimation, which we accomplish by including indicator variables for each probability distribution over treatment intensity represented in the data. Our estimates can be interpreted as the impact of being matched with a marginal would-be apprentice across firms with similar levels of apprentice interest.

We follow the firms in our sample for two years after the placement experiment, with four firm-level follow-up surveys at approximately 3 months, 6 months, 1 year, and 2 years after the start of employment. Our program overlaps in the third and fourth follow up surveys with a study that included cash payments to firm owners with program apprentices. We therefore restrict the main analysis to the short-term effects identified by pooling rounds 1 and 2, and report on the longer-term results in the appendix.

Are labor market frictions limiting firm size in our sample? Indeed, our first main finding is that the placement intervention increased employment in firms assigned apprentices through the randomization. Like most job training and placement programs, apprentice take-up was less than perfect. Firms, however, complied with the program design and did not reject assigned apprentices. We show that about half a program apprentice is reported working for each assigned program apprentice and that total firm size (including program apprentices) increases by slightly more than half a worker. This finding implies two things. First, firms assigned program apprentices did not substitute away from other employment by restricting other hiring or firing existing workers. Second, firms that were not assigned apprentices by the program did not increase their firm size through other means, by retaining existing workers longer than planned or hiring additional employees beyond their baseline firm size.

In the second main result of the paper, we show that the placement experiment (and the apprentice labor inputs brought into firms by the placement experiment) increased both monthly reported revenues and monthly reported profits. Each marginal apprentice assigned to a firm leads to an increase in monthly profits of approximately 10%.[2] Recall that the only subsidy to firms provided by the placement program was in-kind recruitment services. The fact that these services increased both employment and reported profits is evidence of economically significant hiring costs that constrain firm growth and keep employment sub-optimally low in the absence of intervention.

An important limitation of this paper is that the sample of firms we study is selected. Observably, firm recruitment carried out by craft-specific trade associations and geographically concentrated where program apprentices had already been recruited generated firms with more workers at baseline than a representative small firm. In addition, within the sample of firms that showed initial interest in the program, our final sample is composed of firms with larger assets at baseline and higher firm owner ability, as measured by managerial practices and cognitive tests. We find no evidence of treatment effect heterogeneity by firm size, baseline assets, or firm owner ability, suggesting that (within our sample) smaller firms that better match the observables of the modal microenterprise are also labor constrained. Unobservably, the firms that chose to participate may have come to the program with more scope or plans for growth, an important caveat to our findings, which should be interpreted as an existence rather than a prevalence result.

[2] With the noted caveats on longer-run analysis in mind, we show in the appendix that firm size effects persist in the 1 year and 2 year follow-ups, while long-term profit effects are statistical zeros.

Armed with evidence of labor market imperfections in our sample, we scrutinize the existing apprenticeship institution in an effort to characterize these frictions. Perhaps the simplest explanation for the entrance fee is that would-be apprentices are financing their own training. In a setting in which post-training retention is low and most returns come from self-employment rather than wage employment, both general and firm-specific training should be financed by the worker herself (Acemoglu and Pischke, 1998; Frazer, 2006). An alternative (potentially complementary) explanation recognizes that the entrance fee norm in West African apprenticeships is strikingly similar to classic models of bond-posting in the labor market and may instead serve to overcome labor market frictions. In one such model, an entrance fee to gain employment (or a bond redeemable at retirement) serves a disciplinary purpose because workers fired for shirking, malfeasance, or poor performance forfeit the entrance fee upon termination of employment, allowing the market to clear (Becker and Stigler, 1974). A related type of labor market friction concerns imperfect information over worker type. Small firm owners may find it difficult to screen out low-ability workers, may incur recurrent hiring and training costs before the ability of inexperienced would-be apprentices can be observed, and may face informal constraints to firing low-ability workers.

Our study in Ghana began with a series of informal interviews with small firms owners. Later, firm-level baseline surveys included a series of questions meant to quantify the qualitative observations gleaned from these early interviews. What say the firm owners? They nearly universally explain the entrance fee as a mechanism for screening workers. 85% of firms in the baseline survey cite a desire to force apprentices to signal investment in the apprenticeship as the impetus for the fee. The most common colloquialism is that firm owners are looking for apprentices who are "serious", which in this context signifies a combination of ability to learn and motivation. Only 8% of firm owners in the baseline survey mentioned that the fee either finances training supplies or the firm owner's training time, and less than 1% say that the fee finances apprentice wages throughout the apprenticeship.

Building on this qualitative evidence, we develop a stylized model to formalize the labor market for apprentices that exists under the non-intervention entrance fee norm. In our setting and in our model, wages are paid as a proportion of firm revenues and firm revenues depend on worker ability. Firms have imperfect information about worker ability and incur hiring and training costs for each worker they hire. Only high-ability workers can expect a wage large enough to recoup the payment of the entrance fee. The entrance fee norm thus generates a separating equilibrium and allows for non-zero entry-level hiring of high-ability workers despite imperfect information.

Observing a market of this type is unusual, and economists typically argue that explicit entrance fees to gain employment are precluded by limited contracting. Of particular concern are worker liquidity or credit constraints (Dickens et al., 1989), which may be particularly binding in low-income countries where unemployed young people are close to subsistence. In our model, would-be workers vary by both ability and wealth. Missing credit markets to finance the entrance fee cause workers whose ability exceeds fixed hiring and training costs to remain unemployed if they cannot afford to pay the fee, undermining the efficiency of the entrance fee screening technology. We interpret the government intervention as substituting the monetary screening mechanism with a non-monetary, time- and effort-based screening technology, allowing poor high-ability workers to enter employment.

Consistent with a first key prediction of the model, access to the non-monetary screening mechanism differentially predicts entry into apprenticeship for poor applicants (Hardy et al., 2019). In addition, we explore de facto selection under the non-monetary screening mechanism, considering two measures of worker ability: a normalized cognitive ability index including performance on a Ravens tests, a Digits Forward test, an English vocabulary test, and a math word problem test, and a normalized non-cognitive ability index including measures from a Rosenberg self-esteem test and a Rotter Locus of Control test. Consistent with a second key prediction of the model, we find that among initial applicants to the program, cognitive ability predicts completing all necessary steps to enter the match randomization. We find less evidence that our measure of non-cognitive ability proxies for worker productivity in our model. We also discuss other candidate market frictions that could explain our findings, concluding that although we cannot fully rule out alternative explanations, the available empirical evidence is most consistent with our model.

The findings of this paper have potentially important implications for theory and policy. A large empirical literature in development economics has explored constraints to firm growth and interventions that could potentially relax those constraints, allowing small firms to expand. Relaxing capital constraints through lending (Banerjee et al., 2015), cash grants, or in-kind grants (De Mel, McKenzie and Woodruff, 2008) has a positive effect on firm revenues and profits, but fails to generate additional hiring. Business training and mentorship programs have more mixed effects on revenues and profits, but as with cash grant programs, employment effects are mostly nonexistent (Brooks, Donovan and Johnson, 2018; Drexler, Fischer and Schoar, 2014; Karlan, Knight and Udry, 2015). In this paper, we report the results of a rare program that effectively increased firm size, potentially shedding light on some of the key constraints to the growth of small firms.

Our work relates closely to a recent experimental literature exploring labor market frictions faced by jobseekers. Several recent studies have found that creating opportunities for jobseekers to credibly signal quality to prospective employers has positive effects on labor market outcomes (Abebe et al., 2018; Abel, Burger and Piraino, forthcoming; Alfonsi et al., 2019; Bassi and Nansamba, 2019; Carranza et al., 2019; Pallais, 2014). Evidence of non-screening search frictions is more limited and more mixed. Job fairs and matching services appear less effective in improving employment outcomes than signaling technologies (Abebe et al., 2019; Beam, 2016; Groh et al., 2015), but there is evidence that subsidies to ease the financial cost of a job search can affect search behavior and the composition of workers who apply to job postings (Abebe, Caria and Ortiz-Ospina, 2017; Franklin, 2018). All of these studies focus on the supply side and none are designed to test for the impacts of labor market frictions on the firm outcomes of interest in this paper. The findings are, however, consistent with our interpretation that screening frictions in particular are an important impediment to worker-firm matches in low-income countries.

The literature studying labor market frictions on the firm side is much smaller. This paper follows De Mel, McKenzie and Woodruff (2019), who find that wage subsidies to spur hiring in Sri Lankan firms have, on average, no durable effect on firm size. They conclude that employment in small-scale retail and services is consistent with well-functioning labor markets. One key way that our study differs is that our sample is composed primarily of small-scale manufacturing firms, where one might expect the ability of novice workers to be more salient than in retail. Indeed, in the third of their sample that is manufacturing, they find long-term firm size effects consistent with the findings of this paper. Two additional recent studies placed apprentices with small manufacturing and services firms in Uganda and Cote d'Ivoire (Alfonsi et al., 2019; Crepon and Premand, 2019). Consistent with our findings, they document little to no displacement of other workers within the firm. Both papers also show suggestive evidence of positive production impacts, though no statistically significant effects on revenues and profits (perhaps for power reasons). The present paper contributes an existence result of statistically significant profit effects associated with relaxing search and screening frictions. Together, the findings of our study and these other papers suggest that labor market frictions are an important barrier to firm growth in some sectors in low-income countries.

Finally, our work relates to a broader literature on price mechanisms and non-monetary allocative mechanisms in contexts where ability to pay is low, for example in the take-up of preventative health products (see Dupas and Miguel (2017) for a nice summary of the literature) or energy efficient cookstoves (Berkouwer, 2020). A number of field experiments have shown that demand for these types of products is sharply downward sloping and that price mechanisms fail to target those who need them most, resulting in overexclusion (Cohen and Dupas, 2010; Kremer and Miguel, 2007; Ashraf, Berry and Shapiro, 2010; Meredith et al., 2013). In several important cases, such as the use of insecticide treated bednets, free distribution maximizes coverage as price paid is unrelated to usage (Cohen and Dupas, 2010). In others, such as free point-of-use chlorine water treatment and near-sighted eyeglasses for children, a non-monetary ordeal mechanism improves targeting, reducing wastage without reducing overall coverage (Ma et al., 2016; Dupas et al., 2016).

