Baylor University

September 15, 2026

Emily Breza
Harvard University

Davide Viviano
Harvard University

Arun Chandrasekhar
Stanford University

Dear Professors Breza, Viviano, and Chandrasekhar,

I am writing to apply to the NBER Workshop on Development Economics and Networks for Early-Career Researchers, to be held December 5, 2026, in Cambridge, Massachusetts.

I am an Assistant Professor of Economics at Baylor University. I completed my Ph.D. in Agricultural and Resource Economics at the University of California, Davis in August 2024, and am currently within my second year as faculty.

I am submitting my working paper, "Beliefs and the Demand for Employee Training: A Field Experiment with Small Firms in Uganda." (with Andy Brownback, Sarojini Hirshleifer, and Arman Rezaee), which draws on a field experiment with small metal fabrication firms in Kampala, Uganda to study how employers and workers select workers for training, and the role that firm attachments (e.g., family ties and routines) and anticipated skill gains play in that selection process.

My recommender for this application is:

Prof. Travis Lybbert
Department of Agricultural and Resource Economics
University of California, Davis
tlybbert@ucdavis.edu

Please find enclosed my CV, a statement of motivation and research interests, and my working paper. I would welcome the opportunity to present and discuss this work at the workshop.

Contact information:
Isaac Ahimbisibwe
One Bear Place #98003
Waco, TX 76798
Isaac_Ahimbisibwe@baylor.edu
+1 254-710-4162

Sincerely,
Isaac Ahimbisibwe

Isaac Ahimbisibwe

iahimbis.github.io
Isaac_Ahimbisibwe@baylor.edu

One Bear Place #98003
Waco, TX 76798
Office +1254-710-4162

Employment

Assistant Prof. Economics Baylor University, Waco Texas — October 2024-present

Education

Ph.D. in Applied Economics, University of California, Davis — August 2024

MSc. in Applied Economics, Illinois State University — May 2017

Bachelor of Statistics (*Honors*), Makerere University — Jan 2014

References

Prof. Travis Lybbert
Dept. of Agricultural and
Resource Economics
University of California, Davis
tlybbert@ucdavis.edu

Prof. Scott Carrell
Prof. of Economics,
University of Texas, Austin
Prof. Emeritus, Economics
University of California, Davis
secarrell@ucdavis.edu

Prof. Stephen Boucher
Dept. of Agricultural and
Resource Economics
University of California, Davis
srboucher@ucdavis.edu

Research Interests

Development Economics, Economics of Education, Labor Economics, Gender

Research

"The Leaky College Pipeline: Evidence from Administrative Data in Uganda." AEA Papers and Proceedings 116: 206–211. (*with A. Dautheville, L. Fiala, S. Omala, W. Sandholtz)*

"Pipeline vs. choice: the global gender gap in STEM applications" (*with A. Altjmed,G. Artemov, A. Barrios-Fernández, A. Bizopoulou, M. Kaila, J. Liu, R. Megalokonomou, J. Montalbán, C. Neilson, S. Otero, J. Sun, X. Ye*)

"Who Matters in College? Familiarity, Ability Signals, and Student Achievement" (*with Scott E. Carrell, Travis J. Lybbert*)

"Is affirmative Action the Answer? Shrinking the STEM Gender Gap in Higher Education?" (*with S. Chopra*)

"Beliefs and the Demand for Employee Training: A Field Experiment with Small Firms in Uganda." (*with A. Brownback A, S. Hirshleifer, A. Rezaee*)

Work In Progress

Merit vs. Place-Based Quotas: Scholarship Assignment Rules and University Outcomes (*A. Douthaville, L. Fiala , K. Omala, W. Sandholtz*)

Gendered Preferences, Affirmative Action, and Choice of university Major: Evidence from Uganda (*with S. Chopra*)

Thou Must be a Doctor or an Engineer. Unintended Consequences of a Compulsory Science Policy. (*with T. Lallemant, H. J. Opio*)

Returns to College in the Developing World. (with *F. Lenka, O. Kizito, W. Sandholtz*)

Pre-PhD Research

"The contribution of Millennium Development Goals towards improvement in major development indicators, 1990–2015." (2019) *Applied Economics* 44(2), 337-58. (with Rati Ram)

Awards and Grants

2024: Large Research Grant, International Growth Center (PI)

2024: Policy Impacts Early Career Scholars Grant (Co-PI)

2023: Jastro & Shields Graduate Student Research Grant, UC Davis

2023: Travel Award, UC Davis

2021: Jastro & Shields Graduate Student Research Grant, UC Davis

2018: Graduate Education Tution Waiver, UC Davis

2017: Scott Elliott Graduate Academic Excellence Scholarship, ISU

2016: Graduate Education Tuition Waiver, Illinois State University

2010: Government Merit Undergraduate Academic Scholarship

Teaching Experience

Professor, Baylor University

Fall 2024: (Undergraduate) Issues in Economics for Non-Business Majors

Spring 2025: (Undergraduate) Issues in Economics for Non-Business Majors

Fall 2025: (Undergraduate) Issues in Economics for Non-Business Majors

Instructor, University of California, Davis

Summer 2023: (Undergraduate) Applied Summer Research

Teaching Assistant, University of California, Davis

Spring 2019: (Graduate) Demand and Market Analysis

Undergraduate:

- 2020-2021: Microeconomic development Economics
- 2019-2020: Intermediate Microeconomics
- Spring 2023: Managerial Marketing

Invited presentations and conferences

Winter 2026: ASSA Paper Session "Higher Education in Low-Income Countries Paper Session", Philadelphia

Fall 2025: NTxEC Conference, The University of Texas, Arlington

Summer 2025: The Collaborative for Econometrics and Integrated Development Studies Workshop, The University of Notre Damme
Texas Development Workshop, the University of Texas Austin

Winter 2025: The Population Association of America, DC
Midwest Economics Association, Kansas City

Winter 2024: The CSAE Conference on Economic Development, University of Oxford

Fall 2023: The Development Economics Seminar, University of California, Davis
EconNect Africa PhD candidate Online conference
The DIME-KDI School 4th Development Impact conference, The World Bank

Winter 2023: Gianni Agriculture and Resource Economics Student Conference, UC Davis

Summer 2022: Working Group on African Political Economy, CEGA, Kigali

MENTORING AND SERVICE

The Diversity and Inclusion in Research, Education, and Career Training program, UC Davis

- Undergraduate Student cafe coordinator 2021-2022
- Graduate Inclusive Education and Mentoring Training workshop summer 2023

The Graduate Applications International Network - Mentor 2021, 2023

The Bwoga Elementary Library Initiative - Codirector 2022-Present

Other

Languages: English (First language); French (Basic); Runyakitara (Native)

Software: R (intermediate); Stata (advanced); Python (basic) LaTeX(advanced)

Field Work: Uganda

Statement of Motivation and Research Interests

My research sits at the intersection of firm productivity, education access, and gender in developing economies, primarily using data from Uganda and administrative datasets across Sub-Saharan Africa. In "The Leaky College Pipeline" (AEA P&P, 2025) and "Pipeline vs. Choice," I study where students drop out of the path to higher education and why gender gaps in STEM persist despite comparable early ability. In "Who Matters in College?" I look at how familiarity and institutional signals of ability shape student outcomes.

My current working paper, "Beliefs and the Demand for Employee Training: A Field Experiment with Small Firms in Uganda" (with A. Brownback, S. Hirshleifer, and A. Rezaee), focuses on firm-level decisions: a field experiment with 337 small metal fabrication firms, each with at least 4 workers, in Kampala. The paper studies how employers select and workers self-select which workers receive training. Employers weight firm attachment, such as family ties and trust, heavily in these decisions and largely ignore anticipated skill gains, even though their information about skill gains is accurate, while workers do the opposite. Our paper builds on Acemoglu and Pischke (1999) and microfounds firm attachment, showing that it predicts wage compression that makes training in transferable skills profitable while distorting who is selected for it.

I am applying to this workshop to get feedback on this paper's dimensions of social ties, and to learn methods for incorporating network structure more rigorously into my future work on firms and education access in low-income settings. I would also welcome the chance to connect with other early-career researchers working on related questions in African labor markets.

Beliefs and the Demand for Employee Training: A Field Experiment with Small Firms in Uganda*

Isaac Ahimbisibwe[†], Andy Brownback[‡], Sarojini Hirshleifer[§], and Arman Rezaee[¶]

September 15, 2026

Preliminary draft: Please do not cite or circulate

Abstract

Classic models of workforce training predict that the inability of employers to recoup training costs will result in an inefficiently low quantity of training. Labor market frictions can play an important role in enabling a "second best" provision of training. While these frictions may overcome quantity distortions, they create new distortions in selection whereby training may be targeted to workers with relatively lower skill gains from training. We develop a labor market model that produces two main predictions: 1) worker demand for training depends on anticipated skill gains and 2) employer demand for training depends on the ability to extract those skill gains. Using experimental evidence from the Ugandan metalworking sector, we confirm these predictions, finding that employers under-select workers with high anticipated skill gains and over-select workers with strong attachments to the firm while workers self-select narrowly based on skill gains. These misaligned selections do not stem from asymmetric information: we find employer beliefs about skill gains from training are predictive of realized gains.

KEYWORDS: Worker Training, Training Selection, Employer–Worker Conflict, Labor Market Frictions, Developing-country Labor Markets

JEL Classifications: J24, O12, J53, C93

*We gratefully acknowledge that funding for this project was provided by both exploratory and major grants from the Private Enterprise Development in Low-income Countries (PEDL) and exploratory grants from the EASST Initiative at Centre for Effective Global Action (CEGA), and the UK Foreign, Commonwealth & Development Office, awarded through the J-PAL Post-Primary Education Initiative.

[†]Baylor University
[‡]University of Arkansas
[§]University of California, Riverside and CEGA
[¶]University of California, Davis and CEGA

1 Introduction

Firms in developing countries tend to be characterized by low productivity and low levels of human capital (Hsieh and Klenow, 2009; World Bank, 2012). In such settings, worker training plays a special role as both cause and consequence of low productivity and, therefore, serves as a frequent public policy intervention for improvement (McKenzie, 2017). However, such interventions only improve social welfare to the extent that they compensate for underinvestment in training because of market failures. A large theoretical and empirical literature identifies the conditions under which such market failures may arise and inhibit training (Becker, 1962; Acemoglu, 1997; Stern and Ritzen, 2012; Alfonsi et al., 2020; Caicedo, Espinosa and Seibold, 2022).

Classically, Becker (1962), characterizes the incentive compatibility problem for employers to provide general training to their workers. Employers will be unable to recoup the cost of general training by extracting worker skill gains, because those gains follow the worker if they depart for a competitor. Either workers would be required to pay for a fraction of their training, or employers would provide only specific training.

In a series of papers that build on this framework, Acemoglu and Pischke provide a theoretical and empirical basis that accounts for the prevalence of training observed in real-world firms (Acemoglu and Pischke, 1998, 1999*a*,*b*). Specifically, these papers highlight the counterintuitive role played by labor market frictions in facilitating a "second-best" amount of training provision. These frictions produce a distortionary wage compression in the labor market that yields a form of monopsony power for incumbent employers. The ability to extract some share of a worker's marginal revenue product incentivizes the employer to improve the productivity of that worker through training, even general training.

In our paper, we provide further evidence on the primacy of labor market frictions in determining the provision of training. Our model enriches the Acemoglu and Pischke framework to consider a setting where workers have heterogeneous skill gains from training, but the opportunities for training are limited. Under these conditions, the selection of workers for training is critical in determining the overall effectiveness of training. Where Acemoglu and Pischke identify a positive externality of labor market frictions—they make possible employer-sponsored general training—our model shows how these frictions may yet produce a negative externality on the effectiveness of that training. Specifically, when a worker's skill gains from training negatively covary with the labor market frictions the worker faces, employers will forego skill gains from training in order to select workers for training based on how much of the gains from training the employer can extract. This overselection of appropriable workers erodes the welfare gains from Acemoglu and Pischke's

second-best solution.

Our model generates testable hypotheses about which workers employers will select for training and the distortionary effect of their selections on training effectiveness. We predict that employers will make training selections based on their ability to extract the worker's gains from training. This appropriability arises from the worker's attachment to the firm.[1] In our setting, we find that firm attachments are negatively correlated with the worker's gains from training, meaning that employers negatively select based on skill gains.

Our model makes an additional prediction that, all else equal, the diminishing marginal utility of wealth gives employers greater leverage over workers with higher wages. Thus, employers will forego the high-return training of low-skilled low-wage workers in favor of training the more appropriable, high-wage workers. This regressive selection of workers for training produces a distinct negative externality on social welfare and exacerbates within-firm wage inequalities.

We contrast the employer's selection of workers for training with the self-selection decisions of the workers themselves. The workers' objectives more closely resemble the objectives of a social planner with respect to skill gains. Our model predicts that workers will narrowly focus their training selections on skill gains, ignoring or even negatively selecting on the firm attachments that employers focus on.

We study training decisions in the context of small metalworking firms in Kampala, Uganda. Small firms in developing countries tend to operate in competitive and informal labor markets with few exploitable frictions with which to gain leverage over their workers. Moreover, both employers and workers tend to face severe credit constraints. Thus, both Becker (1962) and Acemoglu and Pischke (1999*b*) would predict limited general training to take place endogenously. Indeed, we find this to be true. Only 4% of workers have any training in our baseline data. This indicates that the frictions that do exist—such as a worker's attachment to the firm through relationships, tenure, and trust—are insufficient to overcome the costs of training provision.

We investigate employer demand for worker training through a field experiment in which we offer a free practical skills training program. By providing this comprehensive, 12-week training for free, we can overcome the limitations on training provision while still assessing the distortionary role that labor market frictions may play in determining the selection of workers for training. Our intent-to-treat analysis finds that the training program has a substantial impact on worker productivity. On

[1] Acemoglu and Pischke (1999*b*) explicitly avoid reliance on firm attachments to produce their main results about employer-provided general training. However, with detailed data on firm attachments, we can demonstrate their centrality to employer decision making in our setting. Moreover, our model uses the role of firm attachments as a microfoundation for the wage compression in the labor market that Acemoglu and Pischke (1999*b*) do rely on.

average, treated workers score 0.484 SD higher on our practical skills tests than their counterparts from the control group ($p < 0.001$).

We give employers the opportunity to select one worker from their firm to receive this free training. We elicit the employer's complete ranking of workers based on their desire for them to be trained. Alongside this ranking, we elicit each employer's beliefs about the skills of each of their workers and the skill gains those workers would experience from training.

We find that employers make training selections in the way our model predicts—selecting workers for training based on the worker's attachment to the firm, and their baseline wages and skills. Critically, by selecting on these worker characteristics rather than skill gains from training, employers diminish the productivity gains from training. We show that employers correctly anticipate worker skill gains from training but willingly sacrifice these skill gains in order to over-select based on worker characteristics that yield greater employer leverage. For every 1 SD in additional skill gains an employer anticipates for a given worker, they *reduce* the worker's ranking by 3.6 percentile points ($p < 0.001$). Our data suggest that these selection decisions are not random, nor are they "mistakes," because employer preferences for training closely align with their stated beliefs about the profitability of training each worker.

With respect to firm attachments, we find that employers over-select their own relatives, workers with high anticipated retention, long-tenured workers, and workers they trust ($p < 0.01$ for each). Similarly, with respect to baseline wages and skills, we find that employers over-select high-wage high-skilled workers at the cost of under-selecting the lower-skilled workers with relatively higher returns to training. For every one SD increase in a worker's monthly salary, the employer's selection ranking for that worker increases by 3.4 percentile points ($p \approx 0.017$). Since these characteristics negatively covary with skill gains from training, these selections drive the negative externality on skill gains from training.

We measure worker demand for training using a reverse auction. Our model makes straightforward predictions about the training preferences of workers—they will narrowly prioritize skill gains, thus negatively self-selecting for training based on firm attachments. We find this to be the case. For every one SD in skill gains the employer anticipates, workers increase their demand for training by 0.054 SD ($p \approx 0.084$). This pattern is even stronger with respect to workers' stated beliefs, with every one SD increase in their own anticipated skill gains increasing demand for training by 0.099 SD ($p \approx 0.054$). We find that workers negatively self-select based on two of the measures of firm attachments on which employers positively select: worker tenure and employer trust ($p < 0.05$ for both).

In sum, many of the characteristics that predict employer selection of a worker for training predict the opposite in worker self-selection. We find that employers' selection carry significant costs in terms of social welfare. Employers exacerbate baseline inequalities in wages, tenure, and family connections while sacrificing potential skill gains in order to over-select workers for training based on characteristics that increase the employer's share of productivity gains. Conversely, we find that workers' self-selctions are more closely aligned with social welfare. Workers narrowly self-select into training based on skill gains and optionality.

Using the results of our experiment, we estimate key parameters of our labor market model to demonstrate the relative importance of wage adjustments in the employer training selections. We find that employers' selections for training are highly sensitive to wage adjustment concerns that arise from worker's accumulating transferable skills.

