Chat Over Coffee? Diffusion of Agronomic Practices and Market Spillovers in Rwanda*

Esther Duflo†
Daniel Keniston‡
Tavneet Suri§
Céline Zipfel¶

May 26, 2023

Abstract

Agricultural extension programs often train a few farmers and count on diffusion through social networks for the innovation to spread. However, if markets are imperfectly integrated, this may also inflict negative externalities. In a two-step experiment of an agronomy training program among Rwandan coffee farmers, we first randomize the concentration of trainees at the village level and then randomly select within each village. Knowledge increased, and yields were 6.7% higher for trained farmers. We find no evidence of social diffusion; instead, control households experienced negative spillovers in high treatment concentration areas, likely because of competition for a scarce input, fertilizer.

JEL Classification: O12, Q13, Q16

*We would like to thank Emmanuel Bakirdjian, Jean Louis Uwitonze and Nick Tsivanidis for their exceptional contributions in managing the experiment. We are also grateful to Ali Ahmed, Martin Albouy, Harris Eppsteiner, Gabriella Fleischman, Sara Hernández, Adam Ray and Simon Yaspo for excellent research assistance. Special thanks go to Theodomir Ntasoni, the lead supervisor of our team of enumerators. Theodomir sadly passed away in February 2022; this project would not have happened without his hard work and dedication, we remain ever grateful. We also thank Damien Munyandekwe for sharing his knowledge of coffee agronomy and markets in Rwanda, and his help in transporting over 20,000 paper surveys from Kigali to Nairobi. We gratefully acknowledge financial support from TechnoServe (as part of their funding from the Gates Foundation) and the International Growth Centre.

†MIT
‡Louisiana State University
§MIT Sloan School of Management
¶Stockholm School of Economics

1 Introduction

Given the importance of agriculture to low-income economies, the successful adoption of yield-enhancing technologies is critical for well-being in much of sub-Saharan Africa, where agricultural productivity is increasingly lagging behind compared to the rest of the world (Suri and Udry 2022; FAOSTAT 2022). Agricultural training is poised to play an important role in the diffusion of these technologies (Cole and Fernando 2012) and governments, NGOs and firms spend considerable resources on it.

A common way to organize agricultural training in developing countries is to train a few farmers in each community, and then rely on the organic spread of the information within social networks. This could be an efficient way to use limited training resources if farmers learn from each other. However, a potential downside is that trained farmers could impose a negative externality on untrained farmers, in particular by increasing the demand for scarce inputs in settings where input markets are frequently far from perfect. In this case, direct comparisons of trained and untrained farmers could also overstate the effect of agricultural training by failing to account for those externalities. Another source of potential externalities from an intervention that affects some people but not others is the endogenous re-wiring of social networks (Banerjee et al. 2021): if treated farmers find it advantageous to only socialize with other treated farmers, for example, this could affect the sharing of both information and risk in the community, potentially penalizing the untreated farmers. Conversely, if control farmers seek out treatment farmers, this could accelerate the transmission of any intervention (Comola and Prina 2021).

While a growing literature studies the diffusion of agricultural innovation through social networks, the results of existing experimental studies are mixed, suggesting that diffusion may depend on a variety of factors, including the simplicity of the technology (Chandrasekhar et al. 2022), how novel it really is (Bridle et al. 2019), how profitable it is (Magnan et al. 2015), and the identity of the early adopters.[1] Meanwhile, the potential for negative externalities and network externalities has not, to our knowledge, yet been investigated.

In this paper, we fill this gap, reporting results from an experiment among coffee farmers in Rwanda. We designed a two-stage randomized controlled trial (RCT) and collected detailed social network data to test both for the presence of any transmission of information provided in a farmer training through the social network, and for spillovers (positive or negative) to untrained farmers in the village. We find very little evidence of knowledge diffusion in the social network, and clear evidence of negative spillovers, most likely through the crowding out of inputs in limited supply, including chemical fertilizer, for which there was at the time a very patchy nationwide market, and labor, for which the market is also imperfect. We conclude that the naive comparison between treated and control farmers within a village would have led to a serious overestimate of the benefits of the program.

[1] For recent reviews of this literature, see Suri and Udry (2022) and Caldwell et al. (2019).

The program, designed and conducted by a leading international NGO, was an intensive agronomy training offered to coffee farmers to help them improve their yields. To measure direct and indirect impacts, the design is similar to Crépon et al. (2013). First, we collected interest in the program from coffee farmers. Second, we randomly varied treatment concentration at the village level across 27 villages: in approximately one third of villages, 25% of farmers who signed up were to be treated, in another third, 50% and in the final third, 75%. Finally, the 1,594 farmers who signed up for the program were allocated to treatment and control following the assigned proportions in each village.

Selected farmers received monthly instruction modules for the first year, followed by six refresher modules over the second year. These modules covered nutrition, pest and disease management, weed management, mulching, rejuvenation and pruning, shade, soil and water conservation, and record keeping. Farmers were grouped for the training and picked a lead farmer whose plot was used for demonstrations of the agricultural practices.

We collected ten rounds (including four post-treatment) of data on a wide range of indicators (self-reports and plot audits) to measure the impacts of the training on farmers' knowledge and adoption of the practices as well as on their yields. To measure both diffusion effects and any impact on social connections, we collected social network data before and after the training.

From the household surveys, we construct an index of knowledge and an index of (self-reported) adoption of improved agricultural practices. We also measure use of fertilizer, labor inputs, and yields. From the audit data, we construct an index of adoption of improved agricultural practices, and a measure of leaf nutrition.

The program had significant effects on knowledge and self-reported adoption of the agronomic practices: the training led to a 1.2 standard deviation increase in the knowledge index and a 0.33 standard deviation increase in a self-reported adoption index, although treatment effects on adoption according to the tree audits are insignificant. Treated farmers also use more inputs (fertilizer and labor) than control farmers in their village (0.1 standard deviation higher indexed). After the training, yields are about 6.7% higher in the treatment group compared to control farmers in the same villages. Note, however, that this could be either an overestimate or an underestimate of the actual treatment effect if the program had positive or negative spillovers on the control group. Either of those would violate the SUTVA assumption.

To examine whether these new practices diffuse to those people who treated farmers discuss coffee production with, we exploit the fact that, for a given farmer, conditional on the total number of friends from the original sample that they list as their contacts, the random assignment of the training opportunity generates exogenous variation in the number of contacts treated. Using this variation, we find no evidence of spillovers to friends in the control group on knowledge, adoption, or yields. Using a similar strategy for neighbors, our results suggest that there is no diffusion of the program's effects to geographic neighbors either.

Despite the absence of informational spillovers, there could be externalities on other farmers stemming from other sources, such as the crowding out of limited inputs. To test this hypothesis, we exploit the exogenous variation in village-level shares of treated farmers generated by our experimental design. Even though the number of villages per treatment intensity group is small, there is a clear pattern, which is largely robust to randomization inference confidence intervals: the higher the fraction of treated farmers in the village, the lower the yield. The effects are large: the yields of control farmers are 25% lower in the villages with 75% farmers treated than in the villages with 25% farmers treated. Control farmers in heavily treated villages also use less inputs (chemical fertilizer and labor days per tree), and have poorer leaf health.

This result raises the possibility that the apparent impact of the training on yields that we find when comparing treatment and control farmers is, in part or in full, due to a reduction in the yield of the control group rather than an increase in the yield of the treatment group. To check this, we examine heterogeneity in the difference between treatment and control groups by village treatment concentration. We find that the differences between treatment and control farmers in knowledge and self-reported adoption of the new practices are similar in villages with 25%, 50% and 75% shares of treated farmers. However, the apparent "treatment effect" on yields is actually negative (though insignificant) in villages with few (25%) treated farmers, and becomes more positive in villages with more treated farmers: the yields of treated farmers in villages where 75% are treated are 17% higher than those of control farmers. Overall, this suggests that the apparent positive impact of the program on yields from our within-village analysis is an artifact of the fact that yields went down for control farmers. The training program led to a re-allocation of inputs from untreated farmers to treated farmers, which hurt the yields of untreated farmers without leading to a significant increase in those of treated farmers. This entirely reverses the conclusion that one would have reached by comparing treatment and control farmers within a village.

The fact that the treatment does not increase the yields of treated farmers but decreased those of control farmers suggests that training some farmers but not others may have inadvertently lowered aggregate output by increasing input misallocation. To provide suggestive evidence of this channel, we present estimates of a coffee production function, with labor and fertilizer as the main inputs. A fully robust estimate of the production function is beyond the scope of this paper, but simple descriptive regressions suggest that the production function is indeed concave in how much chemical fertilizer, in particular NPK, and labor is used, and no different across treatment and control farmers. Input re-allocation is thus unlikely to have increased efficiency in aggregate coffee production in this sample.

We end the paper by examining another source of potential spillovers from the intervention affecting some farmers but not others: the fact that the treatment may itself have modified the structure of social networks. We find that treatment farmers indeed made new friends within the treatment group, in particular farmers with whom they attended the training. But this does not come at the cost of friendships in the control group: on net, they just seem to make more friends than their peers in the control group, especially within the sample.

The results in this paper are in some ways specific to the context of this particular program and of the coffee context in Rwanda over the period. Ultimately, although they learned and retained the information, the farmers do not seem to have applied most of the agronomic practices, except for intensifying the use of scarce inputs. The impacts on their crops were not large. This may thus not have been an ideal setting for diffusion. Furthermore, the program took place in a context where NPK was not widely available in the district. The government phased out the direct delivery of fertilizer during our study period, relying on private actors, cooperatives and agro-dealers to take over, but this happened slowly and imperfectly. This likely accentuated the negative externalities. But, as we describe below, the complexities of market-based delivery of fertilizer have roots that make it a fairly general problem in the region. The lesson that any strategy that involves helping some people and not others may backfire when markets are imperfect and poorly integrated is much more general. Furthermore, the fact that this may have been missed without a village-level randomization calls for caution in the evaluation of agricultural extension programs.

2 Background and Program Description

2.1 Context: Rural Rwanda

Coffee is Rwanda's most important export crop, contributing about US$62 million in export earnings per year (NISR 2019). Production is dominated by 500,000 smallholder producers (OCIR-Café 2008). Intensifying coffee production and increasing the sector's productivity were key targets of the government's strategic plan for boosting agricultural development. Rwanda has ideal growing conditions for coffee, but agronomic practices were poor. For example, the national rate of chemical fertilizer consumption per cultivated hectare was 4KG in 2009, below the sub-Saharan African average of 9 to 11 KG per hectare (ROR 2009).

2.2 The Intervention: Agronomy Training Program

The context for our study is a large agronomy training program for small-scale coffee farmers, aimed at improving the health of coffee trees and ultimately yields. TechnoServe, an international agri-business NGO, conducted agricultural training programs in several coffee growing regions in East Africa between 2010 and 2015. This study focuses on the agronomy program in one sub-district in Southern Rwanda, run between February 2010 and October 2011.

The training covered several best practices in coffee growing: tree rejuvenation and pruning; fertilizer use; pest, disease and weed management; mulching; soil and water conservation; optimal shade; and record keeping. Appendix C provides a more detailed description of the practices that were covered and what the NGO's agronomists expected the impacts of each to be. The training sessions took place once a month for eleven months in the first year, and TechnoServe delivered an additional six review sessions the following year. The trainings were conducted with groups of approximately thirty farmers and took place on the plot of a designated "focal farmer". The focal farmers were chosen partly because of the accessibility of their coffee plot but were also meant to be respected members of the local community and have an enthusiasm for learning.

The training itself was conducted by a Farmer Trainer (there were four in total), each of whom supported approximately 10 of these focal farmer groups. These farmer trainers received monthly training from an agronomist for each module, together with lesson plans and activities. They delivered the training to each group on a plot of approximately forty trees (the focal farmer's demonstration plot), with all practical work done by the farmers in the training group.