The remainder of the paper proceeds as follows. In Section 2, we provide additional detail on traditional apprenticeships and the government program. Section 3 lays out our experimental design, describing our data, the randomization, and estimation. Section 4 presents the main results. Section 5 develops our stylized conceptual framework and explore supporting empirical evidence. Section 6 discusses external validity. Section 7 concludes.

I Setting

I.I Apprenticeships in Ghana

Employment in informal sector Ghana is heavily influenced by the apprenticeship system. The emergence and prevalence of apprentices as workers in West Africa is documented in Frazer (2006).[3] Though the apprenticeship institution has a long history throughout West Africa, it is arguably increasing rather than decreasing in importance.[4] The National Industrial Census reports that in 1984, 18% of wage employees in manufacturing were apprentices, while in 2000, 34% of wage employees in manufacturing were apprentices (Sandefur, 2010). These figures are likely understated for small firms, where the vast majority of workers are apprentices. Additionally, while historically the institution tended to function within extended families, modern apprentices are most often hired from outside the extended family. Nearly half of the apprentices observed in our sample firms at baseline were completely unknown to the firm owner before they began their employment relationship. Less than 15% were members of the firm owner's extended family.

[3]The significance of the institution is documented as well in Bas (1989), Boehm (1997), and Birks et al. (1994). Callaway (1964) and King (1977) put apprenticeship in historical context. Mazumdar and Mazaheri (2003) report on survey data from seven countries in Sub-Saharan Africa, where they find that in Ghana and Cote d'Ivoire, over half of manufacturing sector entrepreneurs have completed apprenticeship training.

[4]Apprentices as a proportion of the manufacturing workforce increased dramatically in Ghana in the last thirty years, following liberalization in the eighties and massive expansion in the number of informal sector firms.

Although the system has no centralized rules or regulations, it is characterized by a few widely practiced customs. Most firm owners and their apprentices (or apprentices' families) enter into verbal or written employment and training contracts. The durations vary, with a median of three years, but common contract lengths of one or two years. These agreements generally require the payment of an entrance fee to start the apprenticeship, followed by wages or "chop money" paid throughout the apprenticeship. Wages tend to be quite low, but increase with seniority, and vary with firm revenues, all of which is true in our baseline data. The entrance fee is approximately equivalent to the first three to twelve months of apprentice wages.

At the completion of the apprenticeship, which is marked by the end of a fixed contract duration, by the discretion of the firm owner, or by the apprentice passing an external craftsmanship exam, the apprentice becomes a "master" of their craft. Master workers then transition into one of several roles. They may be retained on a part-time, full-time, or piece rate basis and receive a sharp increase in wages commensurate with their new title as a "paid worker". Alternatively, they may stay on as a "senior apprentice" with a more moderate increase in wages, while they save capital to start their own firm. Others leave the firm, to start their own shop elsewhere, to work as a paid worker at another firm, or to leave the craft entirely.

Apprenticeship training is concentrated in small-scale manufacturing and services, where young people can learn a craft, such as masonry, carpentry, or garment-making. Large firms do, however, employ apprentices and often employ master craftspeople who have already completed apprenticeships at smaller firms. Gender segregation by occupation is nearly universal, though garment making, the most common trade, is done by both men and women. Training often includes basic literacy and numeracy as well as craft skills, and apprentices begin working on actual customer orders almost immediately.

I.II Experimental Program

The program we study in this paper mirrors traditional apprenticeships in several ways, but also includes some key departures from commonly practiced customs. The program was originally conceived as follows: The national government, in an effort to find work and generate skills among unemployed youth too poor to personally finance the entrance fee for apprenticeships, would place apprentices with small firms including a smaller than normal entrance fee payment. The government would also gift each apprentice a toolkit. Concurrently, two other goals were being pursued. First, the national agency wanted to standardize skills testing and certification, which was at the time conducted redundantly across several ministries. Secondly, the agency recognized conceptually that firm owners may exert some ex-post monopsony power over apprentices once an entrance fee had been paid, and extend the length of apprenticeships (and the period of lower wages) unnecessarily. Consequently, the program we study had a stated duration of one year.

Fee payment and program duration were the subject of contentious negotiation between craft-specific trade association leadership and the national agency leading the program during the government's program design period. The conclusion of these negotiations was that no fee payments were made by the government, neither in the first cohort of the program nor in the second cohort of the program, which we study in this paper.[5] Government officials procured and delivered a tiny fraction of toolsets for apprentices, though this also suffered from logistical and political difficulties. Finally, the one year duration conceived originally was likewise essentially abandoned. Though it appears on program documentation, there was no enforcement of that timeline.[6] The program we ultimately study, then, is a recruitment and placement program, where the subsidy to firms comes in the form only of recruitment, de facto screening, and placement of apprentices. The entrance fee mechanism is essentially abolished for this subset of apprentices. Despite the dispute, firm owners continued to be interested in hiring through the program, and the dispute does not appear to have affected training and employment of program apprentices.

II Experimental Design

II.I Sample Recruitment

Our study sample comes from 32 districts around Ghana, a population-weighted random draw from the 100 districts slated to participate in the second year of the national apprentice placement program. The districts include Accra and Kumasi, the two largest cities in Ghana, as well as rural districts in all ten regions. Appendix Figure 1 shows the selected districts.

The program began in August 2012 with the recruitment of apprentices, through local government officials, advertisements publicly posted at the district office and elsewhere in town centers, and via visits to churches and community meetings. The program intended to target economically disadvantaged young people, but did not enforce an income requirement. Apprentices participating

[5] Political and financial considerations unrelated to this evaluation resulted in the second cohort also being the final cohort of the program.

[6] There was, however, a layered experimental study implemented concurrent with the third firm follow up survey, which included skills testing at about 2 years into the program, and which we discuss in the appendix.

in the program were required to submit a formal application to the local government office and attend a short interview with local government officials (generally the district technical training coordinator, another official from the Ghana Education Service, and someone from the local district assembly). Baseline surveys took place concurrently at the site of these interviews and the initial sample of all applicants comes from the pool of people who both submitted a formal application and attended an interview with these district committees.[7]

Following the recruitment of apprentices, which was completed in January 2013, the program experienced a long and unexpected delay.[8] In May of 2013, program activities began again in earnest, with the recruitment of firms and the commencement of district and trade group matching meetings, which continued through September 2013. All 2,360 applicants were invited to attend a relevant matching meeting. Those who attended and provided lists of firms with which they were willing and able to train entered the experiment under study in this paper. Though we cannot characterize entry into the initial pool of applicants, we will use the sample of 2,360 applicants invited to attend matching meetings to provide descriptive evidence on within-sample self-selection. The application process, including the formal application, interview, attendance at group meetings, and the long lag in program roll-out required a non-trivial investment of time and energy from potential apprentices. It is this application process that we interpret and model as a non-monetary screening mechanism.

Firms in the sample were recruited by local government officials and craft-specific trade associations to hire and train the unemployed young people who were the targeted recipients of the program from the perspective of the government. Recruitment of firms took place independently of apprentice recruitment and after the apprentice recipients were chosen, though it was targeted in the sense that local government officials and trade association leadership sought firms that broadly matched the location and trade preference of program apprentices. The program targeted five

[7] The experiment on which we report in this paper was enclosed in a larger randomized controlled trial (Hardy et al., 2019), which randomized over unemployed young people applying to become apprentices. Within the initial pool of 3,928 applicants to the parent project, the research team addressed local political economy concerns (on the advice of our partners in the national agency) by allowing local district committees to hand-pick young people for about 15% of the slots planned for each district. We then randomized over the remainder of applicants. In total, local district committees hand-picked 329 people and 2,031 people were randomly assigned to treatment. Together, these 2,360 young people were invited to continue the application process by attending a matching meeting as described below. These 2,360 applicants to the apprenticeship are the relevant pool for any analysis on worker selection in this paper, because control applicants to the apprenticeship program were not invited to complete the rest of the application procedures. District selection and randomization into the apprenticeship program took place before any firms were recruited and are not the subject of this paper.

[8] The incumbent presidential administration lost the December 2012 election leading to turnover in the leadership of the implementing agency, which was followed by a severe fiscal crisis associated with a drop in the price of gold.

trades: garment making, hair/beauty, welding, carpentry, and blocklaying. In our sample, garment making includes both men and women, hair and beauty is nearly all women, and the construction trades are nearly all men, both among firm owners and apprentices. In general, firms were approached by government officials and trade association leadership and asked if they would be interested in hiring apprentices through the government program.[9] Interested firms were then invited to attend a matching meeting. For the purpose of these meetings, welding, carpentry, and blocklaying were grouped together as skilled construction. It was at these meetings that the research team first enrolled firms in the study, and at these meetings that firm owners participated in the baseline survey.

II.II Placement Intervention

Matching meeting activities began with firm owner registration and briefings that provided firm owners with more detail on the program. In particular, conditional on geographic feasibility and apprentice willingness, apprentices would be randomly allocated. This protocol was acceptable in part because the assignment of apprentices to firms was seen by firm owners as a government benefit, so random placement allowed for arguably fair distribution of that benefit. In addition, firm owners would not have the opportunity to reject program apprentices (because the design sought to ensure a placement for every apprentice). Information on capacity constraints was collected as part of registration, though due to a relatively disperse sample across districts and trades, capacity constraints were never binding (i.e. it turned out no firm owner was randomly assigned more apprentices than she had initially stated she was willing to hire). Firm owners still interested in hiring apprentices through the program then introduced themselves to the gathered group of apprentices, and stated the precise location of their businesses.[10] In addition to a simple introduction and detailed information on location, some firm owners also shared information on their passion for training apprentices, their experience with craft-specific specialties, or the number of apprentices currently under their employ.

Apprentices were then given the opportunity to provide a list of firms with which they would be willing and able to work and train. The instruction was to provide information on firms within

[9] The basis of the communication itself was typically that the government wanted to place poor young people with local firms to work as apprentices, and did not explicitly state that these apprentices would have undergone screening. That interpretation of program effects is the work of the authors and was not built into the conceptualization of the original program design.