Our study contributes to several strands of literature. Broadly, we contribute to the literature on the under-provision of training in general skills (Acemoglu and Pischke, 1998; Becker, 1962; Prendergast, 1993; Acemoglu, 1997). Literature from our specific setting—the informal labor markets of developing countries—finds similarly limited training and knowledge transfer (Oya, Schaefer and Chen, 2022; Liaqat and Nugent, 2015; Schaffner, 2001; Alfonsi et al., 2020; Caicedo, Espinosa and Seibold, 2022). In our paper, we explore the role of employer concerns about retaining trained workers, a focus of recent literature (Brown et al., 2024; Frazer, 2006). By evaluating our practical skills training program, we also add to the set of studies documenting the effectiveness of training programs in the developing world. Vocational programs are a common topic in this literature.[2] In a recent example, Alfonsi et al. (2020) compares employment and earnings impacts of firm-provided training and formal vocational training for Uganda's "disadvantaged" youth starting in the labor market. Dautheville (2025) offers a complementary explanation for limited program effectiveness: micro-entrepreneurs may view self-employment as temporary. Our approach expands this literature through its focus on currently employed workers and the distortions that may arise based on the trainee-selection procedure.

A recent paper by Cefala et al. (2025) similarly explores the under-provision of training in informal labor markets, where the classic appropriability problem arises because trained workers can be poached by neighboring employers. The authors show that tying employer incentives to training provision and tying worker payouts to retention at the training firm can overcome this suboptimal training provision. We complement this study in key ways. First, we demonstrate how appropriability continues to distort training effectiveness through selection even when the underprovision is resolved through free training. Relatedly, we consider heterogeneity in both worker skill gains

[2]See, e.g., McKenzie (2017), Card et al. (2011), Cho et al. (2013), and Hirshleifer et al. (2016). These studies typically find modest impacts of the training programs they study.

from training and the labor market frictions that may allow employers to extract those gains. Next, by collecting employer beliefs about each worker's returns from training, we expand the understanding of employers' training-provision decisions. Finally, our measurement of worker demand for training allows us a complete picture of both sides of the labor market and how their training decisions are misaligned.

2 Background and Setting

Micro, Small, and Medium Enterprises (MSMEs) are a strategic part of Uganda's policy for economic development. The National Development Plan II of Uganda's government underscores the role of MSMEs in wealth and job creation and outlines government objectives to provide institutional support for MSME growth (GoU, 2015*a*). According to Uganda's government, MSMEs constitute over 90 percent of the private sector in Uganda and contribute 18 percent of the GDP (Government of Uganda, 2018). The same report shows that manufacturing-related firms such as welding and steel made up 8 percent of all MSMEs in 2011 and employed 15.4 percent of the workforce in 2012. Statistics from UBOS show that metal and steel fabrication grew over 10 percent from 2011 to 2017, reflecting the dynamic nature and potential of the manufacturing industry in Uganda.

Uganda's metalworking sector operates in an informal labor market where employment relationships are rarely formalized through written contracts (La Porta and Shleifer, 2014). Without formal labor market institutions governing wages or hiring, pay is determined through direct negotiation between employer and worker. Given these informal characteristics, hiring is likely to flow predominantly through family and social networks. Apprenticeship arrangements in particular are often negotiated through an apprentice's family, as documented in Ghana (Frazer, 2006).

Despite their economic importance and potential for growth, Ugandan MSMEs face many challenges. GoU (2015*b*) summarises Uganda's MSME policy and lists access to credit, lack of adequate technical skills, and informality as the top three challenges MSMEs face hindering their growth. More recently, Bassi et al. (2023) show that another limiting factor to growth in the Ugandan metalworking sector is its customer-tailored nature that prevents specialization.

Informal apprenticeship is the dominant training institution in this sector. Across Sub-Saharan Africa, firm owners exercise near-complete control over apprenticeship entry, curriculum, and duration , with apprentices predominantly unpaid (Teal, 2016; Frazer, 2006). Formal vocational training, by contrast, is rare among employed workers in small firms. Uganda's vocational training institutions are well-developed, with Mincerian returns to vocational training exceeding 50 percent

(Alfonsi et al., 2020), yet informal training systems persistently fail to close skills gaps because they reproduce existing skills rather than upgrading them (Johanson and Adams, 2004).

These training failures play out within a sector characterized by very small firms. During our listing exercise, we visited around 2,053 firms in Kampala and neighboring suburbs and found that less than 40 percent of firms employed at least four workers. In a different sample, Bassi et al. (2023) report an average number of six employees in welding firms in Kampala even after oversampling the largest firms. Small welding firms likely face the same headwinds to growth that have been reported in other settings such as suboptimal small-business management practices (McKenzie and Woodruff, 2017).

3 Experimental Design

Our study was conducted and analyzed in accordance with our pre-registration (AEARCTR-0011788) in the AEA RCT Registry.

We conducted a randomized controlled trial to evaluate the impact of our metalworking training program. This study took several stages. First, we identified eligible firms. Then, we surveyed the employers and workers at these firms to determine the two "potential trainees" from each firm. We tested the skills of all potential trainees before randomizing firms to treatment or control groups. For treatment firms, both potential trainees were assigned to training. The potential trainees from control firms remained untrained and were tracked as a counterfactual comparison group. Finally, after the conclusion of the training, we repeated the skills test for all potential trainees.

3.1 Sample Selection: Firm Eligibility

We determined firm eligibility based on key criteria that would ensure a comparable sample with high external validity. Appendix Figure A.1 provides a flowchart of this sample construction.

We first identified our evaluation sample by relying on a brief census survey of all metal fabrication firms located in the suburbs of Kampala. We used this screening survey to physically map firm locations and identify firms with the targeted number of workers (greater than four). We selected this lower bound of eligibility for two reasons: 1) it ensured that employers had experience managing teams of more than one or two workers, and 2) no firm would lose more than 50% of their workforce during our training program. We excluded firms where the employer or workers had previously trained with our expert trainer as we did not want this prior interaction to affect take-up.

We marketed our program to all the qualifying firms. The marketing exercises were designed to increase enrollment in the study. During the marketing visits, we spoke with both employers and workers and invited them to an orientation event. The orientation provided employers and workers with additional information about the training and ensured that they could make informed decisions about the training program and its anticipated benefits. The orientation also filtered firms based on willingness to participate. Those who failed to attend the orientation were excluded unless our research partners determined that they were located sufficiently close to one of our training centers.

3.2 Employer Survey

The employer survey collected data on the employer, the firm, and each of the firm's workers.

At the employer level, we elicited information about demographics, work and training history, and cognitive ability (e.g., Raven's matrices). We collected an incentive-compatible measure of the employer's time preference and an unincentivized measure of risk attitudes.

At the firm level, we collected information on profitability, assets, access to credit, measures of productivity, and personnel practices. Each employer also provided a list of all workers, their skills, their history with the firm, their individual reliability and productivity (as perceived by the employer), the opportunity cost of sending the worker to training, and the employer's expected profit gains from training each worker.

From each employer, we captured a novel, incentive-compatible measure of beliefs about how each worker would benefit from a hypothetical training assignment. To do this, we elicited 1) beliefs about each worker's baseline ability, 2) beliefs about each worker's endline ability after receiving training, and 3) beliefs about each worker's endline ability after not receiving training. These beliefs were captured using a simple elicitation in the style of Charness and Dufwenberg (2006) where employers guessed scores (from 0–100) on the baseline and endline practical skills tests and were paid 10,000 shillings ($\approx$ 2.78 USD) for guesses that were within 10 points of the truth.[3] All employers were informed of the objective scoring rubric for the practical skills test before their guesses. We describe the skills test in detail in Section 3.5. The difference between endline scores with and without training constitutes our prespecified measure of "anticipated skill gains."

[3] For similar belief elicitations, see Abeler, Nosenzo and Raymond (2019) or Madarasz, Danz and Wang (2023). For a more detailed discussion of methodology, see Danz, Vesterlund and Wilson (2022).

3.3 Selection of Potential Trainees

Each firm provided two workers who would serve as potential trainees.

The first potential trainee was determined by the workers' preferences. At each firm, workers cast bids in a reverse auction indicating their "willingness to accept" a stipend for training attendance. Each worker could request an amount between 0 and 120,000 Shillings ($34). The worker who requested the lowest amount was selected as the first potential trainee for the firm. In the event the firm was assigned to treatment, this potential trainee would receive their bid in exchange for attending training, making this first-price reverse auction incentive-compatible.

Since workers' bids reveal their demand for training, we refer to this auction as the elicitation of "worker demand." The worker with the lowest bid—their stipend request—exhibited the highest demand for training.[4]

The second potential trainee was determined by the employer's preferences. On the employer survey, employers revealed their preferences for training by ranking their workers from the most preferred worker to the least preferred worker to receive training. We conducted this elicitation prior to employers learning the outcome of the worker demand elicitation to ensure that their ranking was not affected by its outcome.

The employer's most preferred worker would be selected as the second potential trainee unless 1) they had already been selected as the first potential trainee based on worker demand, or 2) they left the firm prior to treatment assignment. In these cases, we would move down the employer's preference ranking until we found a second potential trainee. This procedure ensured that the employer's entire ranking was incentive-compatible.[5]

3.4 Worker Survey

We surveyed both potential trainees from each firm. From these two workers, we collected a detailed work history and complete wage information including any work done at other firms. Like the employer survey, the worker survey elicited beliefs about the impact of training. For this elicitation, we incentivized workers to guess the average score of all trained workers. To avoid moral hazard,

[4] To conduct this elicitation, we required a sufficient number of workers to be present at the firm. In firms with 4–5 workers, 6–8 workers, or 8+ workers, we required 3, 4, and 5 workers to be present, respectively. Otherwise, we rescheduled for the following day.

[5] The employer's preferred worker and the worker-selected potential trainee were the same person approximately 20 percent of the time.

we did not incentivize our elicitation of beliefs about their own scores on the practical skills tests. We use this to complement our employer-based measure of anticipated skill gains.

3.5 Baseline Practical Skills Tests

We invited all potential trainees to participate in a baseline practical skills test. The skills test involved the fabrication of a common metalworking product: a steel window frame and shutter. We recruited independent assessors from a prestigious and accredited vocational institute in Kampala to evaluate relevant skills (e.g., measuring, cutting, welding, grinding). The assessors rated each worker based on quality, speed, safety, and efficiency when using materials. Scores were based on an objective scoring rubric and ranged from 0 to 100. A score of 100 corresponded to the output of an expert metal fabricator with decades of experience, a score of 50 corresponded to the average metalworker in Kampala, and a score of 0 corresponded to a completely unskilled worker. The assessors scored tests independently without any input from our trainers.

Immediately after completing the practical skills test, workers took a quiz on metalworking knowledge. This quiz evaluated workers on their understanding of the fundamentals of welding such as voltage, electrode size, and workplace safety.

We photographed the window frames and shutters on top of their respective full-sized blueprints. We used these photos to develop a final measure of worker skill: measured deviation from the specified product dimensions.[6]

3.6 Randomization

We randomly assigned eligible firms to the treatment or control group if employers completed the employer survey and both potential trainees completed the worker survey.

We used a stratified randomization design where we assigned firms to one of nine strata. Specifically, we grouped the firms according to which one of the three training locations was most easily accessible using public transportation. We further stratified based on the number of potential trainees from the firm who completed the baseline practical test (zero, one, or two).

Within each stratum, we used a block-randomization design based on Moore and Moore (2013) to group firms based on their total number of workers and the baseline skill levels of the potential

[6] We calculate this by summing all sides of each product (i.e., the perimeter) and comparing the sum to the specified amounts. We aggregate each piece's absolute deviation and divide by the sum of the correct measurements to get a percent deviation. We score a product based on the negative value of this percent deviation.

trainees. This resulted in 33 separate groupings for our final sample of 337 firms (674 potential trainees). In total, we assigned 91 firms (182 potential trainees) to the treatment group, a relatively low treatment assignment rate (27%) reflecting the capacity constraints in our training program.

3.7 Training Program

Our training intervention focused on technical skills for workers at small firms. We developed this technical training program to be scalable and portable. To do so, we identified a metal fabrication expert who, before designing the training, conducted interviews with firms to determine the gaps in their technical knowledge. After the training was initially piloted, we hired a curriculum expert to assist in the development of a detailed, lesson-by-lesson curriculum.

Our training was designed to be accessible to workers of small manufacturing firms who may have limited formal education and low levels of literacy. Publicly available metal fabrication training in Kampala often expects trainees to have completed lower secondary school (O-level) education. In contrast, our training focused on demonstrations and hands-on practice and did not expect trainees to engage with detailed written materials.

The entire training program lasted 168 hours. We made several accommodations to facilitate regular attendance, since much of our low-income target population has high opportunity costs. First, we located the training in temporary facilities near our trainees' workplaces. Second, we conducted training sessions two days a week for six to eight hours a day to ensure workers did not miss more than one-third of their working days per week. Third, all trainees received meals and transport reimbursements on training days.

3.7.1 Endline Practical Skills Tests

Shortly after the training concluded, potential trainees from both the treatment and control groups repeated the practical skills test, this time as the endline practical skills test. They similarly repeated the metalworking quiz. These frames were similarly photographed so that all measures of quality were replicated exactly.

4 Model and Hypotheses

We develop a model that, in the spirit of Acemoglu and Pischke (1999*b*), demonstrates how frictions in an otherwise competitive labor market may make it incentive compatible for employers to provide general training for workers. However, our model features important distinctions that clarify the ways in which the very frictions that make training provision incentive compatible may also distort the allocation of workers to that training.

The first distinction is that, in our setting, training provision faces a quantity restriction. Specifically, employers may only select one worker for training. This is a common constraint for many microenterprises, which would otherwise have to shut down if the majority or entirety of its workforce were attending training.

The second distinction is that training assignment is binary, and the skill gains from training are heterogeneous. There is not an intensive margin for the amount of training a worker may receive. The binary nature of training is common in settings where employers have to outsource worker training because of limited internal capacity to train workers. Heterogeneity in the skill gains from training is a relaxation of the uniform returns to training assumed in many existing models.

The third distinction is that workers have heterogeneous baseline skills and heterogeneous attachments to the firm. These, in consequence, produce heterogeneity in the labor market frictions each worker faces and, therefore the appropriability of their skill gains from training. These are all relaxations of assumptions in prior models.

Finally, when assessing the role of wage compression in motivating the provision of general training, Acemoglu and Pischke (1999*b*) explicitly state that firm attachments need not play a role. Wage compression may be motivated by search costs, asymmetric information about worker skills, firm specific human capital, efficiency wages, or labor market institutions such as minimum wages or union-bargained contracts. In contrast, we will provide direct empirical evidence on the role of firm attachments. Thus, one goal of our model will be to microfound the role firm attachments play in compressing wages in a competitive labor market. To do so, we need only assume concavity of the utility function.

In this context, our model develops hypotheses about employer selection for training. In parallel, we produce hypotheses about worker self-selection for training that we will compare to the hypotheses about employer selection. Our model begins with workers who gain gain utility from two sources: wages and firm attachments. Firm attachments operate similarly to job search frictions, except that they increase with worker tenure. With greater firm attachments, within-firm wage compression

increases because employers do not need to offer fully competitive retention wages to prevent worker turnover. Firms profit from the difference between a worker's market wage (equal to the worker's marginal revenue product) and the retention wages they must offer.

Employers maximize profits for the firm by selecting workers for training based on the increase in extractable profits post-training. In this environment, the correlation between the anticipated skill gains from training and the appropriability of those gains critically determines training effectiveness. When the two are negatively correlated, employer selection for training diminishes productivity gains from training. Employers underselect on skill gains and overselect on firm attachments. Workers, in contrast, are motivated by the potential wage gains from training. This means that worker demand for training will always increase in anticipated skill gains.

4.1 Baseline Wages

Employers maximize the difference between total revenues—an aggregation of the marginal revenue products across all workers—and total costs, which include wages paid to workers.

The competitive labor market provides workers with a wage offer equal to their marginal revenue product. We refer to this marginal revenue product as the worker's "quality" (q_i). After gaining employment at the firm, workers begin to accumulate firm attachments and gain utility from both wages and firm attachments.

Suppose worker attachments to their firm are the only imperfection in an otherwise perfectly competitive labor market. This feature allows a worker's wage (w_i) to fall short of their marginal revenue product (i.e., $w_i \leq q_i$). The profit maximization problem then becomes one in which employers maximize the difference between the marginal revenue product of each worker and the worker's wage, subject to the constraint that all workers must receive at least as much utility from remaining at the firm as they would receive by accepting a wage from the competitive labor market.

Formally, suppose workers have a common utility function that monotonically increases in each of its two arguments: 1) wages (w_i), and 2) "firm attachments" (a_i). Firm attachments may be non-remunerative benefits of maintaining their current employment (e.g., relationships, routines, social capital, etc.).

