



PEDL Research Papers

This research was partly or entirely supported by funding from the research initiative Private Enterprise Development in Low-Income Countries (PEDL), a Foreign, Commonwealth & Development Office (FCDO) funded programme run by the Centre for Economic Policy Research (CEPR).

This is a PEDL Research Paper which emanates from a PEDL funded project. Any views expressed here are those of the author(s) and not those of the programme nor of the affiliated organizations. Although research disseminated by PEDL may include views on policy, the programme itself takes no institutional policy positions.



Regular Article

It takes two: Experimental evidence on the determinants of technology diffusion[☆]Morgan Hardy^{a,*}, Jamie McCasland^b^a New York University Abu Dhabi, United Arab Emirates^b University of British Columbia, Canada

ARTICLE INFO

JEL classification:

D22
L14
O12
O14
O17
O33

Keywords:

Networks
Diffusion
Technology
Demand
Firms

ABSTRACT

This paper reports on an experiment that brings insights from the literature on demand-side determinants of technology adoption to the study of peer-to-peer diffusion. We develop a custom weaving technique and randomly seed training into a real network of garment making firm owners in Ghana. Training leads to limited adoption among trainees, but little to no diffusion to non-trainees. In a second phase, we cross-randomize demand for the technique. Demand shocks increase adoption of the technology in both groups and diffusion to untrained firms, generated by a pattern in which trained firm owners teach approximately 400% more of their peers if they are randomly assigned to the demand intervention. We find no evidence that our main effects are driven by differences in ability (learning-by-doing) or other adoption-based mechanisms. Rather, our findings are most consistent with the demand intervention generating differential willingness to diffuse among potential teachers.

1. Introduction

Firm productivity in low-income countries is both lower on average than in rich countries, and distributed with a far thicker left tail (Hsieh and Klenow, 2009; Bloom and Van Reenan, 2007, 2010). One important source of productivity dispersion is the use of inferior technology, including managerial practices (Syverson, 2011; Bloom et al., 2013). Small firms, where the vast majority of the non-agricultural workforce

is employed, are particularly far from the technology frontier.

This fact has inspired a variety of firm-level training programs aimed at microenterprises. Among garment makers in Ghana, our study population, extension officers are actively engaged in efforts to teach standardized sizing, new styles, and other techniques that the national government believes could increase productivity in small-scale manufacturing. In the economics literature, dozens of studies have evaluated the efficacy of training programs focused on business skills like inventory

[☆] We are grateful to David Atkin, Dan Björkegren, Jing Cai, Andrew Foster, Emily Oster, and Eric Verhoogen, for their detailed reading of and essential commentary on this paper. We also are grateful to Anna Aizer, Arun Chandrasekhar, Kristopher Chatlosh, Pedro Dal Bo, David Glancy, Jeremy Magruder, Isaac Mbiti, Angélica Meinhofer, Sveta Milusheva, Martin Rotemberg, Anja Sautmann, Daniela Scida, Jesse Shapiro, William Violette, Christopher Woodruff, and seminar participants at Brown University, Columbia University, New York University, New York University - Abu Dhabi, University of Zurich, University of Ottawa, Johns Hopkins School of Advanced International Studies, The World Bank, The International Food Policy Research Institute, Northeast Universities Development Consortium (NEUDC), Private Enterprise Development in Low-Income Countries (PEDL) conference at Warwick University, and NBER Summer Institute Development Workshop for taking the time to give helpful comments and suggestions. We also thank Lois Aryee, Robert Obenya, Charles Sefenu, Yani Tyskerud, Innovations for Poverty Action-Ghana, and especially Edna Kobbinah for excellent research assistance in the field, the Hohoe Municipal District Technical and Vocational Education and Training (TVET) coordinators for partnering with us on the design training, and Osman Mutawakil for collaborating on the design of the weaving technique. This research was supported by funding from PEDL, 3ie, USAID, The World Bank, the William and Flora Hewlett Foundation, the Barrett Hazeltine Fellowship for Graduate Research in Entrepreneurship, the Watson Institute for International Studies, and the Population Studies Center at Brown University. All errors are our own. CORRESPONDING AUTHOR: Morgan Hardy, Address: 19 W. 4th Street, New York, NY, 10003, Phone: 212-998-8900, Fax: 212-995-4186, Email: morgan.hardy@nyu.edu.

* Corresponding author.

E-mail address: morgan.hardy@nyu.edu (M. Hardy).

<https://doi.org/10.1016/j.jdevec.2020.102600>

Received 3 March 2019; Received in revised form 10 December 2020; Accepted 22 December 2020

Available online 8 January 2021

0304-3878/© 2021 Elsevier B.V. All rights reserved.

management, marketing, and bookkeeping. Meta-analytic evidence suggests these costly interventions generate only modest benefits, necessitating attention to scale and suggesting a search for impact multipliers (McKenzie and Woodruff (2014); McKenzie (2020)).

This study is among the first to analyze peer-to-peer technology diffusion among non-agricultural microenterprises, a natural candidate for such multiplier effects. Existing studies have documented spillovers of managerial practices and business information when network links between firms are randomized (Fafchamps and Quinn, 2013, 2016; Cai and Szeidl, 2017) and found evidence of social learning about productive technologies in agriculture (Foster and Rosenzweig, 1995; Conley and Udry, 2010; Munshi, 2004; Bandiera and Rasul, 2006). It is not clear, however, that findings about the diffusion of managerial practices would apply to non-managerial technologies nor that findings from tiny farms generalize to tiny firms. In a rare study to test for peer-to-peer diffusion of a non-managerial technology outside agriculture, Atkin et al. (2017b) randomly seed a new input-cost reduction technology in a sample of Pakistani soccer ball producers and find little to no network-based diffusion.

Building on insights from the broader literature on technology adoption, where a growing empirical consensus suggests demand-side factors are key drivers of technology upgrading (Verhoogen, 2020), we pose the hypothesis that demand stimulus paired with training programs can spur peer-to-peer diffusion and multiply any potential productivity effects. Several existing studies use variation in access to export or domestic markets to demonstrate a causal relationship between demand-side factors and firm-level technology upgrading (Bustos, 2011; Atkin et al., 2017), but none to date have tested whether demand can jump-start broader diffusion of new technologies through the network beyond the individual firm. Our primary contribution is to document such a test.

We conduct an experiment using a custom-designed weaving technology in a sample of firm owners that constitute the universe of garment making firms in a mid-size district capital in Ghana. The weaving technique is novel; no firm owner in the sample used the technique prior to the experiment and the steps to complete it differ from any weave commonly in use, making it difficult to backwards engineer. Its simple algorithmic execution was deliberately designed to function like a key or a password, rather than a skill perfected over time, in an effort to focus our study on learning-from-others rather than learning-by-doing. Though we do not claim the technique is necessarily profit or productivity enhancing in the short term, it is widely applicable (gender-neutral, can be added to any garment much like embroidery), relevant (customers in our market research survey cite style innovation as a key driver of organic demand), and well-suited to the moment we study (a period of frequent blackouts; unlike embroidery, this technique does not require electricity).¹ The universal nature of the sample makes it possible to have complete network data across all local garment makers.

The experiment sequentially cross-randomized training in the weaving technology and experimental demand shocks, in which our implementation team randomly placed orders for garments featuring the weaving technology. In the first phase of the experiment, we partnered with the technical training division of the local government to offer training in the technique to 15% of the firm owners in our sample. The technique requires the motor mechanism from a commonly available (and cheap) children's toy and trainees constructed their own tools and kept them following the training. They were also gifted a style calendar featuring the design, a common decor piece in Ghanaian garment making firms and a common gift when attending government-sponsored training. In the six weeks following training, about 13% of firm owners invited to the training adopted the design in a customer order, but there

was little to no diffusion to untrained firms.

In the second phase of the experiment, which takes place about three months after the training, our implementation team cross-randomly attempted to place orders for garments featuring the design with both trained and untrained firms.² Experimental demand was rolled out in waves, to allow for real-time management of project resources. Half of both trained and untrained firms were offered garment orders in the first wave. As completion rates in the first wave were realized, we sequentially offered orders to small random subsets of the remaining sample until resources were exhausted, composing a second experimental demand wave that ultimately included an additional 28% of the sample. Broadly, the infusion of demand shocks appears to have spurred market-level changes in both adoption and diffusion, though we focus naturally on differences between our cross-cut treatment groups.

Our key finding is that among trained firm owners, those who received demand shocks provide more teaching within the network, driving peer-to-peer diffusion of the technology. On the intensive margin, trained firm owners who received demand shocks teach 400% as many of their peers as those who did not. This finding holds for learners who received demand shocks and for learners who did not receive demand shocks. It holds whether we measure an instance of teaching via teacher reports, learner reports, or either. It holds winsorizing the most prolific teachers and estimating using a count model rather than a linear model. It appears to be a quite robust pattern emerging from the experiment.

What explains this finding? A readily available explanation involves learning-by-doing. By their very nature, demand shocks imply differential adoption, differential experience in production, and where intensive margin experience matters, differential skill in production. We note first that our weaving technique was intended to obstruct this mechanism by functioning like a password, rather than a skill to be mastered. In addition, we attempt to rule out this mechanism by searching for evidence of its existence and failing to find it. In the first wave of orders, the size of the order was also randomized, offering a natural experimental test of intensive-margin ability. Among firms in the demand treatment, point estimates on the size of the order are very near zero and never significant in regressions estimating effects on self-reported ability to use the weaving technology and measures of the quality of the weave in delivered garments. Though we are somewhat under-powered to provide precise zeros on the effects of order size on teaching behavior, these point estimates are likewise near zero and never significant.

In the full sample at endline, trained firms with and without experimental orders are about equally likely to self-report ability to produce the weave. We also conducted a long-term survey a year after the experiment and offered single-garment orders to all firms in the sample at that time. Among trained firm owners, there is no statistical difference in their likelihood of continuing to use the weave in non-experimental customer orders and no statistical difference in the quality of the weave as measured in the long-term experimental order.

Other potential explanations of our main finding include differential access to capital and differential visibility to potential learners. We find no evidence that among trained firm owners those with demand were more likely to have access to the necessary capital to produce the weave. Though we do not have direct data on requests for information or skill-transfer, we utilize the lag between the date that orders were placed and the date that firms actually did the work on them to show

² The training team, the implementation team, and the research team intentionally had no overlapping staff, to allow for some potential plausible deniability in the link between these activities. We did not, however, attempt to explicitly employ a secret shopper methodology because we doubted whether it would be possible to effectively conceal the source of a wave of orders that touched more than half of the garment making firms in this small town and its outlying suburbs. Like government training, however, the premise of a wave of orders from a non-governmental organization is not unprecedented, or even particularly obscure in this setting.

¹ See Abeberese et al. (2018) and Hardy and McCasland (2019) for detailed discussions of the power crisis.

that the timing of teaching is dependent on the order date but not on the timing of demand order fulfillment. In other words, trained firm owners taught other garment makers once they had received an order from our implementation team regardless of whether or not they had yet used the weaving technology to actually produce that order. Ultimately, this pattern is inconsistent with any adoption-based explanation.