[10] The formal meeting activities were heavily monitored, though unmonitored communication between participants was also possible.

their craft of interest that were close enough to their homes that they could reach them without incurring large transport costs. However, detailed GPS or other information on firm location and apprentice home location was not available at the time so district officials and research field teams had no ability to enforce that instruction. Consequently, the apprentice-specific firm sets include both geographic feasibility (walkability, generally) and idiosyncratic preference. No minimum or maximum was placed on the number of firms listed and apprentices who listed only one firm were assigned that firm. We refer to apprentices who listed a single firm as degenerate lotteries, and discuss the implications of these lotteries below. Among those who listed at least two firms, the mean and median are both three firms. A small fraction of apprentices were unable to find a walkable firm at the matching meeting; our research team found matches for them later in the process and they are not relevant to the variation for this paper. Anecdotally, we believe the firm sets to be an honest revelation of preferences, where apprentices who listed multiple firms were geographically able and willing to work at all of the listed firms.[11]

Program placement occurred in 28 districts in October 2013, and in the final four districts in January of 2014. Placement involved field staff and local government partners informing apprentices and firms of their placements, and following up to determine whether apprentices had reported to their assigned firms.

II.III Randomization

Within apprentice-specific firm sets, each apprentice was randomly assigned to one of the firms in their set, using a computer generated random number. No re-randomization or stratification beyond individual apprentice was done, and each randomization was independent. If the apprentice only listed a single firm as both geographically feasible and desirable generally, she was assigned to that firm. In this section we discuss our identifying exogenous variation at the firm level and explain how we control for the fact that it is conditional on non-random apprentice interest.[12]

Consider first, a classic RCT, which unconditionally splits a sample randomly into a treatment group and control group, without any stratification. Households in the treatment group receive

[11]The apprentice preference revelation mechanism was not incentive compatible, as that would have required an option in which apprentices do not get a placement. Instead, it focused on ensuring that all apprentices had a feasible placement, as all apprentices invited to the district and trade group meetings were already guaranteed a placement. Note however, that our identification strategy does not in any way depend on an incentive compatible revelation mechanism.

[12]We thank Michael Andersen, Kenneth Chay, Owen Ozier, and Chris Walters for advice on controlling for the randomization and extended comments. This section draws on notation and framing from Duflo, Glennerster, and Kremer (2008).

\$100 in cash; households in the control groups receive \$0. Because the treatment was randomly assigned, the two groups differ in expectation only in their exposure to the treatment. In the potential outcomes framework of Rubin (1974), the difference in observed means post-program ($D = E[Y_i^T|T] - E[Y_i^C|C] = E[Y_i^T - Y_i^C]$) provides an unbiased estimate of the treatment effect. In a regression framework, we simply regress our outcome of interest on an indicator for inclusion in the treatment group ($Y_i = \beta_0 + \beta_1 T + e_i$). Estimated by OLS, β_1 obtains the average treatment effect of receiving \$100.

The first complication in our research design is that the treatment takes multiple values. Suppose we are delivering cash to households in \$100 increments, so some households receive \$0, some households \$100, some households \$200, and so on. Where assignment into treatment groups is unconditionally random, the logic of the binary case extends to the multi-valued case, and the difference in means is an unbiased estimate of the treatment effect. If we are additionally willing to assume linearity in the treatment effect, we can parameterize T to take values 0, 1, and 2, and so on, and estimate the average treatment effect using the same simple equation above. In the present study, the maximum number of possible apprentice assignments in the effective sample is seven. Realized treatment values vary from zero to five apprentices. 44% of the effective sample were assigned zero apprentices, 37% were assigned one apprentice, 14% were assigned two apprentices, and 5% of the sample were assigned three or more apprentices.[13]

The second relevant deviation from the simplest case is that the probability of selection into treatment depends on some observables. With a simple binary treatment, the most common case of conditional randomization is one in which the authors stratify the sample by some observables and randomize within strata. For example, the strata may be cities or schools or genders. If some strata are more oversubscribed than others, the probability of being randomly chosen for the treatment varies by strata. To recover the average treatment effect, one includes indicator variables for each strata in the same simple OLS regression, comparing treatment and control observations within each strata and then taking a weighted average of each effect.

Combining conditional randomization and a multi-valued treatment assignment, we are looking for an observable that perfectly predicts the probability distribution over all treatment values, as strata do over binary treatment assignment in the prior example. Taking our cash example, suppose households in Accra face a probability distribution of \$0, \$100, and \$200 equal to (0.5, 0.25, 0.25)

[13] Given this distribution of realized treatment assignments, sample constraints make it difficult to test for diminishing returns across all values in our data. See the appendix for effects estimated separately for firms assigned zero, one or two apprentices. In the main body of the paper, we assume linearity.

while households in Kumasi face a probability distribution of \$0, \$100, and \$200 equal to (0.3, 0.3, 0.4). Controlling for an indicator variable for city would effectively control for the differences in the probability of treatment faced by households in Accra and households in Kumasi precisely because it is equivalent to an indicator variable for the probability distribution over treatment intensity itself.

In the present study, we observe directly the probability distribution over treatment intensity, because it is fully predetermined by apprentice preferences. For each firm, we know the number of apprentices who listed the firm and the probability that each apprentice would be assigned to that firm. Combining this information across all apprentices yields a probability distribution over all possible treatment assignments. Appendix Tables 1 and 2 list the probability distributions with support in the sample. All cells without any treatment variation are dropped, which implicitly drops both firms that received no interest from apprentices at matching meetings and firms that were listed only in degenerate lotteries. As one example, 85 firms in our sample face a 50% probability of being assigned zero apprentices, a 50% probability of being assigned one apprentice, and 0% probability of being assigned two or more, generated by being listed by a single apprentice who listed two firms. In a more complex example, 9 firms in the sample face a probability distribution of (0.06, 0.25, 0.38, 0.35, 0.06, 0, 0, 0) of being assigned (0, 1, 2, 3, 4, 5, 6, 7) apprentices, generated by being included in the firm sets of four would-be apprentices who each listed two firms.[14] To recover an unbiased average treatment effect, we include in our estimation an indicator variable for each of the probability distributions over treatment intensity delineated in Appendix Tables 1 and 2. Our estimates are thus a weighted average of the treatment effects within each of these cells, as in the city example above.

A final issue to consider is that our research design provides access to workers, not dollar bills. Drawing on familiarity with the observational literature on worker matching, one might be concerned that the characteristics of individual apprentices may bias our estimates, as treatment is literally conditional on apprentice preferences. However, within each probability distribution, the conditionally random nature of the placement experiment guarantees that in expectation worker heterogeneity should be balanced across treatment assignments.[15]

[14] Our notation in the appendix tables for this probability distribution is 2,2,2,2.

[15] In that sense, our study follows experiments that provide inherently heterogeneous treatments like jobs or tutoring or university educations. Apprentice preferences may be of separate academic interest, but are not the subject of this paper. Instead, we use the probability distributions to extract from the apprentice-specific firm sets a simple RCT that generates unbiased estimates by controlling directly for the probability of treatment.

II.IV Data

Data come from four sources: (1) firm baseline surveys, (2) apprentice baseline surveys, (3) apprentice-specific firm sets, and (4) firm-level follow-up surveys conducted at approximately 3 months and 6 months after the start of employment (Hardy and McCasland, 2022). Dropping probability distribution cells with no realized treatment variation, we remain with 755 firms assigned 621 apprentices. The four districts in which program placement occurred in January 2014 (and which account for 86 of the firms in the final sample) were excluded from the first follow-up survey and are only observed in later follow-up data collection. 748 of the 755 sample firms participated in a baseline survey which included personal background, digit span recall, four math questions, capital stock, detailed labor inputs, revenues and profits, managerial aptitude questions, and information on apprenticeship training experiences. 607 of 621 apprentices in the effective sample participated in a baseline survey which included education, training and work background, an adapted measure of the Rotter Locus of Control, an adapted measure of the Rosenberg self-esteem scale, and a series of cognitive tests, including Digit Span Recall, four math questions we developed ourselves, Ravens Matrices Group B, and a fifteen word oral English vocabulary recognition test. The would-be apprentice non-cognitive ability index includes the Rotter and Rosenberg measures; the would-be apprentice cognitive ability index include the four cognitive tests mentioned here. Following Ozier (2018), these indices are the normalized sum of the normalized scores on each of the individual tests.

The first firm follow-up survey targeted 669 firms. Our primary outcome variable, profits, is non-missing in this targeted sample for 88% of the firms assigned zero apprentices, 93% of the firms assigned one apprentice, 91% of the firms assigned two apprentices, and 91% of firms assigned three or more apprentices. These differences are not statistically significant. The second follow-up survey targeted all 755 firms in the effective sample, with non-missing profit observations for 84%, 87%, 94%, and 82% of the sample, respectively, for an overall tracking rate of 86%. Attrition is higher in the second follow-up survey, but also not significantly different by treatment assignment. 95% of the firms in the sample are observed in at least one of the two follow-up surveys. Table 1 presents these results.

Follow-up surveys included revenues, profits, other labor inputs, and detail on program apprentices. The second follow-up also included updated capital stock measures. Note that other labor inputs were measured as number of non-program apprentices and other workers. We did not track individual worker identity beyond the program sample, so we report results on firm size rather than specific information on non-program apprentice retention or hiring. All financial variables, including revenues, profits, assets, and wages have been deflated to April 2013 Ghana Cedis, the first month for which we measure baseline profit values. Our main specifications winsorize the top 0.5% of values for the financial variables with long right tails: revenues, profits, and capital stock. Self-reported revenues and profits follow the question structure of De Mel, McKenzie and Woodruff (2009) and are reported on in more detail in the data appendix.

II.V Summary Statistics and Covariate Balance

In our nationwide sample of 755 small firms, apprentices comprise the vast majority of the workforce. In the 658 firms who have any workers besides the owner at baseline, 78% of the 2,522 workers are apprentices. 50% of the these apprentices had no connection to the firm owner prior to employment and another 38% were unknown to the firm owner but introduced through a relative or friend, underlying that modern apprenticeship is largely an anonymous market activity. The mean monthly wage for an apprentice in our baseline sample is about 24 Ghana Cedis, which at the time of baseline surveys was about 12 US dollars. Garment-makers are the most common trade, we have more female firm owners than male firm owners in the sample, and only about 8% of the sample is registered with the Registrar General (to pay income taxes).