The sub-district in which our study is located comprises 29 villages. Once TechnoServe decided to train in this sub-district, they advertized the program. The farmer trainers were then assigned to visit the villages over a week to register the interested farmers, visiting each village at least twice. In total, 1594 farmers registered interest in the program. Although the program was advertized in all 29 villages in the sub-district, only farmers from 27 of those villages registered to join the program.

3 Experimental Design and Data

3.1 Experimental Design

The 1594 farmers who registered for the program were randomized into a treatment and a control group in two steps. First, we randomly varied treatment concentration at the village level: in approximately one third of villages, 25% of farmers who signed up were to be treated, in another third, 50%, and in the final third, 75%. In the second step, the 1,594 farmers who signed up for the program were allocated to treatment and control following the assigned proportions in each village. 855 farmers were assigned to the treatment group to receive the agronomy training and 739 farmers were assigned to the control.

Farmers were assigned to training groups in their village, or in the village nearest to their location if the number of treatment farmers in their village was less than the minimum size for a training group. In larger villages, treatment farmers were split into two or three groups for training, based on geographical convenience. This split was not randomized. Once assigned to a training group, farmers were expected to remain in the same group throughout the duration of the program.

During the first year of training, farmers in the training villages attended an average of 8 out of the 11 meetings (an attendance rate of about 73%). Attendance tended to be higher in the villages with higher proportions of farmers who were offered training: households in high-density villages attended an average of 8.2 out of 11 meetings, whereas households in lower density villages attended around 7.2. This suggests that farmers were more motivated to attend the training when more members of their community were also invited to attend. Attendance rates of each training session are reported in Appendix Table A2.

3.2 Data

We designed extensive data collection activities over the course of almost three years. In total, we collected ten rounds of survey data, in addition to a census. As we describe in detail in

Appendix D, different modules were asked in different survey waves, and in some rounds we surveyed not just treatment and control households, but all the coffee farmers in the 29 villages of the sub-district in which our sample is located.

In addition to survey data, our enumerators also audited the coffee plots. We were concerned that asking farmers several times about their adoption patterns may result in them erroneously reporting positive adoption simply because they were asked about it repeatedly. For the audits, TechnoServe agronomists trained our field staff to recognize the relevant set of agronomy practices. The field staff were then given an algorithm of which trees (they were to pick five) to inspect on each plot.

To measure diffusion through social networks, we collected the names of farmers to whom the head and spouse in each household talked to about growing coffee, before and after the intervention. We use this data to construct complete social network maps of all the coffee-growing households in the sector, covering over 5,000 households across the 29 villages. We also collected GPS coordinates of plots and households to construct two measures of neighbors: people farmers live close to and people who have coffee plots next to their own coffee plots.

4 Results

As can be seen from Appendix Table A1, we find balance between treatment and control groups across a wide variety of baseline outcomes.[2] We discuss our main results, in turn.

A. Agronomy Training Treatment Effects: Basic Within-village Specification

For our most basic results, we estimate the following specification:

$$y_{ijt} = \alpha + \beta Treat_{ij} + \gamma_j + \delta_t + \epsilon_{ijt} \tag{1}$$

where y_{ijt} is the outcome for household i in village j in survey round t, $Treat_{ij}$ is a dummy variable for whether the household was allocated to the agronomy training program, γ_j are a set of village fixed effects and δ_t are survey round fixed effects.[3] We cluster standard errors at the household level. Appendix Table A3 shows that we had 3% of attrition in our sample between baseline and the final endline, with no evidence of differential attrition by treatment status. We do not include any controls, but our results are robust to controlling for baseline outcomes selected by post-double selection LASSO (see Appendix Table A6).

This specification compares treated and control farmers within the same village. For β to be interpreted as a treatment effect, we would need to assume SUTVA, or the lack of any impact on control households. This would be invalidated if information did diffuse to the control group (in which case the treatment effect on knowledge would be underestimated), or if there

[2] The p-value on the joint F-test for all these outcomes is 0.9998.

[3] We use all survey rounds collected after June 2011 as our endline, namely rounds 6-9 (see Appendix D for details on the module coverage of each survey round).

were negative externalities for some outcomes (in which case the treatment effect would be overestimated). We examine this assumption below.

Table 1 reports β for measures of knowledge and adoption of the agronomic practices, input use, and yields. The first two columns report significant differences between treatment and control farmers on knowledge and self-reported adoption of the agronomic practices: the treatment led to a 1.24 standard deviation (henceforth SD) increase in the knowledge index and a 0.33 SD increase in the self-reported adoption index. This suggests that the treated farmers gained new knowledge about coffee agronomy and report putting this new knowledge into practice.

The tree audits data are an important complement to the self-reports, as they allow us to test whether the treatment group's higher reported adoption is actually visible in practice (column 3), and whether it translates into noticeably healthier-looking trees (column 4). Column 3 shows that the program did not have any effect on the adoption index constructed from the tree audits data. Appendix Table A4, which breaks down this result by component of the index, shows that the treatment effects on observable practices are concentrated in weeding and mulch application. Column 4 provides suggestive evidence that the trees of treatment farmers are better nourished: the audits data reveal a decrease in the index of probabilities that yellow, curling or rusting leaves are observed upon inspection by 0.04 SD. However, this estimate is not statistically significant (the p-value is 0.156).

Column 5 reports treatment effects on input use. Overall, treatment farmers report using significantly more inputs (labor and fertilizer) on their coffee plots, by 0.102 SD. Appendix Table A5 shows that this is due to greater use of NPK (a 20% increase in quantity applied per tree) and labor (in particular, paid labor days per tree are 15% higher). Finally, column 6 suggests that the treatment had an impact on yields, which are about 6.7% higher in the treatment group than in the control.

B. Information Diffusion

We now turn to the diffusion of information through treated farmers' networks. Our analysis focuses on two types of networks: baseline friends and neighbors. We focus on baseline friends in this section, but the results on diffusion through neighbors are similar (Appendix Table A9). The identification strategy exploits the exogenous variation in the number of treatment friends, conditional on the total number of RCT-sample friends.

Table 2 therefore reports results from the following specification, which we run only on the control group sample:

$$y_{ijt} = \alpha + \beta NumTreatFriends_{ij} + \delta NumFriends_{ij} + \gamma_j + \delta_t + \epsilon_{ijt} \tag{2}$$

where y_{ijt} is the outcome for household i in village j in survey round t, $NumTreatFriends_{ij}$ is the number of treatment friends of household i at baseline, and $NumFriends_{ij}$ is their total number of baseline RCT-sample friends (i.e. who entered the treatment lottery). As in Table 1, we cluster standard errors at the household level and we do not include any controls (see Appendix Table A8 for the results where we control for baseline outcomes selected by post-double selection LASSO).

The estimates in columns 1 and 2 report the diffusion effect of treatment friends on the knowledge and self-reported adoption of control farmers. The coefficient estimate in the first row of column 1 shows no diffusion of knowledge about the agronomic practices to the control group. Consistent with this finding, column 2 (self-reported adoption) shows no diffusion of actual practices through farmers' networks, corroborated by the absence of spillovers on leaf health in column 4. Additional results reported in Appendix Table A7 also support the absence of knowledge dissemination from the treatment to the control group, using the control farmers' assessments of whether they learned something new about each of the trained practices from a treatment farmer.

The significantly negative coefficient in column 3 is surprising: it indicates that, for an average farmer in the control group, having one more friend in the treatment group decreases this index of tree audit outcomes by 0.05 SD. This suggests that control farmers end up applying fewer best practices (or applying them less intensively) if they have more links to treatment farmers at baseline. Appendix Table A9 shows very similar results on neighbors of treated farmers, including a negative effect on this outcome.

C. Treatment-Control differences and Spillovers by Treatment Concentration

Next, we exploit the fact that our experimental design also generated exogenous variation in the village share of farmers assigned to treatment. We use this to look at heterogeneity in program impacts by village treatment concentration. The results reported in Table 3 regress the same outcomes as in Tables 1 and 2 on treatment status, indicators for 50% and 75% treatment concentration villages respectively and their interaction with treatment status, controlling for survey round fixed effects. Here, we cluster standard errors at the village level, and since the number of villages is small, we also report randomization inference p-values for the exact null hypothesis. We do not include any controls (see Appendix Table A10 for the results where we control for baseline outcomes selected by post-double selection LASSO).

Panel A shows the aggregate effects of treatment concentration. Panel B shows the interactions of the village-level concentrations with treatment status. The results are strikingly different across the different outcomes. The estimates in Panel B, columns 1 and 2 - the outcomes for which we find the largest average treatment effects in Table 1 - suggest that the program led to differences of similar magnitude between treated farmers and control farmers on agronomic knowledge and self-reported adoption of practices across the three village categories. They also suggest no difference between control farmers' knowledge or self-reported adoption in different types of villages. This confirms the absence of information diffusion within villages that we find when looking at immediate neighbors.

In contrast, the estimates in columns 4-6 of Panel B paint a clear picture of *negative* spillover effects on the control group in 50% and 75% villages. While the average treatment effects reported in columns 4-6 of Table 1 are all positive and significant, columns 4-6 of Table 3 show no positive treatment effects in 25% concentration villages in the top row, lower control group outcomes in both 50% and 75% villages (rows 2 and 3) and positive coefficients on the interaction terms in rows 4 and 5.

The results in columns 4-6 suggest that our treatment effects on input use and yields in Table 1 are driven by worse control group outcomes in higher treatment concentration villages. The treatment effects on knowledge and self-reported adoption, on the other hand, are stable across villages. This indicates that while treated farmers gained agronomic knowledge from the training, they did not implement enough of what they learned to significantly increase their coffee production. In contrast, control farmers in heavily treated villages were negatively impacted.

In principle, a reallocation of labor and fertilizer inputs away from the control group towards the trained farmers might be efficiency-enhancing if the training increased productivity, or if coffee production functions are (locally) convex in inputs. Testing this hypothesis through production function estimation is complicated by the multiple constraints on farmers' input choices (see below) that violate the assumptions of commonly used estimation techniques (Shenoy 2021). To gain some insights, we estimate output as a flexible polynomial of inputs, with heterogeneous parameters across trained and control farmers. Consistent with Table 3, we find no evidence of a productivity increase from training, and returns to labor and fertilizer appear to be concave over the range observed in the data for all types of farmers. (Appendix B contains detailed results.) Thus there is little to suggest that the observed shift of inputs from control to treatment farmers increased aggregate efficiency. Indeed, in Table 3, Panel A, the aggregate impact of treating 50% of farmers or 75% of farmers on yields in column 6 is negative and insignificant (despite the larger number of farmers benefiting from the training).

Taken together, the results in Tables 2 and 3 suggest that the method of training some farmers and hoping the knowledge will spread through their networks does not work in our setting, and that it can even backfire since the control group appears to have been negatively impacted in villages with high treatment intensity.

Why might higher treatment densities lead to lower input use and worse yields for control farmers? One possibility is that greater demand for inputs from treatment farmers raised prices or reduced availability of fertilizer and hired labor. Appendix Table A11 shows that neither the daily wage rates nor NPK prices are significantly higher in 50% and 75% villages, although the point estimates are all positive. However, there is considerable evidence that neither market functions very smoothly. First, agricultural labor markets often do not clear in rural areas in developing countries, so an increase in demand for labor may not have led to an increase in wages (Breza et al. 2021).