We assume that a worker begins any new employment with $a_i = 0$. Thus, a worker considering leaving their current employment must compare the utility they derive from their current wage and attachment to a wage equal to their marginal revenue product, q_i, with zero attachment. We

define the "retention wage" as the wage the incumbent firm must offer to prevent the worker from leaving.

A firm's profit can be approximated by the sum of differences between worker marginal revenue products and wages (represented as the vector, $\mathbf{w} = (w_1, \ldots, w_N)$). The maximization problem is:

$$\max_{\mathbf{w}} \Pi = \max_{\mathbf{w}} \sum_{i=1}^{N} q_i - w_i \tag{1}$$
$$\text{s.t.} \quad U(w_i, a_i) \geq U(q_i, 0) \quad \forall i \in \{1, \ldots, N\}$$

We focus on the partial-equilibrium question of retention wages for existing workers rather than general-equilibrium questions about labor supply, worker recruitment, and firm size. Thus, at baseline, we consider the set of workers and their quality levels fixed, simplifying the profit maximization problem to a wage minimization problem:

$$\min_{\mathbf{w}} \sum_{i=1}^{N} w_i \tag{2}$$
$$\text{s.t.} \quad U(w_i, a_i) \geq U(q_i, 0) \quad \forall i \in \{1, \ldots, N\}$$

Profit maximization causes the constraint to bind. This allows us to implicitly define equilibrium wages as a function of quality and attachment:

$$w_i^* : \quad U(w_i^*(q_i, a_i), a_i) - U(q_i, 0) = 0 \tag{3}$$

While wages remain unknown without further specifying the utility function, we can use the implicit function theorem to derive comparative statics for how a worker's retention wage adjusts to changes

in firm attachment and quality.

$$\frac{\partial U(w_i^*(q_i,a_i),a_i)}{\partial w} \times \frac{\partial w_i^*(q_i,a_i)}{\partial a} + \frac{\partial U(w_i^*(q_i,a_i),a_i)}{\partial a} = 0$$

$$\implies \frac{\partial w_i^*(q_i,a_i)}{\partial a} = -\frac{\frac{\partial U(w_i^*(q_i,a_i),a_i)}{\partial a}}{\frac{\partial U(w_i^*(q_i,a_i),a_i)}{\partial w}} < 0 \tag{4}$$

$$\frac{\partial U(w_i^*(q_i,a_i),a_i)}{\partial w} \times \frac{\partial w_i^*(q_i,a_i)}{\partial q} - \frac{\partial U(q_i,0)}{\partial w} = 0$$

$$\implies \frac{\partial w_i^*(q_i,a_i)}{\partial q} = \frac{\frac{\partial U(q_i,0)}{\partial w}}{\frac{\partial U(w_i^*(q_i,a_i),a_i)}{\partial w}} > 0 \tag{5}$$

Since utility increases in both wages and attachment (i.e., $\frac{\partial U}{\partial a} > 0$ and $\frac{\partial U}{\partial w} > 0$), Equation 4 shows that as a_i rises, the worker's retention wage falls (holding quality fixed). That is, employers can exploit a worker's attachment by offering a lower retention wage.

With fully flexible wages, employers would routinely lower wages to offset a worker's gains in attachment. Research in labor economics finds that such downward adjustments are rare (Altonji and Devereux, 1999; Bewley, 1999).[7] Thus, we assume downward nominal wage rigidity, meaning that a worker's utility strictly increases by remaining at a firm and accumulating attachments.

Next, Equation 5 shows that wages increase in q_i. To further characterize this wage adjustment, suppose that workers' utility functions exhibit diminishing marginal utility of wealth and non-negative complementarity between wages and attachment (i.e., $\frac{\partial^2 U}{\partial w^2} < 0$ and $\frac{\partial^2 U}{\partial w \partial a} \geq 0$). Since we know that employers optimally set wages below quality for attached workers ($w_i^*(q_i,a_i) < q_i \quad \forall a_i > 0$), we have:

$$\frac{\partial U(w_i^*(q_i,a_i),a_i)}{\partial w} \geq \frac{\partial U(w_i^*(q_i,a_i),0)}{\partial w} > \frac{\partial U(q_i,0)}{\partial w}$$

$$\implies 1 > \frac{\frac{\partial U(q_i,0)}{\partial w}}{\frac{\partial U(w_i^*(q_i,a_i),a_i)}{\partial w}} = \frac{\partial w_i^*(q_i,a_i)}{\partial q} \quad \forall a_i > 0$$

That is, for any attached worker, wages grow less than proportionally to a worker's marginal revenue product. The diminishing marginal utility of wealth makes a given wage premium from an outside firm less attractive as quality (and wages) grow. Thus, the gap between the retention wage and the market wage is wider for higher quality workers (holding fixed a_i). This produces the wage compression needed to make training incentive compatible. However, this relationship

[7] In a recent survey, Elsby and Solon (2019) find that downward nominal wage rigidity is the norm internationally, though downward wage adjustments happen more often than previously believed.

between attachment and wage compression foreshadows the distortions they may cause in training selection.

4.2 Employer selection of workers for training

We seek to characterize the process by which employers select workers to receive our free training program. This training program will increase the selected worker's quality (i.e., their marginal revenue product) and, therefore, their market wage. While Equation 5 shows that this will increase the worker's retention wage, it also shows that it will grow more slowly than the worker's quality, increasing the firm's profits. Since training is quantity constrained, employers maximize profits by training the workers whose increases in marginal revenue product (which accrues to the firm) most exceed their increases in wages (which the firm transfers to the worker). In determining this optimal selection, employers anticipate that retention wages will be lower for workers with greater attachment.

Let worker quality be a function of training, $q_i(T_i)$, where T_i is a binary variable equal to 1 if worker i receives training. We assume monotonicity ($q_i(1) \geq q_i(0)$) and differentiability $\exists \frac{dq}{dT}$. Without loss of generality, we explore training selections for a single worker. So, the profit maximization problem for a firm with N workers becomes the selection of a $1 \times N$ vector of binary variables, $\mathbf{T} = \{T_1, \ldots, T_N\}$, with exactly one element equal to 1:

$$\begin{aligned} &\max_{\mathbf{T},\mathbf{w}} \Pi = \max_{\mathbf{T},\mathbf{w}} \sum_{i=1}^{N} q_i(T_i) - w_i \\ &\text{s.t.} \quad U(w_i, a_i) \geq U(q_i(T_i), 0) \quad \forall i \in \{1, ..., N\} \end{aligned} \tag{6}$$

We differentiate per-worker profits with respect to T_i to identify the profitability of training each worker.

$$\begin{aligned} \Pi_i &= q_i(T_i) - w_i^*(q_i(T_i), a_i) \\ \frac{\partial}{\partial T_i}\Pi_i &= \frac{\partial}{\partial T_i}\left[q_i(T_i) - w_i^*(q_i(T_i), a_i)\right] \\ \implies \frac{\partial}{\partial T_i}\Pi_i &= \frac{dq_i}{dT_i} - \frac{\partial w_i^*(q_i(T_i), a_i)}{\partial q} \times \frac{dq_i}{dT_i} \\ \implies \frac{\partial}{\partial T_i}\Pi_i &= \frac{dq_i}{dT_i}\left[1 - \frac{\partial w_i^*(q_i(T_i), a_i)}{\partial q}\right] = \frac{dq_i}{dT_i}\left[1 - \frac{\partial w_i^*(q_i(0), a_i)}{\partial q}\right], \text{ since } T_i = 0 \ \forall i \end{aligned} \tag{7}$$

Equation 7 reveals the two effects of training: a direct effect that increases per-worker profits through increases in worker quality ($\frac{dq_i}{dT_i}$) and an indirect effect that reduces profits through quality-induced wage adjustment ($\frac{\partial w_i^*(q_i,a_i)}{\partial q}$). These are the terms we refer to as "anticipated skill gains" and "wage adjustment," respectively.

Employers seek to maximize anticipated skill gains while minimizing wage adjustments. From Section 4.1, we know that wage adjustment diminishes with both worker quality and firm attachments. Thus, all else equal, employers will profit more by selecting higher-quality workers with more firm attachments. Since wages increase in quality, this means firms will positively select on worker wages as well.

Result: Employers positively select workers for training based on firm attachments, worker quality, worker wage, and anticipated skill gains (wage adjustment held constant).

4.3 Worker self-selection to training

Workers in our study cast competitive bids for the right to receive training. The worker who reveals the highest demand for the program is assigned to training.[8] The worker's utility-maximization problem mirrors the employer selection problem in many ways but is simpler because the incentive-compatibility constraints are unnecessary.

$$\max_{B_i} U(w_i, a_i) =$$
$$\max_{B_i} \left\{ \Pr(B_i \leq \overline{B_{-i}}) \left(U(w_i^*(q_i(1), a_i), a_i) + B_i \right) + \Pr(B_i > \overline{B_{-i}}) \left(U(w_i^*(q_i(0), a_i), a_i) \right) \right\} \tag{8}$$

where, $\overline{B_{-i}} = \min\{B_j\} \quad \forall \quad j \neq i$.

As a first-price auction, the optimal bidding function is a linear transformation of the worker's value-added from training:

$$B_i^* = C \times [U(w^*(q_i(1), a_i), a_i) - U(w^*(q_i(0), a_i), a_i)]$$

where C is a positive constant.[9] Assuming smoothness between the untrained and trained states

[8] In particular, we performed a willingness-to-accept elicitation where workers revealed demand through requesting a smaller stipend to attend training in a first-price sealed-bid auction.

[9] Alternatively, recall that the employer's retention wage will make the worker indifferent with outside employment at market wages. Thus, $U\left(w_i^*(q_i(T_i), a_i), a_i\right) = U\left(q_i(T_i), 0\right)$ and Equation 9 becomes $B_i^* \approx C \times \frac{\partial U(q_i,0)}{\partial w} \times \frac{dq_i}{dT_i}$

allows us to make the approximation:

$$B_i^* \approx C \times \frac{\partial U(w^*(q_i(0), a_i), a_i)}{\partial w} \times \frac{dw^*(q_i(0), a_i)}{dq} \times \frac{dq_i}{dT_i} \tag{9}$$

Thus, worker bids increase with three factors: marginal utility of wages, wage adjustment, and anticipated skill gains.

Result: Workers will positively self-select into training based on anticipated skill gains but negatively self-select on worker quality and worker wage.

Note how worker self-selection on the quality and wages is opposite to the employer's optimal selection. Equation 7 showed how firm profits are sensitive to worker characteristics. We can combine Equations 7, 5, and 9 to show how the alignment of preferences for the two parties is mediated by wage adjustments:

$$\begin{aligned}
\frac{\partial}{\partial T_i}\Pi_i &= \frac{dq_i}{dT} - \frac{\partial w_i^*(q_i(0), a_i)}{\partial q} \times \frac{dq_i}{dT} \\
\frac{\partial}{\partial T_i}\Pi_i &= \frac{B_i^*}{C \times \frac{\partial U(w_i^*(q_i(0),a_i),a_i)}{\partial w} \times \frac{\partial w_i^*(q_i(0),a_i)}{\partial q}} - \frac{B_i^*}{C \times \frac{\partial U(w_i^*(q_i(0),a_i),a_i)}{\partial w}} \qquad (10) \\
\frac{\partial}{\partial T_i}\Pi_i &= B_i^* \times \frac{1 - \frac{\partial w_i^*(q_i(0),a_i)}{\partial q}}{C \times \frac{\partial U(w_i^*(q_i(0),a_i),a_i)}{\partial w} \times \frac{\partial w_i^*(q_i(0),a_i)}{\partial q}} \qquad (11)
\end{aligned}$$

That is, as wage adjustments approach 1 (i.e., proportional), the increase in firm profits from training a given worker falls to zero and is uncoupled from that worker's valuation of the training. While both the worker and employer value anticipated skill gains from the training, they are on opposite sides of the zero-sum conflict over dividing those gains.

5 Data Description, Summary Statistics, and Treatment Balance

In this section, we characterize the data we collected, summarize key characteristics of firms and workers, and perform balance checks to ensure that our randomization maintained the balance of key variables across treatments.

5.1 Data

Our sample includes a total of 337 firms and 1,784 workers (of which 674 are potential trainees). As detailed in Section **??**, employer-level data come from the employer survey. Our worker-level dataset, however, is compiled from several sources: the employer survey, the worker demand elicitation, the worker survey, and the practical skills tests. As a result, the sample size for our worker-level data changes depending on the questions we address. We observe data for nearly all workers when the dataset is derived from either the employer survey or the worker demand elicitation. However, both the worker survey and the practical skills tests gather data only from the subset of potential trainees and are therefore limited and non-randomly selected.

5.2 Summary Statistics

Tables 1 and 2 report summary statistics for our sample at the firm- and worker-levels, respectively. We then report summary statistics for variables directly related to perceptions about training and worker skills in Table 3. Each table is described in detail in the sections that follow.

5.2.1 Firm and employer characteristics

Table 1 presents firm and employer characteristics. The data show that firms in our sample are relatively mature, with an average firm age of 10.35 years. The employers are similarly experienced: 36.50 years old with 14.25 years of metal-fabrication experience, on average. Nonetheless, the firms are dynamic with high growth and high turnover of both workers and products. On average, firms hired 2.3 new workers and lost 1.0 worker to quits in the last 12 months, relative to a mean workforce of 5.42 workers. Average monthly profits of $425 are approximately 2.6 times the average monthly worker earnings in Uganda reported in International Labour Organization (2021).

The average number of workers is comparable to the average size of welding firms in the same setting reported in Bassi et al. (2023) (5.88 workers). The entrepreneurs in our sample are committed—only six percent would leave for wage employment equal to the firm's profits—and optimistic—over 80 percent expect their business to survive at least five years, and they expect their workforce to double in the same time frame.

These entrepreneurs tend to have low levels of formal education (10.35 years) and only 20 percent have completed metal fabrication training. We find evidence of credit constraints in our setting: 25 percent have received a loan of any kind, only three percent have applied for loans to train workers,

and only one percent have ever received such a loan. These features, combined with the high turnover rates make the firm provision of training highly unlikely at baseline (Becker, 1962).

5.2.2 Worker characteristics

Table 2 presents worker characteristics captured from several data sources. Panel A includes attributes of all workers at all firms as reported on the employer survey. Panel B presents characteristics of the potential trainees as reported on the worker survey.

From Panel A, we see that approximately 18% of workers are related to their employer. Workers have an average tenure of 3.5 years at the firm. On average, we find a relatively low likelihood of any individual worker departing based on two subjective metrics: 1) employers report each worker's "likelihood of leaving," which averages to 2.77 out of 10, and 2) employers expect to retain workers at the firm for an average of 19.98 months (out of a maximum of 24 months) after receiving training. The employers reported that they pay workers an average monthly wage of $159 (about 570,000 Ugandan shillings), which is comparable to average wages of Ugandan workers reported in International Labour Organization (2021) .

Panel B reports additional demographic characteristics of workers as collected from all potential trainees. We find that these workers tend to be younger (24 years old) than their employers with far fewer years in metal fabrication (4.5 years). Nearly all workers are male (99%) and have low levels of education (8.3 years), and few have completed metal fabrication training (only four percent). Because of our sample restrictions, all included workers are full-time workers, defined as working at least four days per week at the firm. Still, 24 percent of the workers worked for more than one firm in the month before the survey.

5.2.3 Demand for training, perceptions of training, and beliefs about worker quality

Table 3 summarizes key variables that measure expressions of demand for training, perceptions about the effectiveness of our training, and beliefs about worker quality.

First, Panel A reports the willingness to accept training as elicited from workers. Recall that workers could request a training stipend of up to $33, yet they only requested $14.51, on average ($p < 0.001$ for a two-sided test of differences). That is, on average, workers leave money on the table for the opportunity to attend training, suggesting a general demand for training.

Next, Panel B reports the results of our incentive-compatible elicitation of employer beliefs. We

elicited beliefs about worker quality at baseline and endline (with and without training) on a scale from 0 to 100. Employers believe the baseline quality of their workers is 59.73, well above the midpoint score of 50, which they were told reflected the skills of "the average metalworker in Kampala." Employers believe that worker quality will improve to 70.20 at endline without training. Importantly, they believe that training will accelerate this skill growth, reporting that quality will rise an additional 11 points to 81.1 at endline for trained workers ($p < 0.001$ when tested against endline beliefs without training).

Workers themselves have somewhat more optimistic views of their skills. While the perceptions of own-quality were not incentivized to avoid hedging concerns, they offer a useful contrast to the beliefs of the employers. Panel C shows that potential trainees believe their own quality to be 73.61, on average, at baseline. Workers believe these skills will grow by about 10 points (to 83.96) by the time of the endline test without training. With training, workers believe their endline scores will be nearly perfect, 94.6, on average. When incentivized to guess the average quality of a trained worker, they reported a guess of 75.46, on average, which is close to the employer's estimate of 81.10 in Panel A and also reveals a positive view of the impact of training, on average.