To test systematically for imbalance across firms assigned different numbers of apprentices once we control for non-random apprentice interest (that is, to extract the random element of the design and look for imbalance across experimental groups), we regress firm baseline characteristics on treatment assignment controlling for dummies for all probability distributions over treatment intensity. We also include district and trade fixed effects to mirror our main specification throughout. Each cell in Table 2 comes from a separate regression of the following form:

$$Baseline_i = \beta_0 + \beta_1 * T_i + \gamma_d + \lambda_c + \varphi_p + \epsilon_i \tag{1}$$

where T_i is the number of apprentices assigned to the firm by the randomization, γ_d are district fixed effects, λ_c are trade (craft) fixed effects, and φ_p are probability distribution dummies. Though point estimates are not zero and somewhat imprecisely estimated, none are statistically different from zero. Importantly, we find no evidence of imbalance on labor inputs, capital, revenues, or profits. On these key variables, the randomization procedure achieved conditional balance across treatment assignments.

II.VI Estimation

We have two main outcome groups of interest: (1) labor inputs and firm size, and (2) revenues and profits.[16] As above, our main specification controls for a set of dummies for all probability distributions over treatment intensity, and district and trade fixed effects. We stack data from the first two follow-up rounds as follows:

$$Y_{it} = \beta_0 + \beta_1 * T_i + \eta_t + \gamma_d + \lambda_c + \varphi_p + \alpha * Y_{i0} + \epsilon_{it} \tag{2}$$

where η_t is a round dummy, Y_{i0} is the baseline value of the dependent variable (for outcomes for which this value is captured at baseline), and the coefficient β_1 estimates the intent-to-treat effect. β_1 can be interpreted as the average effect of each assigned apprentice across follow-up rounds, where the effect of each apprentice enters the function linearly. Standard errors are clustered at the district level.

III Results

III.I Take Up and Other Inputs

Take-up requires both that the firm owner accept to train and employ apprentices and that apprentices report to their employment assignments. To our knowledge, only one firm in the study refused to train and employ the apprentice(s) assigned to their firm. Of the 621 apprentices assigned via the random match process, 377 (61%) were reported to be working at their randomly assigned firm in at least one of the two follow up rounds, 52 (8%) were reported to be working at a firm in the study sample other than their assigned firm, 181 (29%) did not report to any firm in the study, and 11 (2%) were not confirmed as their assigned firms attrited from the study.

Table 3 displays regression results. The experimental effect of each assigned apprentice is an

[16] This project was registered with the American Economics Association Randomized Controlled Trial Registry (AEARCTR-0000297), complete with a Pre-Analysis Plan (PAP). The PAP was intended to coalesce ideas on the direction of analysis, and limit both the risks and perception of data mining or specification search. The estimation procedures described in the PAP did not properly control for non-random apprentice interest and were thus abandoned. In addition, we did a poor job in the PAP of grouping hypotheses into families. Consequently, it would be difficult to use these to guide any multiple hypothesis testing adjustments. The spirit of the analysis plan, however, corresponds well with both the early qualitative work that inspired this study and the findings presented in this paper. We focus on a limited set of key outcome variables that were arguably, if not demonstrably, a small set of a priori hypotheses.

implied take up rate of 47% and an estimated increase in overall workforce by 0.58 workers, both significant at the 1% level. Point estimates on other labor and capital inputs are mostly positive, but not statistically different from zero. Across the first two follow up surveys, these findings imply that any impact on revenues and profits are driven by apprentice labor rather than complementary inputs such as capital or firm owner labor supply.

III.II Treatment Effects on Revenues and Profits

Our second main set of results is presented in Table 4. For each additional apprentice assigned, profits increase by about 40 Ghana Cedis, about 10% of the sample mean. Measured in inverse hyperbolic sine, profits increase by 11%, significant at the 1% level.[17] The point estimate on IHS revenues is also statistically different from zero and of a similar magnitude, suggesting an increase in output associated with alleviating labor constraints. Wages in Table 4 are also intent-to-treat; for each assigned apprentice, firms are paying out an additional 12 Ghana Cedis. Given that take-up is about half an apprentice, this corresponds with inflation adjusted wages that match those paid to non-program apprentices observed in sample firms at baseline.

Estimates in Table 5 test for robustness of our main effects to specification choice, with levels in Panel A and IHS in Panel B. Column (1) reproduces the estimates from our main specification. In levels, presented in Panel A, Column (2) presents estimates without winsorizing the top 0.5% of values. This specification is not relevant for IHS profits, an alternative way to improve power for an outcome variable with a long right tail and a data transformation that starts with raw (non-winsorized) levels. Columns (3) and (4) alter the controls, excluding district and trade fixed effects or the baseline value of the dependent variable. Column (5) presents treatment effect estimates using a quantile regression specification. Taken together, Table 5 suggests the point estimates on our key outcome variables are quite stable across a diverse array of specification options, most of which remain significant at traditional levels.

Estimates in Table 6 test for the sensitivity of our main results by sub samples defined by their probability distributions over treatment intensity. The purpose of the table is to show that our point estimates are stable across groups, to assuage concerns that something unusual about certain probability distributions is driving our effects. Columns (1) and (2) present effect estimates for level profits and IHS profits excluding firms whose probability distribution includes any degenerate

[17] Appendix Figure 2 displays distributions of residual profits (estimated by regressing winsorized profits on controls) by treatment assignment. Appendix Figure 3 displays randomization inference treatment effects, where the exact p-value for level profits is 0.09 and the exact p-value for IHS profits is 0.01.

lotteries. Referring to Appendix Tables A1 and A2, the probability distribution (1,2) corresponds to a single degenerate lottery and being ranked by an apprentice who listed two firms. Firms in this probability distribution thus have a 50% chance of being assigned one apprentice and a 50% chance of being assigned two apprentices. This and all other probability distributions with a degenerate lottery are excluded from this sub-sample. Columns (3) and (4) estimate treatment effects separately for firms listed by only one apprentice and for firm listed by more than one apprentice, showing an F test of equality of these treatment effects of 0.58 (for level profits) and 0.25 (for IHS profits), failing to reject equality. The 335 firms listed by a single apprentice more closely resemble a typical RCT in which different strata have a different probability of being assigned a binary treatment. Columns (5) and (6) split the sample in half by the commonness of the relevant probability distribution over treatment intensity. 370 firms, or 49% of the sample, faced one of the six most common probability distributions, each containing more than 30 firms (listed at the top of Appendix Table A1). Again, we fail to reject equality of treatment effects between sub-groups.

IV Interpretation

Large and robust increases in employment and profits in response to recruitment services suggest that firms in our context are labor constrained. To understand why government intervention was necessary to create these worker-firm matches, we develop a stylized asymmetric information model of the labor market for apprentices in West Africa in the absence of intervention, centered around the role of the entrance fee and inspired by qualitative data we discuss in the Appendix.[18] In the model, firms decide whether to incur a fixed cost to hire an individual apprentice and workers decide whether or not to work given an equilibrium wage contract. We then build on the framework to interpret the intervention itself and explain the main findings of this paper.

The model makes a series of simplifications for convenience. Firms are modeled as perfectly competitive; workers are modeled as having discrete ability types; the model is single-period and ignores the dynamic effects of training inputs on worker productivity. In choosing to focus on a firm owner's decision to hire or not hire an individual apprentice, we implicitly assume constant returns to scale over labor inputs. The simplicity however, allows us to focus on the key insight of the model, that imperfect information can constrain hiring and that a monetary or ordeal mechanism

[18]See Appendix Section 5 for a detailed discussion of the qualitative and descriptive evidence from our data that supports our interpretation that screening frictions explain both the entrance fee norm and the efficacy of the program we study.

can be used to screen out low-ability workers.

IV.I Modeling the Entrance Fee Mechanism

Workers are either high-ability θ_H or low-ability θ_L, and their type is known to them. A worker is willing to work if the offered compensation $r_w(\theta) > r_o(\theta)$, the worker's outside option. For simplicity, we assume that the outside option for any ability worker is $r_o(\theta) = 0$ and that workers weakly prefer their outside option, meaning that all workers want to work for any compensation package $r_w(\theta) > 0$. Additionally, workers have an initial wealth endowment of $\gamma \geq 0$ and there is no access to credit. Wealth γ is continuously distributed across workers with some cumulative distribution function F_g.

A worker's contribution to a firm is $Y(\theta) = \theta$. Hiring a worker costs $c > 0$, where $0 \leq \theta_L < c < \theta_H$. Therefore, it is unprofitable for a firm to hire workers with ability θ_L and profitable for a firm to hire workers with ability θ_H. If firms had perfect information about worker types, then θ_L workers would not work and θ_H workers would work for $w_H = \theta_H - c$. Firm owners, however, do not observe ability and make hiring decisions using expected ability $\hat{\theta}$. For simplicity, we assume that $\hat{\theta} < c$ for all workers, meaning that absent a screening mechanism, no hiring would occur.[19]

In the baseline labor market, firms offer a two-part wage contract: would-be apprentices pay an up-front entrance fee and receive positive revenue sharing ($s \in [0, 1]$) once employed. Expected profits are $\hat{\pi} = (1-s)(\hat{\theta}|s, \overline{w}) + \overline{w} - c$, where $\overline{w}$ is the entrance fee paid by the worker to buy into the job. If $\overline{w}$ and s are set such that $s\theta_L \leq \overline{w} < s\theta_H$, then the firm can effectively screen out low-ability workers and hire high-ability workers, who self-select into paying the entrance fee because they can expect a share of revenues large enough to compensate them for the up-front payment.[20]

The status quo absent government intervention is a sophisticated labor market institution that allows firms to hire some high-ability workers despite imperfect information. However, a market failure remains: high-ability would-be apprentices who could generate profits for firms but cannot

[19] These assumptions apply primarily in the anonymous market for non-family workers. Empirically, family members are rarely required to pay an entrance fee, and even close acquaintances or neighbors may also be exempt from the requirement. In these cases, we would presume a few key differences with our model. First, the search and screening costs for family members are likely lower. Secondly, the firm owner likely has better information about the ability of the worker she knows and can therefore choose to employ or not employ her on the basis of that information. Finally, some potential intrahousehold transfers could be enclosed in the employment relationship between family members. Wages in the case of family members would then be a function of both ability and intrahousehold transfers paid as wages. In our baseline data, family members are paid more than non-family members, which we interpret to be the result primarily of intrahousehold transfers paid as wages.