Second, the market for NPK fertilizer was undergoing some changes during the period of our study. At the time of our baseline, NPK was supplied by the government (through OCIR-CAFÉ, which became NAEB in May 2011) to coffee farmers. NPK was then ostensibly free, but the costs of NPK were deducted from the price farmers were paid (the government fixed the price of both NPK and coffee). In 2010-2011, the government gradually phased this out due to organizational capacity constraints (they were struggling to fix prices and supply NPK in large enough quantities) and transitioned to the system still in place today, where the coffee washing stations buy the cherries, but also help farmers with credit and provision of inputs. NPK is thus not free any more, through it is typically provided on credit. This system remains very far from perfect. The cooperatives still rely on government intermediaries to get fertilizer to the villages, while it is their duty to organize credit for the farmers. This change led to a huge drop in NPK use between 2011 and 2012 (from 298 reported users in 2011 across treatment and control farmers, to 126 users in 2012). In addition, since 2007, private agro-dealers (who were only supposed to sell fertilizer for other crops, not coffee) have also sold NPK to coffee farmers, but they only started to trade in our study area in 2012, at the very end of our study. This is because Nyarubaka was the poorest sector of the district until then, making it the least appealing market for agro-dealers.

This narrative illustrates both the specificity of our context, and some general features: fertilizer is heavy, it is imported in most of Africa, and the ports of entry are limited. This means that transport costs are extremely high, and poor remote places would probably be largely cut out of a purely private provision mechanism.[4] More broadly, governments are often in charge of part or all of the fertilizer distribution network, controlling both quantities and prices. In such a context, any program that extolls the value of fertilizer use is likely going to create congestion. Treatment farmers could have become somewhat more enthusiastic about fertilizer use, or they could have been first in line to access the loans that give them access, or to access fertilizer itself. In this third best environment, this ended up being inefficient.

D. Program Impacts on Social Networks

Another source of externalities could have been treatment effects on the social networks themselves. Treatment farmers could have made new friends during the training session, or found it more beneficial to interact with other trained framers, and could have stopped interacting with control farmers. This would have reduced their opportunity to learn, but also potentially risk sharing. Conversely, control farmers could have sought out treatment farmers. In this section we examine the impact of the training groups on farmers' reported social network links to other farmers with whom they discuss coffee. We analyze the total number of friends reported by each farmer, both across different geographical areas and for friends with different treatment statuses.

Table 4 displays the results. Panel A, reporting effects on total friends, shows that treated households gain 0.336 friends who were also chosen for training (a 28% increase over baseline), significantly more than their (insignificant) increase in control group friends. The control group themselves gain neither trained nor control friends overall. Both treated and control households

[4] This is what (Suri 2011) finds: heterogeneity in the availability of seed and fertilizer explains the heterogeneity in the returns to technology.

report fewer friends outside the group of farmers who signed up for the training ("non-sample friends").

If social networks change due to the training groups, the largest impact will be on friends from the same village - we examine this in panel B. Treated households indeed gain most friends within village (0.360), and in panel D we see that virtually all (0.319) of this increase in within-village trained friends can be attributed to additional friends from the same training group. The control group also increases trained friends within village relative to baseline by 0.0988 friends. However, the control group gains a similar number (0.0893) of *control* friends within village as well. Thus it appears that farmers who signed up for the training (both treated and control), a group likely more focused on coffee farming, strengthened intra-village links amongst themselves rather than specifically targeting trained individuals.

This strengthening of within-village links over time comes at the cost of decreasing friendships outside the village. Panel C shows that both treatment and control households dropped outside-village links, with the strongest effects being on contacts with non-sample households outside the village. This effect is statistically the same between treatment and control households, and independent of the fraction of households trained within a village (see appendix Table A12), indicating that it is likely due to a time trend rather than a program effect. Thus the social effect of the training seems to have been largely limited to the relationships formed during the actual training meetings themselves. Conversely, there is no evidence that the control group selectively formed new links with trained individuals to gain access to their increased coffee knowledge.

5 Conclusions

This paper studies whether knowledge about best agronomic practices acquired through training programs can spread through social networks. Our results from an RCT in Rwanda do not support this hypothesis. Although knowledge of best agronomic practices increased among trained farmers, we find no evidence that treatment farmers shared their new knowledge with control households they were socially connected to or lived close to. The fact that control households experienced *negative* spillovers in high treatment concentration areas, along with the null treatment effects we find on observed adoption, suggest that much of the 6.7% higher yields of treatment farmers compared to the control group that we observe at endline may be the result of these negative spillovers, rather than a net gain for the treated farmers.

One possible interpretation of these results is that the treatment group did not experience sufficiently high returns to the taught practices to induce them to encourage control farmers in their information networks to adopt these techniques (Magnan et al. 2015). Indeed, the null effects we find from the tree audits on treatment farmers' own adoption suggest they probably did not undertake enough of these new practices to see a difference.

Another possibility is that information provision alone was not sufficient to unleash the yield-boosting potential of these agronomic practices. There is now growing evidence that the low productivity observed in much of African agriculture is not the result of any one single constraint; rather, different combinations of constraints seem to bind for different farmers (Suri and Udry 2022). The negative externality on NPK use, which is indicative of a shortage of either fertilizer or credit at the village level, suggest that the most important input in the training was not easily accessible by the farmers. Existing studies set in Kenya have shown that intervention packages targeting multiple constraints (e.g. by combining training with financial support, input supply, and marketing assistance) can be effective at increasing adoption of new crops (Ashraf et al. 2009) or fertilizer and improved practices, ultimately increasing yield and profits (Deutschmann et al. 2019). Promising avenues for future research include asking how these multiple constraints interact with information frictions and the complexities of social learning.

What is clear is that in a context where inputs market are imperfect and not well integrated, the strategy of training some farmers and not others creates distortions that may outweigh any gain of the training. Furthermore, it is those distortions that may give the misleading impression that the program is effective. This evidence casts some doubt on a widely practiced strategy.

Table 1: Treatment Effects from Basic Within-Village Specification.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Input quantities index (labor + NPK)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	1.240	0.330	0.019	0.040	0.102	0.065
	[0.075]	[0.041]	[0.030]	[0.028]	[0.034]	[0.036]
	(0.000)	(0.000)	(0.557)	(0.179)	(0.002)	(0.060)
Control mean	-0.000	-0.000	0.000	-0.000	0.000	0.534
R-squared	0.11	0.04	0.02	0.01	0.04	n/a
Observations	4622	4622	47503	47503	6157	6090

Notes: Standard errors clustered by household are in brackets, and randomization inference p-values are in parentheses below those. All specifications control for village and round fixed effects. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix D). For columns (1)-(4), we use data from rounds 6, 8 and 9 only as we did not collect best practices data in round 7. In column (1), we use self-reported knowledge data. In column (2), we use self-reported adoption data. In column (3)-(4), we use tree level audit data. Column (5) is an input quantities index, which is the average of three variables (each standardized by its Control group mean and SD): paid labor days per tree, HH labor days per tree, and KGs of NPK applied per tree. We trim the top 1% of values of this index. In columns (1)-(5), for ease of interpretation, we also normalize each index by its Control group mean and SD. In column (6), we apply a Poisson regression to the outcome, trimming the top 1% of values.

Table 2: Diffusion in the Control Group via Baseline Treatment Friends.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Input quantities index (labor + NPK)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
Number of treatment friends	0.009	0.003	-0.050	0.002	-0.015	-0.013
	[0.027]	[0.026]	[0.021]	[0.019]	[0.026]	[0.029]
Number of sample friends	0.009	0.043	0.040	-0.011	0.002	0.040
	[0.015]	[0.015]	[0.014]	[0.012]	[0.017]	[0.017]
Outcome mean	-0.000	-0.000	0.000	-0.000	0.000	0.534
Mean T friends	1.596	1.596	1.596	1.596	1.596	1.596
Mean tot. friends	2.925	2.925	2.925	2.925	2.925	2.925
R-squared	0.02	0.04	0.03	0.01	0.04	n/a
Observations	2143	2143	21530	21530	2853	2819

Notes: Standard errors clustered by household in brackets. All specifications control for village and round fixed effects. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix D). For columns (1)-(4), we use data from rounds 6, 8 and 9 only as we did not collect best practices data in round 7. In column (1), we use self-reported knowledge data. In column (2), we use self-reported adoption data. In column (3)-(4), we use tree level audit data. Column (5) is an input quantities index, which is the average of three variables (each standardized by its Control group mean and SD): paid labor days per tree, HH labor days per tree, and KGs of NPK applied per tree. We trim the top 1% of values of this index. In columns (1)-(5), for ease of interpretation, we also normalize each index by its Control group mean and SD. In column (6), we apply a Poisson regression to the outcome, trimming the top 1% of values.

Table 3: Treatment Effects interacted with Village Treatment Concentration.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Input quantities index (labor + NPK)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A					
50% T in village	0.194	0.004	0.006	0.059	-0.016	-0.121
	[0.098]	[0.093]	[0.051]	[0.040]	[0.094]	[0.102]
	(0.112)	(0.971)	(0.922)	(0.202)	(0.884)	(0.375)
75% T in village	0.696	0.167	-0.045	-0.015	-0.109	-0.156
	[0.144]	[0.111]	[0.063]	[0.040]	[0.102]	[0.136]
	(0.001)	(0.236)	(0.545)	(0.776)	(0.360)	(0.421)
Control mean, 25% villages	0.056	0.038	0.038	0.006	0.064	0.607
p-value: 50% T= 75% T	0.004	0.055	0.421	0.019	0.312	0.807
R-squared	0.02	0.01	0.00	0.00	0.00	n/a
Observations	4622	4622	47503	47503	6157	6090
	Panel B					
Treatment	1.140	0.301	-0.081	-0.037	0.041	-0.084
	[0.245]	[0.089]	[0.094]	[0.026]	[0.094]	[0.052]
	(0.004)	(0.017)	(0.496)	(0.270)	(0.692)	(0.160)
50% T in village	-0.081	-0.073	-0.014	0.030	-0.057	-0.186
	[0.046]	[0.100]	[0.062]	[0.046]	[0.094]	[0.113]
	(0.156)	(0.533)	(0.826)	(0.577)	(0.604)	(0.191)
75% T in village	-0.109	-0.038	-0.135	-0.082	-0.202	-0.294
	[0.065]	[0.121]	[0.093]	[0.046]	[0.090]	[0.151]
	(0.160)	(0.791)	(0.242)	(0.149)	(0.072)	(0.110)
Treatment X 50% T in village	-0.006	0.007	0.081	0.078	0.061	0.171
	[0.280]	[0.129]	[0.104]	[0.048]	[0.101]	[0.084]
	(0.982)	(0.957)	(0.480)	(0.181)	(0.585)	(0.082)
Treatment X 75% T in village	0.315	0.072	0.174	0.114	0.097	0.238
	[0.286]	[0.100]	[0.102]	[0.036]	[0.118]	[0.089]
	(0.337)	(0.490)	(0.144)	(0.018)	(0.443)	(0.033)
Control mean, 25% villages	0.056	0.038	0.038	0.006	0.064	0.607
p-value: 50% T= 75% T	0.702	0.733	0.153	0.007	0.081	0.513
p-value: Treat x 50% T = Treat x 75% T	0.121	0.538	0.137	0.462	0.655	0.493
p: Treatment + 50% T + Treatment x 50% T=0	0.000	0.037	0.848	0.157	0.637	0.327
p: Treatment + 75% T + Treatment x 75% T=0	0.000	0.005	0.553	0.917	0.558	0.298
R-squared	0.09	0.02	0.00	0.00	0.00	n/a
Observations	4622	4622	47503	47503	6157	6090

Notes: Standard errors clustered by household are in brackets, and randomization inference p-values are in parentheses below these. All specifications control for round fixed effects. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix D). For columns (1)-(4), we use data from rounds 6, 8 and 9 only as we did not collect best practices data in round 7. In column (1), we use self-reported knowledge data. In column (2), we use self-reported adoption data. In column (3)-(4), we use tree level audit data. Column (5) is an input quantities index, which is the average of three variables (each standardized by its Control group mean and SD): paid labor days per tree, HH labor days per tree, and KGs of NPK applied per tree. We trim the top 1% of values of this index. In columns (1)-(5), for ease of interpretation, we also normalize each index by its Control group mean and SD. In column (6), we apply a Poisson regression to the outcome, trimming the top 1% of values.