Rather, we interpret our findings as most consistent with the demand intervention generating differential willingness to diffuse among potential teachers. At the conclusion of the paper, we describe how our demand intervention can be interpreted as reducing the rivalry of the technique, linking to the broader study of competitive incentives in small-scale manufacturing markets. Using heterogeneity by teacher and learner characteristics, we find almost no diffusion from any of our experimental groups to their business-as-usual competitors; our main finding is explained entirely by differential diffusion between trained firms with and without demand to non-competitors.³ In addition, the wedge between trained firm owners with and without demand is mitigated by owners more willing to share rival resources with their peers at baseline. This evidence is consistent with a rivalry interpretation, though we are unable to directly test for it.

Interpreting our main finding as driven by competitive incentives relates back to the literature on social learning in agriculture, where farmers tend to produce tradable cash crops or subsistence products for home consumption, making demand less rival than in small-scale manufacturing or services. [Cai and Szeidl \(2017\)](#) find that non-rival business information is equally likely to diffuse to competitor businesses among large firms in China, while rival business information is less likely to diffuse to competitors. The potential presence of these strategic disincentives to share a new technology highlights another link with the existing literature. While the majority of the literature focuses on the desire of potential adopters to learn a new technology, a few studies have begun to pay heed to the importance of the identity and willingness to teach of incumbent adopters in agriculture ([BenYishay and Mobarak, 2018](#); [Beaman et al., 2018](#)). Our study contributes a similar focus among small-scale manufacturers.

From a policy perspective, our paper highlights a compelling new tool that policymakers interested in increasing the productivity of small-scale non-agricultural businesses may consider. In the presence of a productivity-enhancing training opportunity, it may be possible to magnify impacts by pairing demand shocks with training to spur diffusion of new technologies through business networks. The fact that the demand shocks studied in our paper were interventionist is perhaps a feature more than a bug; they should be replicable and not rely on the generation of organic demand.

The paper proceeds as follows: In Section 2, we describe the project background and the setting. In Section 3, we discuss the research design. Section 4 presents and explores our main results. Section 5 discusses mechanisms, and Section 6 concludes.

2. Project background

2.1. Garment making in Ghana

Small-scale garment making firms are ubiquitous and prolific in Ghana, as in many other parts of Africa and the developing world. The vast majority of their production is bespoke garments for the local market. In many parts of Africa, traditional African wear is worn at weddings, funerals, and special events, “African wear Fridays” are common in government offices and banks, and more modern cuts in

African prints are popular with stylish middle class consumers. Orders for ready-to-wear production, school uniform contracts, and contracted production for export occur, but are a relatively small part of the market (due both to consumer tastes and the fact that uniform sizing is poorly standardized).⁴ Market research by the largest producer of African print fabrics in West Africa predicts growth in consumer spending on bespoke garments, despite local and imported ready-to-wear alternatives ([KPMG, 2014](#)).

We study the garment making industry in Hohoe town and its outlying suburbs, with a total population of 73,641 in 2010. Hohoe town is the capital of Hohoe District, a middle income district by Ghanaian standards, in the Volta Region near the border with Togo. In our own market research survey of nearly 1600 people in our study district, respondents averaged consumption of 3.5 bespoke garments in the last year. Competition in the local market is fierce, and is driven not only by price, but also by fashion style differentiation, including ability to produce the latest trends. Nearly 60% of the respondents in our market research survey cited style/on-trend fashion skills as a primary determinant in choosing a garment maker. Within-industry network relationships in this context are a key resource for firm owners. Appendix Fig. A1 includes detail from our baseline survey on the business resources our sample most frequently utilizes for various needs; within-industry peers are by far the most commonly reported resource for learning about new technology.

2.2. Sample recruitment

Sample recruitment took place in February of 2014, and included a census of all garment making firm owners in Hohoe District. The recruitment strategy began with existing lists of firms procured from trade association leadership, and continued via snowball sampling from there. The final phase of the census included geographic canvassing, in which surveyors covered all roads and alleys in the district searching for commercial storefronts and inquiring with locals in commercial areas after garment making firm owners. The census turned up 1025 active garment making firm owners in the district, with 12% co-owning with one or more other firm owner(s). The unit of measurement throughout is the firm owner, as co-owned firms tend to divide income, variable expenses, and profit streams, and share primarily physical space and physical capital. For logistical reasons, the experiment and all follow-up data were restricted to the subsample of firms located in Hohoe town and its outlying suburbs, a total of 445 garment making firm owners in the initial census. Of these 445, 417 were still operating a garment making business in Hohoe town or a surrounding suburb at the time of the experiment. These 417 make up the main sample for this paper.

3. Research design

3.1. Sharawakil

We collaborated with a designer in Accra to design our own style innovation, which we dubbed Sharawakil.⁵ The design process was iterative, as we wanted a technique that is tricky to backwards engineer but easy to learn if instructed, and relevant for all businesses in our sample. All firm owners who attended the one-day training session easily mastered the technique.

The technique begins with the deconstruction of a toy car, commonly available in nearly all local markets for about 5 Ghana Cedi

³ Concurrent with other data collection, we conducted a market research survey with residents of our study area and used it to create a diversion network: firms that share customers in our market research survey are classified as competitors and we identify, on average, two competitors for each firm in our sample.

⁴ Industrial textile production for export in Ghana has declined sharply in the post-independence period, and more recently due to major electricity shortages. 5000 people were employed in factory-style garment making jobs in Ghana in 2000, down from 25,000 in 1975 ([Quartey, 2006](#)).

⁵ The name of the design is a combination of “shara”, a Hausa word for twisted, and the name of the designer, Mr. Osman Mutawakil.



(a) Toy Car Turned Weaving Tool



(b) Uniquely Identifiable Product



(c) Generalizability of Use

Fig. 1. New Technology - Sharawakil. With permission from the models, this figure depicts photo examples of Sharawakil, the new weaving technique developed for the field experiment. The product is easily distinguishable from others in the market and can be used on wide variety of garments. Photoshoots were conducted at the IPA Ghana Office and in Makola Market, Accra. Models are IPA Ghana staff members and friends and family of the Sharawakil designer, Mr. Osman Mutawakil (pictured top right).

(GhC), or about 2 US dollars at the time of the experiment. Having removed and jerry-rigged the motor, the first step is to use it to spin a bundle of strands of thread, as depicted in Fig. 1(a). The twisted thread is then folded and released, collapsing into a tightly-bound rainbow weave (see Fig. 1(b)) that can be used to adorn garments in a variety of creative ways and in a fashion similar to the widespread use of embroidery. Some potential uses are displayed, with the permission of the models, in Fig. 1(c).

3.2. Randomization design

The random variation exploited in our analysis comes from two experimental phases: random invitations to a government-run Sharawakil training and randomly timed and sized demand offers to complete garments that must be adorned with Sharawakil. In addition, we also conducted a non-randomized “long-run” demand exercise over a year after the experiment, during which each member of our sample had the chance to complete a Shawarakil-adorned garment.

The cross-cut randomizations from Phase 1 and Phase 2 form the four treatment groups of interest in the main analysis of this paper: Control (CC), Demand-only (CD), Training-only (TC), and Training-plus-demand (TD).

Phase 1: Stratified by gender, 15% of firms from the census sample in Hohoe town and its suburbs were randomly selected to receive an invitation to participate in a skills training to learn a new garment design and receive the relevant equipment needed to produce it.⁶ The training was held at the local government technical training offices, where the district technical training coordinator made introductions and acknowledged our training team as the sponsor of the training. The training itself was lead by Mr. Osman Mutawakil (the designer of Sharawakil) who was brought in to teach from Accra. No financial incentives were offered to attend the training. However, all attendees left with the Sharawakil tool (from a toy car they transformed themselves as part of the training) and a design calendar to advertise their new skill to customers (a common practice in Ghana; see Appendix Fig. A2).

Phase 2: Stratified by gender and the skill-training treatment assignment, 50% of the 417 firms in the experimental sample were randomly assigned to the first wave of demand contract offers, which also randomly varied in size. Our implementation team attempted to approach each of these selected firm owners with a contract offer to produce 1, 4, or 10 garments featuring the design for a fixed price of 35 GhC, with 88 firms in the 1 garment group, 88 firms in the 4 garments group, and 33 firms in the 10 garments group.⁷ Firm owners were given three days to decide whether they wanted to accept the contract offer and produce the garments, on the basis of their expected ability to successfully complete the order featuring Sharawakil within two weeks from the offer date. Once they agreed to produce the garment, our implementation team paid 10 GhC in deposit and provided fabric for the garment (standard practice in this industry).

The remaining 50% of the sample was randomly split into subgroups, which was intended to allow for real-time, design-based management of project resources and to introduce a rivalry dimension to the design. On the basis of ongoing order acceptance rates, and chronologically after all wave 1 orders had been offered, two of these subgroups were also sequentially offered orders. Every firm offered a contract in this second wave of orders was offered a contract to produce only one garment. In total, ~22% of the firms in the sample were never targeted with a Phase 2 contract of any size. See Fig. A3 for a breakdown of the experimental design.

Order Process: Experimental demand was both scarce and without replacement and our implementation team emphasized these features of the demand process when placing experimental orders. In addition to a discussion about Sharawakil, our team emphasized that they would make orders until fixed garment needs were fulfilled (all implementation funding exhausted). The process also included a brief explanation of the activity, stating that orders are maximum one per firm, that order size was randomized by lottery, and that the order process does not affect participation in any other activities (namely, the research team's activities in the area, which had started over a year before the experiment).

Long-run Orders: About a year after the experiment, our implemen-

tation team returned to the field and attempted to offer an additional garment order to each of the 417 firms in the experimental sample. 375 firm owners were found and offered contracts. In addition to being universal rather than randomized, these contracts differed from those in the initial experiment in two key ways. Sharawakil was optional and cash-incentivized, and the price was not fixed (so prices paid were the outcome of a bargaining process).⁸

3.3. Data

Fig. A4 visually depicts the data collection timeline. Data for the paper comes from a variety of sources. Pre-experimental data includes the sample recruitment process and associated baseline survey, as well as the first week of a seven week panel that began the week before the training. We measure Sharawakil diffusion in Phase 1 in each of the six weeks of the midline survey following the training and in Phase 2 in an experimental endline that took place shortly after the experiment. Sharawakil adoption is measured at both midline and endline, as well as in a long-run survey conducted about a year after the experiment. We also use data from a market research survey run concurrently with the seven week midline panel, administrative data on attendance at the training and order completion, and Sharawakil quality data measured by the designer of the weave and two other experienced garment makers for both experimental and long-run orders.

Census and Baseline Survey: The census activity included collection of limited demographic and firm-level data. It was followed in July and August of 2014 by an extensive baseline survey. In addition to standard detail about firms and firm owners, the baseline survey included a complete network map within industry and within district for the year preceding the survey. While firm owners were not prompted to confirm or deny acquaintance with each of the more than 1000 other firm owners, the self-reported contact section was designed to capture as many relevant contacts as possible. Firm owners were prompted by category (former employer, former employee or apprentice, trade association co-member, neighbor, close friend in the business, etc.) to report all garment making firm owners in the districts with whom they interacted in any way (even just sharing greetings) in the year preceding the baseline. Once all contacts were revealed, specific network activity over the last year between contacts was collected along the following dimensions: gift and loan giving, skill sharing, labor sharing, equipment sharing, customer referrals, mentorship, outsourcing, and other. When reporting on network interactions in this paper, we restrict network relationships to those within the 417 firms in our final sample and define a connection as existing if either node reports it. Appendix Fig. A5 provides a visualization of the baseline technology diffusion network.

