IT TAKES TWO: EXPERIMENTAL EVIDENCE ON THE DETERMINANTS OF TECHNOLOGY DIFFUSION

Morgan Hardy
New York University - Abu Dhabi

Jamie McCasland
University of British Columbia

September 27, 2016

Abstract

Previous studies of peer-to-peer technology diffusion have primarily focused on the decision of potential adopters. Equally relevant for observed diffusion in many contexts is the willingness of incumbent adopters to actively share technology. We report the results of a field experiment that considers both parties involved in the diffusion process. Specifically, we develop a new weaving technique and randomly seed both training and a limited number of one-time technique-specific contracts in a real network of garment making firm owners in Ghana. We find that firms that need the technology to complete the contract learn it from firms that received training; however, firms selected to receive only training and no contract offer are much less likely to share the technology than those selected to receive both. We document spillover effects in both learning and teaching from baseline technology sharing contacts, but also find that networks display dynamic properties: a large number of firm owners generate new technology sharing contacts in response to the experiment. Teaching patterns to both pre-existing connections and new contacts follow the overall pattern, in which firms selected to receive both training and experimental demand are most likely to share the technology. We interpret these findings as evidence that experimental competition disincentivized diffusion, which we support with additional evidence exploiting random order size and order timing. Finally, we develop a model which conceptualizes observed diffusion patterns as endogenously resulting from both static dyad-specific and dynamic technology-specific costs and benefits to both the potential learner and the potential teacher.

JEL Classifications: D83, D85, D22, L14, O12, O14, O17, O33
Total Word Count: About 10,000

We are grateful to Anna Aizer, David Atkin, Dan Björkegren, Arun Chandrasekhar, Pedro Dal Bo, David Glancy, Andrew Foster, Jeremy Magruder, Isaac Mbiti, Angélica Meinhofer, Sveta Milusheva, Emily Oster, Anja Sautmann, Daniela Scida, Jesse Shapiro, William Violette, Christopher Woodruff, and seminar participants at Brown University, Northeast Universities Development Consortium (NEUDC), Private Enterprise Development in Low-Income Countries (PEDL) conference at Warwick University, NBER Summer Institute Development Workshop, New York University - Abu Dhabi, University of Zurich, University of Ottawa, Johns Hopkins School of Advanced International Studies, The World Bank, and The International Food Policy Research Institute for 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 technical design. 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.

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 Reenan, 2007, 2010). One important source of productivity dispersion is the use of inferior technology and managerial practices (Bloom et al., 2013; Syverson, 2011). Much of the literature on firms in developing countries has focused on (sometimes quite successful) interventions targeting capital, labor, and managerial skill constraints (De Mel, McKenzie and Woodruff, 2008; Hardy and McCasland, 2015; Bruhn, Karlan and Schoar, 2013). While technology upgrading can indeed require capital, skilled workers, and managerial aptitude, basic access to information about a new technology is also a necessary condition to adoption. In our context, small-scale garment making in Ghana, within-industry peers are the most cited source of information about new technology.

While an influential body of research in development economics has analyzed peer-to-peer technology diffusion and specifically learning about new technology in the context of agriculture,[1] relatively little work has been done on learning in manufacturing and industry in low-income countries. There is reason to believe that peer-to-peer diffusion findings from small-scale agriculture, where farmers tend to produce highly tradable cash crops or subsistence agricultural products for home consumption, may not generalize to other contexts. In particular, most small-scale manufacturing firms in developing countries service exclusively local demand, leading to direct competition between neighboring firms. The potential presence of competitive disincentives to share a new technology highlights a more generic point about this literature, namely that it focuses almost exclusively on the adoption decision of the potential new user.[2] However, diffusion of skill-based technologies will in general depend on both the desire of potential adopters to learn the skill and the willingness of incumbent adopters to teach it.

Using only observational data, it is difficult to disentangle the true effects of competition on the willingness of incumbent adopters to share a new technology. For example, in our baseline data, we observe a positive correlation between direct competition over demand and technology

[1] Seminal works include Griliches (1957), Foster and Rosenzweig (1995), Conley and Udry (2010), Munshi (2004), and Bandiera and Rasul (2006).

[2] One notable exception is BenYishay and Mobarak (2015), who consider the importance of the identity of the recipient of initial information seeding, and show an increase in sharing in response to performance incentives.

sharing between firm owners.[3] However, this positive correlation need not necessarily imply that competition itself increases sharing, as there is the potential for confounding factors in the structure of the network and the history of market incentives.

In this paper, we report the results of a field experiment that randomly varied three things: the supply of a new technology, demand for products requiring that technology, and (by virtue of the random timing and rival nature of the demand) a design feature we call experimental competition. We study garment making microenterprises in Ghana, and our sample consists of the universe of garment making firm owners in a mid-size district capital. Made-to-order garments are a staple of Ghanaian culture, making up the majority of clothing worn to weddings, funerals, church, mosque, and holiday events, as well as a sizable share of professional and casual everyday wear. It is typical for new garment styles to become popular periodically and for garment makers to learn these new styles in order to attract and retain customers.[4]

We collaborate with a designer in Accra to design our own style innovation, which we call Sharawakil.[5] The technique involves the use of the motor from a commonly available toy car as a weaving tool to produce a specific pattern in multi-colored thick thread that can be added to augment any garment. The technique was intentionally designed to be very difficult to figure out without being shown. However, we also designed the technique such that, once shown how to construct the weaving tool from the toy car and the correct way of folding the twisted thread so it releases with the correct pattern, it is extremely easy to execute.

We implement a cross-cut randomization in which 15% of the firm owners in our sample were invited to participate in a skills training in this new technique and receive the weaving tool required to produce it. The second randomization, stratified by the first, was implemented several weeks later. It involved a randomly timed rollout of a limited number of randomly sized garment orders featuring Sharawakil, known by firm owners to be without replacement.[6] Firm owners who were not at the training had the opportunity to refuse the order on the grounds of not being able to produce the design, or the option to take the order and find a way to learn the design in the time necessary

[3] Competitors in the baseline market were identified using self-reports of customer diversion and a market research survey which linked firms that shared customers.

[4] In our market research survey, conducted on over 1,600 randomly selected district residents, availability (or lack) of desired styles is the number one reason for choosing (or leaving) specific garment makers.

[5] A mix of the designer's name, Osman Mutawakil, and "shara", a Hausa word for "twisted."

[6] Without replacement means that once a firm owner received an order from IPA, she would not receive a second order regardless of the number of orders completed by herself or others.

to complete the order. Firm owners who attended the training were thus a natural resource.

The limited amount of orders is the key to the experimental competition portion of our design, as groups of firms were invited to complete orders in waves. The use of experimental waves was partly logistical. Because we could not know a priori how many orders would be accepted, it allowed us to manage on-the-spot the pool of resources available for making garment orders. However, It also introduced experimental competition, as unfilled orders would imply an extension of the program into later waves. Firm owners who had already received an order were no longer competing with other firms for these potential future orders featuring Sharawakil, while firm owners who had not yet received an order faced a competitive disincentive to teach others the design.

Throughout the paper, we consider four treatment groups: training only, demand only, both, and neither. The demand only and neither groups enter the market as potential learners of the technology, with the demand only group having a randomly higher benefit to learning. The training only and both groups enter the market as potential teachers of the technology, with the training only group having a randomly higher cost to teaching.

Data collection included a detailed baseline survey, which captured the full network map within the sample along several dimensions, including the sharing of technology, physical capital, workers, customers, orders and cash. The follow-up data includes 8 rounds of information on network activity, as well as firm-level outcomes. The invitation to training and the training itself happened between rounds 1 and 2, while the demand randomization, order offers, and collections happened between rounds 7 and 8. All firm owner level specifications use the panel of network activity, with round and firm fixed effects. Dyadic regressions use all possible dyad pairs in our sample of 417 firm owners.

We find that demand for the product increases learning for potential learners in the peer-to-peer within-industry market for our technology. That is, demand only firms are far more likely to report learning the skill than neither firms. Nearly all of the demand only firms who report learning the skill, and who in fact successfully complete the offered order, report learning the skill from another garment maker in our sample.

Teaching, however, appears to follow a very specific pattern. Firm owners in the both group, who were invited to the training and offered an order contract, do the bulk of the teaching. Firm owners in the training only group, who were invited to the training but not offered an order (as a function of the order acceptance rates in earlier waves), almost never teach the skill.

In our next set of results, we explore the importance of existing firm owners' technology diffusion networks. While our network data is rich in different types of within industry relationships, we focus in the paper on firm dyads in which technology has been shared in the past. If either member of the dyad taught a skill to the other in the year preceding the baseline survey, they are "baseline technology sharing contacts". The vast majority of our firm owners have at least one baseline technology sharing contact within the sample, and the majority have more than one. Overall, diffusion of the new technology is more likely within these already existing social networks. We present evidence of spillovers from baseline contacts with orders (who increase teaching) and baseline contacts with training (who increase learning). Following the direct intention to treat effects, training spillover effects are only experienced by those firm owners' with a trained baseline contact not in competition over future orders (who received both training and an order contract).

In addition to documenting spillover effects, we find that the network appears quite dynamic. We use our 8 rounds of panel data on network activity to examine a common assumption in the literature and in other experimental work that networks are static. While transactions between (pre-experimental) connected peers are more likely, we find that the experiment generated new network relationships. These new network relationships respond to market forces in a pattern that mirrors our other findings. Demand only firms are more likely than others to report learning a skill from a new contact, someone they had not reported interacting with during the entire year preceding the baseline. Similarly, firm owners who were both invited to the training and offered a demand contract report on average nearly one new contact to whom they taught the skill.

We interpret these main empirical findings as evidence that our experimental competition design generated disincentives to teaching. Potential teachers who did not receive an order, and whose potential for a future order depended on the ability of other firm owners to complete current orders, were less willing to teach the technique than those for whom this was no longer a factor. Additional findings exploiting self reports of ability to produce the style, expert assessments of weaving quality, the wave structure of our experimental design, and random variation in the size of the order support our interpretation that the patterns documented above are largely driven by experimental competition.

We find that trained firms with an order are not more likely than trained firms without an order to report an ability to produce Sharawakil. Neither are order only firms with larger orders more likely than order only firms with smaller orders to report such ability. We also see no difference in weaving quality between smaller and larger orders. These findings offer evidence against a learning by doing interpretation of our results.

In addition, teaching the skill is not more likely at the time of production (as proxied by the experimental waves in which a firm is offered an order), but increases upon receiving an order and stays high even after the order has been completed. This finding offers some evidence against an interpretation we term teaching while doing, which would posit that it's easier to teach someone the skill if you are currently using it in production. This finding is also consistent with our interpretation of experimental competition, which becomes irrelevant upon receiving and even after completing the order.

Finally, we find that the incidence of sharing increases discretely upon receiving an order of any size and does not significantly increase with order size for firm owners invited to the training. This finding is consistent with the interpretation of orders removing experimental competition, which does not depend on order size. It also offers evidence against learning by doing (because firm owners with both training and an order who received a larger order are not more likely to teach than those that received an order offer for a single garment), teaching while doing (since larger orders have longer production times), and a third hypothesis we consider: fairness or bitterness. If firm owners not offered an order refuse to teach on the basis of a perceived lack of fairness in the process, then it may also be the case that those firms offered an order of size one are less willing to teach than those offered an order of size 10, owing to bitterness over relative order size. We find no evidence to support that hypothesis.

In the last substantive contribution of the paper, we develop a simple model to formalize a new way of conceptualizing technology diffusion. Unlike much of the previous literature, which conceptualizes diffusion as occurring within a static network structure, we model the network as endogenously resulting from both static and dynamic incentives to both potential parties involved. Technology diffusion is conceptualized as occurring (or not occurring) in a market for technology, in which both potential teachers and potential learners face costs and benefits of participating in a diffusion transaction. These costs and benefits depend both on static pairwise characteristics (e.g. distance, gender, ethnicity), and on dynamic and technology-specific costs and benefits to learning or teaching. These technology-specific incentives, as well as the initial seeding of technology within the industry, can be systematically related to static characteristics, confounding their effects in conventionally observed network data. This can make it difficult to disentangle market based incentives from more static barriers to diffusion, such as ethnicity or gender. Because our experiment exogenously seeded a new technology, and exogenously varied the costs and benefits of teaching and learning, this allows us to disentangle these underlying determinants of diffusion. For example, we see the positive relationship between baseline market competition and technology diffusion disappears when we limit analysis only to diffusion of our randomly seeded and demanded technology. However, we see that homophily in diffusion by gender or ethnicity, albeit attenuated, remains. The model is useful for exploring the external validity of our experimental results within our own context, in addition to thinking about future work that considers the incentives of both parties in peer-to-peer technology diffusion in other contexts.

Our paper contributes to a small but growing literature on network-based technology diffusion in manufacturing in low-income countries. In a pair of papers that randomized network links, Fafchamps and Quinn (2013, 2015) find that new firm-to-firm connections led to limited diffusion of managerial practices, in contrast to cross-sectional evidence from Ethiopia and Sudan that suggests business practices of network-linked firms are actually negatively correlated (Fafchamps and Soderbom, 2013). Atkin et al. (2015) find little network-based diffusion in the context of Pakistani soccer-ball producers, and instead focus on organizational barriers to technology adoption that relate to principal-agent issues between firm owners and their workers. Cai and Szeidl (2016) find observational evidence that information diffuses less between firm owners in direct competition with one another. Other influential explorations of low levels of technological upgrading in manufacturing include Parente and Prescott (1994), who model differences in barriers to technology adoption across countries, Bloom et al. (2013) who find an unwillingness among firm owners to delegate managerial tasks, and Tybout (2000) and Atkin, Khandelwal and Osman (2014) who consider the productivity enhancing benefits of trade.[7] Our paper also relates to work by Bandiera, Bankerey and Rasul (2005) and others on field experiments within and across firms, which is nicely surveyed

[7] Atkin, Khandelwal and Osman (2014) also focus on the demand side in thinking about technology adoption and productivity, something we focus indirectly on in this paper. Other recent work by Ferraz, Finan and de Alencar Szerman (2015) finds persistent effects of winning government procurement contracts, though work on the demand side is otherwise quite limited.

in Bandiera, Bankerey and Rasul (2000).