Table 4: Effect on Social Networks.

	(1) Trained Friends	(2) Control Friends	(3) Non-sample Friends	(4) All Friends	(5) T-C Friend Difference
	Panel A: Friends in All Villages				
Treat × Post	0.336	0.0298	-0.164	0.201	0.306
	[0.0551]	[0.0427]	[0.0458]	[0.0904]	[0.0671]
Control × Post	0.0338	0.0217	-0.150	-0.0947	0.0122
	[0.0499]	[0.0488]	[0.0524]	[0.0937]	[0.0676]
T-C diff. p-value	0.000	0.900	0.837	0.023	0.002
Baseline mean	1.196	1.004	1.567	3.766	0.192
	Panel B: Friends in Own Village				
Treat × Post	0.360	0.0870	-0.00119	0.446	0.273
	[0.0513]	[0.0388]	[0.0407]	[0.0832]	[0.0621]
Control × Post	0.0988	0.0893	0.0392	0.227	0.00947
	[0.0429]	[0.0456]	[0.0416]	[0.0814]	[0.0599]
T-C diff. p-value	0.000	0.969	0.488	0.061	0.002
Baseline mean	1.027	0.831	1.063	2.921	0.195
	Panel C: Friends Outside Own Village				
Treat × Post	-0.0238	-0.0572	-0.163	-0.244	0.0334
	[0.0167]	[0.0165]	[0.0260]	[0.0361]	[0.0224]
Control × Post	-0.0650	-0.0677	-0.189	-0.322	0.00271
	[0.0220]	[0.0164]	[0.0305]	[0.0439]	[0.0247]
T-C diff. p-value	0.137	0.653	0.514	0.171	0.358
Baseline mean	0.169	0.172	0.503	0.845	-0.003
	Panel D: Friends in Same Training Group				
Treat × Post	0.319				
	[0.0438]				
Baseline mean	0.412				

Notes: Standard errors clustered by household in brackets. All specifications control for village fixed effects and for whether the household was selected for training (Treatment/Control status). All columns use household level data from the baseline and round 9 social network surveys. (see Appendix D). Outcome variable in column (1) is the count of HH's friends selected for training. Outcome variable in column (2) is the count of friends who applied for training but were not selected. Outcome variable in column (3) is the count of friends who did not apply for training. Column (4) outcome variable is the sum of (1)+(2)+(3). Column (5) outcome variable is (1)-(2).

References

Ashraf, N., Giné, X., and Karlan, D. (2009). "Finding Missing Markets (and a Disturbing Epilogue): Evidence from an Export Crop Adoption and Marketing Intervention in Kenya". *American Journal of Agricultural Economics*, 91(4):973–990.

Banerjee, A. V., Chandrasekhar, A. G., Duflo, E., and Jackson, M. O. (2021). Changes in Social Network Structure in Response to Exposure to Formal Credit Markets. *NBER Working Paper No. 28365.*

Belloni, A., Chernozhukov, V., and Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies*, 81(2 (287)):608–650.

Breza, E., Kaur, S., and Shamdasani, Y. (2021). Labor rationing. *American Economic Review*, 111(10):3184–3224.

Bridle, Leah amd Magruder, J., McIntosh, C., and Suri, T. (2019). Experimental Insights on the Constraints to Agricultural Technology Adoption.

Caldwell, R., Lambert, R., Magruder, J., McIntosh, C., and Suri, T. (2019). Improving agricultural extension and information services in the developing world. *VoxDev*.

Chandrasekhar, A. G., Duflo, E., Kremer, M., F. Pugliese, J., Robinson, J., and Schilbach, F. (2022). Blue spoons: Sparking communication about appropriate technology use. Working Paper 30423, National Bureau of Economic Research.

Cole, S. and Fernando, A. (2012). The Value of Advice: Evidence from Mobile Phone-Based Agricultural Extension. *SSRN Electronic Journal*.

Comola, M. and Prina, S. (2021). Treatment Effect Accounting for Network Changes. Forthcoming in *Review of Economics and Statistics*.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P. (2013). Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment. *The Quarterly Journal of Economics*, 128(2):531–580.

Deutschmann, J. W., Duru, M., Siegal, K., and Tjernström, E. (2019). Can Smallholder Extension Transform African Agriculture? Working Paper 26054, National Bureau of Economic Research.

FAOSTAT (2022). Technical report, Food and Agriculture Organization of the United Nations. FAOSTAT Statistical Database. [Rome]. Accessed September 22, 2022.

Magnan, N., Spielman, D. J., Lybbert, T. J., and Gulati, K. (2015). Leveling with friends: Social networks and indian farmers' demand for a technology with heterogeneous benefits. *Journal of Development Economics*, 116:223–251.

NISR (2019). National Institute of Statistics of Rwanda, Statistical Yearbook 2019 Edition.

OCIR-Café (2008). Rwanda Industry Review.

ROR (2009). Strategic Plan for the Transformation of Agriculture in Rwanda – Phase II (PSTA II). Final report, Ministry of Agriculture and Animal Resources, Republic of Rwanda.

Shenoy, A. (2021). Estimating the production function under input market frictions. *The Review of Economics and Statistics*, 103(4):666–679.

Suri, T. (2011). Selection and comparative advantage in technology adoption. *Econometrica*, 79(1):159–209.

Suri, T. and Udry, C. (2022). Agricultural Technology in Africa. *Journal of Economic Perspectives*, 36(1):33–56.

World Bank (2007). World Development Report 2008: Agriculture for Development. *World Bank, Washington DC.*

A Additional Tables and Figures

Table A1: Balance Checks.

	Control Mean	Treatment Coeff.	Std Error	P-value
Head, Years of Schooling	3.534	-.061	.146	.678
Female Headed Household	.33	.025	.013	.075
Household Size	5.018	-.037	.141	.794
Average Schooling of Household	3.257	.103	.09	.264
Yield, total KGs per tree	.769	.017	.04	.682
Total Trees	249.892	1.294	11.616	.912
Fraction Unproductive Trees	.312	-.012	.014	.398
Cut Stems	.101	-.009	.014	.546
Book Keeping Done	.025	-.007	.007	.356
Removed Dead Branches	.747	-.043	.018	.023
Removed Suckers	.903	-.018	.011	.131
Removed Weeds	.99	-.005	.004	.256
Applied Compost	.725	.017	.024	.485
Applied NPK	.185	0	.015	.997
Applied Lime	.024	.01	.01	.336
Applied Pesticides	.761	-.022	.019	.264
Applied Mulch	.867	-.022	.016	.179
p-value of joint F-test				0.9998

Notes: All specifications control for village fixed effects. Robust standard errors.

Table A2: Attendance Rates, by Training Session.

Integrated Pest Management	78.0%	[0.109]
Nutrition	77.1%	[0.109]
Harvesting	77.6%	[0.112]
Weeding	72.5%	[0.122]
Mulching	75.7%	[0.093]
Pruning and rejuvenation	74.1%	[0.115]
Pesticide use	74.2%	[0.111]
Composting	71.3%	[0.120]
Erosion Control	73.4%	[0.107]
Shade	70.4%	[0.127]
Nutrition Review	80.2%	[0.110]
Harvesting Nutrition Review	74.6%	[0.112]
Sustainability	78.5%	[0.093]
Composting Review	79.7%	[0.090]
Pruning and Rejuvenation Review	78.2%	[0.096]

Notes: Session-specific average attendance rates and standard deviations across the 38 training groups (spread across 27 villages) that comprise our treatment. Sessions listed in chronological order.

Table A3: Attrition Rates by Treatment Status.

	Control	Treatment	Difference
	Mean [s.d.]	Mean [s.d.]	Coeff. [s.e.]
Survey round	(1)	(2)	(3)
6	0.034	0.034	0.003
	[0.181]	[0.181]	[0.009]
7	0.041	0.036	–0.005
	[0.197]	[0.187]	[0.011]
8	0.035	0.034	0.000
	[0.184]	[0.181]	[0.010]
9	0.034	0.034	0.003
	[0.181]	[0.181]	[0.009]

Notes: Column 1 presents the attrition rate for Control households by endline survey round, Column 2 for Treatment households. Column 3 reports the coefficient from a regression of attrition on a treatment dummy, but also includes village fixed effects. Standard errors clustered at the village level in brackets.

Table A4: Treatment Effects from Within-Village Specification:
Adoption Index and Leaf Health Index Components (Audits Data).

	Dripline is weeded	Tree canopy has mulch	Removed dead branches	Removed branches touching the ground	Opened centers	Removed unwanted suckers	Removed old and dry berries	Tree bark is smoothed	Few signs of leaf rust	Few curled leaves	Few yellow leaves
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Treatment	0.032	0.028	0.001	-0.004	-0.005	-0.000	-0.019	0.004	0.020	0.006	0.008
	[0.015]	[0.014]	[0.013]	[0.007]	[0.013]	[0.013]	[0.013]	[0.006]	[0.012]	[0.010]	[0.009]
	(0.040)	(0.055)	(0.932)	(0.552)	(0.718)	(0.992)	(0.121)	(0.543)	(0.081)	(0.607)	(0.354)
Control mean	0.549	0.583	0.331	0.911	0.434	0.437	0.408	0.041	0.273	0.194	0.152
R-squared	0.060	0.050	0.040	0.040	0.040	0.050	0.010	0.020	0.090	0.040	0.040
Observations	47483	47338	47502	47501	47501	47502	47503	47498	47481	47488	47478

Notes: Standard errors clustered by household in brackets and randomization inference p-values are in parentheses below those. All specifications control for village and round fixed effects. All columns use data from endline rounds 6, 8, 9 at the household-plot-tree-round level. Columns 1-8 constitute the components of the *All Audits* index in column 3 of Tables 1-3. Columns 9-11 are the components of the *Leaf Health* index in column 4 of Tables 1-3.

Table A5: Treatment Effects from Within-Village Specification: Inputs Index Components.

	Household labor days per tree	Paid labor days per tree	Quantity of NPK applied per tree (KGs)
	(1)	(2)	(3)
Treatment	0.035	0.137	0.185
	[0.042]	[0.116]	[0.170]
	(0.439)	(0.254)	(0.329)
Control mean	0.132	0.007	0.004
R-squared	0.000	0.000	0.000
Observations	6157	6157	6157

Notes: Standard errors clustered by household in brackets and randomization inference p-values are in parentheses below those. All specifications are Poisson regressions and control for village and round fixed effects. All columns use data from endline rounds 6, 8, 7, 9 at the household-round level. Columns (1), (2) and (3) are the components of the *Inputs Index* outcome in column (5) of Tables 1-3.