We use two other features of the baseline data in our analysis. In our study district, as in many parts of Ghana, garment makers voluntarily organize themselves into trade associations. The largest is the local chapter of the national Ghana National Tailors and Dressmakers Association (GNTDA), though a second one also has significant membership in our sample. These associations charge small membership fees and offer a range of services, among them access to a network of relatively more powerful and larger garment making firm owners, and organized skill training in new fashion designs and production technologies. The baseline survey collected information on whether a firm owner was a member of a trade association at baseline and to which trade association that person belonged, which we use to measure whether two firm owners in a dyad belong to the same trade association. In addition, we use information we collected on altruism towards other garment making contacts in our analysis. Using a standard non-incentivized dictator

⁶ The training randomization took place several months after the census activities and inadvertently included 29 firms that had permanently closed prior to the start of the experiment, something that our field teams confirmed during experimental field work. Our experimental sample of 417 firms remains balanced in the subset of surviving firms, as described below.

⁷ One garment, four garments and ten garments are equivalent to the median, the 90th percentile, and the 99th percentile of weekly sales in the sample, respectively. Embellished garments range from 20 to 40 GhC in this setting, so our fixed price was chosen to be on the generous end of this range.

⁸ The long-run orders were not anticipated by either the researchers or the experimental subjects, but rather were added to new pilot fieldwork for an unrelated study due to unexpectedly available additional resources.

game, we define altruistic firm owners as those who would give half or more than half of 10 GhC to a hypothetical garment making peer and non-altruistic as those who would give zero or less than half.

Sharawakil Adoption and Diffusion: The seven week midline panel (collected in March and April of 2015) asked indirectly about Sharawakil adoption and diffusion. Creating the weave is a quite different process than other designs in the market so asking respondents to describe any new techniques used in customer orders, learned, or taught and back-coding responses as Sharawakil provides a reasonably precise measure. For this reason, Phase 1 learning is measured using learner reports and Phase 1 teaching using teacher reports, as we do not have dyad-specific measures from which to construct an instance of diffusion as reported by either teacher or learner. All Phase 1 measures sum across all seven weeks of the midline panel, although no adoption or diffusion is reported in the first week (prior to our intervention). No teacher taught more than a single peer in Phase 1, but we report teaching as number of peers taught for consistency throughout the paper. The indirect structure was intended to avoid Hawthorne effects, the concern that we might drive adoption and diffusion simply by asking about it directly.

In the experimental endline (collected directly after the Phase 2 intervention in June and July 2015), we ask directly about adoption, learning, and teaching of Sharawakil, as Hawthorne effects were less of a concern at the conclusion of the experiment. Our primary Phase 2 measure of diffusion defines an instance of dyad-specific teaching when it is reported by either the learner or the teacher, though we also present results disaggregated into learner only and teacher only reports. We also ask directly about working tool ownership and ability to produce Sharawakil at endline, confirming self reports by asking those who say they know how to produce the design to describe the algorithm in detail.

Quality Data: Sharawakil produced through the experiment and for long-run orders by firm owners in our sample was evaluated by the designer of the weave and two other expert garment makers after all orders had been collected. For firm owners who produced more than one garment, one garment was chosen at random and evaluated. An advantage of Sharawakil is that the smoothness and tightness of the weave is a direct result of the correct algorithmic execution of the design, making it a good proxy for knowing the skill. Quality was evaluated on a ten point scale, which we standardize in the analysis.

Market Research Survey: Finally, we also surveyed about 1600 residents of the sample region about their purchasing behaviors and preferences over garment makers, concurrent with the midline surveys. Survey enumerators stood at geographically disperse public points in town, and approached every fifth person who passed their post. The locations were chosen in part to systematically be near clusters of shops in the sample. The primary purpose of the market research survey was to create a diversion network map linking firms that share customers, which serves as our measure of business-as-usual competition between firms. In addition, each market research survey respondent was asked how frequently she buys bespoke garments or uses the services of garment making businesses, how much she has spent on these products and services recently, and what characteristics are most important to her in choosing a garment making business.

4. Results

4.1. Summary statistics

Demographically, a large majority of small-scale garment making firm owners are women (both in Ghana and around Africa), though male-owned firms make up about 23% of our sample and tend to be

both bigger and more profitable. Our sample is 76% Ewe.⁹ Owners have an average of 9 years of schooling, equivalent to a junior high school education, while some have no formal schooling and some have tertiary degrees. The average monthly profits at baseline were 138 GhC, which is equal to approximately 60 USD at the time of the survey. Production technology in these firms typically consists of a mix of hand or foot-crank sewing machines that do not require electricity, and electrically powered embroidery, overlock, and sewing machines. 47% of our sample employ at least one apprentice or other paid worker. Firms employ an average of 1 worker, but the distribution is skewed with a maximum of 15 workers and the median number of workers at 0. Only 17% of our sample firms are registered with any government agency and 22% are members of a trade association.

Table 1 shows balance along major observables across the three treatment and one control groups. With the exception of a lower probability of being Ewe for the demand-only group, a lower probability of being registered for the training-only group, and more years of schooling for the training-plus-demand group, all observables appear balanced.

4.2. Phase 1

Phase 1 of this project was originally intended to be self-contained, allowing for the study of diffusion in a competitive network of garment makers using only exogenous variation in training. That design and that analysis were abandoned when organic demand failed to materialize and fewer than a handful of instances of diffusion were observed.¹⁰ Still, the failed initial experiment is a useful precursor to our primary experimental analysis as it relates directly to the challenges of sparking diffusion of new technologies through networks of this type.

In Panel A of Table 2, we present the basic results from Phase 1 according to the following specification:

$$Y_i = \beta_0 + \beta_T * T_i + \beta_g^* g_i + \epsilon_i \quad (1)$$

where Y_i is the outcome of interest, cumulative across all rounds of the midline survey, g_i is a gender dummy controlling for the randomization strata, and ϵ_i is an error term. The dummy T_i indicates assignment into the training invitation group. The β_T coefficient is the intent-to-treat effect of being randomly assigned to receive a training invitation.

Training compliance is high, at 94%, with 51 of 64 firms in the training group participating in the main training sessions offered at the district technical training office, and another 9 attending mop-up training held a few days later. The primary reason for non-compliance was travel. Given high compliance, we report results as intent-to-treat throughout the paper. The remaining columns of Table 2 present results on Sharawakil adoption and diffusion. 1% of non-trainees and 14% of trainees used the design in customer orders in any of the weeks covered by our midline survey. Trainees taught on average 0.02 of their peers.

⁹ Ewe is the dominant ethnic group in Hohoe. It is a community that straddles the Ghana-Togo border from the coast up to the northern frontier of the Volta Region. Hohoe is also home to small minority ethnic groups, primarily members of various Muslim communities from Northern Ghana, and Twi and Ga people from the Accra area. Ethnic tension in Ghana is relatively minimal, though political parties are divided along ethnic lines and ethnicity is a meaningful proxy for language.

¹⁰ The abandoned original research design for a self-contained Phase 1 was registered with the American Economic Association (AEA) Randomized Controlled Trial Registry under AEARCTR-0000642. We made the decision to shift to an alternative design while in the field and did not update the Pre-Analysis Plan to reflect the new design. All analysis in this paper is thus not pre-specified.

Table 1

Summary Statistics and Covariate Balance. This table reports sample summary statistics and covariate balance between the four main treatment groups used in the analysis of this paper. Columns labeled “mean” give the mean value of each baseline characteristic for all firms in our sample, firms randomized to receive no order-no training (Control), firms randomized to receive training but no order offer (Training Only), firms randomized to receive order offers but no training (Demand Only), and firms randomized to receive both an order offer and training (Training-plus-demand), in that order. Columns labeled “p-val” report the corresponding p-value on the test of equality of each baseline characteristic between each of the three treatment groups (Training Only, Demand Only, Training-plus-demand) and the control group (Control). ***p < 0.01, **p < 0.05, *p < 0.1.

	All mean	Control mean	Training Only mean	(diff = 0) p-val	Demand Only mean	(diff = 0) p-val	Training-plus-demand mean	(diff = 0) p-val
Male	0.23	0.23	0.21	0.92	0.23	1.00	0.24	0.86
Ewe ethnicity	0.76	0.83	0.86	0.78	0.73	0.09*	0.78	0.52
Years of schooling	8.85	8.69	8.21	0.51	8.83	0.66	9.38	0.06*
Ravens	5.63	5.45	6.21	0.28	5.58	0.71	6.00	0.22
Owner age	35.53	33.74	35.14	0.61	36.07	0.07*	35.28	0.33
Within industry/sample degree	15.56	14.37	15.00	0.84	15.75	0.35	16.46	0.36
Firm size (including owner)	1.99	1.93	2.29	0.42	1.99	0.78	2.00	0.82
Has any worker(s) besides owner	0.44	0.44	0.57	0.24	0.46	0.34	0.36	0.66
Revenues (GHC)	197	195	178	0.83	196	0.96	206	0.82
Profits (GHC)	138	144	123	0.72	136	0.68	145	0.99
Assets excl land/building (GHC)	1214	1187	1057	0.79	1256	0.76	1069	0.67
Management practices (of 4)	2.32	2.21	2.21	1.00	2.38	0.20	2.18	0.85
Firm age	9.49	9.06	9.00	0.98	9.66	0.59	9.38	0.83
Trade association member	0.22	0.19	0.14	0.70	0.23	0.41	0.22	0.65
Registered w/any govt agency	0.17	0.19	0.00	0.08*	0.17	0.67	0.18	0.93
Number of Firms	417	75	14		278		50	
F-Test of Joint Significance (pvalue)				.77		.19		.31

Table 2

Intention-to-Treat Effects on Treatment, Adoption, and Diffusion. This table reports intent-to-treat effects of treatment group assignment for Phase 1 (Panel A) and Phase 2 (Panel B) of our experiment on treatment, adoption, and diffusion rates. Training attendance and the successful offer of demand come from administrative data. Adoption is self reported by the firm owner in both phases. Phase 1 diffusion is only teacher reported for number of peers taught and learner reported for learning, due to data limitations. Phase 2 diffusion reflects either teacher or learner reporting. A gender dummy is included to control for the randomization strata. Robust standard errors in parentheses. ***p < 0.01, **p < 0.05, *p < 0.1.

	(1) Compliance	(2) Adoption Used Sharawakil in Customer Order	(3) Diffusion Learned Sharawakil from a Peer	(4) Diffusion # Peers Taught Sharawakil
Panel A - Phase 1				
	<u>Trained</u>			
Training Only or Training-plus-demand	0.94*** (0.03)	0.13*** (0.04)	0.01 (0.02)	0.02 (0.02)
Demand Only or Control Mean	0.00	0.01	0.01	0.00
Observations	417	416	416	416
Panel B - Phase 2				
	<u>Offered Demand</u>			
Demand Only	0.90*** (0.02)	0.51*** (0.04)	0.25*** (0.05)	0.13** (0.05)
Training Only	0.00 (0.00)	0.26** (0.13)	-0.05 (0.09)	0.32* (0.18)
Training-plus-demand	0.98*** (0.02)	0.83*** (0.06)	-0.05 (0.06)	1.21*** (0.28)
Prob > F(Training Only = Training-plus-demand)	0.00***	0.00***	0.94	0.00***
Control Mean	0.00	0.05	0.13	0.07
Observations	417	384	384	384

4.3. Phase 2

We estimate reduced form treatment effects for Phase 2 using the following specification:

$$Y_i = \beta_0 + \beta_{CD} * CD_i + \beta_{TC} * TC_i + \beta_{TD} * TD_i + \beta_g^* g_i + \epsilon_i \quad (2)$$

where Y_i is the outcome of interest, g_i is a gender dummy controlling for the randomization strata, and ϵ_i is an error term. The dummies CD_i , TC_i , and TD_i indicate assignment into the demand-only, training-only, and training-plus-demand treatment groups, respectively. With the pure control group omitted, the β coefficients are the intent-to-treat effects of being assigned each particular treatment group relative to the control.