[20] In a perfectly competitive equilibrium, $s\theta_L = \overline{w}$ and firms raise s and lower $\overline{w}$ until $\hat{\pi} = (1-s)\theta_H + \overline{w} - c = 0$. Plugging $s\theta_L = \overline{w}$ into $\hat{\pi} = (1-s)\hat{\theta} + \overline{w} - c = 0$ we find: $(1-s^*)\theta_H + s^*\theta_L - c = 0 \implies s^* = \frac{\theta_H - c}{\theta_H - \theta_L}$ and $\overline{w}^* = s^*\theta_L = (\frac{\theta_H - c}{\theta_H - \theta_L})\theta_L$

afford to pay the entrance fee (because $\gamma < \overline{w}^*$) remain unemployed.

IV.II Government Intervention

Recall that the government program required applicants to attend several meetings, interviews, and surveys, and continue to show interest in the apprenticeship despite a long and unplanned lag in program roll-out. In our preferred interpretation of the government program, this recruitment process required workers to pay a non-monetary "sweat equity" entrance fee, which allowed for the screening out of low-ability workers without the use of a monetary entrance fee. We call this non-monetary screening cost $\overline{u}$.

Note that unlike $\overline{w}$, $\overline{u}$ is fixed by the program and not determined by a market equilibrium. Absent a clear market equilibrium solution, we assume this non-monetary cost satisfies $s'\theta_L \leq \overline{u} < s'\theta_H$, where s' is the share of revenues paid to program apprentices. The key insight here is that the government program de facto allowed workers to pay the entrance fee in a non-monetary way, drawing out of unemployment that segment of the workforce where $\overline{u} < s'\theta$ (would-be profitable workers with $\theta = \theta_H$) but personal savings $\gamma < \overline{w}^*$. This solves the market failure generated by the combination of the entrance fee screening mechanism and missing credit markets to finance the fee and explains the increase in firm size due to government intervention in this context.

Focusing on the creation of new worker-firm matches, two simple predictions are generated by this interpretation of the government program. First, higher ability applicants to the apprenticeship program should be more likely to overcome the ordeal mechanism and enter the final sample of placements. Secondly, the non-monetary mechanism is open to all applicants regardless of wealth, while the status quo entrance fee mechanism is only open to those applicants who can afford the fee. Consequently, the government program should differentially increase access to apprenticeships for young people from low-wealth households.

To test for the first prediction, we proxy θ with normalized indices of applicant cognitive and non-cognitive ability. Within the initial sample of applicants, we test for whether higher θ applicants are more likely to attend the matching meeting, complete a list of firms, and list more than one firm. The first two of these outcome variables are very similar, as only a small fraction of would-be apprentices were unable to find a geographically feasible match at their relevant matching meeting. The last of these outcome variables is an indicator for generating variation in the match randomization. We find evidence in Table 7 that cognitive ability predicts paying this portion of the non-monetary "sweat equity" entrance fee, with less evidence that non-cognitive ability predicts the same.

To test for the second prediction, we refer to Hardy et al. (2019) Appendix Table 6, which shows that access to the government program differentially increased apprenticeships for poorer applicants. The larger randomized controlled trial in which this study is contained randomized over unemployed young people applying to become apprentices. Young people in the control group could still enter apprenticeships via the traditional entrance fee mechanism (as indeed more than half do), while young people in the treatment group could enter apprenticeships either through the traditional entrance fee mechanism or by paying the non-monetary cost to enter the government program. The first stage on entering an apprenticeship is 50% larger for workers whose baseline asset index is one standard deviation lower than the mean.

IV.III Alternative Explanations

An alternative or complementary explanation for the creation of new worker-firm matches generated by the program is that the matching meetings and the government structure signalled something positive about firm quality to workers. Another way of framing this critique is that firms providing employment and training through the government program could differ from firms providing employment and training through traditional apprenticeships entered under the entrance fee mechanism. We can test for this prediction again using the first stage from Hardy et al. (2019), where Appendix Tables 4 and 5 compare characteristics of apprenticeships undertaken by applicants in the control group to apprenticeships undertaken by applicants in the treatment group (who have access to entry under the ordeal mechanism). There is no evidence of statistical differences in firm size or worker satisfaction with the experience, two potentially important quality metrics. Indeed, the apprenticeships are observably similar except along dimensions explained by the program design (e.g. differences in the entrance fee). Anecdotally, the firms in our sample are precisely those that commonly hire apprentices via the baseline labor market.

One might also be concerned that the non-monetary mechanism is either more stringent or more precise, bringing in higher quality workers than the entrance fee mechanism or reducing uncertainty over type and thus increasing firm owner willingness to hire. Data limitations and our study design preclude a direct comparison of the monetary screening mechanism to the non-monetary screening mechanism. However, evidence against the simplest version of this alternative mechanism comes, as well, from the larger study. Hardy et al. (2019) show in Appendix Table 6 that there is no first stage heterogeneity by cognitive ability, suggesting that applicants in the control group who enter apprenticeships under the entrance fee mechanism do not differ along this dimension from applicants in the treatment group who have access to entry under the non-monetary mechanism.

Further, in our own data, we note that program apprentices receive inflation adjusted wages at follow up similar to those received by non-program apprentices at baseline, within the firms in our sample. Taken together, we fail to find evidence of differences in firm quality by entry mechanism, in worker ability by entry mechanism, and in wages by entry mechanism, suggesting that not only are firms and workers observably similar, but also that we do not have evidence that match quality is improved by the ordeal mechanism.

Finally, our model has focused on screening as the key friction, following the structure of the baseline labor market and qualitative evidence from firm owners who cite screening as the impetus for the entrance fee. However, the placement program also simply eased more basic search costs for employers, by bringing willing applicants to a central location, eliciting preferences, and providing individual placements. The fact that high ability workers are more likely to surpass the hurdles put in place by the ordeal mechanism, as we see in Table 7, suggests that although simple search costs may indeed be important, they are unlikely to explain the entire effect of the program on the creation of new worker-firm matches.

V External Validity

In this section, we discuss the external validity of our findings, focusing on the key issue of sample selection, as well as rent-sharing in our context.

V.I Sample Selection

Sample selection in our study arises from four sources: (1) access to information about the opportunity, (2) firms choosing to participate by attending the matching meeting conditional on receiving an invitation, (3) workers choosing firms with which they were willing and able to work from the pool of firms that attended the matching meeting, and (4) firms that were listed in the feasible sets of apprentices dropping out of the final study sample because their probability distribution cells had no realized treatment variation. The first three of these sources intersect with sample selection generated by geographic feasibility: recruitment intended to target firms where apprentices reside, firms with shorter transit times to district capitals where matching meetings were held may have been more likely to attend, and apprentices who attended matching meetings were able to make geographic feasibility decisions with more detailed private knowledge about transit times and costs from their residence to the location of each business.

Unfortunately, we have relatively little information about the number of firms that were invited to participate. For many, information about the program spread by word of mouth, so estimating how many firms got wind of the opportunity is difficult. Consequently, we also have relatively little information about the proportion or characteristics of *contacted* firms that chose to participate. We can however, pool these two first sources of sample selection, and look for observable differences between firms that attended the matching meeting and all firms, using a census we collected for another study of all self-employed garment making firm owners in one district in our sample (Hohoe District). In addition, across all districts and trades, we can pool the final two sources of sample selection, and look for observable differences between the approximately 1800 firm owners that attended matching meetings and the final sample of 755 firms.

Recruitment of firms into the sample was conducted by local government officials primarily through craft-specific trade associations and in areas where applicants to the apprenticeship program needed firms with which to match. In Hohoe, as in most districts in the sample, applicants to the apprenticeship program were concentrated along the main roads. In Table 8, we present observables for the full census, the sample of firms that attended the matching meeting, and relevant census subsamples. Relative to the full census, matching meeting attendees own larger firms along all measures. They have more workers, more assets, larger revenues, and practice more modern managerial practices. Concentrating on the census subsample most likely to have actually had access to information about the program (those along the main roads whose firm owners were also members of a trade association at the time of the survey), firm size as traditionally measured by the number of workers remains a predictor of participation.

One question of particular interest is how selection relates to expected gains from the program. If firms that are particularly labor constrained are more likely to appear in our sample, then our results are an upper bound. If firms in our sample are selected due to logistical issues with program implementation, but are otherwise representative of the labor constraints of a typical Ghanaian manufacturing microenterprise, then the scope for policy lessons widens. Recent work by Hsieh and Olken (2014), has found that larger firm size is correlated with more severe labor constraints, as measured by average product of labor. A couple data constraints arise in applying their methodology to our sample: the Hohoe census data lacks the wage bill (meaning we cannot calculate value added by summing wages paid and profits) and firm owners in manufacturing and services firms of this type are far more productive than their workers.[21] We adjust their measure of log(value added/worker) by replacing value added with profits less an adjustment for the mean profits in firms in the census with no non-owner labor and explore sample selection by IHS(adjusted profits/worker). Again, by this measure, the firms that attended the matching meeting are far more productive than the entire census of firms. However, in comparing the firms that attended the matching meeting to likely targeted firms, our adjusted measure of labor productivity is remarkably similar.

Across all districts and trades, we can search for observable differences between our final sample and the full set of firms that attended any matching meeting, where sample selection is a function of both worker preferences and the structure of controlling for the randomization. Within district and trade, baseline assets and performance on a firm owner ability index predict entry into our final sample, as shown in Table 9.[22] In this full sample, we have the wage bill, but still face the issue of labor specialization. We therefore adjust standard average product of labor measures by replacing value added with value added less the mean profits for a single person firm and test for selection by IHS(adjusted value added/worker). Here again, this observable measure of labor constraints does not predict entry into the variation for this study.

Taking baseline firm size, baseline assets, and baseline firm owner ability as observables along which our sample is selected, Table 10 presents heterogeneity results by each of these characteristics. We fail to find evidence of heterogeneous treatment effects. Perhaps reassuringly, these point estimates suggest that within our sample, smaller firms and firms led by less-skilled firm owners appear to be labor constrained. Still, we caution the reader that unobservable self-selection remains an area of concern with respect to the external validity of these findings.