Table A6: Treatment Effects from Within-Village Specification, with Post-Double LASSO Selected Controls.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Input quantities index (labor + NPK)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	1.240	0.328	0.019	0.040	0.103	0.068
	[0.100]	[0.049]	[0.034]	[0.022]	[0.038]	[0.038]
	(0.000)	(0.000)	(0.555)	(0.084)	(0.023)	(0.095)
Control mean	-0.000	-0.000	0.000	-0.000	0.000	1.000
Prob > chi2	0.000	0.000	0.941	0.264	0.083	0.000
Number of controls selected	0	4	0	0	1	4
Observations	4622	4622	47503	47503	6157	6090

Notes: Standard errors clustered by household in brackets and randomization inference p-values are in parentheses below those. All specifications control for village and round fixed effects. Same specification as in Table 1 but using post-double selection LASSO (Belloni et al. 2014) to select controls from the list of baseline outcomes in Appendix Table A1. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix D). For columns (1)-(4), we use data from rounds 6, 8 and 9 only as we did not collect best practices data in round 7. In column (1), we use self-reported knowledge data. In column (2), we use self-reported adoption data. In column (3)-(4), we use tree level audit data. Column (5) is an input quantities index, which is the average of three variables (each standardized by its Control group mean and SD): paid labor days per tree, HH labor days per tree, and KGs of NPK applied per tree. We trim the top 1% of values of this index. In columns (1)-(5), for ease of interpretation, we also normalize each index by its Control group mean and SD. In column (6), we apply a Poisson regression to the outcome, trimming the top 1% of values.

Table A7: Learning Spillovers on Control Farmers.

	Learned something new about [practice] from a Treatment farmer									
	Weeding	Fertilizing (manure)	Fertilizing (NPK)	Mulching	Integrated Pest Management	Removal of dead branches	Removal of unwanted suckers	Removal of branches touching the ground	Opening of Centers	Removal of old/dry berries
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treatment friends	0.000	-0.021	-0.014	0.014	-0.003	-0.017	0.013	-0.020	0.002	0.003
	[0.011]	[0.016]	[0.016]	[0.011]	[0.014]	[0.016]	[0.013]	[0.016]	[0.012]	[0.011]
Sample friends	0.003	0.011	0.008	-0.008	0.004	0.007	-0.008	0.009	-0.005	-0.004
	[0.006]	[0.010]	[0.010]	[0.006]	[0.009]	[0.010]	[0.008]	[0.010]	[0.007]	[0.007]
Outcome mean	0.052	0.143	0.140	0.104	0.092	0.099	0.092	0.108	0.088	0.078
R-squared	0.040	0.030	0.060	0.050	0.040	0.040	0.050	0.050	0.040	0.030
Observations	694	694	694	694	694	694	694	694	694	693

Notes: Standard errors clustered by household in brackets and randomization inference p-values are in parentheses below those. All specifications control for total number of friends in sample at baseline and village FE. Control group only. The outcome is constructed from a module collected in the final endline survey (round 9) asking farmers to reflect on how much they have learned about each practice since baseline.

Table A8: Diffusion to Control Group via Baseline Treatment Friends, with Post-Double LASSO Selected Controls.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Input quantities index (labor + NPK)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
Number of treatment friends	-0.003	-0.012	-0.043	0.036	0.004	-0.019
	[0.042]	[0.021]	[0.023]	[0.017]	[0.012]	[0.018]
Number of sample friends	-0.000	0.024	0.023	-0.015	0.008	0.043
	[0.031]	[0.015]	[0.014]	[0.012]	[0.012]	[0.013]
Outcome mean	0.683	0.179	0.001	0.013	0.027	0.537
Mean T friends	1.596	1.596	1.915	1.915	1.597	1.601
Mean tot. friends	2.927	2.927	3.354	3.354	2.928	2.934
Prob > chi2	0.000	0.000	0.000	0.000	0.000	0.000
Number of controls selected	7	8	6	6	7	9
Observations	4622	4622	47503	47503	6157	6090

Notes: Standard errors clustered by household in brackets. All specifications control for village and round fixed effects. Same specification as in Table 2 but using post-double selection LASSO (Belloni et al. 2014) to select controls from the list of baseline outcomes in Appendix Table A1. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix D). For columns (1)-(4), we use data from rounds 6, 8 and 9 only as we did not collect best practices data in round 7. In column (1), we use self-reported knowledge data. In column (2), we use self-reported adoption data. In column (3)-(4), we use tree level audit data. Column (5) is an input quantities index, which is the average of three variables (each standardized by its Control group mean and SD): paid labor days per tree, HH labor days per tree, and KGs of NPK applied per tree. We trim the top 1% of values of this index. In columns (1)-(5), for ease of interpretation, we also normalize each index by its Control group mean and SD. In column (6), we apply a Poisson regression to the outcome, trimming the top 1% of values.

Table A9: Diffusion through Geographic Networks: Neighbors.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Input quantities index (labor + NPK)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Treatment household neighbors					
Num. treatment HH neighbors	0.028	0.027	-0.047	-0.009	-0.022	-0.037
	[0.025]	[0.031]	[0.023]	[0.020]	[0.027]	[0.030]
Num. sample HH neighbors	-0.022	-0.005	0.022	-0.002	0.014	0.016
	[0.015]	[0.017]	[0.015]	[0.013]	[0.017]	[0.019]
Outcome mean	-0.004	-0.002	0.000	0.000	0.003	0.536
R-squared	.02	.03	.02	.01	.04	n/a
Observations	2122	2122	21345	21345	2826	2792
	Panel B: Treatment plot neighbors					
Num. treatment plot neighbors	0.030	0.042	-0.011	0.006	-0.004	-0.017
	[0.019]	[0.018]	[0.014]	[0.012]	[0.017]	[0.020]
Num. sample plot neighbors	-0.013	-0.016	0.004	-0.007	-0.001	0.005
	[0.012]	[0.012]	[0.008]	[0.008]	[0.010]	[0.013]
Outcome mean	-0.001	0.001	-0.001	-0.001	0.002	0.535
R-squared	.02	.04	.02	.01	.04	n/a
Observations	2131	2131	21405	21405	2838	2804

Notes: Standard errors clustered by household in brackets. All specifications control for village and round fixed effects. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix D). For columns (1)-(4), we use data from rounds 6, 8 and 9 only as we did not collect best practices data in round 7. In column (1), we use self-reported knowledge data. In column (2), we use self-reported adoption data. In column (3)-(4), we use tree level audit data. Column (5) is an input quantities index, which is the average of three variables (each standardized by its Control group mean and SD): paid labor days per tree, HH labor days per tree, and KGs of NPK applied per tree. We trim the top 1% of values of this index. In columns (1)-(5), for ease of interpretation, we also normalize each index by its Control group mean and SD. In column (6), we apply a Poisson regression to the outcome, trimming the top 1% of values.

Table A10: Treatment Effects Interacted with Village Treatment Concentration, with Post-Double LASSO Selected Controls.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Input quantities index (labor + NPK)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A			
50% T in village	0.194	0.036	0.006	0.059	-0.015	-0.100
	[0.098]	[0.086]	[0.051]	[0.040]	[0.084]	[0.092]
	(0.106)	(0.760)	(0.921)	(0.193)	(0.880)	(0.405)
75% T in village	0.696	0.166	-0.045	-0.015	-0.077	-0.088
	[0.144]	[0.102]	[0.063]	[0.040]	[0.085]	[0.111]
	(0.000)	(0.205)	(0.550)	(0.775)	(0.451)	(0.553)
Control mean, 25% villages	0.056	0.038	0.038	0.006	0.064	0.607
Prob > chi2	0.000	0.029	0.947	0.037	0.395	0.000
Number of controls selected	0	4	0	0	1	4
Observations	4622	4622	47503	47503	6157	6090
			Panel B			
Treatment	1.140	0.287	-0.089	-0.065	0.065	-0.067
	[0.245]	[0.079]	[0.097]	[0.037]	[0.091]	[0.039]
	(0.002)	(0.021)	(0.471)	(0.159)	(0.539)	(0.160)
50% T in village	-0.081	-0.040	-0.018	0.017	-0.048	-0.159
	[0.046]	[0.100]	[0.062]	[0.046]	[0.084]	[0.097]
	(0.139)	(0.736)	(0.814)	(0.705)	(0.607)	(0.204)
75% T in village	-0.109	-0.054	-0.137	-0.088	-0.157	-0.220
	[0.065]	[0.115]	[0.093]	[0.044]	[0.084]	[0.119]
	(0.142)	(0.699)	(0.254)	(0.106)	(0.104)	(0.139)
Treatment X 50% T in village	-0.006	0.013	0.093	0.119	0.035	0.150
	[0.280]	[0.122]	[0.110]	[0.060]	[0.099]	[0.074]
	(0.985)	(0.924)	(0.473)	(0.098)	(0.752)	(0.087)
Treatment X 75% T in village	0.315	0.102	0.182	0.141	0.063	0.218
	[0.286]	[0.093]	[0.106]	[0.034]	[0.123]	[0.076]
	(0.317)	(0.337)	(0.142)	(0.006)	(0.622)	(0.017)
Control mean, 25% villages	0.056	0.038	0.038	0.006	0.064	0.607
Prob > chi2	0.000	0.000	0.219	0.000	0.036	0.000
Number of controls selected	0	4	1	1	1	5
Observations	4622	4622	47503	47503	6157	6090

Notes: Standard errors clustered by household in brackets. All specifications control for round fixed effects. Same specification as in Table 3 but using post-double selection LASSO (Belloni et al. 2014) to select controls from the list of baseline outcomes in Appendix Table A1. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix D). For columns (1)-(4), we use data from rounds 6, 8 and 9 only as we did not collect best practices data in round 7. In column (1), we use self-reported knowledge data. In column (2), we use self-reported adoption data. In column (3)-(4), we use tree level audit data. Column (5) is an input quantities index, which is the average of three variables (each standardized by its Control group mean and SD): paid labor days per tree, HH labor days per tree, and KGs of NPK applied per tree. We trim the top 1% of values of this index. In columns (1)-(5), for ease of interpretation, we also normalize each index by its Control group mean and SD. In column (6), we apply a Poisson regression to the outcome, trimming the top 1% of values.

Table A11: Coffee Labor Wages and Fertilizer Prices, by Village Treatment Concentration.

	Log daily wage rate		Log NPK price per kg	
	(1)	(2)	(3)	(4)
50% Treatment	0.019	0.027	0.158	0.159
	[0.017]	[0.019]	[0.101]	[0.096]
75% Treatment	0.039	0.033	0.081	0.081
	[0.026]	[0.026]	[0.100]	[0.097]
Outcome mean, 25% villages	6.452	6.438	5.859	5.841
p-value: 50% share=75% share	0.409	0.811	0.237	0.325
R-squared	.18	.16	.01	.01
Observations	2522	2522	277	277

Notes: Standard errors clustered by village in brackets. All RCT-sample farmers who report hiring labor (columns 1-2) or purchasing NPK (columns 3-4) in endline rounds 6-9 are included. Prices are in local currency (RWF). Columns (2) and (4) are weighted by the inverse of the number of observations available per household. Columns (1) and (2): observations are at the household-round-task level. Columns (3) and (4): observations are at the household-round level.

Table A12: Changes in Social Network Friends by Village Treatment Intensity.

	(1) Different village friends	(2) Same village friends	(3) Treated same village friends
Treatment	-0.0366	0.205	0.0644
	[0.0685]	[0.100]	[0.0727]
Treatment × post	-0.221	0.222	0.121
	[0.108]	[0.239]	[0.117]
Control × post	-0.352	0.246	0.0255
	[0.0574]	[0.144]	[0.0579]
50% T × Treatment × post	-0.0827	0.163	0.228
	[0.118]	[0.295]	[0.146]
75% T × Treatment × post	0.0113	0.316	0.300
	[0.111]	[0.291]	[0.142]
50% T × Control × post	0.0583	0.0797	0.142
	[0.0670]	[0.261]	[0.105]
75% T × Control × post	0.0307	-0.249	0.0714
	[0.0690]	[0.266]	[0.167]
25% T village mean	0.857	3.399	0.651
R-squared	0.0590	0.0652	0.214
Observations	3156	3156	3156

Notes: Standard errors clustered by village in brackets. All specifications control for village fixed effects. All columns use household level data from the baseline and round 9 social network surveys (see Appendix D).