In studying small firm growth, we relate to the recent experimental literature that considers interventions aimed at growing both employment and profitability in small firms. For example, De Mel, McKenzie and Woodruff (2008) and Fafchamps et al. (2013) find high rates of return to capital in microenterprises; Hardy and McCasland (2015) find evidence that small firms face large screening costs in hiring; and Bruhn, Karlan and Schoar (2013) study the impacts of business training. Brooks, Donovan and Johnson (2016) find short-run profit effects of randomly pairing young entrepreneurs with older mentors.

The networks literature includes several papers that highlight the potential importance of business networks for businesses and job seekers, though ours focuses more specifically on their role in technology adoption (Greif, 1993; Casella and Rauch, 2002; Barr, 2000, 1998; Patnam, 2011; McMillan and Woodruff, 1999; Munshi, 2011; Beaman and Magruder, 2012). Immorlica, Lucier and Sadler (2014) develop a model particularly relevant to our context, in the sense that they consider incentives to withhold information that coexist with incentives to share information. Banerjee et al. (2012) have experimental variation in rival information about participation in an experimental game, and find evidence of withheld information. Finally, our findings relate to the literature on endogenous network formation (Udry and Conley, 2004).

The paper proceeds as follows: In Section 2, we describe the experimental design. In Section 3, we discuss the data and estimation. Section 4 discusses the results. Section 5 focuses on interpretation of our results. Section 6 presents our model and explores its insights and implications. Section 7 concludes.

2 Experimental Design

2.1 Industry Background

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. 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)[8]. 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). In our own market research survey of nearly 1,600 people in our study district, respondents averaged consumption of 3.5 bespoke garments in the last year. At average garment prices and 2014 average exchange rates, this amounts to about 2% of GDP.

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. 66% of our sample has an electrically powered machine of some kind, while the remainder exclusively use human-powered machines. 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 20% of our sample and tend to be both bigger and more profitable. 40% of our sample (cut about evenly between those with and without electric equipment) employ at least one apprentice or other paid worker. 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 utilized for skill acquisition, to learn about NGO and government opportunities targeting garment makers, to discuss business, and to adapt to highly variable product demand.

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). 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 trainings in new fashion designs and production technologies.

[8]Industrial textile production for export in Ghana has declined sharply in the post-independence period, and more recently due to major electricity shortages. 5,000 people were employed in factory-style garment making jobs in Ghana in 2000, down from 25,000 in 1975 (Quartey, 2006).

2.2 Sample Recruitment

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. The dominant ethnic group is Ewe, a community that straddles the Ghana-Togo border from the coast up to the Northern reaches of the Volta Region.[9]

The sample we aim to study is the universe of garment making firm owners in Hohoe town and its outlying suburbs. Sample recruitment took place in February of 2014, and included a census of all garment making firm owners in all of 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 1,024 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 be able to easily divide income, variable expense, and profit streams, and share primarily physical space and physical capital.

In Hohoe town and surroundings, the February 2014 census identified 445 garment making firm owners. Of these 445, 417 were still operating a business in Hohoe town or a surrounding suburb at the time of the experiment, 95 of whom are men and 322 of whom are women. Table 1 shows the summary statistics for the Hohoe town and surroundings sample which we use for the majority of this paper.

TABLE 1

Our sample is 75% Ewe. 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. Owners have an average of 9 years of school, equivalent to a junior high school education, while some have no formal schooling

[9]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.

and some have tertiary degrees. The average monthly profits at baseline are 138 Ghana Cedis (GHC), which is equal to approximately 60 USD at the time of the survey. Figure 1 shows the distribution of baseline technology sharing degree.

FIGURE 1

If we consider all possibly dyadic relationships, we find in our baseline that the average probability of a skill sharing relationship is .7%. Firm owners report interactions over many dimensions outside of just technology sharing, including working for one another, borrowing each others' equipment, outsourcing subsections of certain orders to one another, discussing prices, exchanging of gifts or loans, discussing business practices, complaining about the continual "lights out" crisis in Ghana or even discussing personal matters. The average probability of any interaction, including something as simple as "shared greetings" is 3.7%.

2.3 Study Design

In this paper we report on a cross-cut randomized controlled trial that randomly seeded both invitation to a specific garment design skills training, and demand for that garment design via direct purchase from the project, without replacement. Stratified by gender, 15% of the 445 firms from the census sample in Hohoe District and its suburbs were randomly invited to participate in a skills training to learn a new garment design and receive the relevant equipment needed to produce it. The training was held at the Ghana Educational Services center in Hohoe. Although IPA was acknowledged as a sponsor of the training, the district coordinator was happy to take credit for inviting and organizing the design training, helping to obscure IPA's interest or involvement with this specific technique.

After the first seven weeks of data collection showed very little take-up (let alone teaching) of the design, we realized we would need to create our own artificial demand to observe diffusion dynamics. 29 firm owners had been confirmed to have left the sample (by leaving the industry or leaving the district permanently) in the period between our census and the start of our experiment, and these firms were thus dropped from the subsequent demand randomization. The remaining 417 firms were then randomly split into two "experimental demand" waves, stratified by gender and the skills training treatment assignment.

All 209 firms in wave 1 were selected to be offered a contract first, though they randomly varied in size. 88 wave 1 firms were selected for a contract to produce one garment, 88 wave 1 firms were selected for a contract to produce four garments, and 33 wave 1 firms were selected for a contract to produce ten garments.[10]

Firms in the second wave were chronologically delayed until the completion of wave 1 and randomly split into three subgroups, intended to allow for real-time design-based management of project resources and to introduce experimental competition. Of these subgroups, two were ultimately offered contracts, on the basis of the completion rates in wave 1 and the first subgroup of wave 2. Every firm offered a contract in wave 2 was offered a contract to produce only one garment. In total, about 20% (89) of the firms in the sample were never offered a contract of any size.

Our order team was given a loose script to follow, collecting a bit of administrative data on previous design knowledge and ability as well as highlighting main points to cover. They were instructed to (1) stress the importance of incorporating Sharawakil for the garment to be accepted, (2) explain that we would continue to make orders depending on resources left after each contract is either completed or rejected, and (3) inform selected garment makers that the contract is a one time offer for a random number of garments.

Order contract offers specified a fixed price of 35 GHC for each shirt or other garment featuring the new design. The price was chosen to reflect a generous price for intricately designed garments. The median prices for a standard (unembellished) slit and kabbah (traditional Ghanaian skirt and top for women) and standard (unembellished) man's shirt are 20 GHC, and 12 GHC respectively, in our baseline survey. However, more intricate designs bring these prices up to as much as 40 GHC. As is typical for this market, we gave a 10GHC deposit for each garment. For orders larger than one garment, we asked for a sample garment and verified the accuracy of the new design before submitting the full deposit.

Figures 2 and 3 show the distribution of network level treatments, with the number of trained contacts in figure 2 and the total number of shirt orders offered to baseline contacts in figure 3. Because those who receive training invitations and orders are random, network level treatments

[10] 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.

are also random, conditional on baseline degree.

FIGURE 2 and **FIGURE 3**

2.4 New Technology

The new design is called Sharawakil, and was commissioned for the project from a somewhat well-known designer in Accra. It involves a technique of twisting thread using a deconstructed common children's toy car, easily found in local markets for around 5GHC at the time of the experiment. Figure 4 depicts the use of the tool in spinning the thread to create Sharawakil. After spinning, the thread is folded using a particular technique before it is released and collapses into the final product, shown in figure 5. The colorful result can and has been applied to garments in many creative ways, an example of which is shown in figure 6.

FIGURE 4, **FIGURE 5** and **FIGURE 6**

This design was created for the sole purpose of this experiment, and has some important features for the purpose of our study: (1) It is very difficult to figure out without being shown how to spin, wrap and release the thread to achieve the correct pattern. (2) Once someone has shown you the technique, it is very easy. Everyone who attended the one-day training had mastered the technique before leaving. (3) It is unisex and versatile in its use and thus should be equally valuable to all kinds of firm owners. (4) It requires no electricity and minimal capital investment and thus it should be equally accessible to all kinds of firm owners. (5) The algorithmic nature of the technique and the results allows for a clear assessment of the quality of the Sharawakil itself, while the freedom in application allows for variation in innovation, allowing us to observe both through analysis of the orders completed.

3 Data and Estimation

3.1 Data

What we refer to in the paper as the baseline network is the listing of all network relationships and network activity among garment making firm owners in our sample, collected in late 2014. The baseline survey targeted the full Hohoe District sample, but we restrict our analysis of baseline network relationships to the 417 firms in our final sample. Data include demographic information, cognitive testing, profits, assets, labor, management, prices, and production breakdown of the most commonly sold products. Also included are general questions about possible substitutes for missing institutions, taxes, and reasons for self-employment, which we took to validate the relevance of networks to the functioning of most of the businesses in our sample.

The network section of the instrument in the baseline survey was designed to create a full within-industry, within-district network map. While firm owners were not prompted to confirm or deny acquaintance with each of the more than 1,000 other firm owners, the self-reported contact section was designed to capture as many relevant contacts as possible. Firm owners were asked whom they know in the district by category (former employer, former employee or apprentice, trade association co-member, neighbor, close friend in the business, etc.). Once all contacts were revealed, 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, and outsourcing. Additionally, these instruments included questions about relative size, experience, work quality, competition, altruism, and trust between the respondent and each network member.

The remainder of the data was collected as weekly follow-up surveys. The first of these took place in March 2015 before the training randomization. The next six took place in the weeks following the training. And the final survey took place in June 2015, after the product demand randomization and purchases. Each survey includes weekly sales and expenses, hours worked by the owner and any other employees, weekly physical outputs in terms of number of garments, and network activity in the preceding week,[11] which mirrors that collected in the baseline survey. These 8 rounds of network activity form the basis of much of our analysis.[12]

In addition to the firm owner level data, we conducted nearly 1,600 customer market research surveys over the course of the study, which attempted to create an additional diversion network mapping competition between firms. Random individuals were approached in public places in

[11]The final survey included two weeks of firm inputs and outputs and asked about the previous month of network activity, generating 9 weeks of firm level data and 8 rounds of network activity data, not including the baseline.

[12]In our data, each type of connection is coded as 1 in a particular round if either firm owner reported the interaction. This is to avoid recall error as we believe firm owners may forget interactions more readily than create them artificially. The results in this paper are qualitatively similar using other definitions of connection (directed or both). However, they are attenuated due to what we believe to be increased measurement error.

Hohoe and its surroundings, systematically near each of the shops in our sample. Each respondent was asked how frequently he or she buys bespoke garments or uses the services of garment making businesses, how much he or she has spent on these services recently, and which businesses he or she frequents, in rank order. In addition, we collected some information on what draws customers to one garment maker or another, or away from their current favored garment maker.

Lastly, three garment making experts (one of whom was the designer) evaluated the quality of all completed orders along four dimensions: (1) quality of the Sharawakil itself, (2) quality of the garment less the Sharawakil, (3) creativity in the use of Sharawakil, and (4) overall estimated cash value of the garment. For firms who were in the 4 garment and 10 garment random treatment groups, we evaluated the single sample garment separately from the remaining 3 or 9 garments, to measure increases or decreases in quality over time.

3.2 Estimation of Direct Effects

We have four main outcome variables of interest: (1) learning the new technology from another garment maker, (2) teaching the new technology to other garment makers, (3) learning a technology from a new technology sharing contact, and (4) teaching the new technology to a new technology sharing contact[13]. Our main specification stacks the 8 rounds of the firm owner level network activity panel, as follows:

$$Y_{it} = \alpha + \beta T_{it} + \eta_t + \varphi_i + \epsilon_{it} \tag{1}$$

where Y_{it} is the outcome of interest, T_{it} is the vector of treatment groups of interest (training only or demand only or both), η_t are round fixed effects, φ_i are firm fixed effects (which control for randomization strata), and ϵ_{it} is an error term. With the pure control group omitted, the coefficients in β are the Intent-to-Treat effects of being assigned each particular treatment group, relative to the control, and are identified from both within round, and within firm variation. Standard errors

[13] An earlier version of this project was registered with the American Economics Association (AEA) Randomized Controlled Trial Registry, complete with a Pre-Analysis Plan (PAP). The PAP was intended to coalesce ideas on the direction of analysis, and limit both the risks and perception of data mining or specification search. The earlier version of the study design did not include random demand for the product, which was inspired partway through implementation both by lower than expected non-experimental demand and skill-sharing in the original design, and field-based insights into what motivates sharing in our context. Consequently, the specifications in this version of the paper do not match those laid out in the PAP. The main hypotheses, however, remain relevant.

are clustered at the firm level.

3.3 Estimation of Spillovers

In the spillovers section of the paper, we estimate:

$$Y_{it} = \alpha + \beta NT_{it} + \omega T_{it} + \eta_t + \varphi_i + \epsilon_{it} \tag{2}$$

where NT_{it} is the vector of baseline technology network treatment (at least one contact invited to training only, at least one contact selected for an order only, at least one contact invited to/selected for both). We continuee to control for round fixed effects, firm fixed effects, and the vector of direct treatment groups. The firm fixed effects control for the baseline degree, making the number of baseline contacts trained random, and allowing for causal identification of β, which in this case can be interpreted as the marginal effect of having at least one baseline network contact in a given treatment group. Standard errors are clustered at the firm level.

