

Working paper

Diffusion of Technologies Within Social Networks

Evidence from a
Coffee Training
Program in Rwanda

Esther Duflo
Tavneet Suri

March 2010

IGC

International
Growth Centre



DIRECTED BY



FUNDED BY



IGC Contract RA-2010-03-001 – Final Report

1. Introduction

Agricultural technologies are key to improving the productivity of the agricultural sector in Sub-Saharan African economies. A large fraction of these populations live in rural areas and derive most of their income from agricultural activities. Previous experiences (such as the Green Revolution in India) have shown the importance of improvements in agricultural productivity in driving structural transformation and it is clear that such productivity improvements are much needed in Sub-Saharan Africa to help the development process. Therefore, understanding the role of agricultural technologies, their diffusion, and the contribution of extension services in improving productivity is essential to understanding effective ways of improving agricultural productivity.

Much as this project is about the diffusion of improved coffee farming practice, the broader questions on the adoption of new technologies are at the heart of models of economic growth. Understanding the drivers of adoption, and in particular the role played by social networks in the special case of coffee, can shed light on technology adoption more generally. There are a number of important studies on this subject (such as Foster and Rosenzweig (1995) and Conley and Udry (2009)), but they are not based on a randomized design, leaving us substantial scope to learn more. In this project, therefore, we focus on the role social networks play in the adoption of new technologies.

To understand these questions, we designed a field experiment in one sector (administrative unit) in Rwanda, in collaboration with the Technoserve (TNS) Coffee Initiative. In this project, we focus on the agronomy training part of the TNS coffee project. As described in detail below, among a set of applicants to the training program, TNS randomly selected the farmers who participated in the agronomy training. Our first objective of the research is to determine the impact of this training on the practice of the targeted farmers. There is not a large literature on the impacts of extension or similar training services. We aim to contribute to that literature. A second objective is to better understand the diffusion of improved agricultural practices (and technologies) through social networks. To do this, we collected detailed social network data to allow us to construct complete network maps for coffee farmers in the sector. To our knowledge, this will produce one of the first complete social network maps in conjunction with a randomized “injection point” for a new technology, and the first in a developing country. In addition, since the intervention was a training program organized in groups of farmers, the intervention itself may have had effects on the social networks themselves. So, we collected midline and endline social network maps to better understand these impacts. This will be the first study to collect data to understand the dynamics in social networks, potentially partly as a result of the field experiment itself.

The rest of this report is structured as follows. Section 2 describes some background on the coffee sector in Rwanda. Section 3 describes the experimental design and Section 4 the data. In Section 5 we discuss some preliminary results for adoption from midline data collected for the research and we

also highlight the additional research to be conducted. In Section 6 we discuss policy implications and in Section 7 we conclude.

2. Background on Coffee in Rwanda

The potential to improve growth are particularly relevant in the context of cash crops, such as coffee. In Rwanda, the coffee sector has enormous potential. According to a Rwanda Ministry of Agriculture national survey in 2002, Rwanda's one billion coffee trees have the potential to generate over \$150 million in annual foreign exchange earnings and when farmer-owned cooperatives sell directly to American buyers, approximately 80% of the total gross sales revenue trickles back to the rural communities. This sector could therefore have the power to drive the rural economy, by creating employment and generating increased revenues and incomes.

However according to the same report, less than 10% of coffee farmers used chemical fertilizers, 14% did not know if fertilizers were useful, and more than 85% did not routinely replace the old trees with new ones. This shows the growth impact that proper agricultural technology could have. If agricultural extension services improve the use of the correct agronomic practices, and if this correct use diffuses from farmer to farmer, the implications for growth could be potentially extremely large.

Coffee is a major export crop for Rwanda and contributes about US\$ 34 million in export earnings a year, about 15% of the country's foreign exchange earnings. The coffee industry in Rwanda is dominated by about 500,000 small-holder coffee producers, with an average of 170 coffee trees each, and the area under coffee is about 29,000 hectares (3% of total cultivated area). Rwanda has ideal growing conditions for coffee (altitude, rainfall and soils), but, given the importance of coffee for the Rwandan economy, the agronomic practices at the farm level and processing are in dire need of improvement. In this context, this project aims to understand two main aspects of the diffusion of agricultural technologies in the coffee sector in Rwanda. We are working with TechnoServe (henceforth TNS), a non-profit organization, who is introducing a four year program to help approximately 180,000 small scale farmers in East Africa boost their income by improving the quality and quantity of coffee through building wet mills for the coffee and training the farmers on agronomy best practices. The project has commenced in Kenya, Rwanda and Tanzania and soon to commence in Ethiopia.