Figure A1: Negative Spillover Effects on the Control Group.

(a) Knowledge index

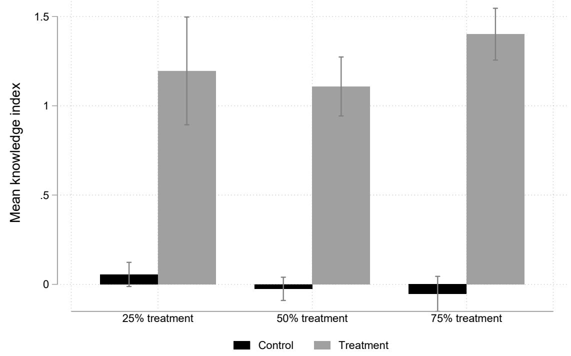

(b) Adoption index (self-reports)

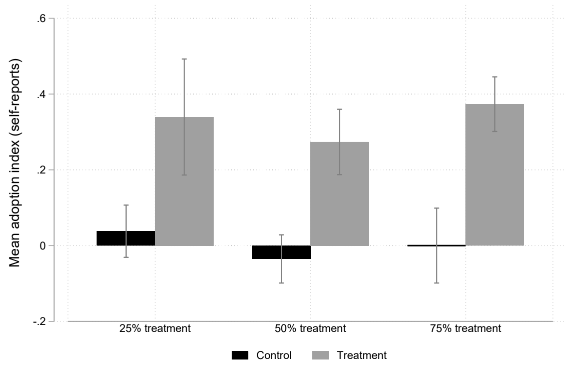

(c) Weeding (tree audits)

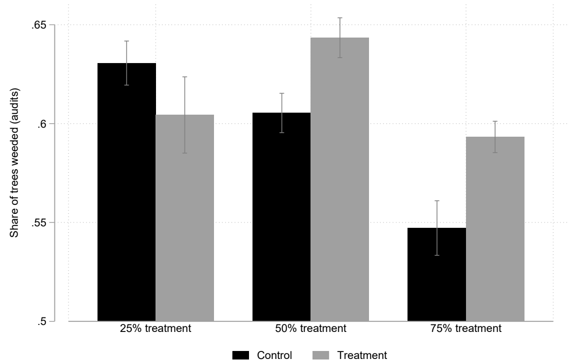

(d) Mulching (tree audits)

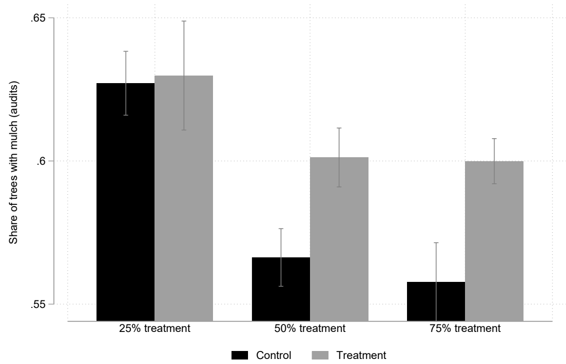

(e) Leaf Health index (tree audits)

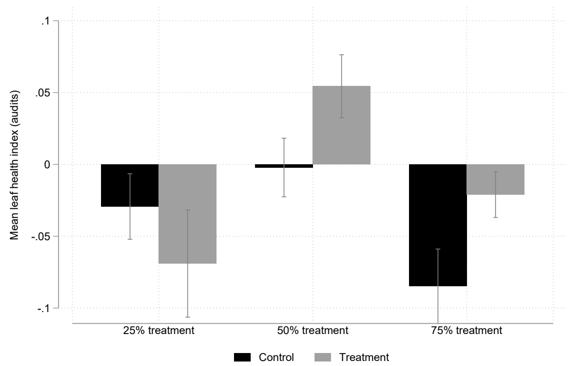

(f) Share of households using NPK

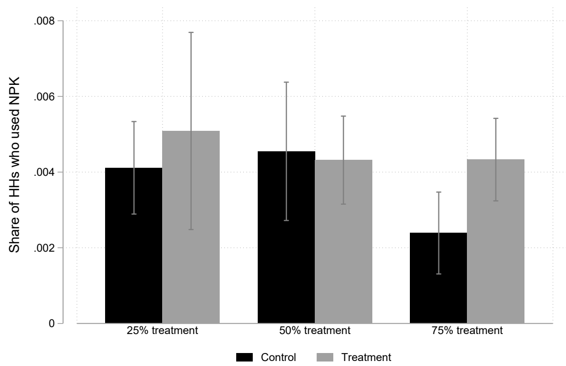

(g) Hired labor days per tree

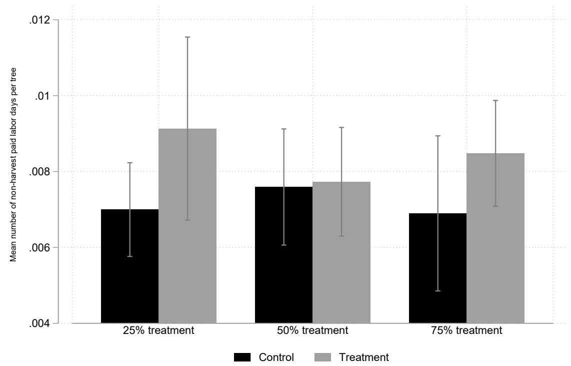

(h) Yields (KG per tree)

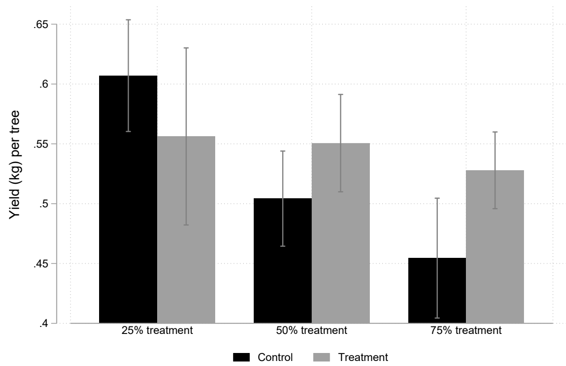

Notes: Outcome means by treatment status and village treatment concentration and 95% confidence intervals constructed from survey rounds 6-9 (pooled). To construct outcomes from tree audits data (weeding, mulching and leaf health index), households are weighted by the inverse of the number of coffee plots they operate in each survey round so as to give each household-round the same weight in the average.

B Agricultural Production Function Estimation

The effects of village-level treatment concentrations in Table 3 suggest a re-allocation of inputs from control to treatment farmers. The impact of these transfers on aggregate output and efficiency is a-priori ambiguous. If the treatment causes an increase in productivity among the trained, then such a re-allocation would increase efficiency by assigning more inputs to the farmers who can use them most productively. Similarly, if coffee production function has locally increasing returns to inputs, then concentrating those inputs in a selected group would also increase aggregate output. Conversely, if production functions are concave in inputs, and the training has no productivity effect, then concentrating more inputs among the trained may increase misallocation. To understand which scenario is most likely, we estimate the parameters of the coffee cherry production function. Since we have few priors about its functional form, we approximate the production function using a flexible polynomial in labor, capital, and number of trees. Since we are also agnostic about the ways in which training might have altered farmers' production functions, we allow its parameters to have two possible values: one for the control and pre-training treatment farmers, and another for the post-training farmers (indexed by t in the equation below):

$$\begin{aligned}
y_{it} = \alpha_t &+ \beta_{1t} Trees_{it} + \beta_{2t} Trees_{it}^2 + \beta_{3t} Labor_{it} + \beta_{4t} Labor_{it}^2 \\
&+ \beta_{5t} NPK_{it} + \beta_{6t} NPK_{it}^2 + \beta_{7t} Trees_{it} Labor_{it} + \beta_{8t} Trees_{it} NPK_{it} \\
&+ \beta_{9t} Labor_{it} NPK_{it} + \beta_{10t} Labor_{ij} NPK_{it} Trees_{it} + \epsilon_{it}
\end{aligned} \tag{3}$$

The estimation of this production function is complicated because the assumptions underlying many structural approaches to productivity estimation are unlikely to be satisfied. As Shenoy (2021) shows, when producers are subject to binding constraints on input use, the single-index assumption upon which control-function approaches to production function estimation rely is no longer satisfied. For Rwandan coffee farmers, who are subject to both financing constraints and limited availability of fertilizer, production functions assuming unlimited access to inputs are likely mis-specified. Furthermore, there is insufficient correlation in input use across seasons to employ a dynamic panel approach in which past input use would serve as an instrument for future input choices. We therefore estimate the production function via OLS.

Results of the production function estimation are displayed in Figure A2 with sub-figure A2a showing the association between NPK and output, and sub-figure A2b showing the association between labor and output. In both cases the relationship is concave, consistent with prior findings about the returns to fertilizer. The slope of the NPK production function appears somewhat steeper for low amounts of fertilizer; however an F-test of the null hypothesis that the slopes are equal for treated/untreated groups fails to reject (p=0.32). The magnitude of predicted output is very similar across treatment groups suggesting that training had little effect on overall productivity, consistent with the lack of evidence of significant adoption of new practices.

Taken together, the evidence from the coffee cherry production function does not support the hypothesis that the reallocation of inputs across farmers increased aggregate productivity.

Figure A2: Coffee Cherry Production Function

(a) Relationship between Coffee Cherry Output and NPK Use

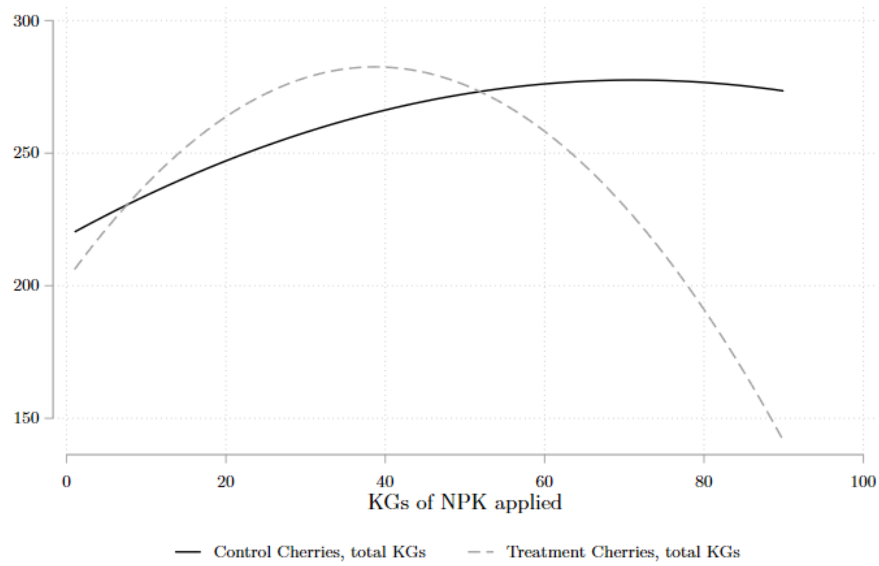

(b) Relationship between Coffee Cherry Output and Labor Use

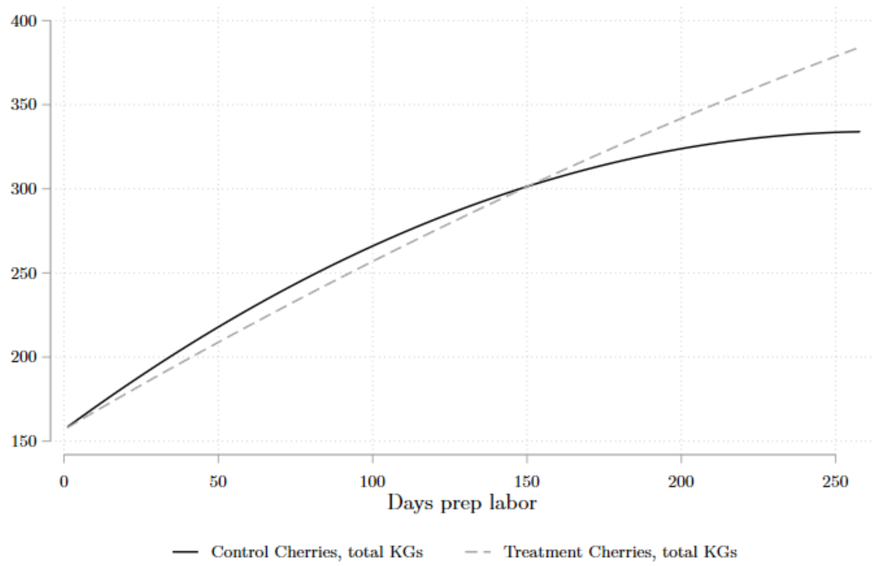

Figures show output of coffee cherries as a function of input quantity, using coefficients estimated from Equation 3. The relationship between output and each input quantity is plotted from 0 to the 99th percentile of that input intensity in the data, with other input quantities set to their mean values.