3.4 Dyadic Level Analysis

In our dyadic specifications that utilize the full network panel, we estimate:

$$Y_{ijt} = \alpha + \beta_1 T_{it} + \beta_2 T_{jt} + \eta_t + \varphi_i + \varphi_j + \epsilon_{it} \tag{3}$$

where Y_{ijt} is a binary variable, equal to one when diffusion of the new technology has occurred from firm i to firm j in round t. T_{it} is the vector of potential teacher treatment groups (i invited to train, selected to receive an order or both) and T_{jt} is the vector of potential learner treatment groups (j invited to train, selected to receive an order or both). φ_i are potential teacher fixed effects and φ_j are potential learner fixed effects, which control for the randomization strata for both firms in each dyad. The vectors β_1 and β_2 here can be be interpreted as the Intent-to-Treat effects of potential teacher treatment groups and potential learner treatment groups, respectively, on technology diffusion from i to j. Standard errors are two-way clustered by potential teacher firm and potential learner firm.

4 Findings

4.1 Balance

Table 1 shows balance along major observables across the four treatment groups. With the exception of a slightly higher probability of being Ewe for the training only group, a slightly lower age for the training only group, and slightly fewer years of schooling for the training and order group, all observables appear balanced. Firm fixed effects are included in all specifications, which should control for any imbalanced constant unobservables.

4.2 Intervention Compliance

For the training treatment, compliance is defined as actually receiving training in the new style, either through attending one of the trainings offered in the district or from a mop-up training held shortly after the main training. For the order treatment, compliance is defined as actually being offered an order for a garment utilizing the new technology (not necessarily accepting the offer). Not every firm owner was able to attend the training, nor were they reachable to receive an order offer during the scope of the project. The largest reason for non-compliance on both accounts was short term travel out of the district.

Figures 7 and 8 show treatment compliance for training and orders, respectively. Over 90% of those randomly selected for training invitations or order receipt were in compliance. No firm owner not randomly selected to receive training or an order offer did so. This is very high compliance and means that reduced form regressions may be considered very similar to treatment on the treated effects.

FIGURE 7 and **FIGURE 8**

4.3 Order Completion and Learning Behavior

Figure 9 shows the probability of completing an order by training treatment assignment for the 298 firm owners who were offered an order contract. Those invited to the training are 17.8% more likely to accept and complete the order, if offered. However, 67.9% of those who were not invited to the training still complete an order, if offered.

FIGURE 9

Figure 10 shows the reported method of completion by training treatment status. The majority of the 211 garment makers who successfully completed a Sharawakil adorned garment reported doing so by weaving it themselves. However, about 27% of this sample report buying Sharawakil from another garment maker.[14] Another 8% did not report an order completion method. Learning and teaching is both the most common way that this technology diffused in the experiment, and the focus of our interest in this paper.

FIGURE 10

Figure 11 depicts the probability of reporting an ability to produce Sharawakil by treatment group. Over 90% of those invited to the training report being able to use the technology. Those firm owners not invited to train have a 43.5% chance of reporting the ability to use the technology if randomly selected to receive an order and a 16% chance if not.

FIGURE 11

Figure 12 depicts the method of learning to produce Sharawakil, by training treatment assignment, for those who reported such an ability. Although a few garment makers reported paying to learn Sharawakil from another garment maker or self-teaching,[15] the overwhelming majority of garment makers reported learning Sharawakil at the government trainings (if assigned to the training treatment groups) or from another garment maker for free (if not).

FIGURE 12

It is this peer-to-peer technology diffusion that is the subject of interest in this paper. Thus, learning and teaching behavior is the focus of our main empirical findings.[16]

[14] Although we did not ask directly from whom each garment maker purchased Sharawakil, only three firm owners reported sales from Sharawakil to other garment makers suggesting that the selling of Sharawakil was relatively rare.

[15] We'd like to note that some of these "self-teachers" came to the training.

[16] In Section 5.2 we pick up the question of outsourcing when addressing the interpretation of our main empirical findings.

4.4 Treatment Effects on Technology Diffusion

Figure 13 plots the probability of learning the new technology from another garment maker by treatment status. Firm owners selected to receive an order but not training have a 38.5% probability of learning the new technology from another garment maker, while those without an order are less than 0.5% likely to learn the new technology from another garment maker.

FIGURE 13

Figure 14 plots the probability of teaching the new technology to another garment maker by treatment status. 54.2% of firms owners selected to receive both an order and training teach the new technology to at least one other garment marker, while only 2.1% of those with training but without an order teach.

FIGURE 14

Because the orders were all placed during the time period covered by 8, the above figures are more illustrative than experimentally informative. In order to truly estimate the treatment effects on technology diffusion, it is important to control for round fixed effects. Table 2 shows the reduced form treatment effects on learning and teaching, controlling for firm and round fixed effects.

TABLE 2

Firm owners not invited to the training, but randomly selected to receive an order are 26.3% more likely to learn the technology from another firm owner than a firm owner who received neither. Those firm owners randomly selected for a training invitation, but not to receive an order are no more likely to teach the technology to another firm owner than a firm owner who was not invited to train and/or received nothing. Firm owners randomly selected for both training and an order are 44.1% more likely to teach the technology to another firm owner.

4.5 Network Treatment Effects on Technology Diffusion

Next, we explore the importance of the pre-existing network. Table 3 shows the reduced form spillover treatment effects from pre-exiting connections on learning and teaching from and to other firm owners.

TABLE 3

Firm owners with at least one baseline technology sharing contact randomly selected to receive both training and an order are 14.5% more likely to learn the new technology from another garment maker. Baseline contacts randomly selected for only an order or training do not increase the likelihood of learning. Firm owners with at least one baseline technology sharing contact selected to receive an order but not invited to the training are 9.9% more likely to teach another firm owner the new technology. These findings show that, although diffusion is more likely through pre-existing contacts, diffusion is still more likely from those in the training and order group than in the training only group.

4.6 Treatment Effects on Network Formation

Spillovers through a pre-existing network have been the subject of interest for many previous networks related experiments. Because we collected many rounds of network panel data, we are also able to explore dynamic changes to the network in response to our experiment. Table 4 shows the reduced form treatment effects on new network formation, controlling for firm and round fixed effects.

TABLE 4

Firm owners not invited to the training, but randomly selected to receive an order are 27.1% more likely to learn the technology from a new technology sharing contact than a firm owner who received neither. Firm owners randomly selected for both training and an order are 86.5% more likely to teach the technology to a new network sharing contact than those who received neither. This table suggests that diffusion occurs both through newly forming connections, not just through pre-existing ones. We also see that newly formed teaching and sharing patterns match the patterns observed for teaching and sharing between to pre-existing contacts.

4.7 Dyadic Analysis

The findings in table 2, table 3, and table 4, at the firm level, are the result of diffusion dynamics actually observed at the dyadic level. We present them first, because they are much easier to understand and interpret than the more complicated dyadic regression table. However, we have included table 5 to illustrate the results of all three tabes and how they fit into the overall story of what has happened in this experiment.

TABLE 5

Table 5 shows that technology diffusion is mainly occurring from those who received both training and an order and to those who received only an order. Additionally, those who received only an order act as vessels for technology diffusion to those who also received only an order. The majority of technology diffusion happens between those firm owners who were contacts already at baseline, from those who received both training and an order to those who did not receive training, even if they did not receive an order (perhaps in anticipation of a potential order). Also note that those who received only training appear more likely to teach the technology to those in need if they are a baseline sharing contact, even though this finding is not quite significant. However, we also see that diffusion still occurs even between those dyads not already contacts at baseline.

5 Interpretation

The previous section clearly shows that technology is diffusing most between those who received orders (from those trained and to those not trained). We interpret these findings as evidence of a negative effect of experimental competition on the willingness of incumbent adopters to share the new technology. Because orders were given without replacement, those who had already received their order were no longer competing over potential future orders. Those who had not yet received an order were in direct competition with potential learners, as each completed order lowered the potential teacher's chance of an eventual order offer (the likelihood that we would reach their experimental wave before exhausting all resources for garment orders). Therefore, receiving an order increases sharing because the loss of market share is no longer adding to the cost of sharing. This interpretation is supported by the intentional simplicity of learning and teaching Sharawakil once both parties are willing.

In this section, we exploit the random timing (over 2 waves) and the random size (1, 4, or 10 garments) of orders to explore the importance of order timing and size on reported Sharawakil learning, teaching, outsourcing, capital investment, and weaving quality. We then discuss the implications of these findings when considering the experimental competition channel against some alternative possible explanations for why teaching might be higher in the training and order group than in the training only group. Although the collinearity of receiving an order with the removal of experimental competition means that no empirical test can perfectly prove that the removal of experimental competition is driving this increase, each test we provide is consistent with what would be expected if competition were the driving channel and inconsistent with any alternative explanation that we have considered to date.

5.1 Order Timing

First, we explore how order timing might relate to the effect of getting an order on Sharawakil diffusion. Recall that orders were ultimately offered over two (barely overlapping) waves. After Wave 1 orders had been accepted or rejected (based on ability to produce or find someone to produce Sharawakil), additional Wave 2 of orders were placed based on remaining resources. After Wave 2, no additional orders were places as resources had been exhausted. Because Wave 1 and Wave 2 both occurred during the time period covered by our final round of panel data, members of each have the same variation in treatment status within the data, moving from control to order only and training only to training and order in the final round. However, limiting our sample by random wave assignments of the potential learner and potential teacher allows us to estimate different treatment effects of receiving an order on diffusion between members in various waves.

Table 6 shows potential teacher and learner treatment effects on sharing at the dyadic level across all experimental rounds. Column 1 shows only those dyads for which i and j were both selected to receive an order in Wave 1, or neither were selected to receive an order. Column 2 shows those dyads for which the potential teacher was selected for an order offer in Wave 1 while the potential learner was selected for Wave 2. Column 3 shows those dyads for which the potential teacher was selected for an order offer in Wave 2, while the potential learner was selected for Wave 1. Column 4 shows only those dyads for which both were selected to receive an order offer in Wave 2. These regressions only include dyads in which firm owners were selected to receive an order size of 1 or none at all to remove the potential for confounding the timing effect with the order size effect explained below.

TABLE 6

From Table 6, we see the highest increase of sharing from Wave 1 teachers to Wave 2 learners, and the lowest occurrence of sharing from Wave 2 teachers to Wave 1 learners. This suggests that something changes upon receiving an order and remains different even after completing the order.

5.2 Order Size

Second, we explore the effect of order size on Sharawakil learning, teaching, outsourcing, capital investment, and weaving quality. Recall that, during Wave 1, firm owners receiving an order were randomized into three order size categories: 1,4, and 10. Table 7 shows the reduced form treatment effects by order size of self-reported ability to produce Sharawakil, learning and sharing across all experimental rounds. These regressions only include those selected to receive an order in Wave 1 or none at all, because the larger order sizes were only given during Wave 1 and including Wave 2 may confound the timing effect with the order size effect.

TABLE 7

From column 1 of table 7, we see a similar pattern as figure 11. A reported ability to produce Sharawakil is driven by either being invited to the training or receiving an order. It is important to note that there is no difference between those invited to the training with and without orders in reporting ability to produce Sharawakil. Also note that there is no increase in reporting an ability for those with larger order sizes in either the trained or untrained groups. From columns 2 and 3, we see that the increase in learning for those not invited to the training and teaching for those invited to the training are both driven by receiving an order of any size, rather than from receiving larger orders. Teaching, for those not invited to the training, only increases when receiving a larger order.

We see a few other interesting results related to the order only group in table 8, which uses itemized sales and expenditure data collected specifically about Sharawakil in the final round. In column 1, the outcome variable is equal to 1 if the firm owner reports positive expenses for outsourcing related to Sharawakil (our proxy for purchasing Sharawakil). In column 2, the outcome variable is equal to 1 if the firm owner reports positive sales from outsourcing related to Sharawakil (our proxy for selling Sharawakil). In column 3, the outcome variable is equal to 1 if the firm owner reports spending money to purchase, rent or repair tools or equipment (most likely the toy car) related to Sharawakil.

TABLE 8

As we observed in figure 10, we see that firm owners in the order only group are more likely to purchase Sharawakil. Purchasing behavior does not increase with order size. We also see that it is those with order of size 10 only, within the order only group, that are more likely to be selling Sharawakil. This result is driven by the fact that the only three firm owners selling Sharawakil are in this group. We also see that it is this same order size of 10 only group that is investing in the Sharawakil tool itself. We do not see the training groups increasing Sharawakil selling, buying or tool investment with the receipt of an order. The tool investment result, in particular, was expected given that every firm owner who attended the training was given a tool car upon completion in order to not confound investment decision with teaching decisions. Because it is easier to teach the technique when one already owns the tool, it is possible that this choice of investment for the order of size 10 only group is contributing to the differences in likelihood of teaching or selling Sharawakil between this group and other order only treatment group members.

In our last table exploiting the random order size, we explore the potential effect of practice producing Sharawakil on quality. Recall that we had each garment purchased as part of this project evaluated by three different garment making experts in Accra, including the Sharawakil designer himself. Table 9 has three columns, one for each independent expert evaluator from Accra, depicting the effect of order size on Sharawakil weaving quality, graded on a forced curve from 1 to 10 by each of our assessors. Column 1 is the designer. Columns 2 and 3 are high ranking officials in the Ghana National Tailors and Dress Makers Association. Because we only have garments for those who received orders, we only include order size dummies in these regressions.

TABLE 9

Although we see a positive relationship between assessed overall garment quality and Sharawakil quality for all assessors, we do not see any effect of order size on Sharawakil quality, suggesting that practice does not make perfect. This is consistent with the intended design of the technique, to be a series of steps that are easily executed once one has found someone willing to share the process.