Panels D and E report the scores for each potential trainee on baseline and endline practical skills tests, respectively. The average practical skills test scores at baseline for the process, product, and overall (i.e., combined) portions were 44.08, 47.63, and 45.92, respectively. Scores on the knowledge assessments at baseline were 36.54%, on average. Finally, at baseline, the measured deviations of the final product from the blueprint were approximately 6.22%. We observed an increase in all scores at endline, pooling treatment and control groups, with process, product, and overall portions rising to 62.00, 50.87, and 56.23, respectively. Knowledge assessment scores rose to 43.84% and measured deviations fell to 3.24%, on average.

In Panel F, we report employer responses to a question about whether they believe that training their workers would improve their profits and why. 82% of employers responded that training would improve profits, while 18% responded that it would not. Employers overwhelmingly believed that profit increases from training would result from improvements in product quality (97%). Other key reasons reported for the increase in profitability are a reduction in wastage (39%) and peer effects where trained workers will train others (29%). Among the minority who did not believe that training would increase profitability, nearly half believed that their product quality was already optimal (48%).

Panel G shows attendance rates for treatment-group workers. Of these, 76% attended at least one day of training, with the average attendance being 10.63 training days (out of a maximum of 23

days).[10] 40% of workers met the requirement for training "completion" by attending at least 70% of the assigned days.[11]

5.3 Balance Tests

As described in Section 3.6, we performed a stratified randomization of firms to the treatment and control groups. Since we stratified our sample based on a limited number of characteristics, it is important to confirm that a broader set of firm-level characteristics remain balanced. We show this in Table 4, where we see that, across a wide range of firm- and employer-level characteristics, there are no statistically significant differences between the treatment and control groups.

Even with balance of firm-level characteristics, randomization does not necessarily ensure that all worker-level characteristics remain balanced, so we demonstrate that in Tables 5 and 6. We first show balance for worker-level demographic characteristics in Table 5. Next, we show balance for worker-level training demand and baseline skills in Table 6. In both cases, we find worker-level characteristics are generally balanced across treatments. A few spurious rejections are expected by chance when comparing 27 different means. In our case, we find that worker wages, worker education, and worker days at the firm per week show marginally significant differences in means across the two groups. While the latter two characteristics are never used in our primary analysis, worker wages do appear as a key source of heterogeneity. Thus, it is reassuring to note that the distributions of wages across the two groups are similar outside of the top few percentiles.

6 Experimental Results

In this section, we present our primary results on the impact of training and the selection of workers for training. We begin by evaluating the causal impact of our training on worker skill gains. We find a large and significant *average* effect of being assigned to training. Thus, the selection of workers to training becomes a key policy-relevant outcome as well as a first-order determinant of the downstream impacts of the program in terms of efficiency and equity.

We evaluate the selection into training by both the employers and workers. We separately consider various worker characteristics on which this selection is based and show how employers' and

[10] Public holidays affected the one training center, reducing the total possible days at that center to 21 days. To compensate, the trainer extended the daily instruction hours to complete the curriculum within the shortened time frame.

[11] This could be any combination of attendance to achieve the 70% mark. Meaning that absences could (and did) occur at the start, middle, or end of the training period.

workers' preferences are broadly misaligned with respect to training selection.

Unless otherwise specified, we perform all analyses using linear regressions with firm-level fixed-effects and standard errors clustered by firm. We normalize many variables for ease of interpretation. In our tables, normalized variables have the suffix "(SDs)."

6.1 Impact of Training

We evaluate our training program using an intent-to-treat design—primarily a comparison of means between the treatment and control groups. This intent-to-treat analysis does not guarantee that any given worker attends the training, but instead incorporates the potential for non-compliance into the estimates. This produces an estimate of the "full-effect" of the policy as implemented.

For all impact analyses, our primary outcome variable is the score on the endline practical skills test. We evaluate this both as a combined "overall score" and decomposed into its two components, a "process score"—the quality of a worker's methods and techniques—and a "product score"—the final result. We supplement practical skills test analysis by estimating the impact of training on 1) scores on an endline quiz of metalworking knowledge, and 2) measured deviations of the final product from the blueprints.

Table 7 presents our evaluation of the impact of training assignment, comparing potential trainees from treatment and control firms using ordinary least squares regression. Columns 1 and 2 report the results of this analysis for overall endline practical skills test scores, without and with baseline covariates, respectively. In both cases, we find large and statistically significant increases from training assignment, 0.484 SD and 0.483 SD, respectively ($p < 0.001$ for both). Columns 3 and 4 show that the overall effect appears to derive more from improvements in the process scores—a 0.529 SD increase ($p < 0.001$)—than from improvements in the final product score—a 0.236 SD increase ($p = 0.056$). These two coefficients are significantly different from each other ($p = 0.014$). In Columns 5 and 6, we see that treatment assignment also produces large and statistically significant increases in endline quiz scores—a 0.556 SD increase ($p < 0.001$)—and measured deviations from the blueprint—a 0.372 SD improvement ($p < 0.001$).

The positive average treatment effects presented in Table 7 may mask important heterogeneity. In Table 8, we divide our sample by whether the potential trainee was selected by the employer or self-selected through the worker demand elicitation. In both subgroups and for all outcome variables, we find positive and significant effects of training assignment, indicating that the benefits to training are broad-based. We do, however, find some suggestion of heterogeneity. The impacts of training assignment on all outcome variables are directionally larger for workers who were self-selected

into training, though none of the differences is significant at the 10% level. Only when comparing measured deviations in Columns 5 and 6 does the difference in treatment effects approach statistical significance. We find that the treatment effect is nearly twice as large for self-selected workers, but this difference is not significant at conventional levels ($p = 0.110$).

We visualize the impact of training attendance in Figure 3, which plots individual skill gains on the vertical axis and training days attended on the horizontal axis. This graph shows how skill gains are similar between the control group and the non-compliers in the treatment group (i.e., those attending zero days of training). However, for the treatment group workers who attend training, the skill gains are positive and increase with additional attendance, suggesting both extensive and intensive margin effects of training assignment.

6.2 Selection

Having established the large benefits of training assignment on worker outcomes, the selection of workers for training becomes a first-order concern in determining both the aggregate effectiveness of worker training and the impact it has on worker welfare. Recall that the respective incentivized demand elicitations each determined one of the firm's two potential trainees. We refer to the monetary value of training elicited from each worker (and normalized across all workers) as "worker demand," and we refer to each worker's percentile rank in the employer elicitation as "employer demand." For employers, we also include their responses to an unincentivized ranking of the profit gains from training each worker ("profitability," hereafter).

Our model identifies three key characteristics that both employers and workers should consider when evaluating training: the anticipated skill gains from training, the worker's attachment to the firm, and the worker's baseline ability (and wages). We evaluate each in sequence below.

6.2.1 Skill Gains from Training

Our model predicts that, on average, both employer and worker demand should increase with worker skill gains from training. For employers, however, this prediction only holds if the wage adjustments after training are constant. If skill gains and wage adjustments negtively covary, then employers must balance a worker's anticipated skill gains from training with their wage adjustment after training. For instance, if the gains from training are highest among low-skilled workers, worker demand should be unambiguously higher for this group, but employer demand may be muted by the large anticipated wage adjustment for this group of workers who face fewer frictions in leaving

for other firms.

As per our pre-specification, we use the difference between employer beliefs about endline scores with and without training as our measure of anticipated skill gains. We normalize this measure within each firm.

Table 9 presents our findings of misaligned demand between employers and workers along the dimension of anticipated skill gains. Panel A explores the impact of anticipated skill gains on several measures of training demand. Columns 1 and 2 show that a one SD increase in anticipated skill gains is associated with 0.036 and 0.073 percentile point decreases in that worker's ranking on the employer demand elicitation and profitability statements ($p < 0.001$ for both). In contrast, Column 3 shows that a one SD increase in anticipated skill gains is associated with a 0.054 SD increase in that worker's demand for training ($p = 0.084$).

To demonstrate robustness to the source of anticipated skill gains, Panel B replicates this analysis using anticipated skill gains stated by the workers during the worker survey. We find patterns of misaligned preferences that are identical in sign and significance, only with differing magnitudes. Employer demand and perceived profitability are both lower for workers with higher anticipated skill gains, while worker demand increases in this dimension.

In sum, while workers prioritize skill gains from training, employer demand is sufficiently distorted by other selection variables that they negatively select on potential skill gains.

6.2.2 Firm Attachments

Our model makes sharp predictions about how firm attachments should affect selection. These characteristics predict which party has more leverage in the zero-sum contest over appropriating the gains from training. Thus, employers should prioritize training workers with greater firm attachments, and workers' demand for training should decrease as their firm attachments increase.

We measure firm attachments using four different variables: whether a worker is a relative of the employer, how much the employer feels they can trust the worker, the worker's tenure with the firm, and how long the employer believes they will retain the worker after training. We put each variable in its own panel of Table 10.

Table 10 shows the importance of these variables in determining employer demand and profitability. Across all panels, Column 1 reports the coefficients from four separate regressions of employer demand on each of the four firm-attachment variables and finds that each increases employer demand to train the worker in question ($p < 0.005$ for each). Similarly, Column 2 reports the coefficients

from four separate regressions of profitability on each of the four firm-attachment variables and finds that they also increase the profitability of training the worker in question ($p < 0.001$ for each). The alignment between employer demand and profitability across all four attachment variables confirms that employer selection is deliberate: employers are not making arbitrary choices but consistently acting on beliefs about which workers are most profitable to train.

In Column 3 we show how workers' preferences conflict with those of their employers. None of the four firm-attachment variables that predicted increases in employer demand are significantly associated with increases in worker demand. Indeed, employer trust and worker tenure are both negatively and significantly associated with worker demand ($p = 0.025$ and $p = 0.011$, respectively). This is not mere indifference: workers demand decreases when they are more attached to the firm in these ways, directly opposing the employer's selection criteria.

Combined with the results of the previous section, we find that employers are willing to sacrifice anticipated skill gains in order to over-select on firm attachment variables that should limit wage adjustments after training. At the same time, workers disregard or negatively self-select on firm attachments, more narrowly basing their demand on anticipated skill gains.

6.2.3 Baseline Skills and Baseline Wages

Similar to firm attachments, our model predicts that employers and workers should have misaligned preferences with respect to how baseline skills and wages affect training demand. Employers should demand training for higher-skilled, higher-wage workers for whom wage adjustments are muted by the diminishing marginal utility of wealth. Workers, on the other hand, should demand training when they are less skilled and have lower wages for the same reason—the marginal benefits of wage gains are highest when wages are low.

In Table 11, we confirm our model's predictions that employers will over-select high-wage workers with high baseline skills. Column 1 shows that a 1,000 USD increase in monthly wages increases employer demand for training by 0.307 percentile points ($p < 0.001$).[12] Columns 2 and 3 show that employer demand also increases significantly with employer beliefs about worker baseline skills and with scores on baseline practical skills test scores. A one SD increase in each metric increases employer demand by 0.105 and 0.095 percentile points, respectively ($p < 0.001$ for both). Columns 4–6 show that employers also anticipate greater profitability from training these workers. We find a pattern of results that is similar in sign, significance, and magnitude to employer

[12] For this analysis, we winsorize the top 1% of wages to address likely data-entry errors. We replicate the analysis using raw wage data in Appendix Table B.3. The results are less precise, but qualitatively unchanged.

demand.

We also confirm our model's predictions with respect to worker demand. We find that worker demand has no significant relationship with baseline wages and is negatively related to baseline skills. Columns 7–9 of Table 11 characterize the misalignment in preferences between employers and workers along these dimensions. Worker demand has a directionally negative relationship with baseline wages and has significant negative relationships with employers' beliefs about baseline skills and scores on baseline practical skills tests ($p \approx 0.037$, and $p \approx 0.002$, respectively).

Note how the selection patterns for employers and workers derive from the same mechanism, yet produce dramatically different outcomes. The employer selection pattern we show in this section is regressive. By leveraging the diminishing marginal utility of wealth to limit wage adjustments, employers consistently direct training toward workers who are already better paid and more skilled, compounding existing within-firm inequality in human capital accumulation. In contrast, worker self-selection favors low-wage, low-skilled workers who seek out training because their marginal utility from the wealth gains from training is high.

7 Robustness Checks

In the previous section, we established our key results with respect to the impact of training and how employers and workers make selections for training. However, much of our analysis relies on novel data collected in a complex environment. Thus, it is important to establish that these metrics accurately capture the preferences or behaviors that we claim. In this section, we perform robustness checks demonstrating the validity of our measures for: 1) elicited employer beliefs, 2) elicited worker demand, 3) worker wages and skills, and 4) employer motivations.

7.1 Robustness Check #1: Predictive Validity of Beliefs

A potential concern with our analysis is that differences in preferences between employers and workers may simply reflect differences in information about the skill gains from training. In this case, employer beliefs may not reflect the truth or may exhibit systematic biases. We find, instead, that our novel measures of employer beliefs about worker skills are highly predictive of objectively measured skills and closely correlated with stated beliefs about skill gains reported by the workers. Recall that we used an incentive-compatible elicitation to capture beliefs about worker skills for three states: at baseline, at endline without training, and at endline with training. We derive beliefs

about anticipated skill gains from training by comparing beliefs about endline worker skills with and without training.

In Table 12, we first test how well employer beliefs about worker skills predict objective measures of worker quality as evaluated through the practical skills tests and metalworking quizzes. Columns 1, 3, & 5 show that employer beliefs strongly predict baseline quality. A one SD increase in beliefs about baseline skills is associated with a 0.623 SD increase in baseline practical skills test score ($p < 0.001$), a 0.212 SD increase in metalworking quiz score ($p = 0.011$), and a 0.342 SD improvement in measured deviations on the practical skills test ($p < 0.001$).

Next, we use our data to rule out the possibility that employers suffer from asymmetric information about worker quality or gains from training. In fact, employer beliefs closely correlate with workers' stated beliefs about their own quality (correlation of 0.386 ($p < 0.001$)). Columns 2, 4, & 6 include worker beliefs about baseline skills to show the limited gains in explanatory power. After conditioning on the employer beliefs, there is no specification in which the worker beliefs are statistically significant, and the R^2 statistic never grows by more than a percentage point.

Finally, we show that employers' beliefs are useful in predicting the impact of training. We construct a measure of employer beliefs about endline skills that accounts for the treatment assignment. This is the employer's belief about endline practical skills test scores with training for treatment and without training for control workers. This measure is standardized with respect to the beliefs about worker skills without training. In Column 7, we show that a one SD increase in our measure of endline beliefs is associated with a 0.250 SD increase in endline practical skills test scores, after controlling for baseline practical skills test scores ($p = 0.013$).

In sum, employer beliefs are accurate, correlate closely with worker beliefs, and predict training outcomes. Therefore the misalignment in selection cannot be attributed to information asymmetry. Employers, like workers, know what skill gains are available. However, unlike workers, employers choose to ignore them.

7.2 Robustness Check #2: Predictive Validity of Worker Demand

Another potential concern with our data is the use of abstract tasks to measure demand—in our case, the use of a reverse auction to elicit worker demand. Here, we demonstrate the validity of our worker demand metric by showing how it captures consequential behaviors towards training.

Table 13 shows that this measure of worker demand is predictive of costly behaviors: the number of training days attended. Column 1 shows that a one SD increase in worker demand for training

is associated with 2.166 more days of training attended, which is a 0.255 SD increase relative to the mean ($p = 0.030$).[13] Moreover, worker demand aligns with their own stated beliefs about the skill gains they would experience from training. Column 2 shows that a one SD increase in worker demand is associated with a 0.110 SD increase in a worker's anticipated skill gains ($p = 0.005$).

The validity of the worker demand measure strengthens the selection results in Section 6.2 showing that differences in employer demand and worker demand reflect actual differences in preferences, not artifacts of the elicitation method.

7.3 Robustness Check #3: Skills, "Productivity," and Worker Wages

One potential concern with our analysis is that we control for skills based on results of our metalworking skills test, which may fail to capture workforce relevant skills. Valid measures of productivity should positively correlate with wages. Thus, to examine the workforce relevance of our measure of skills, we leverage data on worker baseline wages. These data are derived from the employer survey and involve aggregating multiple sources of pay to create a single variable capturing monthly pay from all sources within the firm (salary, piece-rate work, overtime, etc.).

In Section 4, we developed a model of a competitive labor market with workers who gain utility from firm attachments. Under this model, market wages rise proportionally to marginal revenue product, but observed increases in retention wages are limited by firm attachments. However, due to downward nominal wage rigidities, retention wages should always be weakly increasing in worker quality. Additionally, we should see the relationship between quality and wages attenuate for workers with strong firm attachments.

Table 14 confirms the validity of our productivity metrics by demonstrating robust associations with wages.[14] Columns 1–3 show strong relationships between baseline wages and beliefs about quality, worker experience, and scores on our practical skills test, respectively. Column 1 shows that a one SD increase in perceived quality results in \$45.55 higher monthly wages ($p < 0.001$). Column 2 shows that one additional year of experience adds \$8.07 to a worker's monthly wage ($p < 0.001$). Finally, Column 3 shows that a one SD increase in the score on our practical skills test is associated with a \$19.87 increase in monthly wage ($p = 0.043$).