3. Experimental Design

In this project, we study the role of social networks in the diffusion of the agronomy best practices. We do this using a field experiment designed as follows. The field experiment also allows us to understand the impact of this training, or more generally agricultural extension services, on technology adoption and productivity.

TNS Rwanda conducted an agronomy training program for farmers to teach them a number of coffee "best practices": rejuvenation and pruning; nutrition (organic and inorganic depending on the soils); pest, disease and weed management; weed management; mulching; soil and water conservation

(water traps, terracing); shading and record keeping. The experiment was designed by farmers signing up to participate in the training program and then half of these farmers being randomized into a treatment group and half into a control. For the training itself, the participating farmers were allocated to groups of about 30 farmers and elected a focal farmer in each group. Once a month for 11 months, they met at the focal farmer's coffee farm and were (hands on) trained by a TNS staff member on one of these best practices. Training sessions therefore took place once a month for 11 months in the first year of the program (2010) and then an additional 5 review training sessions were done over the course of the second year (2011).

An additional component we built into the design was that in a random third of the villages covered by the agronomy program only 25% of the signed up farmers were allocated to the treatment group, in a second random third 50% were allocated to the treatment and in the final third 75% were allocated to treatment. This created a varying intensity of treatment at the village level, allowing us to better identify and observe the diffusion process.

The last TechnoServe agronomy training was held in October 2011 so we are now in the post-treatment phase of the project, with further data collection scheduled for 2012 in hopes of capturing longer term effects of the program. The next section describes in detail all the survey data that has been collected to date as well as the two remaining rounds of data collection to be completed in 2012.

4. Survey Data Collected

Since we are interested in finding out how information diffuses along the treated farmers' social networks, there were extensive data collection efforts for this project. There were a number of different forms of survey data collected, which we describe here in detail.

The first set of data are detailed and accurate representations of social network maps. For these social network maps and statistics to be accurate, it is important that we not only surveyed the 1,600 farmers in the treatment and control groups of the field experiment, but also all the other coffee farmers in the villages that the treatment and control farmers belong to. This is for two important reasons. First, surveying only farmers in the treatment and control groups would give us only a partial view of the network and information links between farmers. For example, if A is friends with B and B is friends with C and C is friends with D, we will not know that A and D are linked unless we survey B or C. Second, treatment and control farmers were not randomly selected since they applied to the training program. Thus, surveying all of them will not only give us a partial map, but a non-random partial map, which means that the techniques for the "completion" of network maps that statisticians have recently developed cannot easily be used.

To get accurate social network maps, in December 2009, we conducted a Census of more than 5100 (approx.) households in Nyarubaka sector over a two week time span. Of this, 1600 (approx.) households were identified and matched with coffee farmers participating in Agronomy training evaluation and an additional 1300 (approx.) were identified as other coffee-farming households in the sector. The remaining 2200 (approx.) of the visited households were determined as non-coffee farming households and were excluded from the following two survey activities since there was no coffee

farming activity or production to observe. Since we were interested in the diffusion of coffee practices, we decided to collect social network data only on the sample of coffee farmers, all the more since coffee takes many years of investment before the trees produce coffee.

The social network mapping was therefore done for the 3000 (approx.) coffee farming households in the sector. We conducted complete social network surveys for all these 3000 coffee farmers in the sector (administrative unit we are working in). We asked two different questions on links, separately for the household head and for his spouse. The first was just friends and the second was people they spoke to about coffee. The first social network survey was conducted at baseline (late 2009 into early 2010). This was followed by two more rounds of social network surveys in early 2011 and then again in early 2012 (this was just completed). In addition, in early 2012, we collected data on the geographic networks of all the coffee farmers in the sector, by collecting the GPS coordinates of all farmers' houses and plots.

In addition to these maps, a full detailed baseline survey was conducted over December 2009 through January 2010 which collected detailed crop production and profit information (across all crops), a full consumption module, household demographics and assets. This detailed baseline will include only the 1600 treatment and control households. For this sample of treatment and control households, the baseline was followed by eight additional detailed rounds of data collection over two years (with the last two rounds yet to be conducted before the end of 2012). The data was collected at high frequency to minimize error in reporting of coffee yields, inputs (including labor) and sales, given the seasonality in coffee production. In addition to these self reported surveys, we conducted bi-annual best-practice visits to the treatment and control farms to actually observe whether the farmers are indeed adopting the agronomy best practices. These best practice visits are also conducted in the additional "non-sample" households in the sector at the same frequency to monitor diffusion of the practices.