5.3 Considering Alterative Channels

We now consider the implications of these secondary empirical findings for the interpretation of our main empirical finding that receiving and order spurs teaching in the training treatment group. We explain why the removal of experimental competition increasing willingness to teach is the interpretation that is most consistent with our data. We also explain why other possible channels are not.

First, we consider learning by doing. Learning by doing could imply that those who went to the training did not actually feel confident in the skill unless and until they got their own order and practiced. Although we intentionally designed the weaving technique as extremely easy to learn and imitate once a teacher was willing to show the steps in order to minimize the risk of this channel, we can also look for evidence for or against this channel in figure 11, table 7 and table 9. Figure 11 shows no difference in reported ability to produce Sharawakil between those in the training only and the training and order groups. We also do not see a difference between these groups in table 7 nor do we see an increase in weaving quality for larger orders in table 9. We do see a difference in reported ability between the order only group and the control group. However, column 1 in table 7 suggests that this has more to do with the incentive to learn in order to complete the order, rather than an increase in confidence through practice, with the increase in reported ability occurring with any order and not increasing in order size. Our empirical results are inconsistent with this interpretation.

Second, we consider teaching while doing. This alternative explanation posits that the cost of teaching may be lower when one is already doing the skill that must be demonstrated. If teaching while doing is the dominant channel for the reduction in the cost of sharing, we should see that the increase in diffusion is occurring more between dyads that receive orders during the same random wave. Our findings in table 6, that diffusion increases more with early orders to those with later orders, are inconsistent with a teaching while doing hypothesis. Instead they offer evidence for an experimental competition interpretation, in which Wave 1 teachers are no longer concerned about the cost of teaching associated with the loss of potential future orders. Further evidence against teaching while doing is seen in table 7, where we observe no increased likelihood of teaching with larger orders. Larger orders take longer to produce and thus, if teaching while doing were a dominant channel, we might have expected to see those receiving larger orders teaching more.

Another teaching while doing related theory suggests that it is easier to teach if a firm owner already owns the tool required for teaching. However, this is unlikely to drive the main results as we intentionally gave a toy car to every training attendee to avoid this channel. Empirical evidence against this channel is seen in table 8, where we observe no difference in investment between the group invited to train that received an order and the group invited to train that didn't. The only relative increase in likelihood of reporting tool expenses occurs within the untrained and order of size 10 group. The fact that this is also the group more likely to teach and sell Sharawakil amongst those not invited to the training suggests that this might not be an unimportant channel for those not in our training groups, but the fact that there is no difference between those with and without orders in the training group provides evidence against this channel driving the main results in this paper.

Finally, we consider a taste for fairness. Smaller (or no) orders could lead to bitterness or a perceived lack of fairness in order size and thus less teaching. If fairness is based on relative order size, it would be expected that firm owners would be more likely to share when given a larger order. The unimportance of order size for teaching within the training groups observed in table 7 is evidence against fairness as the driving channel for our main findings.

5.4 Experimental Competition

The empirical findings reported in this section are consistent with our preferred interpretation that the removal of experimental competition is what drives sharing. Under this interpretation, we expect to see the significant increase in sharing occur between no orders and any order, rather than with each additional order. Additionally, we would expect to see less diffusion from wave 2 firm owners to wave 1 firm owners, as compared to any other wave combination. These empirical findings, in conjunction with the intended design of the weaving technique and experiment, suggest that competition may be an important barrier to technology diffusion in this context.

6 Model

We conceptualize technology diffusion as occurring (or not occurring) as a transaction in a market for technology. The potential learner is modeled as the potential consumer of a technology and the potential sharer is modeled as the potential supplier of the technology. For trade to occur (technology to diffuse), it must be the case the the potential consumer's reservation price is less than or equal to that of the potential supplier's reservation price. Reservation prices are determined by the consumer's net benefit from learning and the supplier's net benefit of teaching, which are driven both by constant pair-wise characteristics and dynamic market factors.

The learner's benefit and costs are teacher and technology specific. These could be anything, but are most likely related to the profit increases expected from adopting the new technology and the time and effort it takes to find and/or learn from a particular teacher. The teacher's cost is also technology and learner specific, and may include the time or effort it takes to teach a particular learner, as well as any forfeit of market share in the case where the learner is in competition with the teacher over demand for the technology. The price paid to the supplier by the consumer could be anything from real currency, to some other favor or exchange, to a general sense of reciprocity. We abstract from the specifics of the exchange however, as it is the relative value of that price to both the potential consumer and potential supplier that matters for whether technology diffusion will occur.

This framework leaves us with two main lessons. First, the learner's benefit alone will not guarantee technology diffusion. The existence of technology diffusion requires at least one incumbent adopter and one potential adopter to BOTH net benefit from the exchange of that knowledge. Conditional on the existence of such a pair, the particular pattern of technology diffusion will be between specific pairs of incumbents and potential learners for whom the joint net benefit is highest. Second, the mere observation of a technology sharing relationship does not tell us much about its underlying determinants. Recall that the costs and benefits of diffusion between a particular learner and teacher are related both to constant pair-wise characteristics and technology specific factors. Therefore, the resulting diffusion patterns confound the effects of many possible underlying determinants related both to constant pair-wise costs and benefits of diffusion and dynamic technology-specific market factors.

We interpret our field experiment as randomly varying the benefit to the learner (through the placing of orders) as well as the cost to the teacher (through the removal of potential for future orders). This model predicts that potential learners with an order will be more likely to learn and that potential teachers with an order (and thus no potential for future orders) will be more likely to teach. These predictions are consistent with the findings of this paper. Additionally, random variation in the initial adopters of a technology allows us to disentangle these dynamic market incentives from other factors, illuminating forces that may be at play outside of our experimental context.

6.1 Model Set-up

Let N be the set of members of a population. In this case, our population is the set of all firm owners in some particular industry. Let the time-specific $A_t \subset N$ be the incumbent adopters of technology a at time t. Let X be the set vectors of pair-wise characteristic, $\forall(i,j) : i \in N$ and $j \in N$. For example, these vectors may include (i,j) specific distance, relative business size or education level, indicators for matching gender or ethnicity, and any other mostly constant observables or unobservables. For each time, t, $i \notin A_t$ and $j \in A_t$, let i's net benefit of learning be $B_{a,i,j,t} - L_{a,i,j,t} - S_{a,i,j,t}$. Let j's net benefit of teaching be $S_{a,i,j,t} - T_{a,i,j,t}$.

$B_{a,i,j,t} \sim \mathcal{U}(0, 2B_{a,x_{ij}})$ is the benefit to i from learning a from j at time t. $B_{a,i,j,t}$ could be anything, but is most likely related to the profit increases expected from adopting the new technology. The expected value of this benefit varies by the specific elements of $x_{i,j}$, for example an indicator for "speaking the same language", which may effect the expected quality of any technology transfer. The expected benefit also varies by the specific a, with some technologies being particularly useful and others quite useless on average within the the industry.

$L_{a,i,j,t} \sim \mathcal{U}(0, 2L_{a,x_{ij}})$ is the cost incurred by i in learning a from j at time t. The expected value of $L_{a,i,j,t}$ also varies by $x_{i,j}$ and a. Perhaps some particular i has a higher expected cost of learning from j for any technology, due to a large distance between i and j, or perhaps some a may be easier to learn on average than some other technologies.

$T_{a,i,j,t} \sim \mathcal{U}(0, 2T_{a,x_{ij}})$ is the cost incurred by j in teaching a to i at time t. $T_{a,i,j,t}$ may include the time it takes to teach, as well as any forfeit of market share in the case where the learner is in competition with the teacher over demand for the technology. $T_{a,i,j,t}$ has an expected value conditioned both on x and a. It may be that teaching any technology to some i is always expected to be more costly, because of some constant element in $x_{i,j}$. for example, closer firms may compete more over local demand. Additionally, teaching certain technologies may be more costly than others. For example, if some technology is labor enhancing (some method reducing the chance of injury) rather than demand shifting (a new product), this technology may be less costly to share on average within the industry.

$S_{a,i,j,t}$ is the surplus transfer made from i to j for learning a at time t. The particulars of $S_{a,i,j,t}$ could be anything from real currency, to some other favor or broad sense of reciprocity. We abstract from the specifics of the exchange however, as it is the relative value of that transfer to both the potential consumer and potential supplier that matters for whether technology diffusion will occur.

Let $G_{a,t}$ be the graph representing all $i \in N$ and $j \in N$: i learned a from j in t. Note that the typical networks dataset asks respondents about interactions concerning the set of all technologies, Λ over a particular time period, T. Let $G_{\Lambda,T}$ be the set of $(i,j) \in N : (i,j) \in G_{a,t}, \forall t \in T, a \in \Lambda$.

6.2 Conditions for Technology Diffusion

Theorem 1 (It Takes Two):

$G_{a,t} \neq O \iff \exists (i,j) : B_{a,i,j,t} - L_{a,i,j,t} - T_{a,i,j,t} \geq 0$

Proof:

If so, then $\exists (i,j) : B_{a,i,j,t} - L_{a,i,j,t} \geq T_{a,i,j,t}$

$\Rightarrow \exists S_{a,i,j,t} : B_{a,i,j,t} - L_{a,i,j,t} \geq S_{a,i,j,t} \geq TC_{a,i,j,t}$

$\Rightarrow B_{a,i,j,t} - L_{a,i,j,t} - S_{a,i,j,t} \geq 0$ and $S_{a,i,j,t} - TC_{a,i,j,t} \geq 0$

If not, then $\forall (i,j), B_{a,i,j,t} - L_{a,i,j,t} < T_{a,i,j,t} \Rightarrow$

$\nexists S_{a,i,j,t} : S_{a,i,j,t} - T_{a,i,j,t} \geq 0$ and $B_{a,i,j,t} - L_{a,i,j,t} - S_{a,i,j,t} \geq 0$

Intuitively, what this theorem states is that the overall surplus from the transfer of technology a from j to i at time t must be positive. Otherwise, there is no surplus transfer that would make the transaction worthwhile for both parties involved. In the case of peer-to-peer technology diffusion, it must be the case that both parties net benefit.

Theorem 2 (Path of Least Resistance):

$g_{ij} \in G_{a,t} = 1 \iff B_{a,i,j,t} - L_{a,i,j,t} - T_{a,i,j,t} \geq 0$

and $B_{a,i,j,t} - L_{a,i,j,t} - T_{a,i,j,t} \geq B_{a,i,k,t} - L_{a,i,k,t} - T_{a,i,k,t}, \forall k$

Proof:

If so, then $B_{a,i,j,t} - L_{a,i,j,t} - T_{a,i,j,t} \geq B_{a,i,k,t} - L_{a,i,k,t} - T_{a,i,k,t}, \forall k$

$\Rightarrow (B_{a,i,j,t} - L_{a,i,j,t}) - (B_{a,i,k,t} - L_{a,i,k,t}) \geq T_{a,i,j,t} - T_{a,i,k,t}, \forall k$

$\Rightarrow \forall S_{a,i,k,t} : S_{a,i,k,t} - T_{a,i,k,t} \geq 0,$

$\exists S_{a,i,j,t} : S_{a,i,j,t} - T_{a,i,j,t} \geq 0 \cap B_{a,i,j,t} - L_{a,i,j,t} - S_{a,i,j,t} \geq B_{a,i,k,t} - L_{a,i,k,t} - S_{a,i,k,t}$

If not, then $\exists (i,k) : B_{a,i,j,t} - L_{a,i,j,t} - T_{a,i,j,t} < B_{a,i,k,t} - L_{a,i,k,t} - T_{a,i,k,t}$

$\Rightarrow (B_{a,i,j,t} - L_{a,i,j,t}) - (B_{a,i,k,t} - L_{a,i,k,t}) < T_{a,i,j,t} - T_{a,i,k,t}$

$\Rightarrow \exists S_{a,i,k,t} : S_{a,i,k,t} - T_{a,i,k,t} \geq 0 \cap B_{a,i,j,t} - L_{a,i,j,t} - S_{a,i,j,t} < B_{a,i,k,t} - L_{a,i,k,t} - S_{a,i,k,t}$

Intuitively, what this theorem states it that, conditional on the overall surplus from the transfer of technology a being positive for at least one j, i pair at time t, the exact j from whom i will choose to learn will be the one for whom the overall surplus is highest. Otherwise, some other k teacher can offer to teach i for a lower surplus transfer and still benefit from the exchange.

Taken together, Theorems 1 and 2 imply that $G_{\Lambda,T}$ will be determined (with some noise) by both constant industry factors, X, as well as dynamic market factors, Λ, during T.

6.3 Experimental Predictions

In our experiment, we create a new technology a and essentially randomize N into 4 distinct subsets:

- β_{cs} receives training only
- β_{oc} receives an order only
- β_{os} receives both training and an order
- β_{cc} receives neither

We conceptualize β_{cs} and β_{os} as $\subset A_t$ and β_{oc} and β_{cc} as $\subset A_t^c$, where A_t^c is the compliment of A_t. For some $i \in A_t^c$, receiving an order involving a increases $B_{a,i,j,t}$ by $\beta_1, \forall j$. For some $j \in A_t$, receiving an order involving a lowers $T_{a,i,j,t}$ by $\beta_2, \forall i$. Let $P_{a,i,j,t}$ be the probability that $g_{ij} \neq 0 \in G_{a,t}$.