[13] Greater worker demand increases the likelihood of training assignment through the auction, so we restrict our focus to the employer-selected workers in the treatment group.

[14] As with the analysis in Section 6.2.3, we winsorize the top 1% of wages to address likely data-entry errors. We replicate the analysis using raw wage data in Appendix Table B.4. The results are less precise, but qualitatively unchanged.

Importantly, we verify our model's assertion that the relationship between quality and wages will be attenuated for workers with strong attachments to the firm. In Column 4 Table 14, we replicate the analysis of Column 1 after restricting our sample to only the relatives of the employer. We find a smaller and only marginally statistically significant relationship between wages and perceived quality. When we restrict the sample to non-relatives in Column 5, we find that this relationship grows in magnitude and statistical significance ($p < 0.001$). We note that these two coefficients are not significantly different from one another.

7.4 Robustness Check #4: Alternative Motivations for Training Selections

We demonstrated that employer demand for training is not motivated primarily by anticipated skill gains for workers, but instead by the ability to extract those gains. However, it may be the case that the wedge between anticipated skill gains and appropriability is motivated by more practical concerns about workforce development. For instance, trained workers may be expected to pass on their knowledge to other workers. To address these alternatives, we use data on employer and worker stated motivations for training selections. We captured these data by asking for the motivations behind training selections on the employer and worker surveys.

We report motivations for training in Table 15. When asked why they chose their selected worker, 22% of employers reported that they selected the highest quality worker, 20% selected a worker they perceive to be trustworthy, 18% selected a worker they thought deserved to receive training, and 13% selected a worker they thought would be good at teaching their peers. 27% of employers reported one of a variety of other reasons, none of which surpassed a 10% reporting rate by employers. Interestingly, only 9% of the employers selected a worker based on their anticipated skill gains from training.

We measured two different possible motivations for training among workers: wage increases and improvements in wastage. We found that over 80% of workers believed that their wages would increase post-training. Additionally, we found that workers anticipated an average post-training reduction in wastage of 7.31% relative to baseline.

8 Structural Model: Wage Setting and Worker Selection for Training

We propose a basic structural model to capture the way that wage adjustments and skill gains compete to determine training demand for both employers and workers. We care in particular about the role of firm attachments in driving a wedge between employer and worker demand. In order

to estimate this model, we must make certain assumptions about the utility function of the workers and the mappings between the data we observe and the model parameters.

Suppose workers possess the following quasi-linear utility with diminishing marginal utility of wages:

$$U(w, a) = w^{\rho} + \tau \times a$$

where $\rho \in [0, 1]$ captures risk aversion and $\tau \in [0, 1]$ captures the weight given to firm attachments.

Under this functional form, the worker's (binding) incentive compatibility constraint from Equation 3 yields:

$$w^*(q_i, a_i)^{\rho_i} + \tau \times a_i = q_i^{\rho_i} \Rightarrow w^*(q_i, a_i) = \left(q_i^{\rho_i} - \tau \times a_i\right)^{1/\rho_i} \tag{12}$$

We derive the wage adjustment term by differentiating Equation 12 with respect to q_i:

$$\frac{\partial w_i^*(q_i, a_i)}{\partial q} = q_i^{\rho_i - 1} \times \left(q_i^{\rho_i} - \tau \times a_i\right)^{\frac{1-\rho_i}{\rho_i}} \tag{13}$$

We previously derived employer profit gains from training a given worker in Equation 7. Plugging in our wage adjustment function, we get:

$$\frac{\partial}{\partial T_i}\Pi_i = \frac{dq_i}{dT}\left[1 - q_i^{\rho_i - 1}\left(q_i^{\rho_i} - \tau \times a_i\right)^{\frac{1-\rho_i}{\rho_i}}\right] \tag{14}$$

Similarly, we derived the worker demand for training in Equation 9. Again, plugging in our functional form assumptions yield:

$$\begin{aligned} B_i^* &\approx C \times \frac{\partial U(w^*(q_i, a_i), a_i)}{\partial w} \times \frac{dw^*(q_i, a_i)}{dq} \times \frac{dq_i}{dT} \\ \Rightarrow B_i^* &= C \times \rho_i \times w_i^{\rho_i - 1}\left(q_i^{\rho_i - 1}\left(q_i^{\rho_i} - \tau \times a_i\right)^{\frac{1-\rho_i}{\rho_i}}\right) \times \frac{dq_i}{dT} \end{aligned} \tag{15}$$

where we can replace $w_i^*(q_i, a_i)$ with w_i in the first term because we are at equilibrium.[15]

Given the data we have collected on worker wages (w_i), attachments (a_i), and risk aversion (ρ_i), we are able to estimate τ directly from baseline wages. This provides a complete parameterization

[15] As in Footnote 9, we can simplify this equation by leveraging the binding incentive compatibility constraint to get ($B_i^* = C \times \rho_i \times q_i^{\rho_i - 1} \times \frac{dq_i}{dT}$). Given that we observe quality, wages, and attachments directly, estimating this simpler model is more dependent on measurement error in our estimates of ρ_i and $\frac{dq_i}{dT}$, while the fuller model diversifies our sources of variance.

of our model and allows us to test for the fit of the model in predicting employer and worker demand.

8.1 Graphical Model

We graphically illustrate the wage-setting problem under our proposed utility function to build intuition for how worker and employer preferences for training will diverge. Specifically, we present a set of indifference curves that capture the way in which employers may leverage firm attachments to minimize wage adjustments. This, in turn, helps predict which workers are most desirable for an employer to send for training, because the same leverage that suppresses baseline wages can be used to suppress wage adjustments from training.

In Figures 1 and 2, we present indifference curves for the utility function $U(w_i, a_i) = w^{1/2} + (3/2) \times a_i$ at various levels of utility. Figure 1 presents the situation of a worker with firm attachment $a_i = 0.2$, while Figure 2 presents a similar situation for a worker with firm attachment $a_i = 0.4$. Consider Figure 1 and suppose that this worker possesses skills commensurate with a market wage of $w_i = 1$. Departing for a new firm at this wage (thus, resetting firm attachment to $a_i = 0$) would place the worker's utility along the blue indifference curve. To retain this worker, the incumbent firm would need to just match this utility by providing a retention wage of approximately $w_i = 0.45$. Now, consider Figure 2 and suppose this more-attached worker were also more skilled, with a market wage of approximately $w_i = 1.8$. Because of the greater attachment, this worker would require no higher of a retention wage than the lower-skilled, less-attached worker.

Recall that a worker's market wage is determined by their marginal revenue product and notice that changes in market wages need not increase retention wages for sufficiently attached workers. For these attached workers, improvements in marginal revenue product through training may be

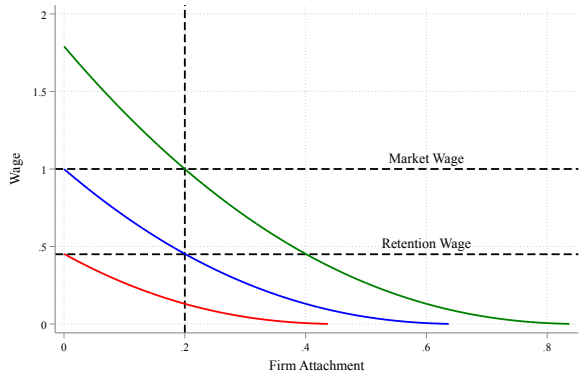


Figure 1: Retention wage (low attachment)

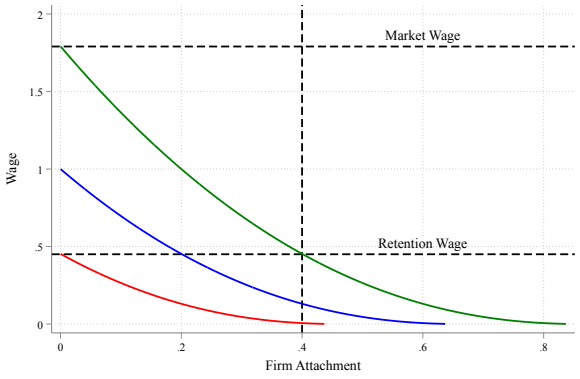


Figure 2: Retention wage (high attachment)

partially or fully captured by the firm. Thus, employer demand for training will be critically dependent on anticipated wage adjustments that are, in turn, dependent on firm attachments.

8.2 Estimating the Structural Parameters and Testing the Model Fit

In our data, we observe all of the key features of our model—wages, skills, anticipated skill gains, attachment to the firm, and risk preferences. We use these to 1) estimate the relative influence of firm attachments in determining the selection of workers for training, 2) document the explanatory power of our model in capturing the selection behaviors of employers and workers, and 3) provide simulated policy analysis showing training selections under different counterfactuals.

We first estimate the role of firm attachments in determining employer demand for training a given worker. Recall that, in our model, the profits associated with training a given worker are determined as a function of quality (q_i), firm attachments (a_i), anticipated skill gains ($\frac{dq_i}{dT_i}$), risk aversion (ρ_i), and the weight on attachments (τ):

$$\frac{\partial}{\partial T_i}\Pi_i = \frac{dq_i}{dT}\left[1 - q_i^{\rho_i - 1}\left(q_i^{\rho_i} - \tau \times a_i\right)^{\frac{1-\rho_i}{\rho_i}}\right] \tag{14}$$

We observe measures of each variable except for τ. For quality, we use the employer's beliefs about a worker's baseline skill and rescale it such that $q_i \in [0, 1]$. For firm attachments, we use a principal components analysis to extract one component from the following survey questions about the worker: 1) are they related to the employer, 2) did the employer know them before hiring them, 3) how much does the employer trust them, 4) how long has the worker worked at this firm, 5) how likely is this worker to leave the firm, and 6) how many months does the employer expect the worker to remain at the firm after training. This attachment variable is measured as a percentile between zero and one. For anticipated skill gains, we use the elicited value from the employer survey for each worker normalized within firm. For risk aversion, we also use principal components analysis to extract one variable out of two survey questions about the worker's self-reported willingness to take risks. This is also rescaled to fall between zero and one.

We perform a non-linear least squares estimation of τ to establish the role of firm attachments in determining employer demand for training a given worker. Alongside τ, we estimate a new scale factor, S'. To do so, we regress the observed employer ranking of each worker for selection to training (in percentile terms) on the profit function of Equation 14 after substituting in our rescaled values of quality, firm attachments, anticipated skill gains, and risk aversion. Since our risk aversion parameter is missing for a large fraction of the data, we impute it using predicted values from a

linear regression of observed risk aversion on worker characteristics.

Panel A of Table 16 presents our estimates. We find a substantial effect of firm attachments with $\hat{\tau} = 0.128$ ($p < 0.001$). Our scale parameter estimate, $S' = 0.982$, simply ensures that we can match the profit function to the ranking percentile, which falls between zero and one.

We measure the fit of this model in Panel B. We estimate the profitability of training a given worker based on Equation 14 using observed data on quality, firm attachments, anticipated skill gains, risk aversion, and the estimated τ parameter from Panel A. We regress the observed employer ranking of each worker for selection to training (in percentile terms) on these respective estimates of profit. Panel B of Table 16 shows that, while our estimated profits have limited explanatory power, they nonetheless positively and significantly predict employer demand.

We also break down the estimated profit into its two key components: anticipated gains from training and wage adjustment. These are reflected in the two terms of the expanded profit function: $\frac{\partial}{\partial T_i}\Pi_i = \frac{dq_i}{dT} - \frac{dq_i}{dT}\left[\beta(q_i, a_i, \rho_i)\right]$ where the wage adjustment function $\beta(q_i, a_i, \rho_i)$ is equal to $q_i^{\rho_i - 1}\left(q_i^{\rho_i} - \tau \times a_i\right)^{\frac{1-\rho_i}{\rho_i}}$. We find that the predicted impact of each term matches our model, with selection for training increasing in anticipated skill gains and decreasing in the combination of skill gains and wage adjustments. That is, while skill gains are valuable to an employer, all else equal, they become less valuable when wages are likely to adjust substantially in response to skill gains.

To illustrate the importance that employers place on the wage adjustment dimension, we plot the individual-level estimates of wage adjustments for workers who were and were not selected for training by their employers. Figure 4 presents the CDFs for these worker-level estimates across the two groups. We find that employers over-select workers with low estimated wage adjustments. While the distributions do not have a first-order stochastic dominance relationship, they are significantly different based on Kolmogorov–Smirnov tests ($p < 0.001$).

9 Conclusion

Ineffective worker training is a critical part of the vicious cycle of low-skilled workers, low wages, and low quality goods that characterizes much of the developing world. In this paper, we provide a theoretical model describing how firm attachments impose negative externalities on the effectiveness of such training. In an environment where training in available and free, our model predicts that employers and workers will use opposing criteria to select who should receive training. Using a field experiment, we find that the model's predictions are accurate. Employers willingly forego skill

gains from training in favor of selecting workers based on firm attachments, over-selecting relatives, long-tenured workers, trusted workers, and workers they can retain after training. Employers' selections are regressively targeted, sending workers to training who have higher wages and higher skills at baseline. Workers, on the other hand, either ignore or negatively self-select based on firm attachments. Workers more narrowly self-select based on skill gains from training, with the highest demand coming from low-skilled workers.

Importantly, we demonstrate dramatic skill gains from our training program. Workers assigned to training experience skill gains of nearly one-half of a standard deviation. These increases are corroborated across multiple skill measures and operate at both the extensive and intensive margins. Therefore, the misalignment in selection criteria for employers and workers becomes a critical determinant of training effectiveness.

Based on our results, we conclude that policies to increase workforce training need to be cautious about employer influence over training assignment and attendance. Employer preferences distort the assignment of training and may limit its impact given their incentive to select workers from whom they can extract a larger fraction of the productivity gains. On the other hand, we find that workers reliably self-select for training based on anticipated skill gains. Thus, we expect that aggregate productivity will improve the most when training opportunities and incentives are targeted to workers. An additional benefit of worker-targeted training incentives is its progressive nature—self-selected potential trainees have lower wages and are earlier in their careers than those selected by the employer. While employer selections exacerbate within-firm inequality, worker selections ameliorate it.

References

Abeler, Johannes, Daniele Nosenzo, and Collin Raymond. 2019. "Preferences for truth-telling." *Econometrica*, 87(4): 1115–1153.

Acemoglu, Daron. 1997. "Training and Innovation in an Imperfect Labour Market." *The Review of Economic Studies*, 64(3): 445–464.

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

Acemoglu, Daron, and Jorn-Steffen Pischke. 1999*a*. "Beyond Becker: Training in imperfect labour markets." *The economic journal*, 109(453): 112–142.

Acemoglu, Daron, and Jörn-Steffen Pischke. 1999*b*. "The structure of wages and investment in general training." *Journal of political economy*, 107(3): 539–572.

Alfonsi, Livia, Oriana Bandiera, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali. 2020. "Tackling Youth Unemployment: Evidence From a Labor Market Experiment in Uganda." *Econometrica*, 88(6): 2369–2414.

Altonji, Josepth, and Paul J Devereux. 1999. "The extent and consequences of downward nominal wage rigidity."

Bassi, Vittorio, Jung Hyuk Lee, Alessandra Peter, Tommaso Porzio, Ritwika Sen, and Esau Tugume. 2023. "Self-Employment Within the Firm." National Bureau of Economic Research Working Paper 31740.

Becker, Gary S. 1962. "Investment in Human Capital: A Theoretical Analysis." *Journal of Political Economy*, 70(5, Part 2): 9–49.

Bewley, Truman F. 1999. *Why Wages Don't Fall during a Recession.* Harvard University Press.

Brown, Gabriel, Morgan Hardy, Isaac Mbiti, Jamie McCasland, and Isabelle Salcher. 2024. "Can Financial Incentives to Firms Improve Apprenticeship Training? Experimental Evidence from Ghana." *American Economic Review: Insights*, 6(1): 120–36.

Caicedo, Santiago, Miguel Espinosa, and Arthur Seibold. 2022. "Unwilling to train?—firm responses to the colombian apprenticeship regulation." *Econometrica*, 90(2): 507–550.

Card, David, Pablo Ibarrarán, Ferdinando Regalia, David Rosas-Shady, and Yuri Soares. 2011. "The Labor Market Impacts of Youth Training in the Dominican Republic." *Journal of Labor Economics*, 29(2): 267–300.

Cefala, Luisa, Pedro Naso, Michel Ndayikeza, and Nicholas Swanson. 2025. "Under-training by employers in informal labor markets: Evidence from Burundi." Working paper.

Charness, Gary, and Martin Dufwenberg. 2006. "Promises and partnership." *Econometrica*, 74(6): 1579–1601.