Finally, in 2011 we tried to improve the collection of yield data by giving farmers very basic scales along with a calendar where the farmers record their daily harvests of coffee. We check and pick up the calendars at regular intervals. This was done only in the treatment-control sample. We are not collecting detailed coffee harvest data for the "non-sample" farmers as for that sample we are only interested in the diffusion process.

Since the start of the study, therefore, eight rounds of survey data have been collected for our sample households (approx. 1600 farmers) and four rounds have been carried out on an additional 1300 (approx.) "non-sample" coffee farmers in the area of our study. These additional 1300 farmers are the sample that allow us to track the diffusion of the agronomy training best practices through social networks. In terms of future data collection, there are still two rounds of detailed production data to be collected on the treatment-control sample of households. The yield-recording calendars will continue to be distributed and collected every two months through December 2012 to the treatment-control households, and at least one more round of the Best-Practices survey will be carried out on each of the coffee farmers (all 3000) to further assess the adoption and diffusion of the agronomy training techniques.

5. Preliminary Results

In this section, we present some preliminary results from the data collected so far, in particular, using data from the baseline and the first six rounds of data collection (which spanned over the 2010 and 2011 coffee cycles). This section summarizes results achieved to date.

We first looked at attendance data based on the TechnoServe attendance sheets that are filled out by farmer trainers at each meeting to compute attendance rates. Looking at the attendance data from 11 trainings which took place between February 2010 and May 2011 (excluding the 4 review sessions), we find that the average number of meetings attended by at least one member of treatment households is around 8, which gives us an attendance rate of about 73%. An interesting finding is that attendance tends to be higher in villages with higher treatment-to-control ratios: households in 75% treatment villages attended around 8.2 out of 11 meetings, whereas households in 50% treatment villages attended around 7.9 meetings and households in 25% treatment villages attended around 7.2. This suggests that there might be a “critical mass” effect of the trainings – farmers who live in villages where more of their neighbors attend the trainings might be incited to attend themselves – so we might also expect diffusion of agronomic practices, if any, to be higher in villages with higher treatment-to-control ratios. This will be one of the main areas of investigation that forthcoming analysis will focus on.

Data on coffee harvests used for this analysis came from two sources. First, we used farmer responses to survey questions on harvest quantities from the baseline survey in December 2009 (when farmers were asked to recall the whole year's production in a single survey) to the bimonthly monitoring survey in July 2011. This covers both 2010 and 2011 coffee harvest seasons. Second, we used the data from daily harvest calendars filled out by farmers using scales and calendar sheets provided by TechnoServe. These cover only the 2011 season since these measuring tools were only distributed to the sample households in January 2011. All regressions controlled for village and year fixed effects as well as baseline differences between treatment and control. Controlling for household fixed effects was also included in the analysis but this did not affect the results much.

The results suggest that the training led to a roughly 10% increase in harvests reported in 2010 (both for in-season harvests and for all harvests) though the estimate is only significant at around the 15-20% level – which is likely to be a power issue and should be rectified by the ongoing rounds of data collection. Both the 2011 rounds of self-reported data and the calendar data showed no evidence of any effect on harvests in 2011, but this can likely be attributed to the fact that 2011 was a down year for coffee (and low coffee production might have discouraged treatment farmers from adopting time-consuming and costly practices for this coffee season if the quantities harvested were initially substantially lower than over the previous harvest cycle, thus making any investment into coffee for that year less attractive).

Interestingly, there is some evidence that there may be an effect of the 2011 training on those farmers whose harvests were the lowest in 2009 (at baseline), the year before the program began. Results from interacting the 2011 treatment dummy variable (i.e. indicating whether the household was in the treatment group and whether the year was 2011) and baseline yields suggest that a 1% lower pre-program harvest is associated with an increase in the effect of training of 12.5 percentage points. While this relationship appears absent for 2010, it suggests that the effects of the training may not be

equal for all farmers – perhaps the program works best on farmers whose pre-treatment yields were the lowest because they had the least knowledge of best practices initially. Heterogeneity of treatment effects will be further investigated in the upcoming rounds of analysis in 2012.