Prediction 1: ($\uparrow$ Benefit of Learning $\uparrow$ Learning)

$P_{a,i,j,t} > P_{a,c,j,t}, \forall j \in A, i \in \beta_{oc}, c \in \beta_{cc}$

Proof:

Because the assignment of $i \in \beta_{oc}, c \in \beta_{cc}$ is random, their expected distributions no longer depend on X, but differ only in the mean shift of B_a by β_1. Therefore:

$P_{a,i,j,t} > P_{a,c,j,t} \iff$

$P(B_{a,i,j,t} - L_{a,i,j,t} - T_{a,i,j,t} \geq 0) > P(B_{a,c,j,t} - L_{a,c,j,t} - T_{a,c,j,t} \geq 0) \iff$

$B_a - L_a - T_a + \beta_1 > B_a - L_a - T_a \iff$

$\beta_1 \geq 0$

Prediction 2: ($\downarrow$ Cost of Teaching $\uparrow$ Teaching)

$P_{a,i,j,t} > P_{a,i,k,t}, \forall i \notin A, j \in \beta_{os}, k \in \beta_{cs}$

Proof:

Because the assignment of $j \in \beta_{os}, k \in \beta_{cs}$ is random, their expected distributions no longer depend on X, but differ only in the mean shift of T_a by β_2. Therefore:

$P_{a,i,j,t} > P_{a,i,k,t} \iff$

$P(B_{a,i,j,t} - L_{a,i,j,t} - T_{a,i,j,t} \geq B_{a,i,k,t} - L_{a,i,k,t} - T_{a,i,k,t})$

$> P(B_{a,i,k,t} - L_{a,i,k,t} - T_{a,i,k,t} \geq B_{a,i,j,t} - L_{a,i,j,t} - T_{a,i,j,t}) \iff$

$B_a - L_a - (T_a - \beta_2) > B_a - L_a - T_a \iff$

$\beta_2 \geq 0$

Prediction 3: (The Network has Memory)

$P_{a,i,j,t} > P_{a,c,k,t}, \forall (i,j) : g_{ij} \neq 0 \in G_{\Lambda,t-1}, (c,k) : g_{ck} = 0 \in G_{\Lambda,t-1}$

Proof:

Because the assignment of N to A and A^c is random, membership to either set should be orthogonal to X. Thus, if a randomly chosen $(i, j) \in G_{\Lambda,t-1}, (c, k) \notin G_{\Lambda,t-1}$, then knowing nothing else about these pairs, it can be expected that:

$$B_{a,x_{i,j}} - L_{a,x_{i,j}} - T_{a,x_{i,j}} \geq B_{a,x_{c,k}} - L_{a,x_{c,k}} - T_{a,x_{c,k}} \iff$$
$$P(B_{a,i,j,t} - L_{a,i,j,t} - T_{a,i,j,t} \geq B_{a,c,k,t} - L_{a,c,k,t} - T_{a,c,k,t})$$
$$> P(B_{a,c,k,t} - L_{a,c,k,t} - T_{a,c,k,t} \geq B_{a,i,j,t} - L_{a,i,j,t} - T_{a,i,j,t}) \iff$$
$$P_{a,i,j,t} > P_{a,c,k,t}$$

Taken together, these three predictions suggest that the diffusion of our new technology should occur mostly from $j \in \beta_{os}$ and mostly to $i \in \beta_{oc}$. They suggest that this diffusion is more likely to occur if $i, j : g_{ij} \neq 0 \in G_{\Lambda,t-1}$, but that $g_{ij} \neq 0 \in G_{\Lambda,t-1}$ is not necessary for diffusion to occur between i and j.

6.4 Unpacking Observed Diffusion

Taking together our model and experiment, we can begin to think about unpacking the determinants of baseline diffusion. The experiment has shown the importance of demand and competition in driving (or inhibiting) diffusion. Table 10 shows the relationship between various potential dyadic characteristics in X and both baseline diffusion and the diffusion of our new technology.

TABLE 10

Because the seeding of our new technology was random, we are able to both observe the effects of various characteristics on sharing decoupled from real world patterns of technology seeding and, by comparing column 2 with column 1, we can say something about those real world patterns.

The only dyadic characteristic associated with higher sharing of the new technology is distance. Being within 500 feet of one another is associated with both higher baseline technology diffusion and higher sharing of our new technology. Specifically, baseline sharing and the sharing of our new technology are higher by about 4 times the average.

Competition has a positive relationship with baseline diffusion. Specifically, being listed as competitors by our market research survey is associated with a 1.9% increase in the likelihood of sharing, or an increase of 3 times the average. However, baseline competition has little to no relationship with the diffusion of Sharawakil. Similarly, management practices has a positive and significant relationship to sharing at baseline, but loses significance for sharing of Sharawakil.

We see significant, but lower, relationships between gender, ethnicity, business size and baseline sharing that remain relatively constant with sharing of Sharawakil. This suggests that the observed relationship between diffusion and these dyadic characteristics is due to more static market frictions or preferences than the result of systematically related market based incentives.

7 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 sharing in the context of garment makers in Ghana. We designed a new weaving technology and randomly varied direct training (direct access to the technology), benefits of adoption (demand for garments featuring the technology), and the existence of competition over potential future orders in a real network of garment making firm owners. We find that the majority of this new technology's diffusion occurs from those with direct training who are not in competition over future orders to those without direct training with a higher benefit of adoption. Those with direct training who remained in competition over future orders were no more likely to share than those without direct training.

We interpret our findings as evidence that competition is an important barrier to technology diffusion in the context of small-scale manufacturing in low-income countries. We develop a model designed to frame our understanding of the findings as the result of internal market incentives to both learn and share, with diffusion being the final result of a learner's overall benefit being higher than a teacher's overall cost. Taken together, the findings and the model suggest that productivity-increasing technologies may fail to diffuse, in part because the incentives of incumbent adopters may not encourage active sharing.

References

Atkin, David, Amit K Khandelwal, and Adam Osman. 2014. "Exporting and Firm Performance: Evidence from a Randomized Trial." NBER Working Paper 20690.

Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit K Khandelwal, and Eric Verhoogen. 2015. "Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan." NBER Working Paper 21417.

Bandiera, Oriana, and Imran Rasul. 2006. "Social Networks and Technology Adoption in Northern Mozambique." *Economic Journal*, 116(514): 869–902.

Bandiera, Oriana, Iwan Bankerey, and Imran Rasul. 2000. "Field Experiments with Firms." *Journal of Economic Perspectives*, 25(3): 63–82.

Bandiera, Oriana, Iwan Bankerey, and Imran Rasul. 2005. "Social Preferences and the Response to Incentives: Evidence from Personnel Data." *Quarterly Journal of Economics*, 120(3): 917–962.

Banerjee, Abhijit, Emily Breza, Arun Chandrasekhar, Esther Duflo, and Matthew Jackson. 2012. "Come Play with Me: Experimental Evidence of Information Diffusion about Rival Goods." Working Paper.

Barr, Abigail. 1998. "Enterprise performance and the functional diversity of social capital." CSAE Working Paper.

Barr, Abigail. 2000. "Social Capital and technical information flows in the Ghanaian manufacturing sector." *Oxford Economic Papers*, 52(3): 539–559.

Beaman, Lori, and Jeremy Magruder. 2012. "Who Gets the Job Referral? Evidence from a Social Networks Experiment." *American Economic Review*, 102(7): 3574–3593.

BenYishay, Ariel, and A Mushfiq Mobarak. 2015. "Social Learning and Incentives for Experimentation and Communication?" Working Paper.

Bloom, Nicholas, and John Van Reenan. 2007. "Measuring and Explaining Differences in Management Practices Across Countries." *Quarterly Journal of Economics*, 122(4): 1351–1408.

Bloom, Nicholas, and John Van Reenan. 2010. "Why do Management Practices Differ Across Firms and Countries?" *Journal of Economic Perspectives*, 24(1): 203–224.

Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. 2013. "Does Management Matter? Evidence from India." *Quarterly Journal of Economics*, 128 (1): 1–51.

Brooks, Wyatt, Kevin Donovan, and Terence R Johnson. 2016. "The Dynamics of Inter-Firm Skill Transmission among Kenyan Microenterprises." Working Paper.

Bruhn, Miriam, Dean S Karlan, and Antoinette Schoar. 2013. "The Impact of Consulting Services on Small and Medium Enterprises: Evidence from a Randomized Trial in Mexico." WB Policy Research Working Paper 6508.

Cai, Jing, and Adam Szeidl. 2016. "Interfirm Relationships and Business Performance." Working Paper.

Casella, Alessandra, and James E Rauch. 2002. "Anonymous market and group ties in international trade." *Journal of International Economics*, 58(1): 19–47.

Conley, Timothy, and Christopher Udry. 2010. "Learning About a New Technology: Pineapple in Ghana." *American Economic Review*, 100(1): 35–69.

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

Fafchamps, Marcel, and Mans Soderbom. 2013. "Network Proximity and Business Practices in African Manufacturing." *World Bank Economic Review*, 28(1): 99–129.

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

Fafchamps, Marcel, and Simon Quinn. 2015. "Networks and Manufacturing Firms in Africa: Results from a Randomized Field Experiment." Working Paper.

Fafchamps, Marcel, David McKenzie, Simon Quinn, and Chris Woodruff. 2013. "Microenterprise Growth and the Flypaper Effect: Evidence from a Randomized Experiment in Ghana." *Journal of Development Economics*, 106: 211–226.

Ferraz, Claudio, Fred Finan, and Dimitri Joe de Alencar Szerman. 2015. "Procuring Firm Growth: The Effects of Government Purchases on Firm Dynamics." NBER Working Paper 21219.

Foster, Andrew, and Mark Rosenzweig. 1995. "Learning by doing and learning from others: Human capital and technical change in agriculture." *Journal of Political Economy*, 103: 1176–1209.

Greif, Avner. 1993. "Contract enforceability and economic institutions in early trade: The Maghribi traders' coalition." *American Economic Review*, 525–548.

Griliches, Zvi. 1957. "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica*, 25(4): 501–522.

Hardy, Morgan, and Jamie McCasland. 2015. "Are Small Firms Labor Constrained? Experimental Evidence From Ghana." Working Paper.

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

Immorlica, Nicole, Brendan Lucier, and Evan Sadler. 2014. "Sharing Rival Information." Working Paper.

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

McMillan, J, and Christopher Woodruff. 1999. "Interfirm Relationships and Informal Credit in Vietnam." *Quarterly Journal of Economics*, 114(4): 1285–1320.

Munshi, Kaivan. 2004. "Social Learning in a Heterogeneous Population: Technology Adoption in the Indian Green Revolution." *Journal of Development Economics*, 73: 185–213.

Munshi, Kaivan. 2011. "Strength in Numbers: Networks as a Solution to Occupational Traps." *Review of Economic Studies*, 78: 1069–1101.

Parente, S, and E Prescott. 1994. "Barriers to Technology Adoption and Development." *Journal of Political Economy*, 298–321.

Patnam, M. 2011. "Corporate Networks and Peer Effects in Firm Policies." Working Paper.

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.*

Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature*, 49(2): 326–365.

Tybout, James. 2000. "Manufacturing Firms in Developing Countries: How Well Do They Do, and Why?" *Journal of Economic Literature*, 38(1): 11–44.

Udry, Christopher, and Timothy Conley. 2004. "Social Networks in Ghana." Working Paper.

Figure 1: **Baseline Degree**

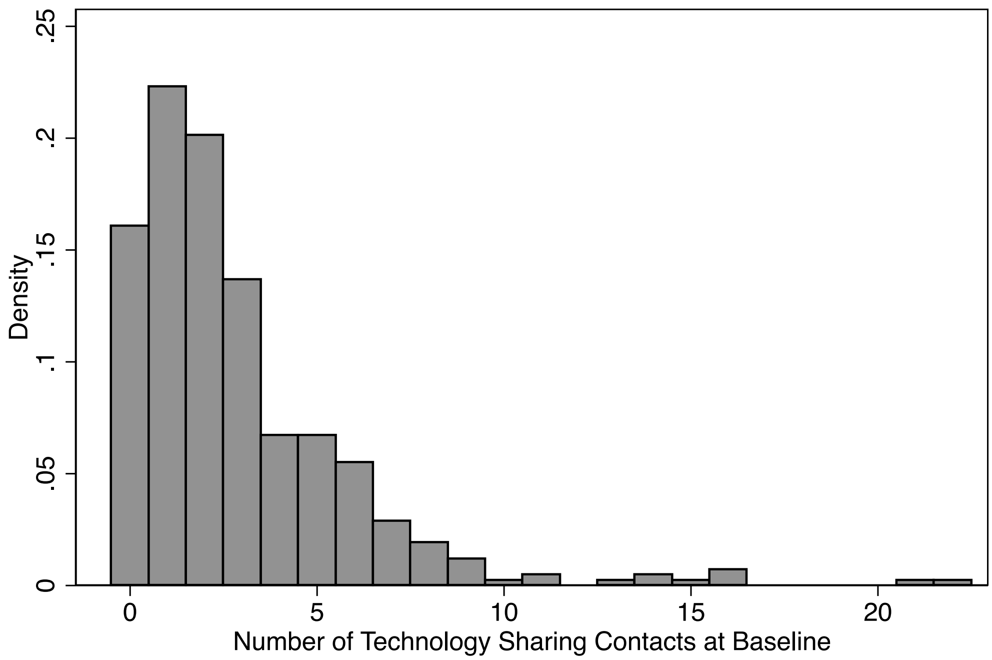

Note: A firm owner is another firm owner's technology sharing contact at baseline if either reports teaching/learning a business related technique or skill to/from the other in the year prior to the baseline.