We define firms as demand treatment compliers if our implementation team successfully found them, spoke with firm owners, and offered to place a garment order contract. Compliance was 90% and 98% of firms in the demand-only and training-plus-demand groups, respectively. Again, we should expect intent-to-treat effects to be close to treatment-on-the-treated. Not all compliers chose to accept the order contract and produce the garments. 67.9% of demand-only compliers accept and complete a Sharawakil-adorned garment order, displaying ingenuity in the face of demand for new products. However, Phase 1 treatment assignment matters, as training-plus-demand compliers are 17.8% more likely to accept and complete an order than demand-only compliers.

Table 3

Robustness of Intention-to-Treat Effects on Diffusion. This table reports intent-to-treat effects of treatment group assignment for Phase 2 diffusion (number of peers taught), varying learner subgroups, measurement, and specification choices. Columns 1 and 2 show diffusion effects by learner treatment group assignment. Columns 3 and 4 show diffusion effects under alternative choices of measurement. Columns 5 and 6 show estimates for diffusion effects with winsorized measures of number of peers taught. Finally, Column 6 shows estimated effects using poisson estimation count model. In columns 1, 2, 5, 6, and 7, diffusion reflects either teacher or learner reporting. A gender dummy is included to control for the randomization strata. Robust standard errors in parentheses. ***p < 0.01, **p < 0.05, *p < 0.1.

	(1) Demand Only or Training- plus-demand	(2) Training Only or Control	(3) Learner Reported	(4) Teacher Reported	(5) Winsorized 1%	(6) Winsorized 5%	(7) Poisson
Demand Only	0.10** (0.05)	0.02** (0.01)	0.10* (0.05)	0.04*** (0.01)	0.11** (0.05)	0.09** (0.04)	1.07** (0.53)
Training Only	0.32* (0.18)	-0.00 (0.00)	0.24 (0.17)	0.08 (0.07)	0.32* (0.18)	0.32* (0.18)	1.77*** (0.66)
Training-plus-demand	1.12*** (0.26)	0.08** (0.04)	0.96*** (0.27)	0.29*** (0.10)	1.02*** (0.19)	0.77*** (0.13)	2.97*** (0.53)
Prob > F(Training Only = Training-plus-demand)	0.01**	0.04**	0.02**	0.09*	0.01***	0.04**	
Prob > Chi2(Training Only = Training-plus-demand)							0.02**
Control Mean	0.07	0.00	0.07	0.00	0.07	0.07	0.07
Observations	384	384	384	384	384	384	384

Column (2) presents the impact of group assignment on adoption by Phase 2 follow-up. Here adoption is again defined as having used the weave in a customer order (in the last one month). In Phase 2, this definition is inclusive of our experimental demand. The first thing to note is that overall adoption is higher. The Phase 2 endline, from which this data is drawn, was completed directly after the demand intervention period which was a couple months after the final midline survey, which could explain part of the increase in the adoption of the weave. Another plausible explanation is that the exogenous demand shocks shifted market-level tastes for using the weave (from customers or garment makers themselves). In either case, 5% of control group firms are using the weave by the Phase 2 follow-up. Each of the three treatment groups is more likely to have adopted the weave than the control group, which the highest adoption rates in the training-plus-demand group (as anticipated).

In Column (3) we explore learning about the new technology via peer-to-peer diffusion. About 13% of firms in the control group learned through another garment maker, while that number is much higher in the demand-only group. The wedge between adoption in the demand-only group and diffusion through the network to the demand-only group is explained by the fact that 32.65% of demand-only group adopters report buying Sharawakil from another garment maker to affix to their garments. Fig. A6 Sub-figure A depicts the method by which firm owners who adopted reported acquiring Sharawakil. Only three firm owners reported sales from Sharawakil to other garment makers suggesting that the selling of Sharawakil was relatively rare. Nearly all of those who learned from another garment maker did so for free (see Fig. A6 Sub-figure B). Less than 3% of our sample report reverse-engineering the technique themselves.

Finally, Column (4) of Table 2 presents our main finding. Firms in the training-plus-demand group teach about 400% more of their peers. This difference amounts to training about one extra garment making peer for each trained firm owner randomly selected to receive a demand shock. The F-stat on the difference between the coefficients on teaching among training-only and training-plus-demand firms is significant at the 1% level.

We explore the robustness of this finding in Table 3. The first two columns vary the identity of the learner, restricting learners to those who received demand (the demand-only and training-plus-demand groups) and those who did not receive demand (the training-only and control groups). Few learners come from either trained groups, so this is effectively splitting learners between the demand-only and control groups. Note that there are 328 total potential learners in Col-

umn (1) and 89 total potential learners in Column (2). If something strange about the experimental structure of demand is driven by both the learner and the teacher receiving experimental demand orders, we should expect our main finding to only hold when learners also received demand. However, the coefficients on teaching for the training-only and training-plus-demand groups are significantly different from each other regardless of the identity of the learner, suggesting that demand is acting on the teacher independently of its effects on the learner's desire to adopt the new technology. Point estimates are larger in Column (1), both because there are more potential learners in that set and because learners in that set are more motivated to learn.

The next two columns explore robustness of the main finding to differences in measurement. In our primary specification with Phase 2 endline data, an instance of teaching is recorded if either the learner or the teacher reported the connection (and was able to name the garment maker on the other side of that connection). This specification choice depends on the belief that recall error is important and that firm owners are less likely to misreport additional (erroneous) instances of teaching than they are to forget to report an instance of teaching. A worry with this decision is that surveyor demand effects could drive training-plus-demand firms to differentially misreport additional instance of teaching because they may view these as socially desirable. Here we split the outcomes and report findings for learner-reported and teacher-reported instances separately. Some instances of teaching are reported by both the learner and the teacher and thus appear in both Columns (3) and (4), but instances of teaching are more frequently reported by only the learner than only the teacher. In both columns, training-plus-demand firms teach significantly more of their peers than training-only firms, and the ratios between the two point estimates remain at about 400%. However, it appears recall error is more prominent among teachers than it is among learners, an intuitive feature of the data for two reasons. Learners report a maximum of a single teacher while in order to have a completely matching network map, teachers would have to report (in many cases) more than one learner. In addition, learning a new skill (and recalling the person who taught one that skill) is arguably a more salient experience than teaching a skill. We find no evidence that surveyor demand effects among training-plus-demand firms are driving our main results.

Columns 5 and 6 check for the sensitivity of our main results to outliers, using winsorized measures of number of peers taught at the 1% and 5% level and finding no sensitivity of the difference in teaching between training-plus-demand and training-only firms to choice of

Table 4

Intensive Margin Effects of Random Order Size. This table reports intent-to-treat effects of increasing order size on Phase 2 order completion, production ability, and diffusion for firm owners randomized to receive demand. Number of experimental orders completed comes from our administrative data. Ability is self-reported by the owner and verbally verified (through description of the algorithm) by our enumeration staff. Quality measures come from expert evaluations of Sharawakil quality on completed garments (for larger order sizes, one garment is randomly selected for review) using a 10 point rubric. Diffusion reflects either teacher or learner reporting. Orders were given out over two waves. Wave 1 contained orders of size 1,4 and 10. Wave 2 contained orders only of size 1. Wave 2 firm owners are dropped from analysis in all columns except column 3 to not confound order size and timing effects. A gender dummy is included to control for the randomization strata. Robust standard errors in parentheses. ***p < 0.01, **p < 0.05, *p < 0.1.

	(1) Number Phase 2 Garments Completed	(2) Able to Produce Sharawakil	(3) Z-Score Sharawakil Quality	(4) # Peers Taught Sharawakil
Panel A - Saturated				
Order Size Four	1.61*** (0.22)	0.02 (0.07)	-0.04 (0.18)	0.01 (0.18)
Order Size Ten	6.32*** (0.80)	0.02 (0.09)	-0.23 (0.22)	0.37 (0.24)
Prob > F(Size Four = Size Ten)	0.00***	0.98	0.39	0.11
Panel B - Linear				
Number of Extra Shirts (0,3,9)	0.70*** (0.09)	0.00 (0.01)	-0.03 (0.02)	0.04 (0.03)
Demand Only and Order Size One Mean	0.63	0.49	0.09	0.20
Training-plus-demand and Order Size One Mean	0.77	1.00	0.01	1.38
Observations	209	198	126	198

winsorization.¹¹ Finally, column 7 reports coefficients from a Poisson regression, testing for sensitivity to the choice of linearity in our main specification. Again, the p-value for the Chi2 test of equality of number of peers taught between the training-plus-demand and training-only groups comes in at 0.02. The key pattern of interest is remarkably stable.

5. Mechanisms

An initial intuitive explanation for why demand increases teaching is that demand drives adoption by potential teachers. Adoption has the potential to increase an owner’s ability to use a technology through learning by doing. Adoption can increase complimentary capital investment (through tool purchases). In a dense real-world network like the one we study, adoption may increase visibility to potential learners in their search for a teacher. We argue in this section that the design of the weave, the structure of our experiment, and the empirical evidence do not support these adoption-based mechanisms to explain our main findings. Finally we explore alternative mechanisms and conclude that the limited available evidence is consistent with an interpretation that the experiment differentially manipulated the degree to which the skill was rival across treatment groups.

5.1. Evidence on adoption-based diffusion

Three key pieces of the experimental design were intended to blunt common barriers to technology diffusion and mute adoption-based channels. Sharawakil was custom-built to function like a password or a secret rather than a skill perfected with practice. Anecdotaly, firm owners easily mastered the technique in the government-training and we have no instances of firm owners reporting difficulty with its execution. The capital necessary to produce the weave was intentionally cheap and provided for free to all firm owners who attended a training session. All firm owners who attended a session also received a calendar, a simple signal to others that the firm owner may have attended the training.

Turning to empirical evidence, we begin by exploiting the fact that order size was randomized in the first wave of demand offers to search for intensive margin effects. Table 4 excludes firms outside the demand treatments and estimates the intensive margin effects of larger orders using the following saturated and linear specifications:

$$Y_i = \beta_0 + \beta_{D4} * D4_i + \beta_{D10} * D10_i + \beta_{TD} * TD_i + \beta_g^* g_i + \epsilon_i \tag{3}$$

$$Y_i = \beta_0 + \beta_{ADD} * ADD_i + \beta_{TD} * TD_i + \beta_g * g_i + \epsilon_i \tag{4}$$

where Y_i is the outcome of interest, g_i is a gender dummy, and ϵ_i is an error term. β_{D4} and β_{D10} are the estimated impact of receiving an order size of four and ten, respectively, relative to the omitted group receiving order size one. β_{ADD} is the estimated linear increase in Y_i with each additional order over the base order size of one. Note that Wave 2 firm owners are not included as Wave 2 included only order sizes of 1.