V.II Monopsony

An additional external validity concern relates to the baseline labor market. As we note above, observing a labor market in which aspiring workers explicitly pay an entrance fee (or post a refundable bond) to enter employment is quite unusual. In the case of a non-refundable monetary or non-monetary entrance fee, the structure of the implicit contract suggests that ex-post monopsony power is likely to arise once an aspiring worker has entered an apprenticeship. Once the fee or

[21] See Jensen and Miller for another example of labor specialization in small-scale manufacturing in a low-income country, where the majority of highly-skilled work is conducted by the firm owner, even as firm size increases.

[22] The ability index is a normalized sum of the normalized scores on a Digits Forward test, a four question math test, and the five managerial skills questions.

the "sweat equity" is paid, shifting to another employer would require a repayment of the initial screening fee, generating downward pressure on worker wages at the existing employer. In addition, eyeballing the wages paid (about 25 GHC per month) against the profits gains for firm owners suggests that a large share of worker product accrues to firms. Finally, a back-of-the-envelope exercise to estimate labor supply elasticity (which approaches infinity where employers are price takers in the labor market and is closer to zero where monopsony employers face upward sloping labor supply) generates estimates across districts that range from about 0.5 to about 5.[23] In other words, these estimates suggest a fairly monopsonistic labor market.

On the other hand, the firms in our study are small, so the classic idea of a monopoly employer seems out of place. In addition, anecdotal settings in which employees put in "sweat equity" with a particular employer (e.g. novice work, internships, student research assistantships, etc.) are exceedingly common. The literature on monopsony across settings suggests that a large share of worker product accrues to firms rather generalizably. Studying the German apprenticeship system, Acemoglu and Pischke (1998) argue that ex-post monopsony power arises where workers receive on-the-job training and signal ability *during* the apprenticeship. Dal Bo, Finan and Rossi (2013) use randomized wage levels across Mexico to estimate a labor supply elasticity of about 3. Kline et al. (2018) estimate that workers capture about 30 percent of the surplus generated by an employer's successful patent. Dube et al. (2020) estimate elasticities in the range of 0.1 in online labor markets and Azar, Berry and Marinescu (2019) use job vacancies to estimate market-level elasticities of about 0.6 across United States local labor markets. Taken together, one way to place our findings in context is to recognize that there may be downward wage pressure in this setting, but that is also true of many settings.

VI Conclusion

A large literature tests for barriers to small firm growth. This literature has primarily focused on capital, credit, and managerial constraints. Our study provides the first experimental evidence that some small firms face labor constraints.

This evidence of labor constraints in our setting compels a reconsideration of the common assumption that all small firms in developing countries face a frictionless labor market characterized

[23] We do not have exogenous variation in wages or productivity, so these estimates are necessarily quite rough. We take our adjusted measure of the average product of labor and use it to proxy for marginal revenue product in a labor market version of the Lerner Index, $((MRP - W)/W)$ which equals $1/\epsilon$, and solve for labor supply elasticity at the district level.

by an unlimited supply of unproductive workers. The justification for modeling firms in this way comes primarily from the idea that larger firms are subject to more stringent regulations and wage premia and therefore face much higher hiring costs. This line of thinking, however, misses the fact that large firms have the ability and capacity to put significant resources into recruitment and screening of potential workers. Consequently, they have access to both a larger pool and a more complex mechanism by which to screen workers. The second half of this paper argues that small firms in our context rely on the sophisticated but not fully efficient screening mechanism embedded in the widespread practice of charging novice would-be employees an entrance fee to enter employment.

Using the results from a field experiment which randomly gave firms access to worker recruitment services, we show that small firms offered workers grew relative to those not offered workers through the program. In addition, we show that the marginal revenue product of labor (even when that labor is unemployed young people not productively employed elsewhere) is positive and quite large. It appears there is substantial room for small firms to profitably grow.

More work remains to be done to better understand small firms and labor markets in developing countries. This paper presents evidence for one type of labor market friction constraining employment in small firms, but its limitations leave further empirical tests as future work. Understanding whether this existence result is generalizeable to a broader population of firms is a key area for future research. In addition, the findings in this paper suggest that some labor market institutions for vacancy posting and non-monetary screening are missing. Seeking further policy options to address these market failures is a way forward.

References

Abebe, Girum, Marcel Fafchamps, Michael Koelle, and Simon Quinn. 2019. "Learning Management Through Matching: A Field Experiment Using Mechanism Design." *Working Paper.*

Abebe, Girum, Stefano Caria, and Esteban Ortiz-Ospina. 2017. "The Selection of Talent: Experimental and Structural Evidence from Ethiopia." *Working Paper.*

Abebe, Girum, Stefano Caria, Marcel Fafchamps, Paolo Falco, Simon Franklin, and Simon Quinn. 2018. "Anonymity or Distance? Job Search and Labour Market Exclusion in a Growing African City." *Working paper.*

Abel, Martin, Rulof Burger, and Patrizio Piraino. forthcoming. "The Value of Reference Letters: Experimental Evidence from South Africa." *American Economic Journal: Applied Economics.*

Acemoglu, Daron, and Jörn-Steffen Pischke. 1998. "Why Do Firms Train? Theory and Evidence." *The Quarterly Journal of Economics*, 113(1): 79–119.

Alfonsi, Livia, Oriana Bandiera, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, Anna Vitali, et al. 2019. "Tackling Youth Unemployment: Evidence from a Labor Market Experiment in Uganda." *Working paper.*

Ashraf, Nava, Jim Berry, and Jesse Shapiro. 2010. "Can Higher Prices Stimulate Product Use? Evidence from a Field Experiment in Zambia." *American Economic Review.*

Azar, Jose, Steven Berry, and Ioana Elena Marinescu. 2019. "Estimating Labor Market Power." *Working Paper.*

Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan. 2015. "The Miracle of Microfinance? Evidence from a Randomized Evaluation." *American Economic Journal: Applied Economics*, 7(1): 22–53.

Bas, Daniel. 1989. "On-the-job Training in Africa." *International Labor Review*, 128: 485–496.

Bassi, Vittorio, and Aisha Nansamba. 2019. "Screening and Signaling Non-Cognitive Skills: Experimental Evidence from Uganda." *Working Paper.*

Beam, Emily A. 2016. "Do Job Fairs Matter? Experimental Evidence on the Impact of Job-Fair Attendance." *Journal of Development Economics*, 120: 32–40.

Becker, Gary S, and George J Stigler. 1974. "Law Enforcement, Malfeasance, and Compensation of Enforcers." *The Journal of Legal Studies*, 3(1): 1–18.

Berkouwer, Susanna. 2020. "Credit, attention, and externalities in the adoption of energy efficient technologies by low-income households." *Working Paper.*

Birks, Stace, Fred Fluitman, Xavier Oudin, and Clive Sinclair. 1994. "Skills Acquisition in Microenterprises: Evidence from West Africa." *OECD*, Paris.

Boehm, Ullrich. 1997. "Human Resource Development in African Small and Microenterprises: The Role of Apprenticeship." *In: Bass, H.H., et al. African Development Perspectives Yearbook 1996. Regional Perspectives on Labor and Employment.*

Brooks, Wyatt, Kevin Donovan, and Terence R Johnson. 2018. "Mentors or Teachers? Microenterprise Training in Kenya." *American Economic Journal: Applied Economics*, 10(4): 196–221.

Callaway, Archibald. 1964. "Nigeria's Indigenous Education: The Apprenticeship System." *University of Ife, Journal of African Studies*, 62–79.

Carranza, Eliana, Robert Garlick, Kate Orkin, and Neil Rankin. 2019. "Job Search and Hiring with Two-Sided Limited Information About Workseekers' Skills." *Working Paper.*

Cohen, Jessica, and Pascaline Dupas. 2010. "Free Distribution or Cost-Sharing? Evidence from a randomized malaria experiment." *Quarterly Journal of Economics*, 125: 1–45.

Crepon, Bruno, and Patrick Premand. 2019. "Direct and Indirect Effects of Subsidized Dual Apprenticeships." *Working Paper.*

Dal Bo, Ernesto, Frederico Finan, and Martin A. Rossi. 2013. "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service." *Quarterly Journal of Economics.*

De Mel, Suresh, David McKenzie, and Christopher Woodruff. 2008. "Returns to Capital in Microenterprises: Evidence from a Field Experiment." *Quarterly Journal of Economics*, 123(4): 1329–1372.

De Mel, Suresh, David McKenzie, and Christopher Woodruff. 2009. "Measuring Microenterprise Profits: Must We Ask How The Sausage Is Made?" *Journal of Development Economics*, 88, pages 19-31.

De Mel, Suresh, David McKenzie, and Christopher Woodruff. 2019. "Labor Drops: Experimental Evidence on the Return to Additional Labor in Microenterprises." *American Economic Journal: Applied Economics*, 11(1): 202–235.

Dickens, William T, Lawrence F Katz, Kevin Lang, and Lawrence H Summers. 1989. "Employee Crime and The Monitoring Puzzle." *Journal of labor economics*, 7(3): 331–347.

Drexler, Alejandro, Greg Fischer, and Antoinette Schoar. 2014. "Keeping it Simple: Financial Literacy and Rules of Thumb." *American Economic Journal: Applied Economics*, 6(2): 1–31.

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddarth Suri. 2020. "Monopsony in Online Labor Markets." *American Economic Review: Insights.*

Dupas, Pascaline, and Edward Miguel. 2017. "Impacts and Determinants of Health Levels in Low-Income Countries." *Handbook of Field Experiments.*

Dupas, Pascaline, Vivian Hoffman, Michael Kremer, and Alix Peterson Zwane. 2016. "Targeting health subsidies through a non-price mechanism: A Randomized Controlled Trial in Kenya." *Science*, 889–895.

Franklin, Simon. 2018. "Location, Search Costs and Youth Unemployment: Experimental Evidence from Transport Subsidies." *The Economic Journal*, 128(614): 2353–2379.

Frazer, Garth. 2006. "Learning the Master's Trade: Apprenticeship and Human Capital in Ghana." *Journal of Development Economics*, 81: 259–298.

Garcia, Marito, and Jean Fares. 2008. "Youth in Africa's Labor Markets." *World Bank Publications*, The World Bank: number 6578.

Groh, Matthew, David McKenzie, Nour Shammout, and Tara Vishwanath. 2015. "Testing The Importance of Search Frictions and Matching Through a Randomized Experiment in Jordan." *IZA Journal of Labor Economics*, 4(1): 7.