We also looked at the adoption of the agronomy best practices that the treatment farmers were trained in and compared adoption across treatment and control farmers using two rounds of post-treatment “best practices” survey data, the first sample survey (April/May 2010) and the sixth sample survey (June/July 2011). The best-practices surveys consist of visits to each and every of the farmers’ coffee plots and included thorough plot and tree inspections. These surveys are administered by trained surveyors to observe and document the present status of farming activities of mulching, weeding, pruning, pest control, and other observable measurements to determine what level of the “best practices” taught by TechnoServe are being implemented. Baseline data was also included in this analysis when possible, but comparability with monthly survey data is limited due to the self-reported nature of the baseline data - enumerators did not perform plot or tree inspections in the baseline survey. An additional round of plot and tree inspections was recently carried out in January and February 2012, as part of the eighth sample survey, but this data is currently in the process of being cleaned before it can be used for further analysis.

The most highly significant treatment effects to be detected using these three datasets concerned bookkeeping habits and keeping a compost heap: treatment farmers were more likely to keep record books than control farmers by 82.6 percentage points in the first sample survey, and more likely by 71.5 percentage points in the sixth sample survey. Treatment farmers were also more likely by 7.2 percentage points to keep a compost heap in the First Monthly, and by 8.8 percentage points in the sixth sample survey.

Our preliminary results also suggest that awareness of coffee agronomy best practices (in particular, awareness of which fertilizers are particularly effective for coffee and integrated pest management methods) is higher among treatment farmers, which shows that notions of what *can* be done to improve yields are assimilated by the farmers who attend the trainings. Treated households are statistically more aware of most pest control methods in first sample survey (average awareness is higher for the treatment group by 7.4 percentage points) and all practices in the sixth sample survey (where average awareness is higher by 5.9 percentage points). Treated households are also more likely to report NPK as a good fertilizer for coffee (4.8 percentage points more likely in the first sample survey and 5.4 percentage points in the sixth), as well as zinc/boron (4 percentage points more likely to be reported as a good fertilizer in the sixth sample survey).

Now, treatment effects on *actual* application of the practices are somewhat less consistently detected and vary across the different types of practices. The data does show that such effects are likely to exist for Integrated Pest Management practices: the first sample survey self-reported data suggests that treatment effects exist for several pest control methods including beneficial insects (5.7 percentage points more likely to be used as a pest control method), and keeping trees fed/healthy (10.7 percentage points more likely). Treatment households are also 7.4 percentage points more likely to use a TechnoServe method for pest control in the sixth sample survey. We find significant negative treatment effect for pesticide spraying (-11.7 percentage points in the First Monthly, and -6.6 percentage points in the sixth sample survey) – which shows again that farmers do retain some of the methods that they are

trained in since TechnoServe encourages farmers to use other more cost-efficient and ecologically-friendly pest control methods and keep pesticide spraying as a last resort for pest control. Significant treatment effects exist for several methods including removing old and dry berries (3.2 percentage points more likely in first sample survey – though this effect is not detected in the sixth sample survey) and squashing *Antestia* (3.2 percentage points more likely to be reported as a used pest control method in first sample survey and 3.8 percentage points in sixth). *Antestia* is one of the pests that might require chemical spraying to be eliminated, according to the TechnoServe trainings, and we do find that treated households using pesticides are significantly more likely than control households to report using them to combat *Antestia* (11.8 percentage points more likely in first sample survey and 7.2 percentage points more likely in sixth).

These results suggest that trainings have not only been effective in making farmers retain information on the most effective pest control methods, but that the knowledge imparted to the trained farmers also triggered some behavioral changes: treatment farmers are more likely to withhold from applying any unnecessary practices that might do more harm than good to their trees' productive ability (such as pesticide spraying), and they are more efficient and strategic in their use of pest control methods to target specific pests.

Significant treatment effects were also detected for tree nutrition and mulching. Tree inspections data reveals that trees of treated households receive better nutrition (significantly so in the first sample survey): leaves of treatment trees are less likely to exhibit yellowing (-2.7 percentage points in the first sample survey) and curling (-1.1 percentage points in the first sample survey), and treatment trees are significantly more likely to show signs of mulching (4.9 percentage points in the first sample survey). In addition, we find significantly lower incidence of leaf rust (-3.9 percentage points in the first sample survey) on the treatment farmers' trees. None of these effects are significant in the sixth sample survey however. Nonetheless, trees of treated households are significantly more likely to be weeded (by 3.6 percentage points) in the sixth sample survey.

On the other hand, we find no effects of the trainings on the adoption of pruning techniques in either the first or the sixth sample surveys. New measures were added to the existing set of tree inspections criteria which should increase precision and might help detect any treatment effects on pruning techniques, if any, in upcoming rounds of analysis.

All in all, the data suggests that, though the trainings do make the farmers more knowledgeable about best agronomic practices, there is a gap between awareness of the practices and their actual implementation, as farmers are more likely to adopt practices that require the least effort (such as withholding from spraying pesticides) over practices that are distinctively more time-consuming and "unpleasant" tasks such as pruning.