We first confirm that firm owners with larger orders produced more garments. Anecdotaly, no firm completed only a portion of the garments ordered, so the linear effect here comes from the fact that not all wave 1 orders were accepted and completed. Column (2) presents self-reported ability to produce Sharawakil. The data collection here required that the firm owner say they knew how to produce the weave and then described the algorithm for producing the weave, with survey enumerators confirming the firm owner had the algorithm correct. Point estimates in both linear and saturated specifications are very near zero. In Column (3), we present quality data from garments produced in Phase 2. Outcomes in Column (3) are presented in standard deviations and point estimates are near zero and never significant.

In the final column of Table 4, we test for whether order size affects our main outcome of interest. Point estimates are positive but never significant. The confidence intervals are wide, however, and we are somewhat underpowered to rule out small effect sizes. Only 15% of the firms that received demand orders in Wave 1 were offered orders of size ten; the 95% confidence interval excludes effects larger than 0.36 from three additional garments, which is small relative to our main effect sizes. In Appendix Table A1 we show that point estimates in Panel A on order size ten are driven not by firms in the training-plus-demand group, but by firms in the demand-only group. Demand-only firms with order size ten spend more on capital purchases, repairs, and rentals than demand-only firms with smaller order sizes, and the intensive margin effect on

¹¹ Fig. A7 shows distribution of number taught by treatment group.

Table 5

Intention-to-Treat Effects on Ability, Capital, and Future Adoption. This table reports intent-to-treat effects of treatment group assignment on firm owner's ability to produce Sharawakil (column 1), the quality of Sharawakil production (column 2), ownership of a working Sharawakil tool (column 3), and the future use of Sharawakil in customer orders during the year following the random demand shocks (column 4). Ability is self-reported by the owner and verbally verified (through description of the algorithm) by our enumeration staff. Quality measures come from expert evaluations of Sharawakil quality on completed garments from the long-run follow-up using a 10 point rubric. Ownership of a working Sharawakil tool was self-reported and visually verified by our enumeration staff. Sharawakil use in the year following our demand shocks was collected during a later survey round on the same sampling frame over one year after the completion of this experiment. A gender dummy is included to control for the randomization strata. Robust standard errors in parentheses. ***p < 0.01, **p < 0.05, *p < 0.1.

	(1) Able to Produce Sharawakil	(2) Has Working Sharawakil Tool	(3) Future Use of Sharawakil in Customer Order	(4) Z-Score Sharawakil Quality
Demand Only	0.28*** (0.06)	0.11*** (0.03)	0.11*** (0.04)	-0.07 (0.18)
Training Only	0.74*** (0.09)	0.75*** (0.12)	0.32** (0.14)	0.13 (0.40)
Training-plus-demand	0.76*** (0.06)	0.73*** (0.06)	0.38*** (0.08)	0.23 (0.25)
Prob > F(Training Only = Training-plus-demand)	0.86	0.89	0.70	0.82
Control Mean	0.18	0.02	0.06	0.00
Observations	384	384	375	255

teaching does not survive controlling for differential expenditures on the Sharawakil-specific tool.

Turning back to the full sample, Table 5 tests directly for differential reported ability, capital access, persistence in adoption, and weave quality between training-only and training-plus-demand firms. Point estimates on self-reported ability and owning a working tool are almost identical and not statistically different from each other in the training-only and training-plus-demand groups. Future use of Sharawakil in a customer order measures adoption in the year following the experiment, excluding any orders placed by the implementation team during the experiment or at the one year follow-up. Point estimates in the training-only and training-plus-demand groups are not statistically different from each other. The final measure is Sharawakil quality in the one-year follow-up orders among firms that completed the garments, where we find no treatment effects relative to control of any experimental group.¹² Together, this table is inconsistent with a learning-by-doing or capital complementarity explanation of our main findings.

Finally, we provide summary evidence on the set of adoption-based mechanisms using the timing of instances of teaching, relative to the timing of order placement, order production, and order collection. Administrative records catalogue the order placement and collection dates, while the Phase 2 endline asked respondents about the date of learning Sharawakil and the date of the first use of Sharawakil in a customer order, whether it be for our order or a market customer. Fig. 2 presents weighted scatter plots of these relative dates. The line in each sub-figure represents the 45° line, i.e. the same date. Bubbles above the line in each sub-figure signify that the event on the y-axis occurred after the event on the x-axis. Bubbles below the line in each sub-figure signify that the event on the y-axis occurred before the event on the x-axis. Bubbles directly on the line signify that both events happened on the same date.

Sub-figures (a) and (b) provide a sanity check for what we can learn from these figures.¹³ Sub-figure (a) shows that, although most firm owners first adopted the technology (first used Sharawakil for any cus-

tom order) only after receiving our order, this use is far from immediate and often delayed by a week or two. In Sub-figure (b), we see that all respondents adopted the technology before our team collected their garments, as adoption is a requisite to collecting a garment adorned with Sharawakil.

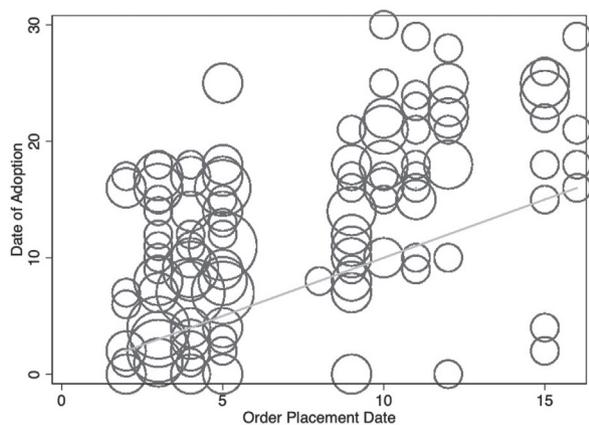
Sub-figures (c) and (d) plot the actual timing of diffusion of Sharawakil by order placement timing and the timing of first use. The x-axis depicts information from the teacher in a Sharawakil-diffusion teacher-learner pair on the teacher's date of order receipt and date of first use, while the y-axis depicts the date that the learner in that teacher-learner pair reported learning the skill. A larger share of the mass is above the "same date" line in Sub-figure (c), signifying that among training-plus-demand and demand-only teachers (firms must be in one of the demand groups to have an order placement date, and must have taught someone to have a diffusion date), teaching was more likely to happen after the teacher received an order from the experiment. In Sub-figure (d), there is no pattern of the bubbles appearing above or below the "same date" line. This is inconsistent with a learning-by-doing or a visibility story, which would predict that a larger share of the mass would appear above the "same date" line (i.e. that teaching would happen after teachers had first used Sharawakil on a garment). In fact, Sub-figure (d) is inconsistent with any adoption-based mechanism of demand driving diffusion. Absent functional differences between trained firms owners who did and did not receive random demand shocks, our findings are most consistent with the demand intervention generating differential willingness to diffuse among potential teachers.

5.2. Interpretation

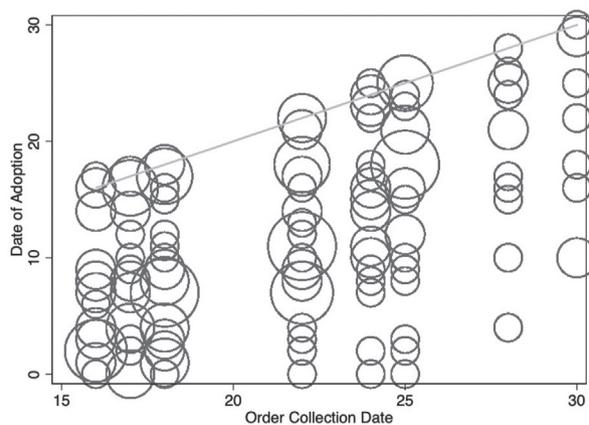
Unlike in many agricultural markets, where farmers tend to produce highly tradable cash crops or subsistence products for home consumption, small-scale manufacturers in low-income countries service almost exclusively local demand, leading to direct competition between firms in the same market. An intuitive explanation for unwillingness to diffuse new technologies in real-world interactions among small-scale manufacturers is that firm owners recognize that skills may be rival and strategically choose to keep their peers in the dark. Using competing firms matched in our market research survey, we can explore heterogeneity in teaching behavior by whether the potential learner is a competitor or not in the normal course of business. In Table 6 column

¹² Note, we detect no differential selection into adorning garments with Sharawakil for these long-run orders between training-plus-demand and training-only firms.

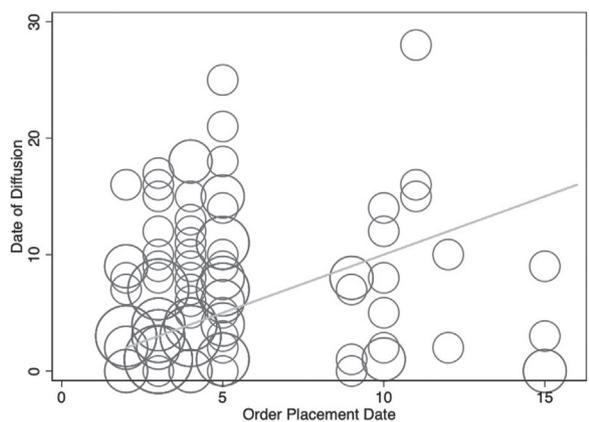
¹³ Note that we can see three timing clusters emerging, associated with our randomized order timing: Wave 1 and the two pieces of the second wave.



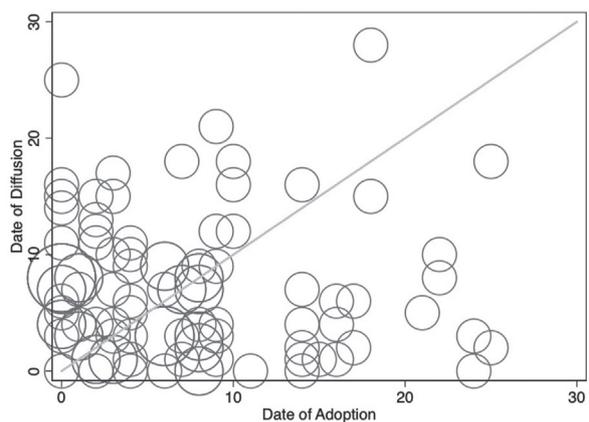
(a) Order Placement and Adoption Date



(b) Order Collection and Adoption Date



(c) Order Placement and Diffusion Date



(d) Adoption and Diffusion Date

Fig. 2. Observing Demand, Adoption, and Diffusion Timing. This figure depicts observational evidence on the relationship between the timing of order placement, order collection, Sharawakil adoption, and Sharawakil diffusion. Sub-figure (a) depicts a weighted scatter plot of the administratively-recorded date of order placement and respondent-reported date of first use of Sharawakil for a customer order (which are primarily experimental orders). Sub-figure (b) depicts a weighted scatter plot of administratively-recorded date of Sharawakil-adorned garment collection and respondent-reported date of first use. Sub-figure (c) depicts a weighted scatterplot of learner-reported date of learning and administratively-recorded date of teacher order placement for learner-teacher dyads. Sub-figure (d) depicts a weighted scatterplot of that same learner-reported date of learning by teacher-reported date of teacher’s first use of Sharawakil. The line in each sub-figure represents the 45° line, i.e. the same date. Bubbles above the line in each sub-figure signify that the event on the y-axis occurred after the event on the x-axis. Bubbles below the line in each sub-figure signify that the event on the y-axis occurred before the event on the x-axis. Bubbles directly on the line signify that both events happened on the same date.