Hardy, Morgan, and Jamie McCasland. 2014. "Apprentice Labor and Firm Output: Evidence from a random match experiment." *AEA RCT Registry. March 7.*, https://doi.org/10.1257/rct.297-1.1.

Hardy, Morgan, and Jamie McCasland. 2022. "Data and Code for: "Are Small Firms Labor Constrained? Experimental Evidence from Ghana"." *Inter-University Consortium for Political and Social Research, Ann Arbor, MI. openicpsr-155121*, http://doi.org/10.3886/E155121V1.

Hardy, Morgan, Isaac Mbiti, Jamie McCasland, and Isabelle Salcher. 2019. "The Apprenticeship-to-Work Transition." *World Bank Policy Research Working Paper 8851.*

Hsieh, Chang-Tai, and Benjamin Olken. 2014. "The Missing "Missing Middle"." *Journal of Economic Perspectives*, American Economic Association: vol. 28(3), pages 89–109, Summer.

Iacovone, Leonardo, Vijaya Ramachandran, and Martin Schmidt. 2014. "Stunted Growth: Why Don't African Firms Create More Jobs?" *Working Paper 353*, Center for Global Development.

ILO. 2013. "Global Employment Trends for Youth." *Working Paper, International Labor Organization.*

Karlan, Dean, Ryan Knight, and Christopher Udry. 2015. "Consulting and Capital Experiments with Microenterprise Tailors in Ghana." *Journal of Economic Behavior & Organization*, 118: 281–302.

King, Kenneth. 1977. "The African Artisan: Education and the Informal Sector in Kenya." Heinemann, London.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar. 2018. "Who Profits from Patents? Rent-Sharing at Innovative Firms." *NBER Working paper.*

Kremer, Michael, and Edward Miguel. 2007. "The Illusion of Sustainability." *Quarterly Journal of Economics*, 1007–1065.

Ma, Xiaochen, Sean Sylvia, Matthew Boswell, and Scott Rozelle. 2016. "Ordeal Mechanisms and Training in the Provision of Subsidized Products in Developing Countries." *Working Paper.*

Mazumdar, Dipak, and Ata Mazaheri. 2003. "The African Manufacturing Firm." Routledge, London.

Meredith, Jennifer, Jonathan Robinson, Sarah Walker, and Bruce Wydick. 2013. "Keeping the doctor away: Experimental evidence on investment in preventative health products." *Journal of Development Economics*, 196–210.

Ozier, Owen. 2018. "Exploiting Externalities to Estimate The Long-term Effects of Early Childhood Deworming." *American Economic Journal: Applied Economics*, 10(3): 235–62.

Pallais, Amanda. 2014. "Inefficient Hiring in Entry-Level Labor Markets." *American Economic Review*, 104(11): 3565–99.

Sandefur, Justin. 2010. "On the Evolution of the Firm Size Distribution in an African Economy." *CSAE Working Paper Series 2010-05*, Centre for the Study of African Economies, University of Oxford.

Table 1: **Attrition**

	(1) =1 if Surveyed in Round 1	(2) =1 if Surveyed in Round 2	(3) =1 if Surveyed in Round 1 or Round 2	(4) =1 if Profits Non-Missing in Round 1	(5) =1 if Profits Non-Missing in Round 2	(6) =1 if Profits Non-Missing in Round 1 or Round 2
Panel A: No Controls						
Treatment Apprentices	0.02	0.01	0.00	0.01	0.02	0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Panel B: With Controls						
Treatment Apprentices	0.01	0.01	0.00	-0.00	0.01	0.00
	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)	(0.01)
Observations	669	755	755	669	755	755
Mean of Dep Variable T=0	0.92	0.88	0.97	0.88	0.84	0.95

Notes: Regressions in Panel A exclude all controls and show raw differences in participation in surveys by treatment, where treatment is defined as number of apprentices assigned to the firm. Regressions in Panel B include district and trade fixed effects and dummies for all probability distributions over treatment intensity, mirroring our preferred specification throughout. The mean of the dependent variable is reported for firms assigned zero apprentices. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2: **Covariate Balance**

	(1) Has Any Worker(s)	(2) Total Workforce	(3) Apprentices	(4) Paid Workers	(5) Owner Hrs/Week
Treatment Apprentices	0.00	0.09	-0.05	0.08	0.04
	(0.02)	(0.16)	(0.14)	(0.06)	(1.07)
Observations	755	755	755	755	747
R^2	0.19	0.30	0.26	0.44	0.31
Mean of Dep Variable T=0	0.86	2.98	2.31	0.52	57.65

	(6) Revenues (GHC)	(7) IHS Revenes	(8) Profits (GHC)	(9) IHS Profits	(10) Capital (GHC)
Treatment Apprentices	68.34	0.06	26.39	0.09	-376.63
	(50.41)	(0.07)	(27.68)	(0.07)	(370.27)
Observations	742	742	743	743	744
R^2	0.43	0.35	0.33	0.31	0.34
Mean of Dep Variable T=0	613.07	6.46	295.80	5.66	4325.97

	(11) IHS Capital	(12) Bank Account	(13) Reg w/ Reg General	(14) Reg w/ Dist Assembly	(15) Firm Age
Num Treatment Apprentices	-0.06	0.00	-0.02	0.03	0.23
	(0.06)	(0.03)	(0.02)	(0.03)	(0.42)
Observations	744	747	747	747	747
R^2	0.41	0.22	0.29	0.28	0.29
Mean of Dep Variable T=0	8.53	0.65	0.09	0.34	11.72

	(16) Management Skills (of 5)	(17) Female Owner	(18) Years Schooling	(19) Math (of 4)	(20) Digit Span (of 14)
Num Treatment Apprentices	0.03	0.01	0.19	0.05	0.04
	(0.09)	(0.02)	(0.19)	(0.05)	(0.15)
Observations	742	755	747	745	748
R^2	0.33	0.76	0.23	0.17	0.23
Mean of Dep Variable T=0	2.59	0.69	9.10	2.62	6.94

Notes: Each coefficient is from a separate regression of baseline firm-level covariates on treatment assignment, district and trade fixed effects, and dummies for all probability distributions over treatment intensity, mirroring our preferred specification throughout. All 755 firm owners were registered for a matching meeting and provided information on their gender and employees at registration. 748 of 755 firm owners completed a baseline survey. *Paid Workers* is a Ghanaian colloquialism for workers who have already completed an apprenticeship, though both apprentices and paid workers receive wages. Profits are self-reports of all sales less all expenses (including the wage bill) in the reported month. Profits, sales, and capital stock are in April 2013 Ghana Cedis, when 1 US dollar was equivalent to 1.95 Ghana Cedis. The top 0.5% of profit, sales, and capital stock observations have been winsorized in level specifications. An F test of the joint significance of all 20 covariates in predicting treatment yields of p-value of 0.64. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: **Treatment Effects on Labor and Capital Inputs**

	Take Up		Other Labor Inputs			Capital	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Program Apprentices	Total Workforce	Other Apprentices	Paid Workers	Firm Owner Hours/Week	Capital Stock (GHC)	IHS Capital Stock
Treatment Apprentices	0.47***	0.58***	0.15	0.07	-0.23	171.95	0.05
	(0.04)	(0.13)	(0.11)	(0.06)	(0.80)	(242.01)	(0.04)
Observations	1315	1315	1315	1315	1312	674	674
Mean of Dep Variable T=0	0.06	3.18	2.42	0.55	53.47	2712.45	8.10

Notes: Regressions include round fixed effects, district and trade fixed effects, and dummies for each probability distribution. Program apprentices are apprentices placed with these firms by the experimental program. Total workforce includes program apprentices and all other non-owner labor. *Paid Workers* is a Ghanaian colloquialism for workers who have already completed an apprenticeship, though both apprentices and paid workers receive wages. Capital Stock was only collected in the second follow-up survey, excludes land and buildings, and is in April 2013 Ghana Cedis, when 1 US dollar was equivalent to 1.95 Ghana Cedis. The top 0.5% of capital stock observations have been winsorized in the level specification. Columns (2) through (7) include baseline values of the dependent variable. Standard errors in parentheses are clustered at the district level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: **Treatment Effects on Revenues and Profits**

	(1) Profits (GHC)	(2) IHS Profits	(3) Revenues (GHC)	(4) IHS Revenues	(5) Program Apprentice Wages (GHC)
Panel A					
	Primary Specification				
Treatment Apprentices	40.54***	0.11***	50.92	0.09**	11.82***
	(12.00)	(0.04)	(30.06)	(0.04)	(2.41)
Panel B					
	With Additional Baseline Controls				
Treatment Apprentices	41.27***	0.12***	53.80*	0.08**	11.82***
	(12.06)	(0.04)	(31.43)	(0.03)	(2.34)
Observations	1257	1257	1257	1257	1257
Mean of Dep Variable T=0	401.08	6.12	736.24	6.68	1.13

Notes: Regressions include round fixed effects, district and trade fixed effects, and dummies for each probability distribution. In both panels A and B, Columns (1) through (4) include baseline values of the dependent variable. In Panel B, we use a LASSO estimator and post-double selection to choose additional baseline controls from those presented in Table 2. The procedure selects the baseline number of apprentices at the firm for each of the first four outcome variables, and additionally firm owner years of schooling for the outcome variable in Column (4). It selects no additional controls in Column (5). Profits are self-reports of all sales less all expenses (including the wage bill) in the reported month. Profits, sales, and wages are in April 2013 Ghana Cedis, when 1 US dollar was equivalent to 1.95 Ghana Cedis. The top 0.5% of profit and sales observations have been winsorized in level specifications. Program Apprentice Wages are all wages paid to apprentices placed with these firms by the experimental program. Standard errors in parentheses are clustered at the district level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: **Robustness to Specification Choice**

Panel A

	Dependent Variable: Profits (GHC)				
	(1)	(2)	(3)	(4)	(5)
	Preferred Specification	Without Winsorizing	Alternative Controls		Quantile (Median)
Treatment Apprentices	40.54***	32.71	32.38**	44.17***	26.24**
	(12.00)	(20.48)	(13.53)	(14.63)	(13.23)
Winsorizing top 0.5% of profits	YES	NO	YES	YES	YES
Including District and Trade FEs	YES	YES	NO	YES	YES
Including Baseline Value of Dep Variable	YES	YES	YES	NO	YES
Observations	1257	1257	1257	1257	1257
Mean of Dep Variable	401.08	401.08	401.08	401.08	401.08