C TechnoServe's Agronomy Best Practices

The agronomy training program covered the following eight basic modules:

- Rejuvenation and pruning to produce new and productive wood. A multi-stem un-capped system was promoted.

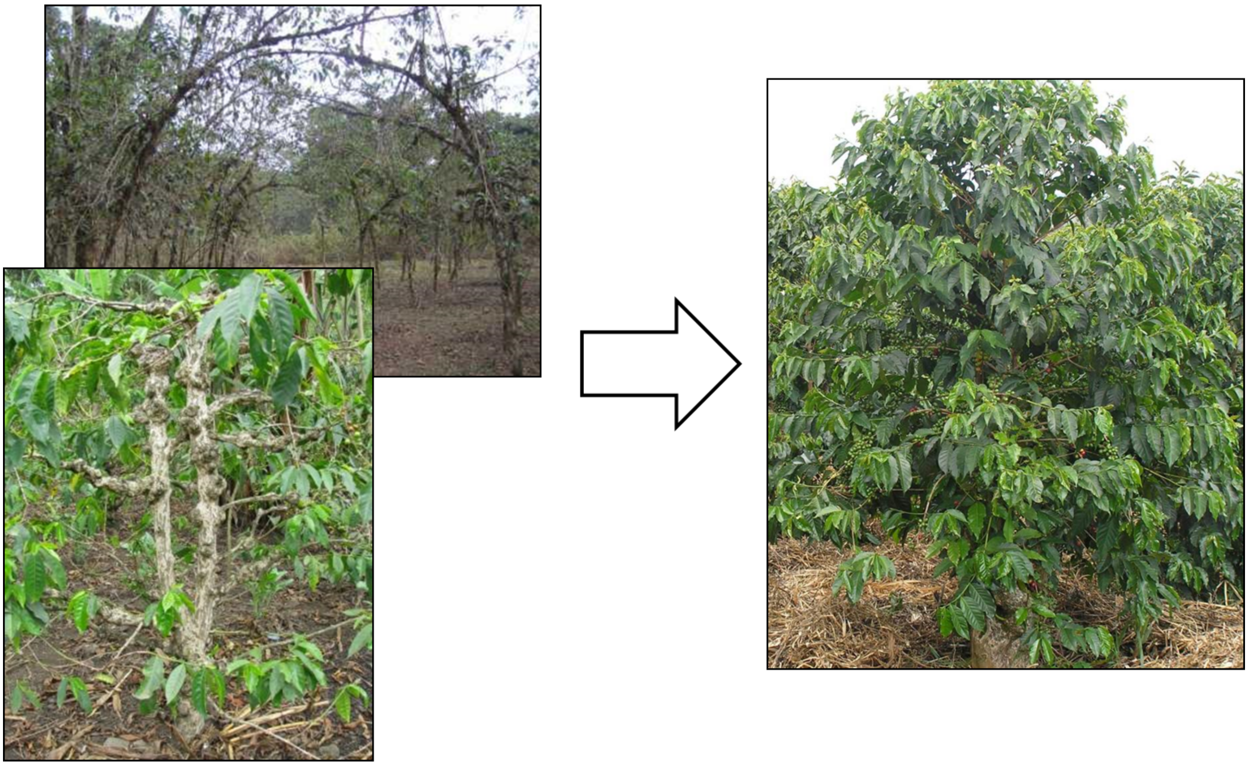

- Nutrition: a balanced nutritional program based on organic and inorganic additives, with the exact requirements determined by soil analysis. It included homemade compost, and may include liming, inorganic fertilizers and foliar feeds depending on the requirements of the area.

- Integrated pest management: multiple techniques to manage pests and diseases, such as correct nutrition, tree management, biological control, traps etc. Selective pesticides used as a last resort, but safe use of pesticides promoted.

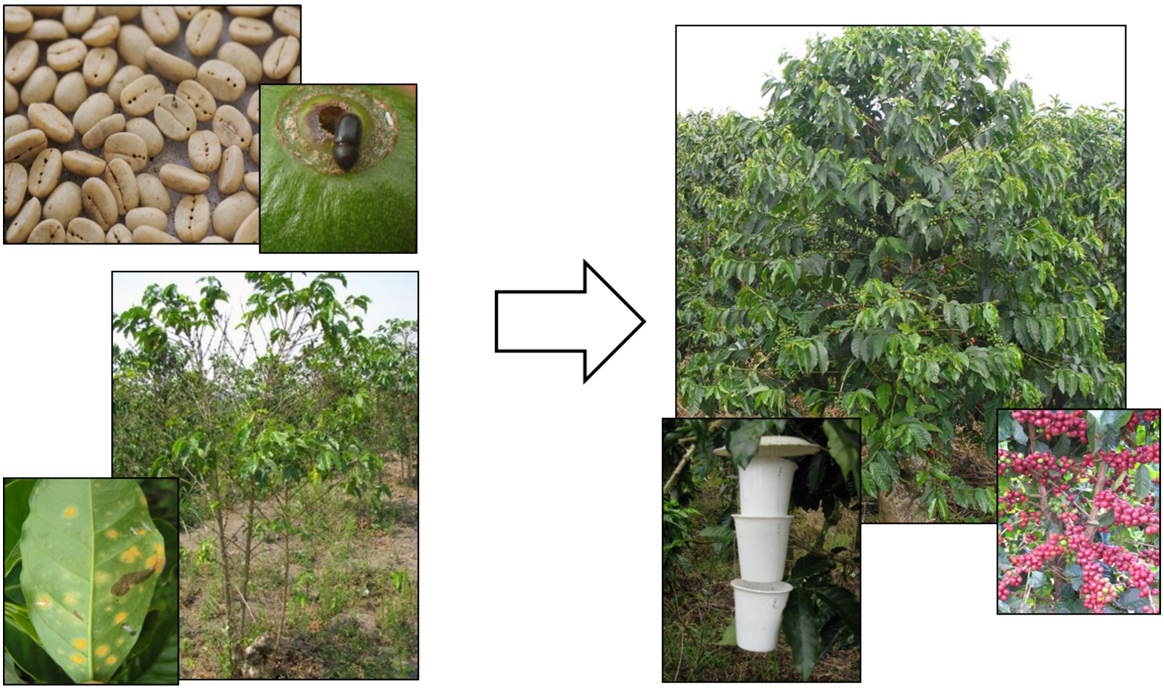

- Weed control: management of weeds through mulching, hand weeding and/or cover crops.

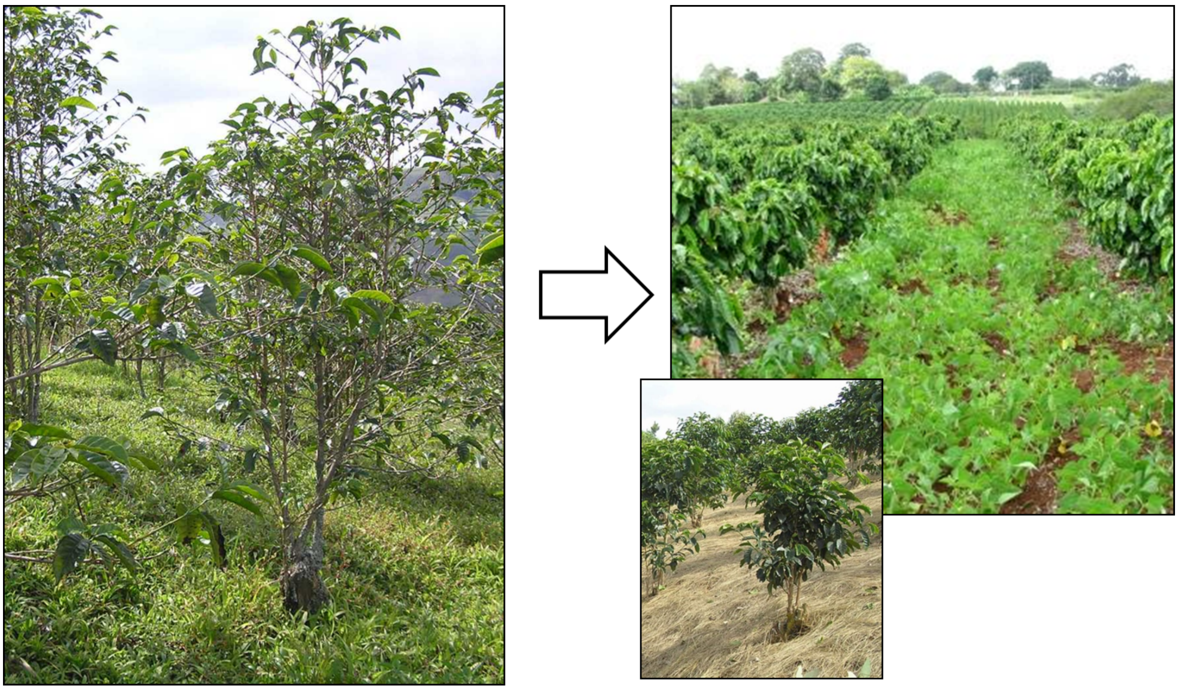

- Mulching: techniques to conserve moisture, add organic matter, and control soil erosion.

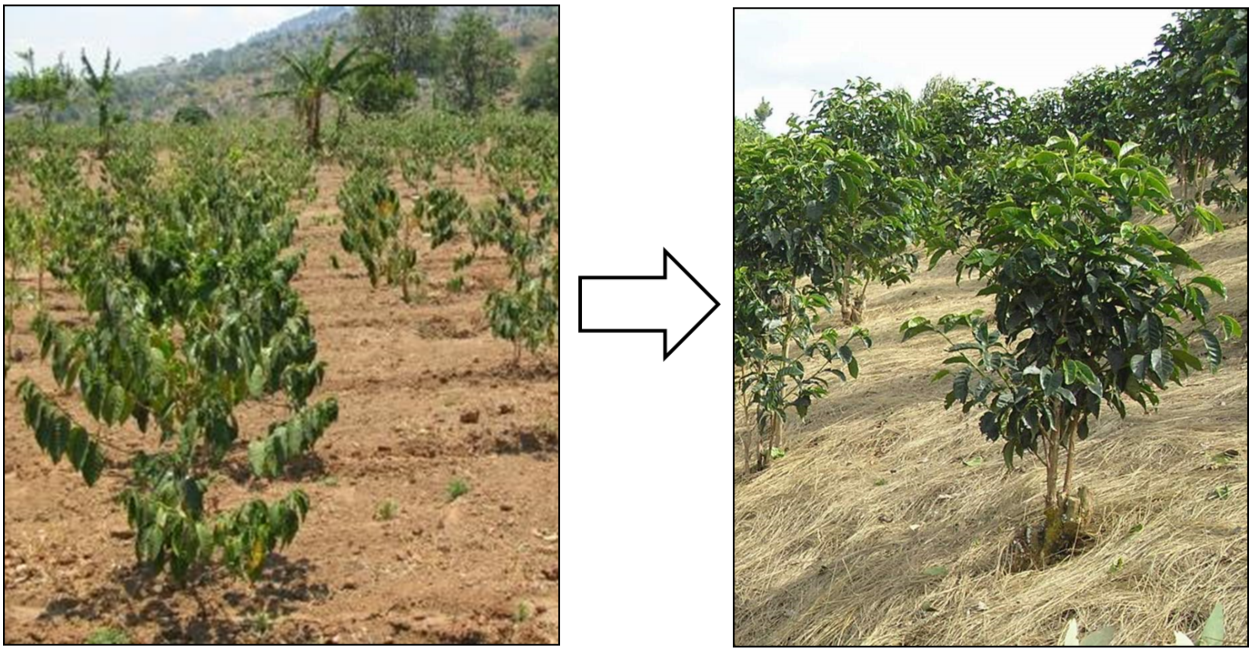

- Soil and water conservation: use of a number of techniques such as mulching, terracing and water traps to control soil erosion and maintain soil fertility. Encourage the management of water resources through conservation zones.