(1) and (2), we show that our main finding is explained entirely by diffusion to non-competitors. We observe almost no diffusion from any treatment group to their business-as-usual competitors, suggesting that baseline competition may have inhibited diffusion of our exogenously-seeded technology.

In addition to studying business-as-usual competition in the network, the demand intervention was designed to randomly attenuate the rivalry of Sharawakil-making skills for training-plus-demand firm owners. Absent a secret shopper structure, the orders placed by our implementation team were pitched (and implemented) as follows: demand was limited to a fixed number of garments (driven by resources), demand shocks were rolled out in a random order, and each firm would receive at most one purchase request (of a random size). This structure implies that trained firm owners with demand shocks, having already received their one-time purchase order, no longer faced experimental competition with potential learners. Trained firm owners without demand shocks would weaken their chance of receiving an experimental order by teaching their peers, a form of rivalry that may have limited diffusion by this group and which mirrors most closely competition in the non-experimental market. In other words, the demand interven-

tion generated an explicitly less-rival demand environment for training-plus-demand firms, potentially increasing the willingness of these firm owners to diffuse technology.

Absent data from firm owners on their decision-making processes, we are unable to test directly for whether experimentally-induced differences in demand rivalry drive our key finding. In Columns (3) and (4) we explore heterogeneity by baseline propensity to share a rival resource with a garment-making peer: dictator games are a direct measure of willingness to share in the face of such rivalry. The wedge between training-only and training-plus-demand firms is attenuated among altruistic firm owners, where the relative rivalry of Sharawakil-making skills may have constituted a less-severe disincentive. Among non-altruistic firm owners, training-only firms that face business-as-usual rivalry respond by teaching almost none of their peers. Training-plus-demand firms, in contrast, behave similarly across baseline levels of altruism, consistent with the argument that our research design shut down experimental rivalry of Sharawakil-making skills for firms assigned to the demand intervention treatment. Note that one could also argue that some sense of fairness contributed to this pattern. While nailing down competition as the sole strategic consideration is a chal-

Table 6

Intention-to-Treat Effects on Diffusion by Teacher and Learner Characteristics. This table reports intent-to-treat effects of treatment group assignment on diffusion to various learner and teacher types. Columns 1 and 2 show diffusion patterns to new and old contacts, respectively. Columns 3 and 4 consider those within and outside of the teacher's trade association (if applicable). Columns 5 and 6 include those learners who were and were not matched as competitors to the teacher using our market research survey. Columns 7 and 8 show teaching by teacher baseline responses to hypothetical dictator games about their network contacts. Diffusion reflects either teacher or learner reporting. A gender dummy is included to control for the randomization strata. Robust standard errors in parentheses. ***p < 0.01, **p < 0.05, *p < 0.1.

	(1) Competitor Learner		(4) Altruistic Teacher		(6) Pre-Existing Teaching		(8) Pre-Existing Contact		(10) Ta Co-Members	
	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Demand Only	0.12** (0.05)	0.00 (0.00)	0.07 (0.10)	0.17*** (0.04)	0.07 (0.04)	0.06*** (0.02)	0.03 (0.03)	0.10*** (0.04)	0.12*** (0.04)	0.01 (0.03)
Training Only	0.24 (0.17)	0.08 (0.07)	0.09 (0.21)	0.52** (0.25)	0.17 (0.17)	0.15 (0.10)	0.12 (0.10)	0.20* (0.12)	0.34* (0.18)	-0.02 (0.02)
Training-plus-demand	1.21*** (0.28)	-0.00 (0.00)	1.13*** (0.40)	1.30*** (0.43)	0.85*** (0.24)	0.35*** (0.08)	0.32** (0.13)	0.88*** (0.18)	1.08*** (0.25)	0.13* (0.07)
Prob > F(Training Only = Training-plus-demand)	0.00***	0.30	0.01***	0.11	0.02**	0.11	0.22	0.00***	0.01**	0.04**
Control Mean	0.07	0.00	0.01	0.00	0.07	0.00	0.03	0.03	0.05	0.05
Avg. # Possible	414.00	2.02	417	417	413.45	2.54	399.44	16.53	403.94	9.28
Observations	384	384	177	206	384	384	384	384	383	383

lenge, we interpret this finding as suggestive evidence in a real network, in a quasi-normal business transaction, of strategic unwillingness to share technology with industry peers.

If this interpretation is correct, a natural next question would be what mitigates these competitive disincentives to share between small firm peers. We observe more connections and more preexisting technology diffusion between business-as-usual competitors at baseline, a fact that could derive from some kind of cooperative sharing equilibrium but also potentially stem merely from the endogenous seeding of non-experimental technologies into the network. Columns (5) through (10) explore heterogeneity by preexisting network relationship. We find little difference in our key pattern across these subgroups, suggesting that any prior cooperative contracts do not mitigate the wedge in teaching behavior created by varying rivalry of Sharawakil-making skills.

A final important caveat when taking evidence from this paper to the policy space is that our experiment involved a single buyer, in a one-shot fashion, making purchases from a large share of the network. Balancing authenticity and identification, we attempted to naturalize the experiment as much as possible while implementing a specific experimental design. The fact that the weave is still in use a year later suggests that experimental demand jump-started organic demand as well, perhaps offering reassurance that our findings may be applicable to policy applications that vary many of the features of our experimental setting.

6. Conclusion

Increases in firm productivity are the backbone of economic growth. Understanding how and when technology upgrading occurs is thus a central challenge for academics and policymakers interested in combating poverty. As a conceivably scalable alternative to direct intervention, network-based technology diffusion presents both an opportunity and a puzzle. Why do we observe some peer-to-peer technology diffusion within industry networks, but not full access to new technologies across

the board? What market incentives and barriers drive the observed pattern?

In this paper, we report the results of a field experiment designed to study technology diffusion in Ghana's garment making sector. We designed a new weaving technique and randomly varied initial training in a real network of garment making firm owners, finding little adoption and diffusion before the introduction of experimenter demand for the technique. Upon the arrival of experimental demand, we see an increase in both adoption and diffusion of the design. We find that, although training renders incumbent adopters equally able to produce the new weaving technique regardless of experimental demand, the majority of teaching comes from those who received demand shocks. The policy applications of this insight are important, as it could be applied to the design of other training programs aiming to use training to spur technology upgrading in small-scale manufacturing firms.

Our findings also suggest that the incentives of both potential learners and potential teachers are important to consider in understanding the determinants of technology diffusion. In particular, contexts in which demand for new technologies is rival may exhibit diffusion dynamics that are different from those where demand is non-rival. We leave further exploration of this insight to future research on peer-to-peer diffusion.

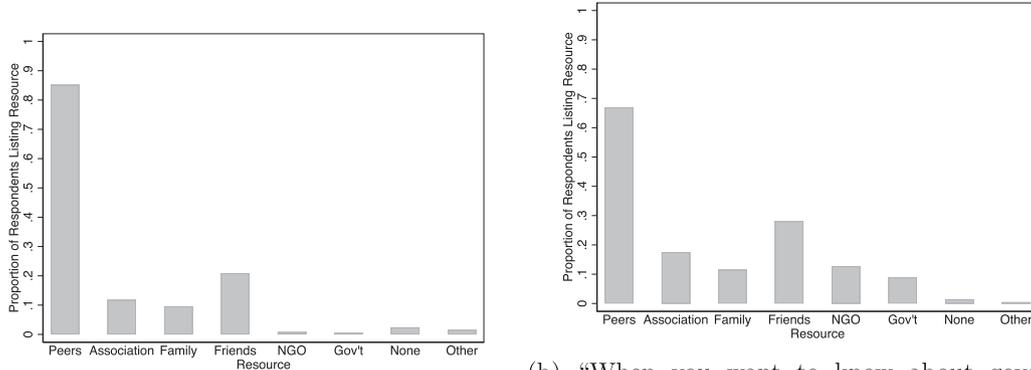
Author statement

Morgan Hardy: Conceptualization, Methodology, Software, Analysis, Investigation, Resources, Funding Acquisition, Writing. Jamie McCasland: Conceptualization, Methodology, Software, Analysis, Investigation, Resources, Funding Acquisition, Writing.

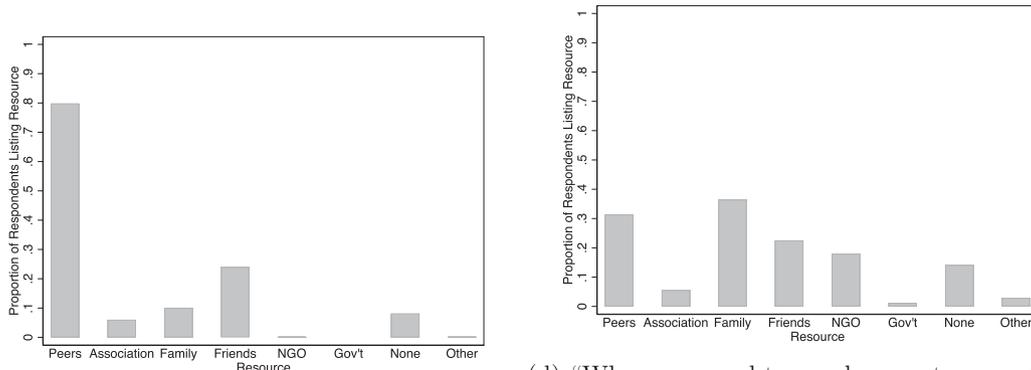
Data availability

Data will be made available on request.

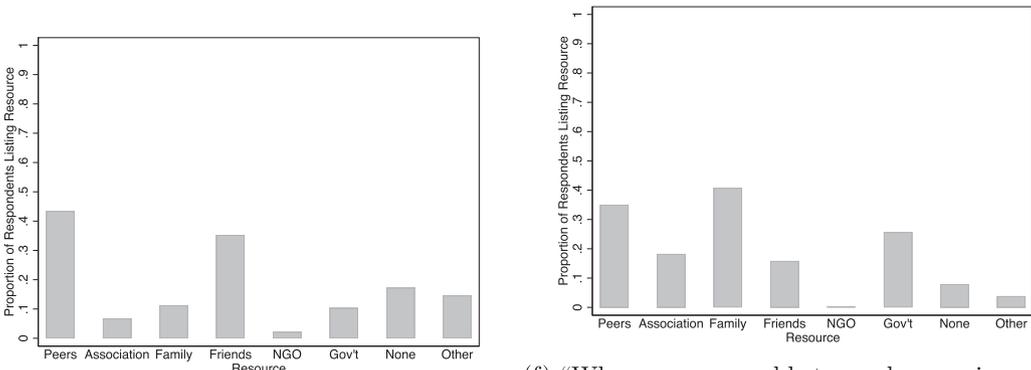
Appendix



(a) "When you want to learn a new technology or skill relevant to your business?" (b) "When you want to know about government/NGO programs and other opportunities relevant to your business."



(c) "When your current labor force isn't sufficient to handle your work flow." (d) "When you need to purchase or to use a tool or machine you do not already have for your business."