Panel B

	Dependent Variable: Inverse Hyperbolic Sine Profits			
	(1)	(2)	(3)	(4)
	Preferred Specification	Alternative Controls		Quantile (Median)
Treatment Apprentices	0.11***	0.10**	0.13***	0.11**
	(0.04)	(0.04)	(0.04)	(0.05)
Including District and Trade FEs	YES	NO	YES	YES
Including Baseline Value of Dep Variable	YES	YES	NO	YES
Observations	1257	1257	1257	1257
Mean of Dep Variable	6.12	6.12	6.12	6.12

Notes: Preferred specifications in Panels A and B replicate those presented in Table 4, including round fixed effects, district and trade fixed effects, dummies for each probability distribution, and baseline values of the dependent variable. Profits are self-reports of all sales less all expenses (including the wage bill) in the reported month, reported in April 2013 Ghana Cedis, when 1 US dollar was equivalent to 1.95 Ghana Cedis. In Panel A Column (1), the top 0.5% of profit observations have been winsorized. The remaining columns alter one specification choice at a time. Column (2) removes winsorizing, using raw self-reported profits, a specification not relevant to inverse hyperbolic sine transformation as the inverse hyperbolic sine transformation already uses non-winsorized profits. Column (3) removes district and trade dummies. Column (4) excludes the baseline value of the dependent variable. The final column in both panels displays quantile regression estimates at the median. In Panel A, the quantile estimates use non-winsorized raw profits. Standard errors in parentheses are clustered at the district level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: **Sample Sensitivity**

	Excluding Firms with Any Degenerate Lotteries		Splitting Sample by Probability Distribution Type			
	(1)	(2)	(3)	(4)	(5)	(6)
	Profits (GHC)	IHS Profits	Profits (GHC)	IHS Profits	Profits (GHC)	IHS Profits
Treatment Apprentices	35.53**	0.11**				
	(13.31)	(0.04)				
Treatment Apps X Listed by One Apprentice			19.86	0.20**		
			(41.98)	(0.09)		
Treatment Apps X Listed by More than One Apprentice			47.48***	0.09**		
			(14.70)	(0.04)		
Treatment Apps X Six Most Common Probability Distributions					38.77	0.16**
					(28.61)	(0.06)
Treatment Apps X Less Common Probability Distributions					42.03**	0.07
					(18.54)	(0.05)
F test of equality of treatment effects			0.58	0.25	0.94	0.27
Observations	1120	1120	1257	1257	1257	1257
Mean of Dep Variable	401.08	6.12	401.08	6.12	401.08	6.12

Notes: Regressions include round fixed effects, district and trade fixed effects, dummies for each probability distribution, and baseline values of the dependent variable. Profits are self-reports of all sales less all expenses (including the wage bill) in the reported month. Profits are in April 2013 Ghana Cedis, when 1 US dollar was equivalent to 1.95 Ghana Cedis. The top 0.5% of profit observations have been winsorized. Columns (1) and (2) exclude firms with probability distributions that contain any degenerate lotteries, where a degenerate lottery is defined as an apprentice listing only a single firm. This excludes 82 firms but does little to alter the point estimates. Columns (3) through (6) split the sample to test for sensitivity of the estimates to peculiarities in the probability distributions. Columns (3) and (4) consider firms listed by a single apprentice, a subset of 335 firms that more closely resembles a typical experimental setup with a binary treatment indicator. 420 firms were listed by more than one apprentice, and thus had some positive probability of being assigned more than one apprentice. Columns (5) and (6) split the sample in half: 370 firms, or 49% of the sample, faced one of the six most common probability distributions. Point estimates are statistically indistinguishable from each other and qualitatively of similar magnitude across subgroups of the sample. Standard errors in parentheses are clustered at the district level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: **Predictors of Program Applicant Selection**

	(1) Attended MM	(2) Ranked Any Firm	(3) Ranked >1
Cognitive Index z-score	0.03**	0.02*	0.04***
	(0.01)	(0.01)	(0.01)
Non-Cognitive Index z-score	0.02	0.02*	-0.01
	(0.01)	(0.01)	(0.01)
Observations	2288	2288	2288
Mean of Dep Variable	0.56	0.51	0.32

Notes: *Attended MM* takes a value of one for any apprentice who attended a matching meeting. *Ranked Any Firm* takes a value of one for those who completed a list of firms with which they were willing and able to work. *Ranked >1* takes a value of one for those apprentices who listed more than one firm and thus are the source of random variation in apprentices matched with firms in our sample. Columns (1) through (3) include all applicants who were invited to attend a relevant matching meeting. Of 2,360 treatment applicants, 2,288 completed an apprentice baseline survey.

Table 8: **Firm Sample Selection**

	(1) Matching Meeting Pooled	(2) Rest of Census	(3) Diff	(4) Census Located on Main Roads	(5) Diff	(6) Main Roads Trade Assoc Member	(7) Diff
Total Workers	2.54	0.60	1.95***	0.87	1.67***	1.68	0.86**
Non-program Apprentices	2.31	0.45	1.87***	0.62	1.70***	1.19	1.13***
Paid Workers	0.20	0.04	0.16***	0.06	0.14***	0.15	0.05
Revenues (GHC)	272.14	118.97	153.17***	138.80	133.34***	218.85	53.29
Profits (GHC)	185.57	81.85	103.72***	97.32	88.25***	157.85	27.72
Capital Stock (GHC)	1476.97	624.39	852.58***	838.86	638.11***	1383.45	93.52
Firm Owner Female	0.86	0.78	0.08	0.77	0.08	0.77	0.08
Firm Owner Years Schooling	9.09	8.94	0.15	8.85	0.23	9.51	-0.42
Management Skills (of 5)	2.54	1.94	0.60***	2.35	0.19	2.75	-0.21
IHS Adjusted Profit Per Worker	2.29	-0.41	2.70***	0.27	2.01***	2.27	0.01
Observations	35	991		501		75	

Notes: This table compares the census of all garment making firm owners in Hohoe District (collected for another project) to the garment making firm owner sample for this study from this district. The program sample for this study was recruited in areas with concentrations of applicants to the apprenticeship program, which in this district were concentrated along the main roads, including Hohoe Town (the district capital), Santrokofi, and Gbi on the N2 highway and Likpe Bata, Likpe Mate, and and Likpe Bala on the highway that runs to Badou, Togo. The primary mechanism by which government officials recruited firms was through craft-specific trade associations. We therefore present samples of the full district census (including quite rural areas), those firms along the main roads, and those firms along the main roads whose firm owners were also members of a trade association at the time of the survey. Profits are self-reports of all sales less all expenses (including the wage bill) in the reported month. Profits, sales, and capital stock are in April 2013 Ghana Cedis, when 1 US dollar was equivalent to 1.95 Ghana Cedis. The top 0.5% of profit, sales, and capital stock observations have been winsorized. *Paid Workers* is a Ghanaian colloquialism for workers who have already completed an apprenticeship, though both apprentices and paid workers receive wages. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: **Firm Sample Selection - Matching Meetings**

	=1 if Firm in Final Sample	
	(1)	(2)
Baseline Total Workforce	0.01	0.02
	(0.00)	(0.02)
Baseline Apprentices	0.01	-0.02
	(0.00)	(0.02)
Baseline Paid workers	0.00	-0.03
	(0.01)	(0.03)
Baseline Revenues (1000 GHC)	0.02	0.01
	(0.02)	(0.03)
Baseline Profits (1000 GHC)	0.03	-0.00
	(0.03)	(0.05)
Baseline Assets (1000 GHC)	0.01**	0.00*
	(0.00)	(0.00)
Firm Owner Female	-0.03	-0.04
	(0.04)	(0.05)
Firm Owner Years Schooling	0.01	0.00
	(0.00)	(0.00)
Firm Owner Ability/Skill Index (z-score)	0.03***	0.03**
	(0.01)	(0.01)
IHS Adjusted Avg Product Labor	0.00	-0.00
	(0.00)	(0.00)
Observations	1832	1832

Notes: Each coefficient estimate in Column (1) comes from a separate regression. Column (2) coefficients come from a single regression that includes all covariates. Regressions include district by trade fixed effects. Firm owner ability index is a normalized sum of the normalized scores on a Digits Forward test, a four question math test, and five managerial skills questions. *IHS Adjusted Avg Product Labor* takes value added (profits plus the wagebill) and subtracts the mean profits for a single person firm to adjust for the fact that firm owners are more productive than workers. It then divides by the number of workers (inclusive of the owner) and takes the inverse hyperbolic sine of this adjusted value added per worker measure. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: **Firm Heterogeneity**

	Program Apprentices			Profits (GHC)		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment Apprentices	0.41***	0.52***	0.41***	87.80**	55.70**	46.35**
	(0.07)	(0.05)	(0.06)	(39.30)	(24.90)	(22.15)
Baseline Total Workforce	-0.01			57.91***		
	(0.01)			(20.60)		
Baseline Total Workforce X Treatment Apps	0.02			-14.14		
	(0.01)			(11.13)		
Baseline Assets Above Median		-0.01			74.77	
		(0.04)			(44.88)	
Baseline Assets Above Median X Treatment Apps		-0.12***			-41.49	
		(0.04)			(47.20)	
Firm Owner Ability Above Median			0.03			89.74*
			(0.06)			(46.49)
Firm Owner Ability Above Median X Treatment Apps			0.14			-32.90
			(0.09)			(46.60)
Observations	1315	1298	1290	1257	1244	1236

Notes: Regressions include round fixed effects, district and trade fixed effects, and dummies for each probability distribution. Columns (4) through (6) include baseline values of the dependent variable. Profits are self-reports of all sales less all expenses (including the wage bill) in the reported month in April 2013 Ghana Cedis, when 1 US dollar was equivalent to 1.95 Ghana Cedis. The top 0.5% of profit observations have been winsorized. Baseline Total Workforce is the number of employees (including paid workers and apprentices) at the firm at baseline excluding the owner; 12% of firms have zero workers at baseline. Standard errors in parentheses are clustered at the district level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$