Cho, Yoonyoung, Davie Kalomba, Ahmed Mushfiq Mobarak, and Victor Orozco. 2013. *Gender Differences in the Effects of Vocational Training: Constraints on Women and Drop-Out Behavior.* The World Bank.

Danz, David, Lise Vesterlund, and Alistair J Wilson. 2022. "Belief elicitation and behavioral incentive compatibility." *American Economic Review*, 112(9): 2851–2883.

Dautheville, Adrien. 2025. "Job-Seeking Entrepreneurs." Unpublished. Available at https://drive.google.com/file/d/1xXBviUDuIbPkDSmrsT_WGmXciWaRzjok/view.

Elsby, Michael WL, and Gary Solon. 2019. "How prevalent is downward rigidity in nominal wages? International evidence from payroll records and pay slips." *Journal of Economic Perspectives*, 33(3): 185–201.

Frazer, Garth. 2006. "Learning the master's trade: Apprenticeship and human capital in Ghana." *Journal of Development Economics*, 81(2): 259–298.

Ghanem, Dalia, Sarojini Hirshleifer, and Karen Ortiz-Becerra. 2023. "Testing attrition bias in field experiments." *Journal of Human resources*.

GoU. 2015*a*. "Second National Development Plan (NDPII) 2015/16 – 2019/20." Government of Uganda, National Planning Authority. Uganda Vision 2040: "A Transformed Ugandan Society from a Peasant to a Modern and Prosperous Country within 30 years".

GoU. 2015*b*. "Uganda Micro, Small and Medium Enterprise (MSME) Policy: Sustainable MSMEs for Wealth Creation and Socio-Economic Transformation." Ministry of Trade, Industry and Cooperatives (MTIC), Kampala, Uganda.

Government of Uganda. 2018. "Trade and Industry Sector Statistical Abstract." Ministry of Trade, Industry and Cooperatives (MTIC), Kampala, Uganda.

Hirshleifer, Sarojini, David McKenzie, Rita Almeida, and Cristobal Ridao-Cano. 2016. "The Impact of Vocational Training for the Unemployed: Experimental Evidence from Turkey." *The Economic Journal*, 126(597): 2115–2146.

Hsieh, Chang-Tai, and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India*." *The Quarterly Journal of Economics*, 124(4): 1403–1448.

International Labour Organization. 2021. "ILOSTAT Country Profile: Uganda." *https://ilostat.ilo.org/data/country-profiles/uga/*, Accessed March 2026.

Johanson, Richard K., and Arvil V. Adams. 2004. *Skills Development in Sub-Saharan Africa. World Bank Regional and Sectoral Studies*, Washington, DC:World Bank.

La Porta, Rafael, and Andrei Shleifer. 2014. "Informality and Development." *Journal of Economic Perspectives*, 28(3): 109–26.

Liaqat, Zara, and Jeffrey B Nugent. 2015. "Under-provision of private training by MENA firms: what to Do about It?" *IZA Journal of Labor & Development*, 4: 1–29.

Madarasz, Kristof, David Danz, and Stephanie Wang. 2023. "Projective Thinking: Model, Evidence, and Applications."

McKenzie, David. 2017. "How Effective Are Active Labor Market Policies in Developing Countries? A Critical Review of Recent Evidence." *The World Bank Research Observer*, 32(2): 127–154.

McKenzie, David, and Christopher Woodruff. 2017. "Business practices in small firms in developing countries." *Management Science*, 63(9): 2967–2981.

Moore, Ryan T., and Sally A. Moore. 2013. "Blocking for Sequential Political Experiments." *Political Analysis*, 21(4): 507–523.

Oya, Carlos, Florian Schaefer, and Wei Chen. 2022. "Management Workforce Localization and Skill Development in Ethiopia's Light Manufacturing." IDCEA Research Report, SOAS, University of London.

Prendergast, Canice. 1993. "The Role of Promotion in Inducing Specific Human Capital Acquisition." *The Quarterly Journal of Economics*, 108(2): 523–534.

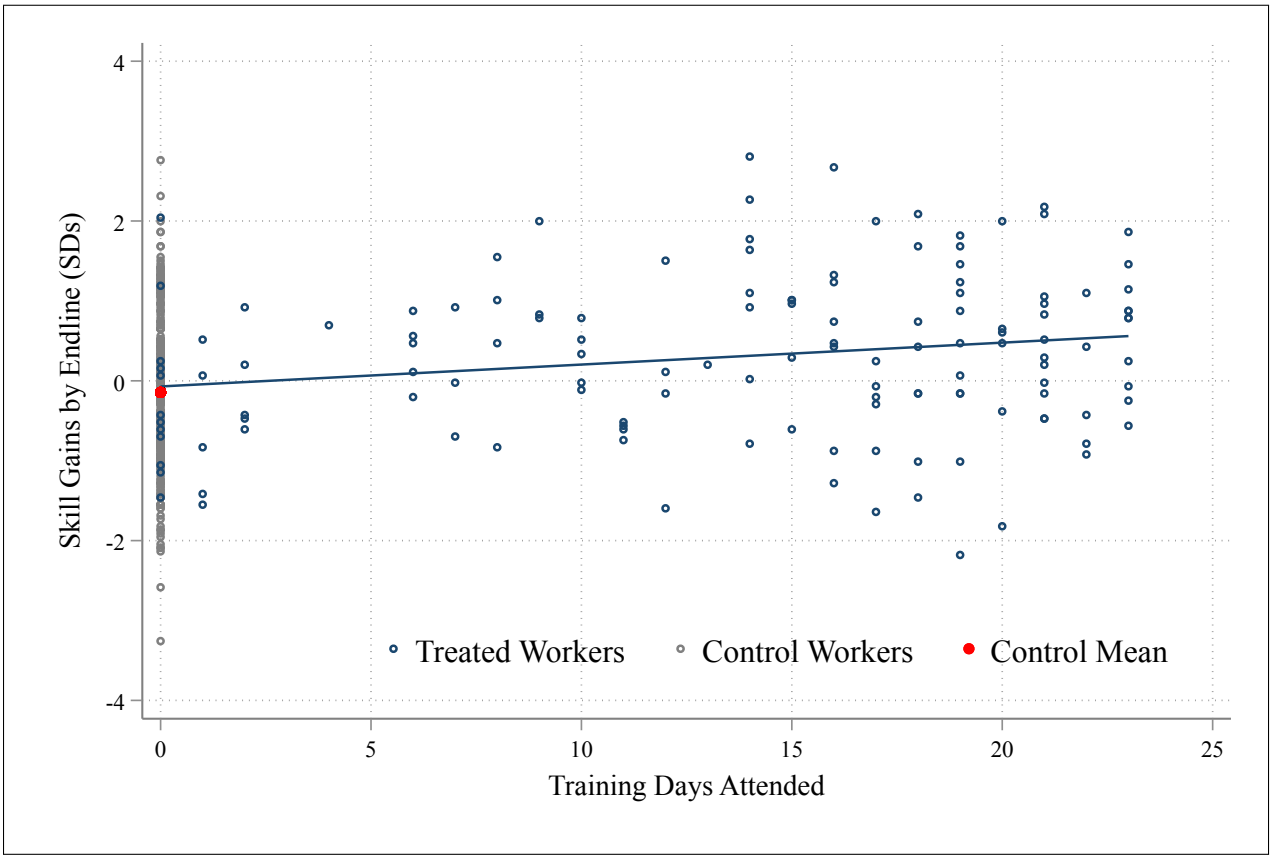
Schaffner, Julie Anderson. 2001. "Job stability in developing and developed countries: Evidence from Colombia and the United States." *Economic Development and Cultural Change*, 49(3): 511–535.

Stern, David, and Jozef MM Ritzen. 2012. *Market failure in training?: New economic analysis and evidence on training of adult employees.* Springer Science & Business Media.

Teal, Francis. 2016. "Are apprenticeships beneficial in sub-Saharan Africa?" *IZA World of Labor*, 2016: 268.

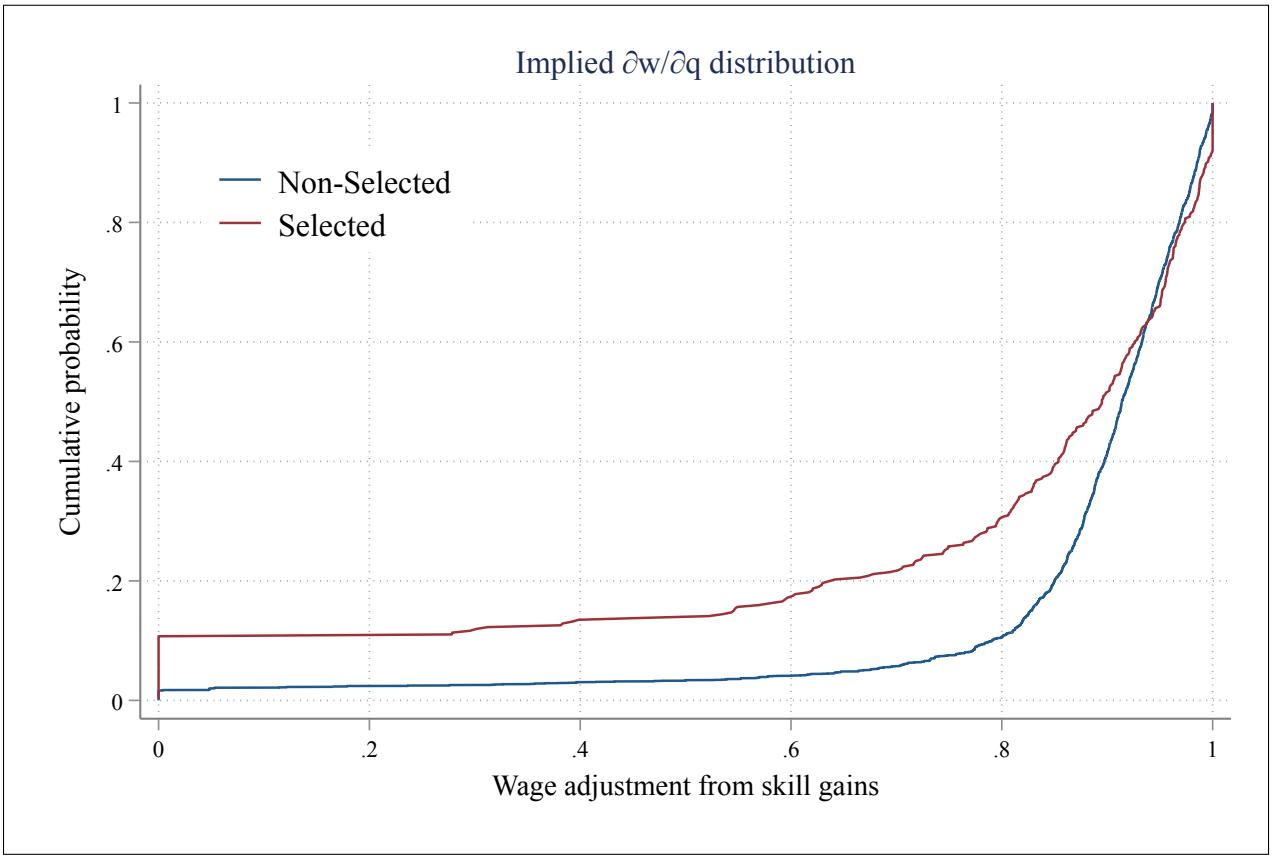
World Bank. 2012. *World Development Report 2013: Jobs.* Washington, DC:World Bank.

Figure 3: Training Attendance Increases Skill Gain.



The X-axis shows the number of days attended by workers assigned to the treatment group, while the y-axis is the skill gain from training. Additionally, the figure plots the line of best fit.

Figure 4: Estimated Wage Adjustment from Training for Selected and Non-Selected Workers



The x-axis is the worker-level estimate of how much wages will adjust after training. Workers are divided based on whether the employer selected them to receive training (i.e. ranked them first on their preference ranking) or not.

Table 1: Employer/Firm-level Descriptive Statistics

	N	Mean	SD	Min	Max
Age	336	36.5	8.9	19	73
Years in metal fabrication	337	14.3	7.9	0	49
Completed formal metal fabrication training	275	0.2	0.4	0	1
Education (years)	336	10.4	4.2	0	22
Would shut down business for a job offer (wage=profits)	337	0.1	0.2	0	1
Firm age	330	10.3	7.9	0	51
Number of employees	337	5.4	1.9	4	18
New employees hired in the last 12 months	336	2.3	3.4	0	30
Employees who have quit in the last 12 months	337	1.0	1.7	0	12
Introduced new products in the last 12 months	337	0.5	0.5	0	1
Profits in the last 30 days (USD)	337	425.1	592.2	0	3,333
Revenue in the last 30 days (USD)	337	1976.0	2590.4	0	12,500
Employees specialize on tasks	337	0.3	0.5	0	1
Tendency to avoid risks	337	0.4	0.5	0	1
Very Likely still in business in 5 years	337	0.8	0.4	0	1
Expected number of employees in 5 years	333	11.4	8.3	2	70
Expected number of similar firms (out of 10) in 5 years	333	6.9	2.6	1	10
Ever received a loan	337	0.2	0.4	0	1
Ever applied for a loan to train employees	337	0.0	0.2	0	1
Ever received for a loan to train employees	337	0.0	0.1	0	1
Change wages after working	337	0.9	0.3	0	1

Notes: Data are from the employer survey. Unless specified, the variable is binary whenever the min is zero and the max is one, such as "Completed formal metal fabrication training". Ever received loan refers to any other type of loan irrespective of training loan. We interviewed businesses if they have ever received any loan from any commercial bank, microfinance institution, savings cooperative, or other certified institutions separately from if they have ever applied or received a loan to train their workers specifically. Four employers did not have expectation about the number of workers and number of businesses that will be in business in five years. Seven employers did not know how long the firm had been open. A coding error limited the sample size for the question about completing formal training.

Table 2: Worker-level Descriptive Statistics

	N	Mean	SD	Min	Max
Panel A: Employer Survey					
Worker Related to Employer	1,784	0.2	0.4	0	1
Worker Tenure in Years	1,784	3.4	3.1	0	34
Likelihood of leaving (0–10)	1,782	2.8	2.5	0	10
Months (out of 24) expected to stay after training	1,739	20.0	7.0	1	24
Monthly wage (USD)	1,617	159.1	167.5	0	889
Panel B: Worker Survey					
Age	674	24.1	5.2	18	53
Female	674	0.0	0.1	0	1
Years in metal fabrication	673	4.5	3.8	0	23
Years at firm	573	4.5	5.0	0	50
Education	673	8.3	3.1	0	22
Completed vocational training	674	0.0	0.2	0	1
Ever under apprentice at firm	674	0.8	0.4	0	1
Days at the firm per week	674	5.8	1.3	0	7
Days at the firm last month	673	22.6	7.4	0	30
Worked for other firm last month	604	0.2	0.4	0	1
Shares about job opportunities with different firm people	674	0.5	0.5	0	1

Notes: Data in Panel A are from the employer survey and cover all workers in a firm. Data in Panel B are from the worker survey.
An average education of 8.34 years is equivalent to completing one and a half years of secondary education. Some workers did not remember when they started at the firm. Ever under apprentice at the firm includes both employees who were currently at the firm under apprenticeship or started and completed an apprenticeship at the firm (with out pay). We asked workers if they share/discuss job opportunities with workers from different firms.

Table 3: Descriptive Statistics: Training Demand, Worker Quality, and Training Impact

	N	Mean	SD	Min	Max
Panel A: Worker Demand for Training					
Willingness to accept training (USD)	1,351	14.5	10.0	0	33
Panel B: Employer Beliefs about Worker Quality					
Perceived baseline quality (0-100)	1,782	59.7	24.6	0	100
Perceived endline quality if untrained (0-100)	1,783	70.2	22.4	0	100
Perceived endline quality if trained (0-100)	1,778	81.1	18.9	0	100
Panel C: Worker Beliefs about Own Quality					
Perceived baseline quality (0-100)	674	73.6	17.6	13	100
Perceived endline quality if untrained (0-100)	674	84.0	14.2	20	100
Perceived endline quality if trained (0-100)	673	94.6	8.3	50	100
Anticipated average quality of a trained worker (0-100)	672	75.5	14.5	30	100
Panel D: Baseline Worker Practical Skills Test Scores					
Process score (0-100)	537	44.1	21.9	5	95
Product score (0-100)	537	47.6	13.1	0	80
Practical score (0-100)	537	45.9	14.6	2	81
Knowledge score (%)	598	36.5	7.9	16	64
Measured deviation of final product (%)	531	6.2	11.9	0	100
Panel E: Endline Worker Practical Skills Test Scores					
Process score (0-100)	448	62.0	26.7	8	98
Product score (0-100)	448	50.9	15.0	0	93
Practical score (0-100)	448	56.2	17.7	7	94
Knowledge score (%)	552	43.8	7.1	26	71
Measured deviation of final product (%)	445	3.2	5.4	0	30
Panel F: Effect of training on profits					
Training will increase profits	337	0.8	0.4	0	1
Because of increased product quality	276	1.0	0.2	0	1
Because of reduced wastage	276	0.4	0.5	0	1
Because trained workers will train others	276	0.3	0.5	0	1
Training will not increase profits	337	0.2	0.4	0	1
Because of increased wages	61	0.0	0.2	0	1
Because product quality already optimal	61	0.5	0.5	0	1
Because training happens at firm	61	0.2	0.4	0	1
Panel G: Worker Training Attendance					
Attended training at least one day	182	0.8	0.4	0	1
Training days completed	182	10.6	8.5	0	23
Completed training (> 70% attendance)	182	0.4	0.5	0	1

Notes: Our sample size in Panel A falls short of that in Panel B because we only elicited WTA bids from workers who were present at the firm to prevent coordination (see Section 3.3 for details). Data in Panel B were collected from the employer survey and cover all workers from all firms; Panel C data come from the worker survey. Data in Panels D and E come from the baseline and endline practical skills tests, respectively. Measured deviations are evaluated as the negative percent deviation, so that larger numbers indicate greater accuracy. Data in Panel F were collected on the employer survey. For employers who indicated that training would (or would not) increase profits, we included a follow-up question to gather their reasons. A total of 276 owners believed that training would increase profits, while 61 did not. Panel G presents attendance statistics for potential trainees from the treatment group only. Training completion corresponds to attending at least 70% of the course days.