5.1 Upcoming Analysis

Here we outline some of the forthcoming analysis as the majority of it will be carried out over the remainder of 2012. Extensive social network and diffusion analysis is scheduled for the remainder of 2012. Once this third round of mapping is complete, we will be able to use these maps to help us determine what practices are spilling over to those farmers not directly benefitting from the training

(the non-sample farmers) as well as track diffusion of these practices (if any) from the treatment to the control farmers. This will require using the best practices data from the First, Sixth and Eighth Monthly surveys (and the social network maps corresponding to these dates). We will study how variation in farmers' "connectedness" to other coffee producers tends to affect their own adoption of coffee best practices.

We will also use the GPS data we collected for each farmer's house and coffee plots to detect any diffusion of practices on farmers living close to plots of "adopters" – as we might expect people living close to plots where the effects of certain practices are visibly evident to be inspired to adopt these same practices as well.

We will also study the diffusion from the treatment to the control group to potentially explain the rise in adoption that we see among the control group for some practices. Finally, we will compare diffusion rates across 25% treatment villages, 50% treatment villages and 75% treatment villages – as this variation in treatment intensity was designed to help detect diffusion better by introducing some variation in the number of "injection points" for the knowledge of best practices across our clusters.

As suggested by the preliminary yield analysis, there might be some heterogeneity in the effects of the treatment (other than that related to their level of connectivity within social network maps), which we will further investigate in the upcoming months. We will test for any heterogeneity among treatment farmers' adoption of best practices using variation in farmers' wealth (we also have data on farmers' non-agricultural sources of revenue), plot size, number of coffee trees, and pre-treatment yields.

We will also measure whether adoption varies with the fraction of the time and land that farmers allocate to coffee – one might expect that the farmers who are the most dependent on coffee for their livelihood might also be keener to improve their coffee yields and spend more time applying inputs and practices on their coffee plots. For this we will use the data we have collected on a bi-yearly basis since 2009 (including the baseline) on crop production and profit information across all crops.

This is important, since if the information diffuses rapidly, targeting a few farmers may be sufficient to induce large changes in farming practices. This context is a particularly interesting one to learn about diffusion, as there are interesting sources of variation in the type of technologies that farmers are trained in. Some of the coffee best practices result in outcomes that are visibly evident to farmers in the area (i.e. your coffee trees just look much better/healthier) whereas others are less visible. So, we may expect to see differences in the relative speeds of diffusion of these various practices (and some diffusion maybe even outside the group of farmers a trained farmer usually talks to, i.e. in geographical space).

7. Policy Implications

This research has extremely important implications for policy. Understanding the adoption of new, productivity enhancing technologies is extremely important for policy. For a partner like TNS, it is important for the design of their program to understand the diffusion process. If there is a lot of diffusion, then TNS can train a few farmers and expect that others will adopt and learn from the trained

farmers. However, if there is little diffusion, then TNS is better off concentrating their program over small geographic spaces to get the largest impact. The additional important policy implications that will emerge from this research deal with the impacts of extension services in a developing economy. Given the magnitude of resources that governments spend on extension services of this kind, it is important to understand whether they have effects, how large and how broad given that there may be diffusion of the training. We hope to be able to answer all these questions over the course of the next several months with the data collected in this project.

8. Conclusions and Timeline

We expect final results for the impacts of the training program to be available after the last round of data collection at the end of 2012 for the treatment and control sample of farmers. In addition, we expect to have some results on the adoption process and diffusion from the eighth survey sample that was just completed in March 2012 (the data is still being cleaned and processed). This is the first survey post the TNS program (after both the training sessions as well as the review sessions) and so will give us an idea of the magnitude of the adoption response. In addition, we expect one more round to adoption data collection in the Fall of 2012. Finally, we expect to be able to look at the dynamics in social networks over the next few months using the eight survey round as that collected the third and final social network map.

The International Growth Centre (IGC) aims to promote sustainable growth in developing countries by providing demand-led policy advice based on frontier research.

Find out more about our work on our website
www.theigc.org

For media or communications enquiries, please contact
mail@theigc.org

Subscribe to our newsletter and topic updates
www.theigc.org/newsletter

Follow us on Twitter
[@the_igc](https://twitter.com/the_igc)

Contact us
International Growth Centre,
London School of Economic and Political Science,
Houghton Street,
London WC2A 2AE

IGC

**International
Growth Centre**

DIRECTED BY



FUNDED BY



Designed by soapbox.co.uk