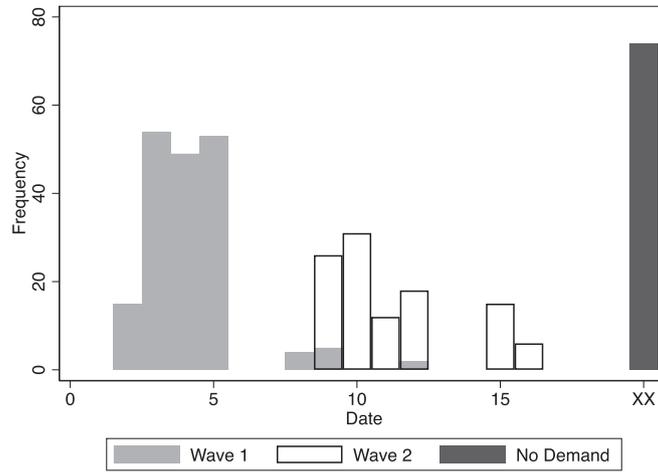


(e) "When you are seeking more customers or contracts for your business." (f) "When you are unable to resolve a serious dispute on your own, involving your business. For example, a theft or breach of contract?"

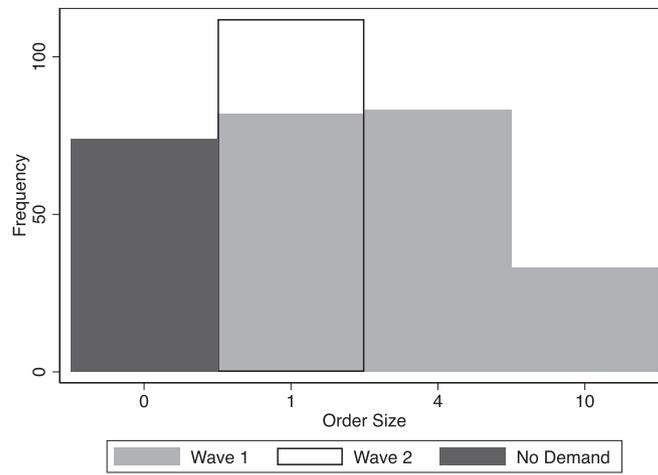
Fig. A1 Available Resources for Business Needs. This figure visually depicts baseline responses to the question "We are interested in what resources are available to help you with various business relevant needs. I will read to you various things you may need as a business owner and you will tell me which of the following is your most important resource?".



Fig. A2 Sharawakil Design Calendar. With the permission of the models, this figure depicts the design calendar for Sharawakil that was given to any firm owner who was trained to display in their shop/advertise their new product. Photoshoots were conducted at the IPA Ghana Office and in Makola Market, Accra. Models are IPA Ghana staff members and friends and family of the Sharawakil designer, Mr. Osman Mutawakil, pictured top right.



(a) Date of Order Placement



(b) Size of Order Placed

Fig. A3 Demand Size and Timing. This figure depicts frequency histograms of size and timing by wave.

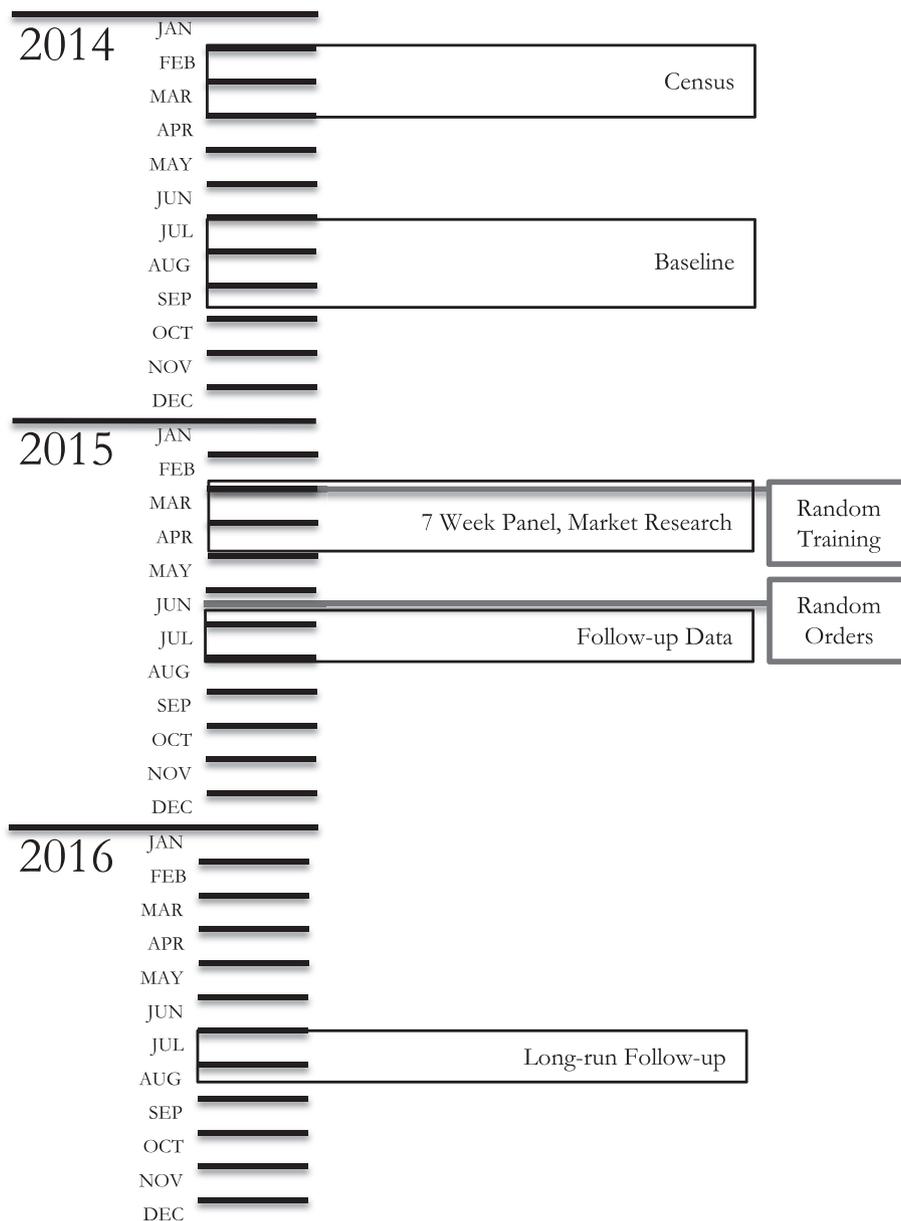


Fig. A4 Project Timeline. This figure depicts the data collection timeline for the project. Firm owners were identified during the census and network relationships over the previous year were measured during the baseline, with firm and owner characteristics measured across both. Weekly network interactions and new technique adoption were documented during the seven week panel that began a week before the training. Order contracts were offered and completed and network interactions documented during the month covered by the experimental endline. A contract offer of one garment was attempted for the entire sample one year following the experiment, for which Sharawakil inclusion was not required, but cash-incentivized.

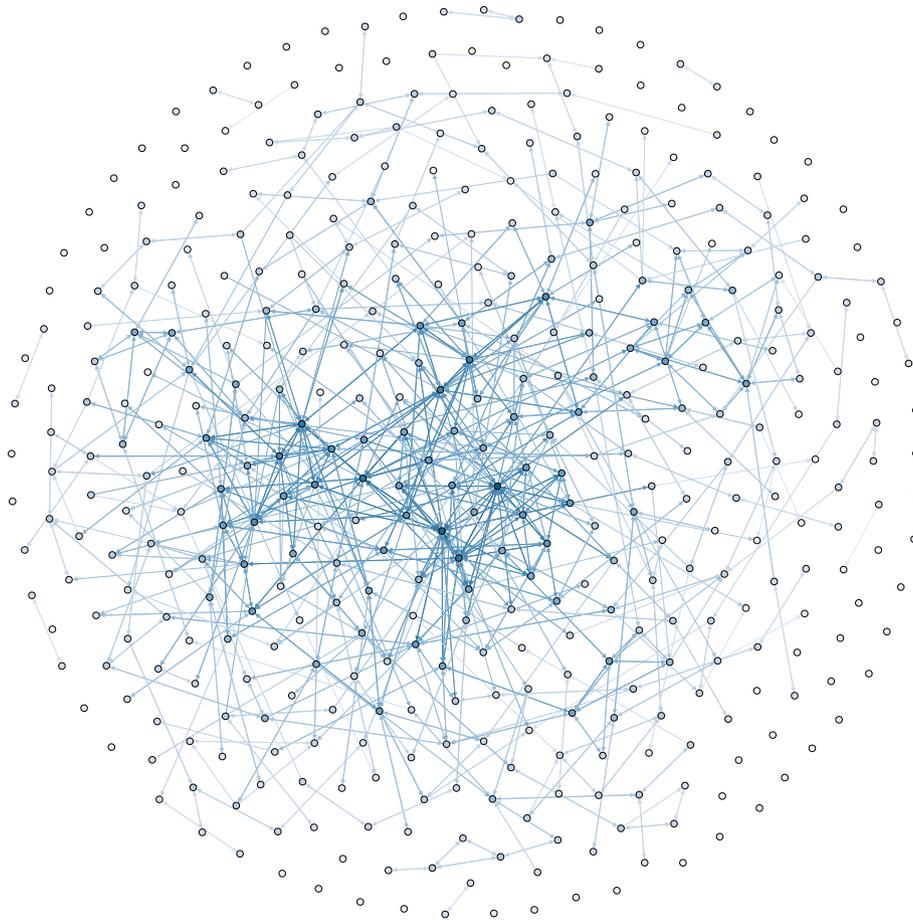
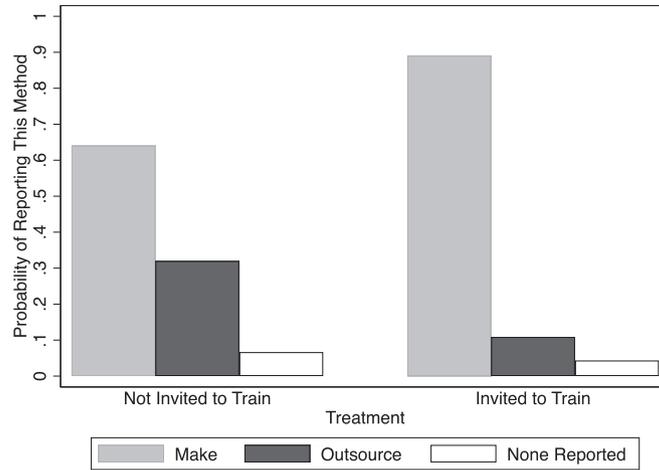
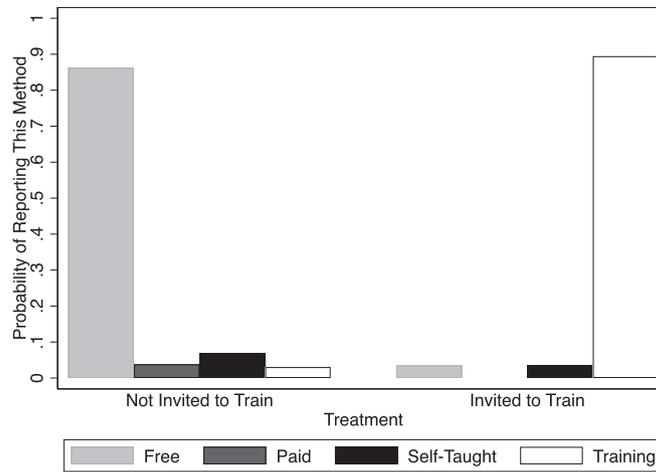


Fig. A5 Technology Diffusion at Baseline. This figure displays the directed graph of technology (skill/technique) diffusion reported at baseline, with arrows on each edge indicating the direction of diffusion. An edge appears between two nodes if either member of that dyad reported diffusion occurring during the year preceding the baseline. The nodes are placed using a “Fruchterman-Reingold” algorithm that attracts connected nodes and repels unconnected nodes. More profitable nodes are darker.