Table 4: Employer/Firm-level Balance Table

	Control	Treatment	P value (Diff in means)
Age	36.7	36.0	0.72
	(8.9)	(8.8)	
Years in metal fabrication	14.4	13.9	0.96
	(7.8)	(8.1)	
Completed formal metal fabrication training	0.2	0.2	0.30
	(0.4)	(0.4)	
Education (years)	10.2	10.8	0.12
	(4.1)	(4.3)	
Would shut down business if offered a job (wage = profits)	0.1	0.0	0.47
	(0.2)	(0.2)	
Firm age	10.5	10.0	0.83
	(8.0)	(7.8)	
Number of workers	5.5	5.2	0.62
	(2.0)	(1.6)	
New employees hired in the last 12 months	2.2	2.4	0.82
	(3.6)	(2.8)	
Employees who have quit in the last 12 months	1.0	1.1	0.77
	(1.7)	(1.9)	
Introduced new products in the last 12 months	0.5	0.4	0.36
	(0.5)	(0.5)	
Profits in the last 30 days (USD)	421.1	411.6	0.69
	(544.1)	(610.9)	
Revenue in the last 30 days (USD)	1970.5	1908.5	0.58
	(2468.4)	(2612.7)	
Employees specialize on tasks	0.3	0.3	0.56
	(0.5)	(0.5)	
Often takes risks	0.4	0.4	0.88
	(0.5)	(0.5)	
Very likely still in business in 5 years	0.8	0.9	0.31
	(0.4)	(0.3)	
Expected number of employees in 5 years	11.7	10.8	0.40
	(8.3)	(8.0)	
Expected number of similar firms (out of 10) in 5 years	7.0	6.6	0.21
	(2.6)	(2.6)	
Ever applied for a loan to train employees	0.0	0.0	0.47
	(0.2)	(0.1)	
Ever received a loan	0.2	0.3	0.21
	(0.4)	(0.5)	
Change wages after working	0.9	0.9	0.86
	(0.3)	(0.3)	
Firms	246	91	

Notes: This table shows the balance of variables on the employer survey between the treatment and control groups. The first two columns present the means with standard deviations in parentheses. The final column reports the p-values from the tests of equality between the treatment and control group means.

Table 5: Worker-level Balance Table (Demographics)

	Control	Treatment	P value (Diff in means)
Panel A: Data from Employer Survey			
Worker is a relative of employer	0.2	0.2	0.19
	(0.4)	(0.4)	
Worker Tenure in Years	3.6	3.1	0.49
	(3.2)	(2.7)	
Worker likelihood of leaving	2.7	2.9	0.31
	(2.5)	(2.5)	
Months (out of 24) worker is expected to stay after training	19.9	20.2	0.62
	(7.1)	(6.8)	
Worker monthly wage (USD)	168.9	132.0	0.05
	(177.6)	(131.2)	
Panel B: Data from Worker Survey			
Age	24.1	23.8	0.41
	(5.3)	(4.5)	
Years in metal fabrication	6.5	5.8	0.37
	(4.6)	(3.9)	
Years at the firm	3.4	3.1	0.67
	(3.0)	(2.4)	
Education	8.2	8.8	0.06
	(3.1)	(3.1)	
Completed vocational training	0.0	0.0	0.63
	(0.2)	(0.2)	
Ever under apprentice at firm	0.8	0.8	0.56
	(0.4)	(0.4)	
Days at the firm per week	5.8	5.7	0.05
	(1.3)	(1.2)	
Days at the firm last month	22.7	22.3	0.48
	(7.3)	(7.3)	
Worked for another firm last month	0.2	0.2	0.58
	(0.4)	(0.4)	
Discusses job opportunities with people from other firms	0.5	0.5	0.34
	(0.5)	(0.5)	

Notes: This table shows the balance of variables from employer and worker surveys between the treatment and control groups. In Columns 1 & 2, we present the means with standard deviations in parentheses. Column 4 reports the p-values from the tests of equality between the treatment and control group means. These tests control for strata fixed effects, and standard errors are clustered at the firm level

Table 6: Worker-level Balance Table (Training Demand and Skills)

	Control	Treatment	P value (Diff in means)
Panel A: Worker Demand for Training			
Willingness to accept training (USD)	14.5	14.8	0.57
	(9.8)	(10.6)	
Panel B: Employer Beliefs about Worker Quality			
Perceived baseline quality (0-100)	59.4	60.6	0.62
	(24.9)	(23.8)	
Perceived endline quality if untrained (0-100)	69.7	71.6	0.27
	(22.7)	(21.6)	
Perceived endline quality if trained (0-100)	80.6	82.6	0.17
	(19.3)	(17.7)	
Panel C: Worker Beliefs about Own Quality			
Perceived baseline quality (0-100)	74.0	72.5	0.22
	(17.8)	(17.2)	
Perceived endline quality if untrained (0-100)	84.1	83.6	0.52
	(14.5)	(13.4)	
Perceived endline quality if trained (0-100)	94.8	94.2	0.33
	(8.5)	(8.0)	
Panel D: Baseline Worker Practical Skills Test Scores			
Process score (0-100)	0.4	0.4	0.94
	(0.2)	(0.2)	
Product score (0-100)	0.5	0.5	0.35
	(0.1)	(0.1)	
Overall score (0-100)	0.5	0.5	0.71
	(0.1)	(0.1)	
Knowledge score (%)	0.4	0.4	0.83
	(0.1)	(0.1)	
Measured deviation of final product (%)	0.1	0.1	0.61
	(0.1)	(0.1)	

Notes: This table shows the balance of the baseline variables related to beliefs about quality and scores from practical tests. In Columns 1 & 2, we present the means with standard deviations in parentheses. Column 4 reports the p-values from the tests of equality between the treatment and control group means. These tests control for strata fixed effects, and standard errors are clustered at the firm level.

Table 7: Impact of Assignment to Training on Endline Skills (ITT)

	Endline Practical Skills Test Scores (SDs)				Endline	Endline Measured
	Overall Score		Process	Product	Quiz	Deviations
Treatment	0.484***	0.483***	0.529***	0.236*	0.556***	0.372***
	(0.11)	(0.10)	(0.10)	(0.12)	(0.12)	(0.08)
Baseline Practical Skills Test Score (SDs)		0.535***				
		(0.06)				
Baseline Score Missing		1.040***				
		(0.28)				
Control Mean	0.539	0.539	0.579	0.501	0.430	0.039
R-Squared	0.218	0.351	0.227	0.127	0.151	0.131
Workers	448	448	448	448	552	445

Column 2 controls for baseline practical skills test scores and accounts for missing data by imputing those scores as the minimum value. "Baseline Score Missing" is an indicator variable equal to one if the baseline practical skills test score is missing.
To calculate "Measured Deviations," we take the negative absolute value of the deviations of the measurements of the perimeter of the final product from the blueprint and standardized it. Thus, all values are denominated in standard deviations and positive values indicate greater accuracy.
All columns report results of ordinary least squares regressions with indicator variables for randomization strata.
Standard errors clustered at the firm level in parentheses.
Standardized baseline test scores are z-scores computed using the sample mean and standard deviation, while standardized endline test scores are z-scores computed using the mean and standard deviation of the control group.
* p<0.10, ** p<0.05, *** p<0.01

Table 8: Impact of Assignment to Training Split by Employer-Selected and Worker-Selected

	Practical Skills Test Scores (SDs)		Quiz		Measured Deviations	
Selected By:	Employer	Worker	Employer	Worker	Employer	Worker
Treatment	0.398***	0.488***	0.639***	0.646***	0.241**	0.429***
	(0.142)	(0.134)	(0.165)	(0.153)	(0.113)	(0.112)
R-squared	0.270	0.292	0.229	0.203	0.233	0.218
Workers	232	220	280	273	230	217

Columns 1,3,5 estimate treatment effects for employer-selected potential trainees, while Columns 2,4,6 estimate treatment effects for worker-selected potential trainees.
All columns report results of ordinary least squares regressions with indicator variables for randomization strata.
Standard errors clustered at the firm level in parentheses.
Standardized practical test scores are z-scores computed using the mean and standard deviation of the control group.
* p<0.10, ** p<0.05, *** p<0.01

Table 9: Impact of Anticipated Skill Gains on Training Selection

	Employer Demand	Profitability	Worker Demand
Panel A: Anticipated Skill Gains Elicited from the Employer			
Anticipated Skill Gains (SDs)	-0.036***	-0.073***	0.054*
	(0.01)	(0.01)	(0.03)
R-Squared	0.010	0.043	0.003
Workers	1,778	1,726	1,347
Panel B: Anticipated Skill Gains Stated by the Worker			
Anticipated Skill Gains (SDs)	-0.071***	-0.058***	0.099*
	(0.02)	(0.02)	(0.05)
R-Squared	0.053	0.035	0.012
Workers	670	647	642

Fixed-effects regressions with firm-level fixed-effects.
Standard errors clustered at the firm level in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

Table 10: Impact of Firm Attachments on Training Selection

	Employer Demand	Profitability	Worker Demand (SDs)
Relative of Employer	0.115***	0.120***	0.069
	(0.02)	(0.03)	(0.08)
R-squared	0.017	0.018	0.001
Workers	1,784	1,732	1,351
Employer's Trust	0.077***	0.108***	-0.064**
	(0.01)	(0.01)	(0.03)
R-squared	0.063	0.127	0.006
Workers	1,784	1,732	1,351
Worker Tenure (Years)	0.009***	0.020***	-0.028**
	(0.00)	(0.00)	(0.01)
R-squared	0.007	0.037	0.008
Workers	1,784	1,732	1,351
Retention after Training (months)	0.017***	0.012***	0.007
	(0.00)	(0.00)	(0.01)
R-squared	0.057	0.032	0.001
Workers	1,739	1,691	1,319

Each row presents three separate fixed-effects regressions with firm-level fixed-effects.
Standard errors clustered at the firm level in parentheses.
Employer Demand and Profitability are measured as percentiles. Worker Demand is the willingness to accept training assignment standardized across all workers in the study.
"Relative of Employer" is an indicator variable equal to one if the worker is related to the employer. Trust is measured on a 0-10 Likert scale asking for the extent to which the owner can depend on the worker to do assigned work reliably. Worker tenure is measured in years and captures the time the worker has been associated with the current firm. Retention after training measures the number of months that the employer believes a worker would remain at the firm after receiving training.
Sample sizes change because of missing data for certain variables.
* p<0.10, ** p<0.05, *** p<0.01

Table 11: Impact of Baseline Wages and Skills on Training Demand (Winsorized Wages)

	Employer Demand			Profitability			Worker Demand (SDs)		
Worker Baseline Monthly	0.307***			0.457***			-0.157		
Wages (in 1000 USD)	(0.07)			(0.08)			(0.26)		
Employer Belief about		0.105***			0.141***			-0.069**	
Baseline Skill (SDs)		(0.01)			(0.01)			(0.03)	
Baseline Practical Skills			0.095***			0.098***			-0.188***
Test Score (SDs)			(0.02)			(0.02)			(0.06)
R-squared	0.018	0.093	0.083	0.040	0.171	0.085	0.000	0.005	0.036
Workers	1,620	1,782	534	1,572	1,730	515	1,233	1,349	516

Fixed-effects regressions with firm-level fixed-effects.
Standard errors clustered at the firm level in parentheses.
The top 1% of wages are winsorized (i.e. replaced with the 99th percentile wage).
Standardized baseline practical test scores are z-scores computed using the mean and standard deviation of the control group.
* p<0.10, ** p<0.05, *** p<0.01

Table 12: Beliefs about Worker Quality

	Baseline (SDs)						Endline (SDs)
	Practical Skills Test		Quiz		Measured Deviation		Practical Skills Test
Employer Belief about	0.623***	0.595***	0.212**	0.216**	0.342***	0.354***	
Baseline Skill (SDs)	(0.06)	(0.07)	(0.08)	(0.09)	(0.09)	(0.09)	
Worker Belief about		0.063		-0.008		0.011	
Baseline Skill (SDs)		(0.06)		(0.07)		(0.05)	
Baseline Practical							0.306***
SkillsTest Score (SDs)							(0.10)
Employer Beliefs about							0.250**
Endline Skills (SDs)							(0.10)
R-squared	0.291	0.295	0.026	0.026	0.149	0.158	0.202
Workers	536	536	597	597	530	529	373

Fixed-effects regressions with firm-level fixed-effects.
Standard errors clustered at the firm level in parentheses.
Measured deviations are evaluated as the negative of the percent deviation (in absolute value) so that larger numbers indicate greater accuracy.
Standardized baseline practical test scores are z-scores computed using the sample mean and standard deviation, while Standardized endline test scores are z-scores computed using the control group's mean and standard deviation.
Employer beliefs about endline skills are conditional on treatment assignment. For treatment group workers, this is the employer's belief about worker skills with training. For control group workers, this is the employer's belief about worker skills without training.
* p<0.10, ** p<0.05, *** p<0.01

Table 13: Worker Demand and Training Attendance

	Training Days Completed	Worker Anticipated Skill Gains (SDs)
Worker Demand (SDs)	2.166**	0.110***
	(0.98)	(0.04)
R-squared	0.063	0.010
Workers	83	642

Ordinary least squares regressions.
Standard errors clustered at the firm level in parentheses.
* p<0.10, ** p<0.05, *** p<0.01

Table 14: Wage Response to Worker Quality (Winsorized Wages)

	Worker Monthly Wage				
	Full Sample			Relatives	Non-Relatives
Employer Belief about Baseline Skill (SDs)	45.548***			36.136*	45.006***
	(4.82)			(20.93)	(4.67)
Years in metal fabrication		8.070***			
		(1.98)			
Baseline score STD			19.868**		
			(9.79)		
R-squared	0.090	0.055	0.020	0.039	0.089
Workers	1,640	623	503	275	1,365

Fixed-effects regressions with firm-level fixed-effects.
Standard errors clustered at the firm level in parentheses.
The top 1% of wages are winsorized to address likely data-entry errors (i.e. wages above the 99th percentile are replaced with the wage at the 99th percentile). Standardized baseline test scores are z-scores computed using the sample mean and standard deviation.
* p<0.10, ** p<0.05, *** p<0.01

Table 15: Descriptive Statistics: Training Motivations for Employers and Workers

	N	Mean	SD	Min	Max
Employer Survey: Why did you select the worker you selected?					
Selected worker is highest quality	337	0.22	0.42	0	1
Selected worker is trustworthy	337	0.20	0.40	0	1
Selected worker deserves training more	337	0.18	0.38	0	1
Selected worker will teach peers better	337	0.13	0.34	0	1
Selected worker because of other reasons	337	0.27	0.44	0	1
Worker Survey: Gains from training					
Perceived wastage at baseline	664	18.3	19.1	0	180
Perceived wastage if trained	662	10.5	15.7	0	153
Perceived duration (hrs) to make a window at baseline	664	2.7	4.1	0	50
Perceived duration (hrs) to make a window if trained	663	1.6	3.0	0	48
Improvements in wages	664	0.8	0.4	0	1

Notes: Other reasons for the selection of a worker by employers include anticipated gain in skill, anticipated (un)separation from the firm, current skill levels being too low to perform any tasks, the worker's family relationship with the employer, or the worker not having any jobs at the firm, for which the incidence is lower than 10%. Perceived wastage is calculated as the percentage of material a worker wastes out of the total material provided to make a window valued at $42. Three workers reported their wastage level as greater than the value of the window. Perceived duration is the self-reported time a worker estimates they would take to make a window, while 'Improvements in wages' is a dummy variable equal to one if a worker anticipates their wage will increase.