Figure 2: **Network Training Treatment**

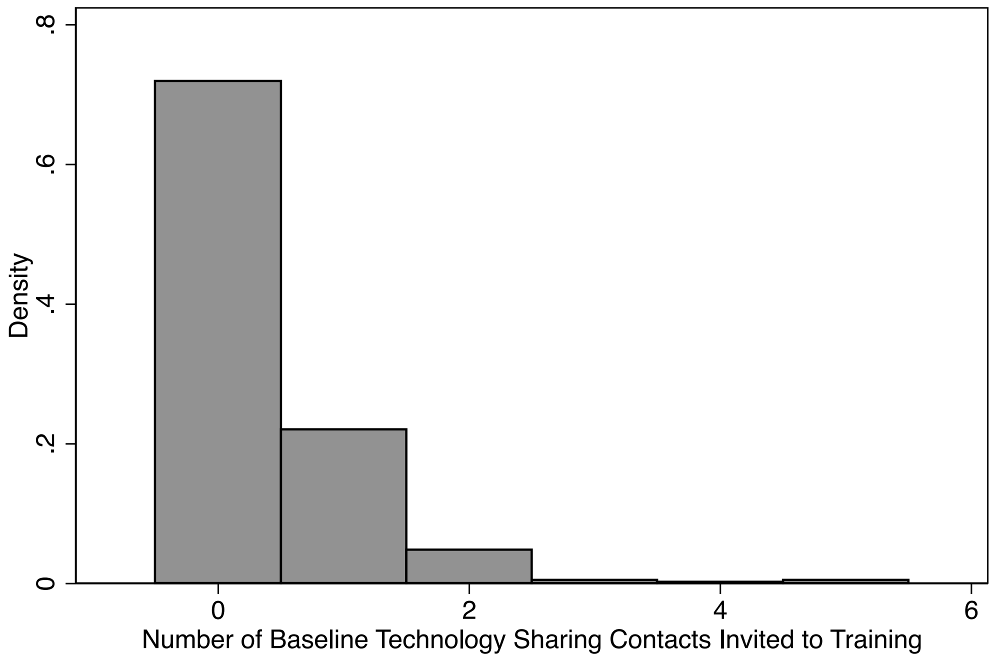

Note: A firm owner is another firm owner's technology sharing contact at baseline if either reports teaching/learning a business related technique or skill to/from the other in the year prior to the baseline.

Figure 3: **Network Order Treatment**

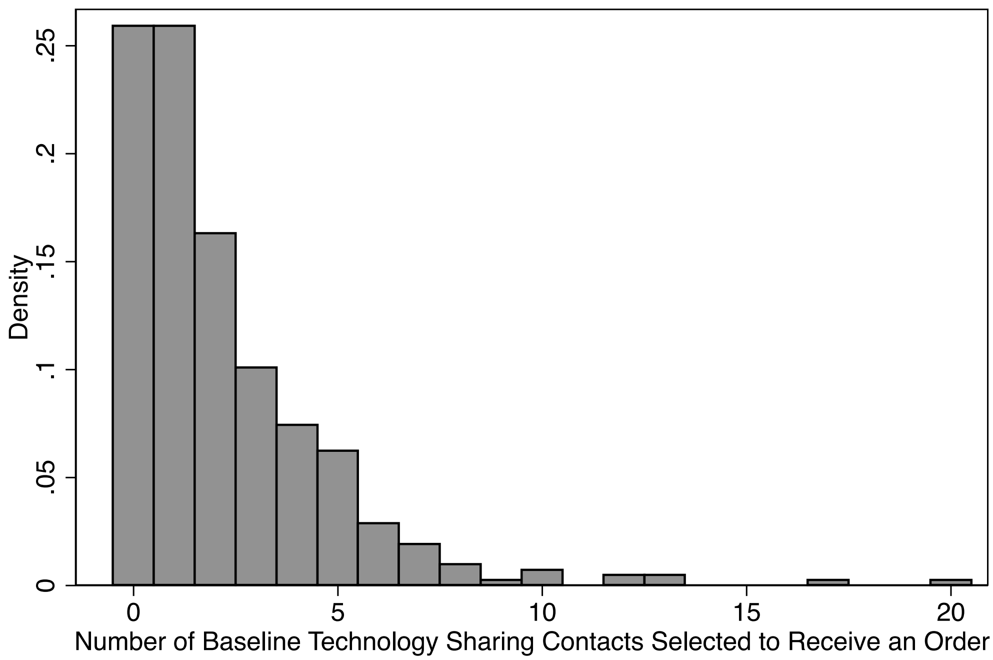

Note: A firm owner is another firm owner's technology sharing contact at baseline if either reports teaching/learning a business related technique or skill to/from the other in the year prior to the baseline.