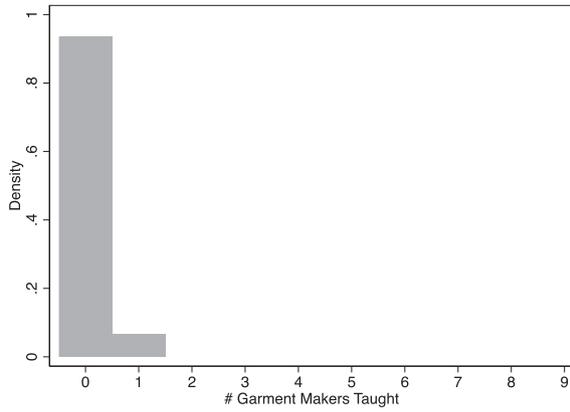


(a) Method, Conditional on Adoption

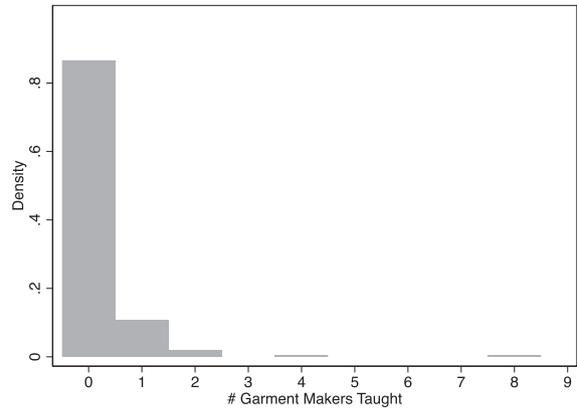


(b) Method, Conditional on Learning

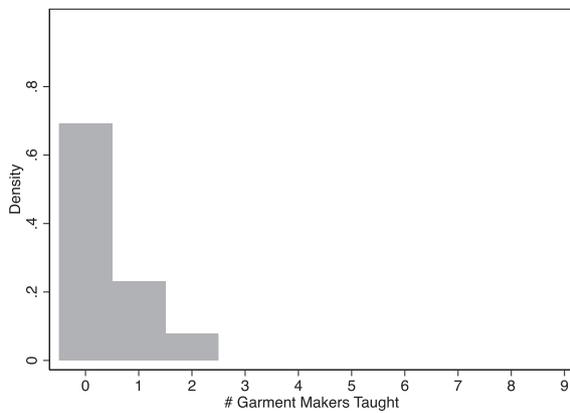
Fig. A6 Sharawakil Learning and Adoption Method. This figure displays conditional methods of learning and adopting the new technique by Phase 1 treatment group. Method were self-reported during the experiment follow-up data collection activity.



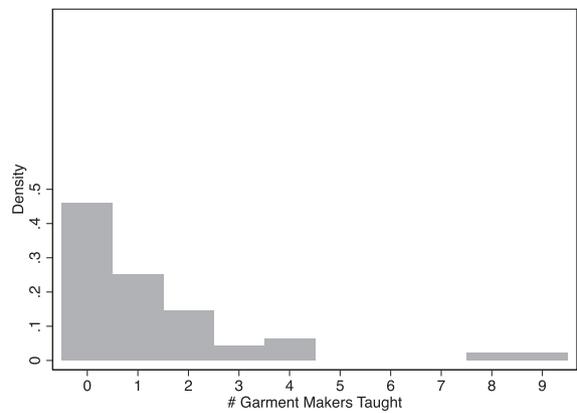
(a) Control



(b) Demand Only



(c) Training Only



(d) Training and Demand

Fig. A7 Number of Other Garment Makers Taught by Treatment Group. This figure depicts density histograms of number of other garment makers taught during the order period by each of the four treatment groups. Diffusion reflects either teacher or learner reporting.

Table A1

Heterogeneous Intensive Margin Effects of Random Order Size. This table reports heterogeneity by training group assignment in intent-to-treat effects of increasing order size on Phase 2 order completion, production ability, and diffusion for firm owners randomized to receive demand. Number of experimental orders completed comes from our administrative data. Ability is self-reported by the owner and verbally verified (through description of the algorithm) by our enumeration staff. Quality measures come from expert evaluations of Sharawakil quality on completed garments (for larger order sizes, one garment is randomly selected for review) using a 10 point rubric. Diffusion reflects either teacher or learner reporting. Ownership of a working Sharawakil tool was self-reported and visually verified by our enumeration staff. Spending on the tool through the market is a dummy indicative of reporting any positive expenditure on Sharawakil related capital (rentals, repairs, purchases) during the order period. Orders were given out over two waves. Wave 1 contained orders of size 1,4 and 10. Wave 2 contained orders only of size 1. Wave 2 firm owners are dropped from analysis in all columns except column 3 to not confound order size and timing effects. A gender dummy is included to control for the randomization strata. Robust standard errors in parentheses. ***p < 0.01, **p < 0.05, *p < 0.1.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Number Phase 2 Garments Completed	Able to Produce Sharawakil	Z-Score Sharawakil Quality	# Peers Taught Sharawakil	Own Working Sharawakil Tool	Spending on Tool Through Market	# Peers Taught Sharawakil
Panel A - Saturated							
Order Size Four	1.49*** (0.24)	0.03 (0.09)	-0.10 (0.22)	0.01 (0.13)	-0.02 (0.06)	0.05 (0.05)	-0.03 (0.14)
Order Size Ten	6.08*** (0.92)	0.03 (0.12)	-0.24 (0.25)	0.39* (0.22)	0.05 (0.09)	0.28*** (0.10)	0.19 (0.24)
Order Size Four * Training-plus-demand	0.81 (0.58)	-0.03 (0.09)	0.28 (0.39)	0.02 (0.95)	-0.24 (0.17)	0.04 (0.17)	-0.01 (0.96)
Order Size Ten * Training-plus-demand	1.43 (1.78)	-0.02 (0.12)	0.05 (0.53)	-0.11 (0.97)	-0.13 (0.20)	-0.26 (0.21)	0.07 (0.97)
Spending on Tool Through Market							0.70** (0.33)
Prob > F(Size Four = Size Ten)	0.00***	0.97	0.57	0.05	0.47	0.03	0.19
Panel B - Linear							
Number of Extra Shirts (0,3,9)	0.67*** (0.10)	0.00 (0.01)	-0.03 (0.03)	0.04* (0.02)	0.00 (0.01)	0.03*** (0.01)	0.02 (0.03)
Number of Extra Shirts (0,3,9) * Training-plus-demand	0.17 (0.19)	-0.00 (0.01)	0.01 (0.06)	-0.01 (0.11)	-0.02 (0.02)	-0.03 (0.02)	0.01 (0.11)
Spending on Tool Through Market							0.70** (0.33)
Demand Only and Order Size One Mean	0.63	0.49	0.09	0.20	0.49	0.09	0.20
Training-plus-demand and Order Size One Mean	0.77	1.00	0.01	1.38	1.00	0.01	1.38
Observations	209	198	126	198	198	198	198

References

Abeberese, Ama Baafra, Ackah, Charles Godfred, Asuming, Patrick Opoku, 2018. Productivity losses and firm responses to electricity shortages: evidence from Ghana. *World Bank Econ. Rev.*

Atkin, David, Amit, K Khandelwal, Adam, Osman, 2017a. Exporting and firm performance: evidence from a randomized experiment. *Q. J. Econ.* 132 (2), 551–615.

Atkin, David, Chaudhry, Azam, Chaudry, Shmyla, Amit, K Khandelwal, Verhoogen, Eric, 2017b. Organizational barriers to technology adoption: evidence from soccer-ball producers in Pakistan. *Q. J. Econ.* 132 (3), 1101–1164.

Bandiera, Oriana, Rasul, Imran, 2006. Social networks and technology adoption in northern Mozambique. *Econ. J.* 116 (514), 869–902.

Beaman, Lori, BenYishay, Ariel, Magruder, Jeremy, Ahmed, Mushfiq Mobarak, 2018. Can Network Theory-Based Targeting Increase Technology Adoption? Working Paper.

BenYishay, Ariel, Mobarak, A Mushfiq, 2018. Social learning and incentives for experimentation and communication. *Rev. Econ. Stud.*

Bloom, Nicholas, Van Reenan, John, 2007. Measuring and explaining differences in management practices across countries. *Q. J. Econ.* 122 (4), 1351–1408.

Bloom, Nicholas, Van Reenan, John, 2010. Why do management practices differ across firms and countries? *J. Econ. Perspect.* 24 (1), 203–224.

Bloom, Nicholas, Benn, Eifert, Mahajan, Aprajit, McKenzie, David, Roberts, John, 2013. Does management matter? Evidence from India. *Q. J. Econ.* 128 (1), 1–51.

Bustos, Paula, 2011. Trade liberalization, exports, and technology upgrading: evidence on the impact of MERCOSUR on argentinian firms. *Am. Econ. Rev.* 101 (1), 304–340.

Cai, Jing, Szeidl, Adam, 2017. Interfirm relationships and business performance. *Q. J. Econ.* 133 (3), 1229–1282.

Conley, Timothy, Udry, Christopher, 2010. Learning about a new technology: pineapple in Ghana. *Am. Econ. Rev.* 100 (1), 35–69.

Fafchamps, Marcel, Quinn, Simon, 2013. Social Networks and Business Practices: Evidence from a Randomized Experiment with Microentrepreneurs. Working Paper.

Fafchamps, Marcel, Quinn, Simon, 2016. Networks and manufacturing firms in Africa: results from a randomized field experiment. *World Bank Econ. Rev.* 32 (3), 656–675.

Foster, Andrew, Rosenzweig, Mark, 1995. Learning by doing and learning from others: human capital and technical change in agriculture. *J. Polit. Econ.* 103, 1176–1209.

Hardy, Morgan, McCasland, Jamie, 2019. Lights off, lights on: the effects of electricity shortages on small firms. *World Bank Econ. Rev.*

Hsieh, Chang-Tai, Klenow, Peter, 2009. Misallocation and manufacturing TFP in China and India. *Q. J. Econ.* 124 (4), 1403–1448.

KPMG, 2014. Material Gains. *High Growth Markets Magazine*. September.

McKenzie, David, 2020. Small Business Training to Improve Management Practices in Developing Countries: Re-assessing the Evidence for “Training Doesn’t Work”. Working Paper.

McKenzie, David, Woodruff, Christopher, 2014. What are we learning from business training and entrepreneurship evaluations around the developing world? *World Bank Res. Obs.* 29 (1), 48–82.

Munshi, Kaivan, 2004. Social learning in a heterogeneous population: technology adoption in the Indian green revolution. *J. Dev. Econ.* 73, 185–213.

Quartey, Peter, 2006. The Textiles and Clothing Industry in Ghana. *Book Chapter in The Future of the Textile and Clothing Industry in Sub-Saharan Africa*.

Syversen, Chad, 2011. What determines productivity? *J. Econ. Lit.* 49 (2), 326–365.

Verhoogen, Eric, 2020. Firm-Level Upgrading in Developing Countries. Working Paper.