Table 16: Structural Estimates of Weight Given to Firm Attachments and Scale Parameters

	Estimate	SE	95% Conf. Interval	R^2	Workers
Panel A: Worker-Selection Equation					
τ	0.128***	0.021	[0.087 , 0.169]	0.4091	1730
Panel B: Model-Fit (Parameters from Worker-Selection Equation)					
Estimated Profits	0.309***	0.064	[0.183 , 0.436]	0.018	1730
Skill Gains	0.246***	0.067	[0.115 , 0.378]	0.033	1730
Skill Gains×Wage Adjust	-0.406***	0.065	[-0.534 , -0.278]		

In Panel A, estimates are derived from non-linear least-squares regressions. We imputed the risk aversion parameter based on the predicted values of observed risk aversion parameters on worker characteristics: worker tenure and worker training history
In Panel B, estimates are derived from linear regressions with firm-level fixed effects. Estimated profits are derived using the τ parameter from Panel A and imputed risk aversion parameters.
All panels report standard errors clustered at the firm level.
* p<0.10, ** p<0.05, *** p<0.01

A Appendix A

A.1 Timeline and Sample Construction

Our research and data collection process began with an in-person census of metal fabrication firms in the Kampala metropolitan areas coupled with marketing. Our two-person team (enumerator and marketer) visited all firms in our study area, identified those that meet the criteria, and introduced the program. Our marketers were students from nearby vocational schools with the ability to clearly explain the nature of the training, including potential skills and benefits, with the purpose of recruiting businesses in the study area into our training program. We gave business owners an opportunity to opt in by applying.

One month later, business owners who applied and their workers were invited for a 3-4 hour orientation where they interacted with our implementing partner and accredited trainers. The trainers provided a high-level introduction to the curriculum and other training details, such as facilitation on training days. Both listing and orientation were crucial in identifying our evaluation sample by identifying firms that met the number of employees requirement or were located at the training centers.

Figure A.1 highlights the workflow of the sample construction. Every firm that attended the orientation had more information about our training program, making the owner likely to buy into our training program. We also considered the firms located close to the training centers even though they did not attend orientation, as these would incur a small cost (in terms of travel time). As mentioned, our training program covered the financial cost eliminating credit constraints that would affect employee training. We randomized at the firm level after conducting the post-baseline survey. As the figure shows, two employees are selected at each firm through incentive-compatible demand elicitation with the workers or owner elicitation for worker selection. We identify spillover workers as any other workers at the treatment firm whom the owner did not select for training or did not have the highest value for training.

Figure A.1 illustrates our sample selection procedure:

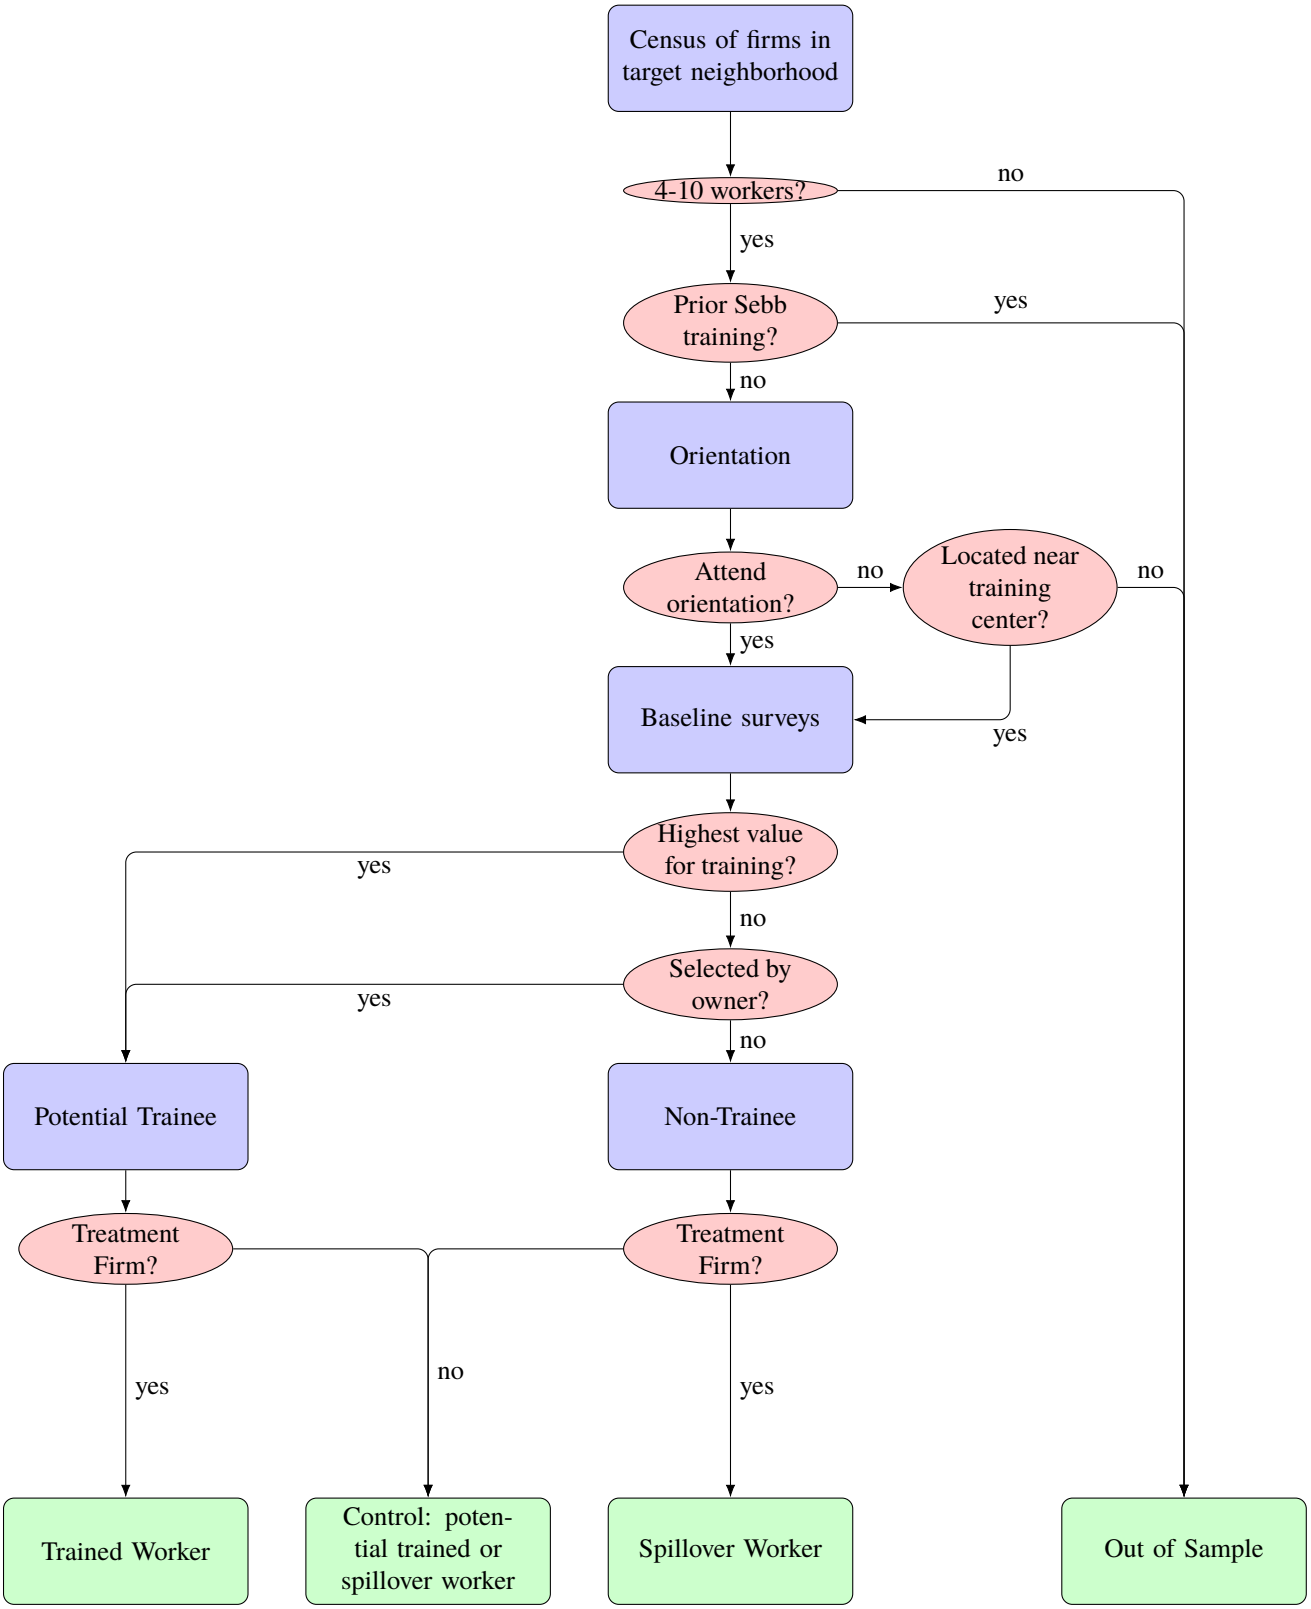

Census of firms in target neighborhood
4-10 workers?
no
yes
Prior Sebb training?
yes
no
Orientation
Attend orientation?
no
Located near training center?
no
yes
yes
Baseline surveys
Highest value for training?
yes
no
Selected by owner?
yes
no
Potential Trainee
Non-Trainee
Treatment Firm?
Treatment Firm?
yes
no
yes
Trained Worker
Control: potential trained or spillover worker
Spillover Worker
Out of Sample

Table A.1: Project Timeline

Oct 2022 - Jan 2023	**Listing and marketing**
	• Introduced our program to the owners and all workers that present at the firm.
	• Identified firms that would qualify to meet the number of employees collected, location, and contact information.
	• Allowed firms an opportunity to apply to participate in our training program.
Feb 2023	**Orientation**
	• Our expert trainer introduced the curriculum to owners.
Jul - Nov 2023	**Baseline survey**
	• Collected demographics from both owner and workers.
	• Incentive-compatible exercises with the owner about perceived worker quality.
	• Collected Owner selection for training.
	• Collected worker demand for training, and thus, worker selection for training.
Aug - Nov 2023	**Baseline practical tests**
	• Objective measure for worker quality at baseline for all workers in our evaluation sample.
Jan 2024	**Treatment assignment**
Feb-May 2024	**Training**
	• Every cohort trains for two days a week and 6-8 hours per day.
May -June 2024	**Endline practical tests**
	• Objective measure for worker quality at endline for all workers in our evaluation sample.

We invited workers from each location whenever we completed the baseline survey from that location as enumerators moved to another location during the baseline survey, which explains the overlap between baseline surveys and baseline practical tests.

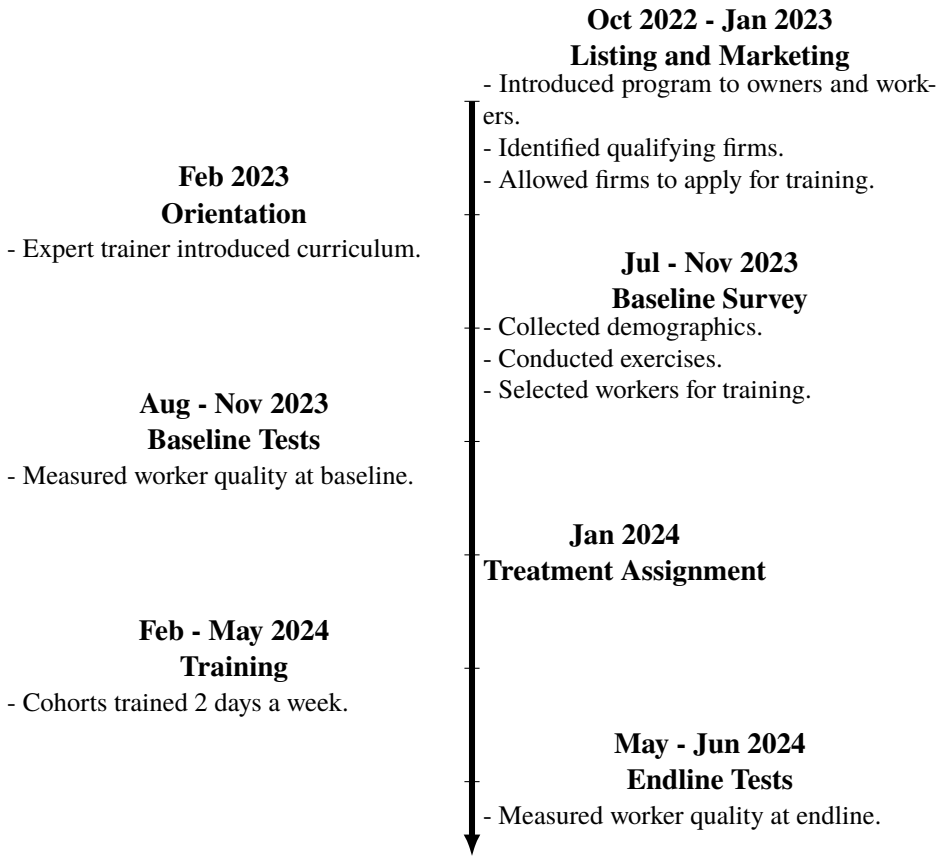


Figure A.1: Project Timeline

A.2 Attrition

Table A.2: Baseline and Endline Practical Test Completion

	Has endline	No endline	Total
All 674 potential trainees			
Has baseline	414	123	537
No baseline	34	103	137
Total	448	226	674
Control (n=492)			
Has baseline	280	86	366
No baseline	33	93	126
Total	313	179	492
Treated (n=182)			
Has baseline	134	37	171
No baseline	1	10	11
Total	135	47	182

Notes: The sample consists of 674 potential trainees across 337 firms. Treated workers complete the endline practical test at a higher rate than control workers (74.2 percent versus 63.6 percent). We fail to reject the null hypothesis of no attrition bias using the regression-based tests of Ghanem, Hirshleifer, and Ortiz-Becerra (2023), with strata fixed effects as appropriate for experiments with more than ten strata (IVal-R p=0.674; IVal-P p=0.280).

Of the 674 potential trainees, 537 completed the baseline practical test and 448 completed the endline practical test, for an overall endline completion rate of 66.5 percent. Table A.2 shows the joint distribution of baseline and endline test completion by treatment status. Treated workers complete the endline at a higher rate than control workers (74.2 percent versus 63.6 percent), a difference of 10.6 percentage points (p=0.021).

To assess whether this differential attrition threatens internal validity, we apply the regression-based attrition tests suggested by Ghanem, Hirshleifer and Ortiz-Becerra (2023). These tests examine whether the baseline practical test score, the baseline measure of our main outcome, is balanced across the four treatment-by-response subgroups. Since our randomization used 33 strata, we implement the strata fixed effects specification, which is appropriate when the number of strata exceeds ten (Ghanem, Hirshleifer and Ortiz-Becerra, 2023, footnote 63). We fail to reject the null hypothesis of internal validity for respondents (IVal-R, p=0.674) and for the full study population (IVal-P, p=0.280). We conclude that differential attrition does not bias our estimates.

B Appendix B: Additional analyses

Table B.3: Impact of Baseline Wages and Skills on Training Demand

	Employer Demand			Profitability			Worker Demand (SDs)		
Worker Baseline Monthly Wages (in 1000 USD)	0.181*** (0.06)			0.246*** (0.08)			0.024 (0.22)		
Employer Belief about Baseline Skill (SDs)		0.105*** (0.01)			0.141*** (0.01)			-0.069** (0.03)	
Baseline Practical Skills Test Score (SDs)			0.095*** (0.02)			0.098*** (0.02)			-0.188*** (0.06)
R-squared	0.011	0.093	0.083	0.021	0.171	0.085	0.000	0.005	0.036
Workers	1,620	1,782	534	1,572	1,730	515	1,233	1,349	516

Fixed-effects regressions with firm-level fixed-effects.
Standard errors clustered at the firm level in parentheses.
Standardized baseline practical test scores are z-scores computed using the mean and standard deviation of the control group.
* p<0.10, ** p<0.05, *** p<0.01

Table B.4: Wage Response to Worker Quality

	Worker Monthly Wage				
	Full Sample			Relatives	Non-Relatives
Employer Belief about Baseline Skill (SDs)	48.622*** (6.65)			21.971 (28.08)	48.472*** (5.70)
Years in metal fabrication		9.028*** (2.55)			
Baseline Practical Skills Test Score (SDs)			18.764* (10.14)		
R-squared	0.058	0.052	0.017	0.007	0.058
Workers	1,640	623	503	275	1,365

Fixed-effects regressions with firm-level fixed-effects.
Standardized baseline test scores are z-scores computed using the sample mean and standard deviation.
Standard errors clustered at the firm level in parentheses.
* p<0.10, ** p<0.05, *** p<0.01