Figure 4: **Design Sample 1**

Figure 5: **Design Sample 2**

Figure 6: **Design Sample 3**

Figure 7: **Training Treatment Compliance**

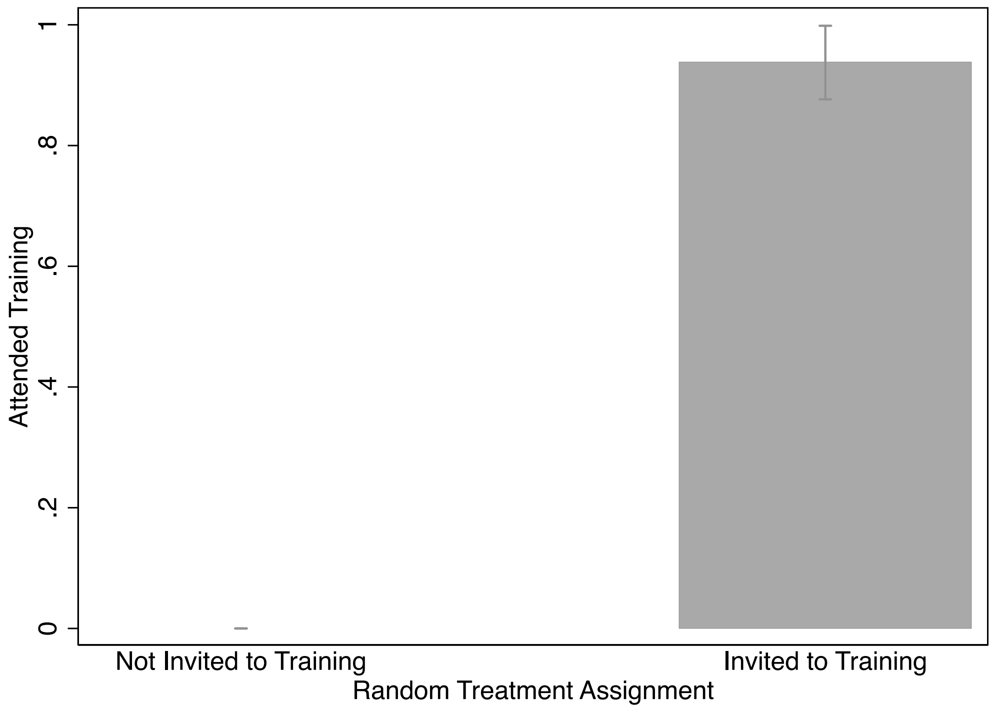

Figure 8: **Order Treatment Compliance**

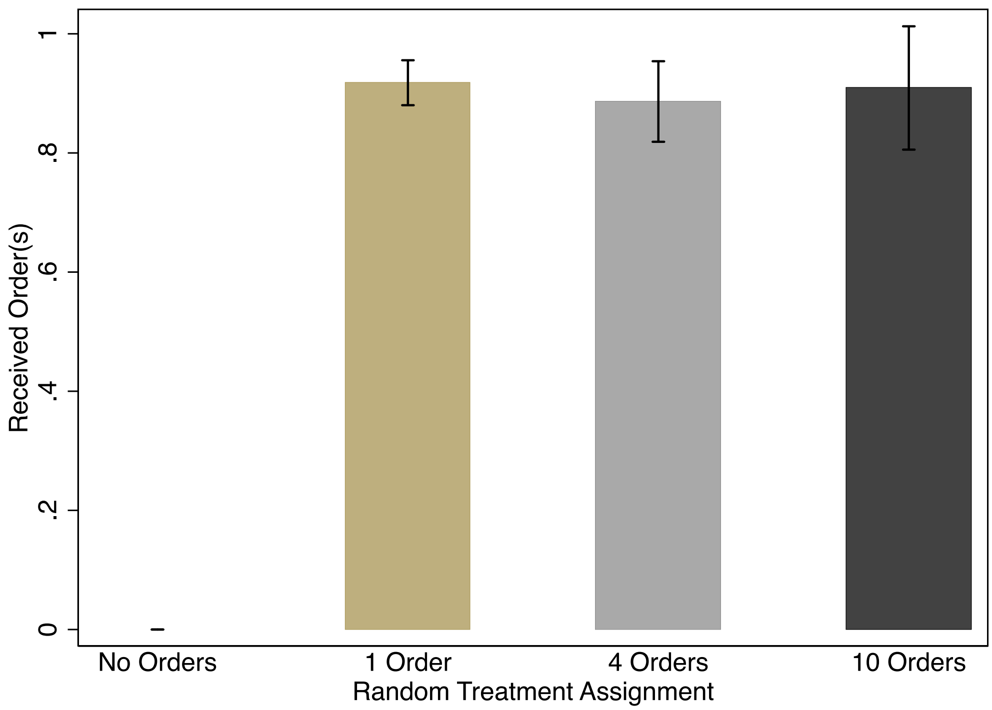

Figure 9: **Order Completion Conditional on Order Request**

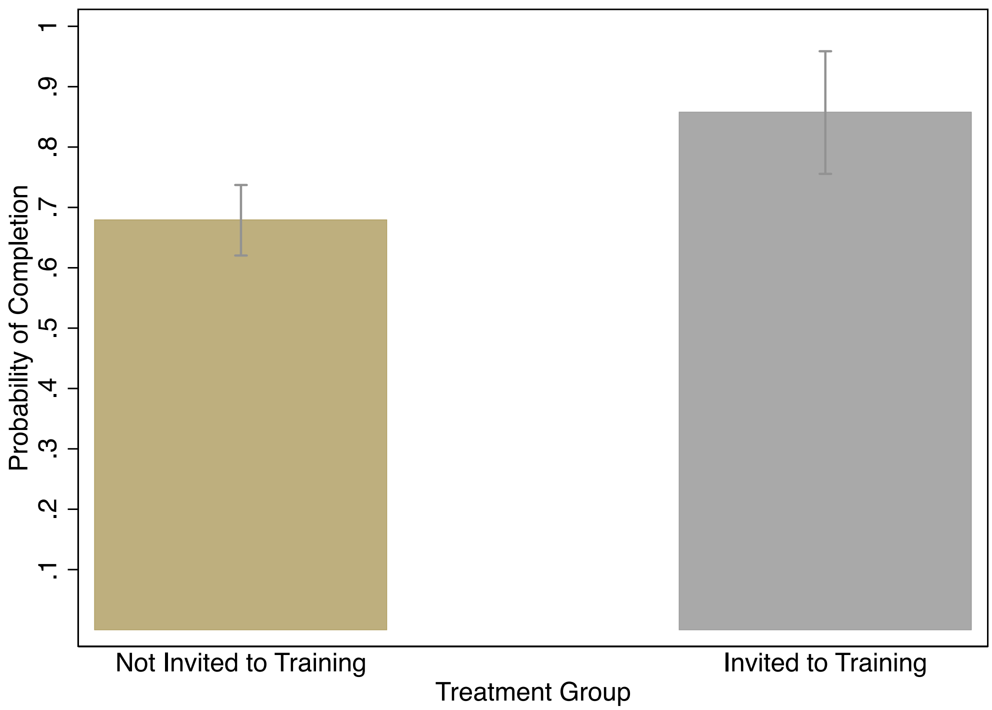

Figure 10: **Reported Completion Method (Conditional on Order Completion)**

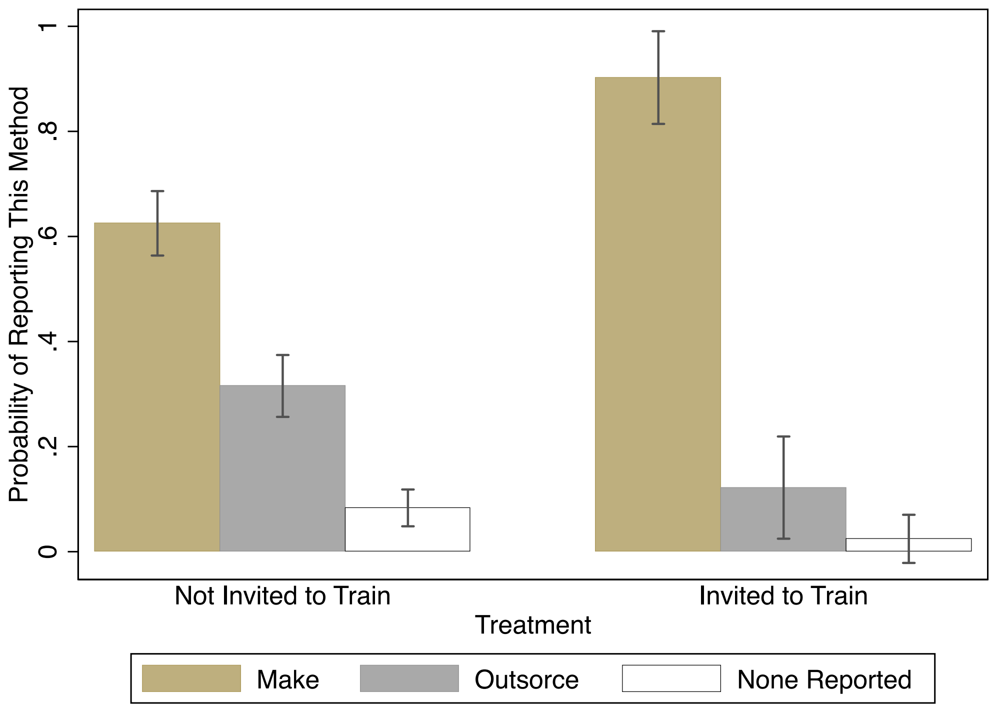

Figure 11: **Reported Ability to Produce Sharawakil by Treatment Group**

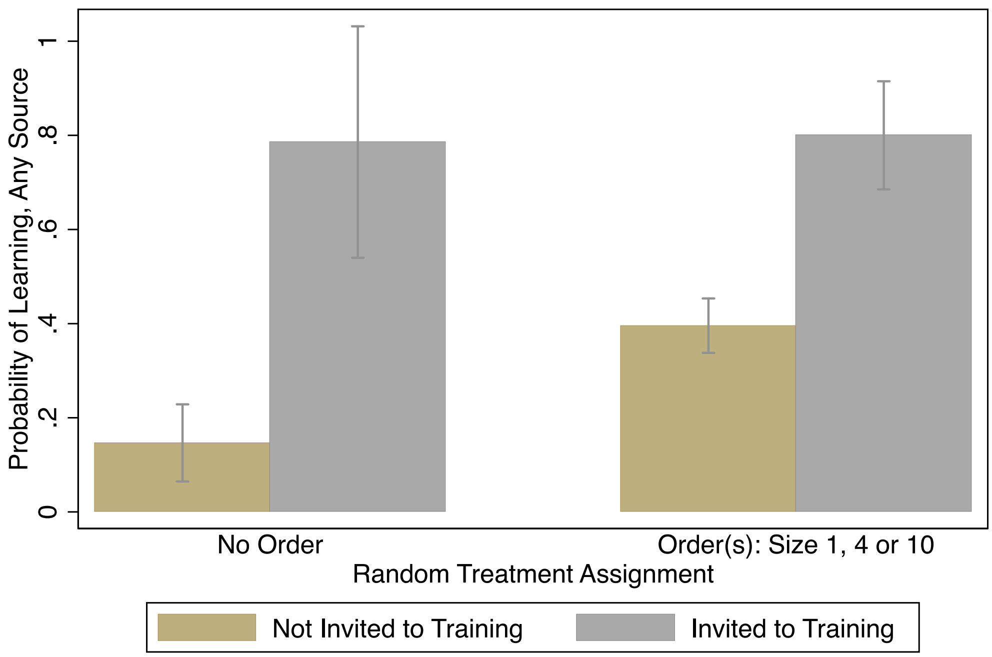

Note: A firm owner is considered to have learned the technology if they've ever reported learning or being able to produce it in any round of data collection, regardless of the learning source.

Figure 12: **Reported Learning Method (Conditional on Reported Ability)**

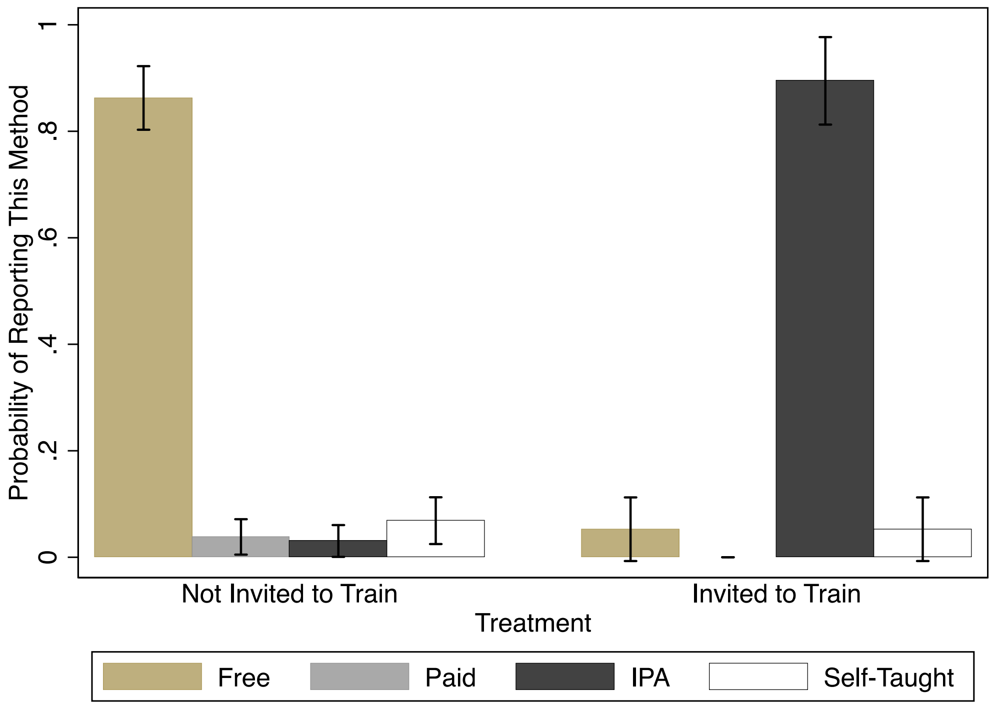

Figure 13: **Learning From Another Firm Owner by Treatment Group**

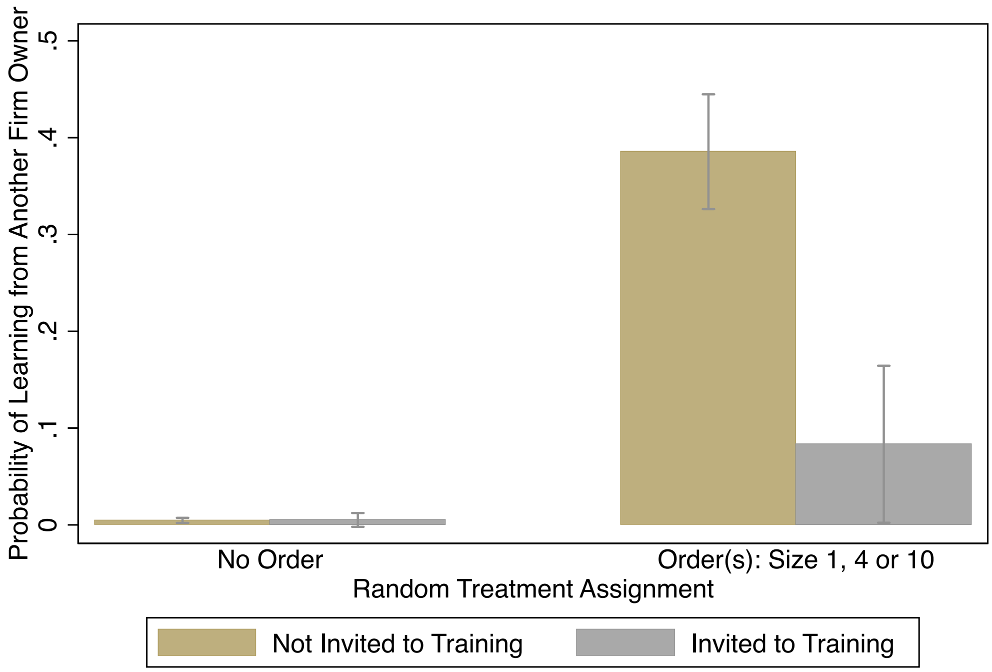

Note: Includes 8 rounds. Training occurred in round 2. Orders occurred in round 8. Learning is 1 if either the firm owner or any other respondent reports that firm owner learning during that round.

Figure 14: **Teaching To Another Firm Owner by Treatment Group**

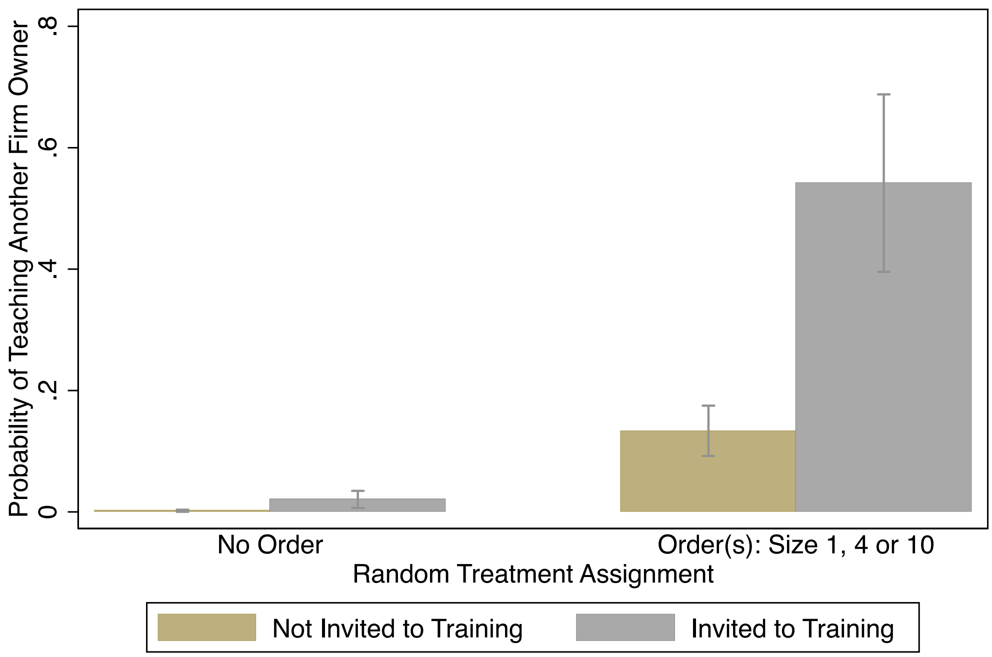

Note: Includes 8 rounds. Training occurred in round 2. Orders occurred in round 8. Teaching is 1 if either the firm owner or any other respondent reports that firm owner teaching during that round.

Table 1: **Summary Statistics and Covariate Balance**

Columns labeled "mean" give the mean value for all firms in our sample, firms randomized to receive no order-no training (cc), firms randomized to receive training but no order offer (cs), firms randomized to receive order offers but no training (oc), and firms randomized to receive both an order offer and training (cs), in that order. Columns 4-5, 7-8, and 10-11 show the difference between the mean in the control group and the three treatment groups, with the corresponding p-value on the test of equality. *** p<0.01, ** p<0.05, * p<0.1.

	All Firms	cc	cs	cc-cs			cc-oc		os	cc-os	
	mean	*mean*	*mean*	*diff*	*p-val*	*mean*	*diff*	*p-val*	*mean*	*diff*	*p-val*
Male	0.23	0.23	0.21	0.01	0.92	0.23	0.00	0.99	0.24	-0.01	0.86
Ewe ethnicity	0.76	0.83	0.86	-0.03	0.78	0.73	0.10	0.09*	0.78	0.05	0.52
Years schooling	8.85	8.69	8.21	0.47	0.51	8.83	-0.14	0.66	9.38	-0.69	0.06*
Ravens score (of 12)	5.63	5.45	6.21	-0.76	0.28	5.58	-0.13	0.71	6.00	-0.55	0.22
Owner age	35.53	33.74	35.14	-1.40	0.61	36.07	-2.33	0.07*	35.28	-1.54	0.33
Within industry/sample degree	15.56	14.37	15.00	-0.63	0.84	15.75	-1.38	0.35	16.46	-2.09	0.36
Firm size (including owner)	1.99	1.93	2.29	-0.36	0.42	1.99	-0.06	0.78	2.00	-0.07	0.82
Has any worker(s) besides owner	0.47	0.44	0.57	-0.13	0.37	0.50	-0.06	0.38	0.36	0.08	0.38
Revenues (GHC)	197	195	178	17	0.83	196	-2	0.96	206	-11	0.82
Profits (GHC)	138	144	123	21	0.72	136	8	0.68	145	-0	0.99
Assets excl land/building (GHC)	1214	1187	1057	130	0.79	1256	-69	0.76	1069	118	0.67
Management practices (of 4)	2.32	2.21	2.21	-0.00	0.99	2.38	-0.16	0.20	2.18	0.03	0.85
Firm age	9.49	9.06	9.00	0.06	0.98	9.66	-0.60	0.59	9.38	-0.32	0.83
Trade association member	0.22	0.19	0.14	0.04	0.69	0.23	-0.05	0.41	0.22	-0.03	0.65
Registered w/any govt agency	0.17	0.19	0.00	0.19	0.08*	0.17	0.02	0.68	0.18	0.01	0.93
Number of Firms	417	75	14			278			50		
F Stat of Joint Sig					0.77			0.19			0.31

Table 2: **Reduced Form Treatment Effects on Technology Diffusion**

This table reports reduced form effects of being randomly invited to a training in and/or selected to receive an order for a new technology on the sharing of that technology between firm owners over the 8 rounds of data collection. The training was held in the period covered by the second round. The orders occurred during the period covered by the final round. Orders were given out on a randomized rolling basis until all resources had been exhausted. Once a firm owner received an order, they could not receive another. If a firm owner had not yet received an order, they were eligible to receive one if enough of the existing orders were not met by their peers. A firm owner is considered to be a teacher or learner if either firm owner reports an interaction in a given round. Standard errors are clustered at the firm level. *** p<0.01, ** p<0.05, * p<0.1.

	(1) Learned Technology from Another Garment Maker	(2) Taught Technology to Another Garment Maker
Randomly Selected to Receive:		
Order Only	0.263***	0.0255
	(0.0526)	(0.0453)
Training Only	0.00657	0.0122
	(0.00866)	(0.00810)
Order and Training	-0.0436	0.441***
	(0.0591)	(0.0860)
Average Across All Rounds	0.0365	0.0227
Observations	3,263	3,263
R-squared	0.407	0.349
Firm FE	YES	YES
Round FE	YES	YES

Table 3: **Reduced Form Spillover Effects on Technology Diffusion**

This table reports reduced form effects of a firm owners' technology sharing contact being randomly invited to a training in and/or selected to receive an order for a new technology on the sharing of that technology between firm owners over the 8 rounds of data collection. Another firm owner is a technology sharing contact if either owner reported sharing technology during the year prior to the start of the study. The training was held in the period covered by the second round. The orders occurred during the period covered by the final round. Orders were given out on a randomized rolling basis until all resources had been exhausted. Once a firm owner received an order, they could not receive another. If a firm owner had not yet received an order, they were eligible to receive one if enough of the existing orders were not met by their peers. A firm owner is considered to be a teacher or learner if either firm owner reports an interaction in a given round. Standard errors are clustered at the firm level. *** p<0.01, ** p<0.05, * p<0.1.

	(1)	(2)
	Learned Technology from Another Garment Maker	Taught Technology to Another Garment Maker
At Least One Baseline Contact Randomly Selected to Receive:		
Order Only	0.0203	0.0994***
	(0.0529)	(0.0377)
Training Only	-0.00401	-0.0112*
	(0.0153)	(0.00640)
Order and Training	0.145**	0.00385
	(0.0605)	(0.0440)
Firm Owner Randomly Selected to Receive:		
Order Only	0.261***	0.0292
	(0.0512)	(0.0446)
Training Only	0.00651	0.0106
	(0.00916)	(0.00786)
Order and Training	-0.0162	0.448***
	(0.0583)	(0.0859)
Average Across All Rounds	0.0365	0.0227
Observations	3,263	3,263
R-squared	0.419	0.358
Firm FE	YES	YES
Round FE	YES	YES

Table 4: **Reduced Form Treatment Effects on Network Formation**

This table reports reduced form effects of being randomly invited to a training in and/or selected to receive an order for a new technology on the formation of new network connections between firm owners over the 8 rounds of data collection. A network connection is "new", if teaching/learning to/from that contact has never been mentioned in a previous round by either the firm owner or that contact. The training was held in the period covered by the second round. The orders occurred during the period covered by the final round. Orders were given out on a randomized rolling basis until all resources had been exhausted. Once a firm owner received an order, they could not receive another. If a firm owner had not yet received an order, they were eligible to receive one if enough of the existing orders were not met by their peers. A firm owner is considered to be a teacher or learner if either firm owner reports an interaction in a given round. Standard errors are clustered at the firm level. *** p<0.01, ** p<0.05, * p<0.1.

	(1) Learned a Technology from a New Contact	(2) Taught a Technology to a New Contact
Randomly Selected to Receive:		
Order Only	0.271***	0.0421
	(0.0453)	(0.0635)
Training Only	0.0116	-0.000879
	(0.0219)	(0.0276)
Order and Training	0.0504	0.865***
	(0.0679)	(0.272)
Average Across All Rounds	0.0429	0.0429
Observations	3,263	3,263
R-squared	0.283	0.269
Firm FE	YES	YES
Round FE	YES	YES

Table 5: **Reduced Form Treatment Effects on Dyadic Level Technology Diffusion**

This table reports the effects of potential teachers and/or potential learners being randomly invited to a training in and/or selected to receive an order for a new technology on dyadic level technology diffusion over the 8 rounds of data collection. Column 1 shows the basic effects of each teacher/learner treatment status on the full sample. Columns 2 includes only dyads that are not technology sharing contacts at baseline, while column 3 includes only the opposite. The training was held in the period covered by the second round. The orders occurred during the period covered by the final round. Orders were given out on a randomized rolling basis until all resources had been exhausted. Once a firm owner received an order, they could not receive another. If a firm owner had not yet received an order, they were eligible to receive one if enough of the existing orders were not met by their peers. Technology diffused between a potential teacher and learner in a given round, if either respondent reported teaching or learning in that round. Standard errors are two-way clustered at the teacher and learner level. *** p<0.01, ** p<0.05, * p<0.1.

	(1)	(2)	(3)
Dyads Included in Regression:	All	Not Technology Sharing Contacts at Baseline	Technology Sharing Contacts at Baseline
	New Technology Diffused	New Technology Diffused	New Technology Diffused
Teacher RandomlySelected to Receive:			
Order and Training	0.00303***	0.00212***	0.13860***
	(0.00071)	(0.00060)	(0.03303)
Order Only	0.00018	0.00009	0.01071
	(0.00015)	(0.00013)	(0.00829)
Training Only	0.00002	-0.00000	0.00449**
	(0.00003)	(0.00002)	(0.00219)
Learner Randomly Selected to Receive:			
Order and Training	-0.00014	0.00002	-0.00574
	(0.00016)	(0.00011)	(0.01292)
Order Only	0.00070***	0.00063***	0.01331
	(0.00017)	(0.00016)	(0.01210)
Training Only	0.00002	0.00002	0.00026
	(0.00002)	(0.00002)	(0.00117)
Average of New Technology Difussion	0.00009	0.00006	0.00430
Average of Any Technology Diffusion	0.00017	0.00010	0.01065
Observations	1,328,400	1,319,102	9,298
R-squared	0.003	0.002	0.146
Teacher FE	YES	YES	YES
Learner FE	YES	YES	YES
Round FE	YES	YES	YES

Table 6: **Reduced Form Treatment Effects on Technology Diffusion by Order Timing**

This table explores the importance of order timing in the effects of potential teachers and/or potential learners being randomly invited to a training in and/or selected to receive an order for a new technology on dyadic level technology diffusion over the 8 rounds of data collection. The training was held in the period covered by the second round. The orders occurred during the period covered by the final round. Orders were given out on a randomized rolling basis until all resources had been exhausted. Once a firm owner received an order, they could not receive another. If a firm owner had not yet received an order, they were eligible to receive one if enough of the existing orders were not met by their peers. 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. Therefore, to remove any size effect, these regression include only dyads in which owners were selected to receive an order of size 1 or not selected for an order at all. Technology diffused between a potential teacher and learner in a given round, if either respondent reported teaching or learning in that round. Standard errors are two-way clustered at the teacher and learner level. *** p<0.01, ** p<0.05, * p<0.1.

Dyads Included in Regression:	(1)	(2)	(3)	(4)
Teacher Wave, Learner Wave	Wave 1, Wave 1	Wave 1, Wave 2	Wave 2, Wave 1	Wave 2, Wave 2
	New Technology Diffused	New Technology Diffused	New Technology Diffused	New Technology Diffused
Teacher RandomlySelected to Receive:				
Order and Training	0.00229	0.00257*	0.00167**	0.00202**
	(0.00156)	(0.00145)	(0.00080)	(0.00097)
Order Only	0.00028	-0.00000	-0.00020	-0.00005
	(0.00061)	(0.00030)	(0.00033)	(0.00025)
Training Only	0.00026***	0.00006	0.00013**	0.00001
	(0.00009)	(0.00005)	(0.00006)	(0.00003)
Learner Randomly Selected to Receive:				
Order and Training	-0.00038	-0.00045	0.00018	0.00017
	(0.00026)	(0.00030)	(0.00049)	(0.00038)
Order Only	0.00053	0.00022	0.00032	0.00043
	(0.00046)	(0.00037)	(0.00029)	(0.00027)
Training Only	-0.00001	0.00001	-0.00002*	0.00002
	(0.00004)	(0.00003)	(0.00001)	(0.00003)
Average of New Technology Difussion	0.00010	0.00007	0.00005	0.0000517
Average of Any Technology Diffusion	0.00019	0.00012	0.00014	0.0001145
Observations	177,834	221,248	221,248	270,730
R-squared	0.004	0.003	0.003	0.003
Teacher FE	YES	YES	YES	YES
Learner FE	YES	YES	YES	YES
Round FE	YES	YES	YES	YES

Table 7: **Reduced Form Treatment Effects on Technology Diffusion by Order Size**

This table reports reduced form effects of being randomly invited to a training in and/or selected to receive an order for a random number of garments utilizing a new technology on the sharing of that technology between firm owners over the 8 rounds of data collection. The training was held in the period covered by the second round. The orders occurred during the period covered by the final round. Orders were given out on a randomized rolling basis until all resources had been exhausted. Once a firm owner received an order, they could not receive another. If a firm owner had not yet received an order, they were eligible to receive one if enough of the existing orders were not met by their peers. 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. Therefore, to remove any timing effect, these regression include only firm owners selected to receive an order in Wave 1 or never received an order. A firm owner is considered to be a teacher or learner if either firm owner reports an interaction in a given round. Standard errors are clustered at the firm level. *** p<0.01, ** p<0.05, * p<0.1.

	(1)	(2)	(3)
	Reported Ability to Produce Sharawakil	Learned Technology from Another Garment Maker	Taught Technology to Another Garment Maker
Randomly Selected to Receive:			
Order Only, Any	0.338***	0.325***	-0.00570
	(0.0770)	(0.0767)	(0.0556)
Order Only, Size 4	-0.0366	0.00309	0.0806
	(0.0896)	(0.0904)	(0.0629)
Order Only, Size 10	-0.104	-0.113	0.269**
	(0.118)	(0.117)	(0.107)
Training Only	0.684***	0.00587	0.0133
	(0.0641)	(0.0114)	(0.0110)
Order and Training, Any	0.699***	-0.0393	0.344**
	(0.105)	(0.0905)	(0.154)
Order and Training, Size 4	-0.0829	-0.0277	0.165
	(0.110)	(0.0951)	(0.213)
Order and Training, Size 10	-0.0971	0.0897	0.227
	(0.105)	(0.182)	(0.255)
Average Across All Rounds	0.1461845	0.0364695	0.0226785
Observations	2,324	2,324	2,324
R-squared	0.764	0.436	0.373
Firm FE	YES	YES	YES
Round FE	YES	YES	YES

Table 8: **Direct Training and Order Size Effects on Outsourcing and Capital Investment**

This table reports reduced form effects of being randomly invited to a training in and/or selected to receive an order for a random number of garments utilizing a new technology on technology specific outsourcing and capital investment behavior. Because technology specific sales and expenditure were only included in the final round, this table only includes one round of data. Gender is included to control for the randomization strata. *** p<0.01, ** p<0.05, * p<0.1.

	(1) Outsourcing Buying Sharawakil	(2) Outsourcing Selling Sharawakil	(3) Investing Sharawakil Tool
Randomly Selected to Receive:			
Order Only, Any	0.0914**	0.00621	0.0559
	(0.0401)	(0.0131)	(0.0443)
Order Only, Size 4	0.0494	-0.00610	0.0677
	(0.0380)	(0.0124)	(0.0419)
Order Only, Size 10	-0.0545	0.0680***	0.297***
	(0.0555)	(0.0181)	(0.0613)
Training Only	-5.99e-05	0.000142	-0.0175
	(0.0817)	(0.0266)	(0.0902)
Order and Training, Any	0.0999*	0.000131	0.0826
	(0.0596)	(0.0194)	(0.0658)
Order and Training, Size 4	-0.0166	-7.52e-05	0.151
	(0.0913)	(0.0297)	(0.101)
Order and Training, Size 10	-0.0998	-0.000451	0.0702
	(0.120)	(0.0390)	(0.132)
Gender	-0.00190	0.00451	-0.0357
	(0.0319)	(0.0104)	(0.0352)
Observations	384	384	384
R-squared	0.078125	0.0078125	0.1041667

Table 9: **Direct Training and Order Size Effects on Quality**

This table reports reduced form effects of being randomly invited to a training in and/or selected to receive an order for a random number of garments utilizing a new technology on the sharing of that technology between firm owners on assessed order quality from three separate garment making experts. Both the garment and the weaving technique, Sharawakil, on the garment were given assessments on a curve from 1 to 10. Gender is included to control for the randomization strata. *** p<0.01, ** p<0.05, * p<0.1.

	(1) Assessor 1 Sharawakil Quality (1-10)	(2) Assessor 2 Sharawakil Quality (1-10)	(3) Assessor 3 Sharawakil Quality (1-10)
Randomly Selected to Receive:			
Order Only, Size 4	0.318	-0.00494	-0.0107
	(0.394)	(0.285)	(0.306)
Order Only, Size 10	-0.243	-0.604	0.163
	(0.537)	(0.390)	(0.418)
Order and Training, Any	-0.183	0.339	0.257
	(0.453)	(0.329)	(0.353)
Order and Training, Size 4	-0.186	-0.331	0.514
	(0.785)	(0.569)	(0.613)
Order and Training, Size 10	-0.741	-0.282	-0.461
	(1.031)	(0.750)	(0.803)
Garment Quality (1-10)	0.784***	0.425***	0.270**
	(0.0977)	(0.0981)	(0.124)
Gender	0.329	-0.287	-0.272
	(0.377)	(0.285)	(0.262)
Observations	211	211	211
R-squared	0.314	0.111	0.043

Table 10: **Experimental and External Determinants of Technology Diffusion**

This table reports correlations between baseline dyadic level characteristics and all technology sharing at baseline, in Column 1, and sharing of our newly developed technology during the orders round, in Column 2. Because the seeding of training in and demand for the new technology were random, any persistent relationship between a dyadic characteristic and sharing in Column 2 exists as a result of something unrelated to the supply and demand of technology. Standard errors are two-way clustered at the teacher and learner level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

	(1) Any Technology Diffusion in the Year Prior to Experiment	(2) Sharawakil Diffusion During Experiment
Dyadic Characteristics		
Competitors (Market Research Survey)	0.01916***	0.00081
	(0.00650)	(0.00157)
Neighbors (< 500')	0.02366***	0.00309***
	(0.00200)	(0.00050)
Same Gender	0.00430***	0.00035**
	(0.00052)	(0.00017)
Same Ethnicity	0.00202***	0.00033**
	(0.00068)	(0.00015)
Same Size (Either No Workers or Workers)	0.00127***	0.00027*
	(0.00040)	(0.00014)
Matching Management Practices (>2/5)	0.00102**	0.00020
	(0.00049)	(0.00017)
Average Diffusion	0.0062	0.0008
Observations	139,886	139,886
R-squared	0.015	0.002