- Shade: use of the correct level (20-40%) of shade to reduce tree stress, conserve moisture, increase organic matter and increase biodiversity.

- Record keeping: maintenance of records of inputs, outputs, profit and loss in a record book.

The schedule of the modules covered in the training was as follows:

- February 2010: Record Keeping
- March 2010: Integrated Pest Management (IPM)
- April 2010: Coffee Nutrition
- May 2010: Coffee Harvesting
- June 2010: Weed control
- July 2010: Mulching
- August 2010: Pruning and Rejuvenation
- September 2010: Safe use of pesticides
- October 2010: Composting
- November 2010: Erosion control
- December 2010: Coffee Shade Management
- January 2011 to October 2011: Review

D Data Details

D1. Survey Data

From December 2009 to October 2012, we conducted ten rounds of surveys. These surveys mostly focused on the 1,594 farmers who were part of the experiment: the RCT-sample farmers. However, given the social networks focus of the study, we wanted to map the full social networks of the RCT-sample farmers. Therefore, alongside the baseline in December 2009, we also conducted a full census of the 5,198 farming households in all 29 villages of the the sub-district, including many who had not signed up for the study. Out of these, we focused on the 57% who had grown or harvested coffee in the year prior to the census, given the training program was targeted to coffee alone and it takes five years for coffee trees to grow once they have been planted. This meant that there were an additional 1,327 coffee farmers in the sub-district who did not register for the agronomy training program. Throughout, we refer to these farmers as the non-RCT-sample farmers, implying they grow coffee and live in the same sector as the RCT-sample farmers but are neither treatment nor control farmers for the agronomy program.

The data collection was split into modules that covered different aspects of the household's behavior. The modules covered household demographics, detailed plot level data for coffee as well as all other crops (including harvests, sales, labor and other inputs), coffee plot performance, coffee farming activities and practices, a consumption module, household finance and social networks for the household head and the spouse. In the social networks module, we asked both the household head and spouse who their friends were (with no limit on the number that could be listed). In addition, we asked which of these friends grew coffee and which they spoke to about coffee. Throughout the paper, we define friends as being "coffee friends", the friends that respondents in the sample report talking to about coffee.

Not every round of data collection covered the same modules and not every module in a given round covered all farmers. We collected fewer rounds of data for the farmers in the non-RCT-sample. Appendix D5 and D6 show a schedule of which modules were collected in which rounds, separately for the farmers in the RCT-sample and for those in the non-RCT-sample, as well as the timing of each survey wave. The first nine rounds of surveys took place every 2-3 months over the course of the program (recall that the training was run monthly between February 2010 and October 2011), and the tenth and final round took place in September-October 2012.

D2. Audit Data

One of our adoption measures was constructed from plot and tree inspections data, collected using plot and tree audits. Field staff visited each coffee plot of all the coffee farmers in the sub-district and inspected five trees, looking for signs of adoption of the agronomic practices covered in the training. The enumerators were given specific instructions on how to pick the five trees on each plot. They were instructed to start at the corner of the coffee plot closest to the farmer's house and walk in the direction of the opposite corner. They were then to inspect the second tree into the field, walk to the middle of the field and inspect the tree in the middle. Starting from the middle, the field staff was to walk towards the other two corners of the field and inspect the second tree in each direction. The field staff was then to walk back to the middle of the field and continue on the original path and inspect the second to last tree in the field. For each tree, the field staff would also note the GPS coordinates of each tree. Different variables were collected at different levels (the household level, the plot level and the tree level):

- Household level: We collected data on two practices that were also observed by the field staff, in particular
  1. whether the household kept record books
  2. whether the household has a compost heap
- Plot level: we collected data on three practices at the plot level, in particular
  1. whether the farmer had used any methods to control for soil erosion (such as using stabilizing grasses, water traps, etc.)
  2. whether there were any shade tress on the plot
  3. whether the farmer had grown other crops among the coffee
- Tree level: the audit data covered twelve different practices for each of the five trees per plot that were inspected. The practices were:
  1. whether the tree had any antestia (an insect)
  2. how many leaves were yellowing
  3. how many leaves were curling
  4. use of mulch
  5. evidence of weeding
  6. evidence of rejuvenation
  7. evidence of pruning
  8. evidence of integrated pest management
  9. whether the tree had any berry borers (an insect)
  10. evidence of damage from white borers
  11. evidence of scales or mealy bugs or mould
  12. signs of leaf rust

D3. Weigh Scale Data

Starting in March 2011, we distributed weigh scales to all the farmers in the RCT-sample for them to keep accurate counts of their coffee harvests. The bulk of the coffee harvest arrives in May and June. Starting in March 2011 and through June 2012, every month we distributed a yield calendar to the farmers in the RCT-sample for them to record daily harvests for that month. An example of a yield calendar is shown in Appendix D7. The farmers were given instructions on how to use the weigh scale and how to record their coffee harvests on the calendars. At the end of every month, we collected up the calendars and distributed new ones for the following month.

D4. Best Practices Indices

For outcomes related to best agronomic practices, we create the following indices (all standardized with respect to the mean and variance of the outcome in the control group):

1. **Knowledge**: this index averages fifteen standardized measures of what a farmer knows. It includes whether the farmer knows each of the ten methods used to control insects, pests and other diseases and how they should be used and whether the farmer knows each of five different fertilizers that should be used.

2. **Adoption**: this index is the mean of nine standardized measures of which practices the farmer adopts. Importantly, these are collected using survey questions, and therefore measure *self-reported* adoption, as opposed to *observed* adoption. It includes whether the farmer adopted each of eight methods used to control insects, pests and other diseases, and whether the farmers kept a compost heap. We do not include the indicator of whether farmers kept record books here because the farmer trainers checked these at every session, so they were well kept throughout the study period. Record-keeping is also not an agronomic practice as such, and thus does does not have a counterpart in the tree audits data.

3. **All audits**: this index is the mean of nine standardized measures of what the farmer adopts as per the observed tree audits. This index includes two measures of whether the farmer uses integrated pest management (whether old and dry berries are removed, whether the bark is smoothed or banded to control white borer), whether the farmer used mulch, whether the tree was weeded, whether there are signs of rejuvenation, and four measures of pruning (removal of dead branches, removal of branches touching the ground, removal of crossing branches, removal of unwanted suckers).

4. **Leaf health**: this index is the mean of three standardized measures of leaf health from the tree audits. It includes: whether there are signs of the leaves yellowing, whether the leaves are curling and whether there are signs of the leaves rusting. We changed the sign of the variable so that any increase in the index would indicate an improvement in tree nutrition (i.e. a decrease in the prevalence of leaf defects).

D5. RCT-Sample Modules

The timing of the ten survey rounds of surveys was as follows:

1. December 2009 to January 2010
2. April 2010 to May 2010
3. July 2010
4. September 2010 to October 2010
5. November 2010
6. January 2011 to February 2011
7. June 2011 to July 2011
8. October 2011
9. January 2012 to February 2012
10. September 2012 to October 2012

SECTION	1	2	3	4	5	6	7	8	9	10
Cover Page	X	X	X	X	X	X	X	X	X	X
Consent	X	X	X	X	X	X	X	X	X	X
HH Roster	X	X		X	X	X	X	X	X	X
Plot Roster		X		X	X	X	X	X	X	X
Household-Level Sections										
HH Member Demographic Characteristics	X									
HH Characteristics		X				X			X	
HH Characteristics (Extended)										X
Group Memberships		X						X		
Crop Inventory	X	X		X			X	X		
Plot Questionnaire										
Long Season	X			X				X		
Short Season	X	X					X			
Other	X			X			X	X		
Labor Activities										
Long Season	X			X				X		
Short Season	X	X					X			
Crop Harvests and Sales										
Long Season	X			X				X		
Short Season	X	X					X			
Coffee Activites										
Coffee Plot Details	X	X	X							
General Household Coffee										
A. Coffee Plot Performance/Future	X									
B. Training	X									
C. Cooperative Membership	X									
D. Coffee Farming Activities/Practices	X									
Coffee Delivery	X									
Coffee Module										
Labor Activities for Coffee	X	X	X	X	X	X	X	X	X	X
Coffee Harvests		X	X	X	X	X	X	X	X	X
Coffee Sales	X	X	X	X	X	X	X	X	X	X
Coffee Inputs		X	X	X	X	X	X	X	X	X
Best Practices Schedule/Training Attendance								X	X	X
Consumption Module	X			X				X		
Household Finance										
Decisionmaking - Use of Money	X			X			X			
Bank Holdings/Savings/Debts/Credits	X			X			X			
Remittances	X									
Bank Account Location			X							
Gifts				X			X			
Non-Agricultural Income and Credit		X		X	X	X	X	X	X	
Social Networks	X	X				X			X	X
Best Practices Module: Audits										
Coffee Plot Measurements and ID		X					X		X	X
Best Practices Sheet		X					X		X	X
Tree and Plot Inspection		X					X		X	X
Feedback on Training/Improvement of Knowledge of BP										X
Barriers to Adoption of Best Practices										X

D6. Non-RCT-Sample Modules

SECTION	**2**	**3**	**6**	**7**	**9**	**10**
Cover Page	X	X	X	X	X	X
Consent	X	X	X	X	X	X
HH Roster				X	X	X
Plot Roster				X	X	X
Household-Level Sections						
HH Demographics (Basic)	X					
HH Characteristics (Extended)						X
Group Memberships	X					
Coffee Activites						
General Household Coffee						
A. Coffee Plot Performance/Future	X					
B. Training	X					
C. Coffee Farming Activities/Practices	X					
Coffee Module						
Labor Activities for Coffee			X	X	X	X
Coffee Harvests			X	X	X	X
Coffee Sales	X		X	X	X	X
Coffee Inputs			X	X	X	X
Best Practices Schedule/Training Attendance					X	X
Consumption Module						
Household Finance						
Decisionmaking - Use of Money						
Bank Holdings/Savings/Debts/Credits						
Remittances						
Bank Account Location		X				
Gifts						
Non-Agricultural Income and Credit						
Social Networks	X		X		X	X
Best Practices Module: Audits						
Coffee Plot Measurements and ID		X		X	X	X
Best Practices Sheet		X		X	X	X
Tree and Plot Inspection		X		X	X	X
Feedback on Training/Improvement of Knowledge of BP						X
Barriers to Adoption of Best Practices						X

D7. Yield Calendars

HHID: VILLAGE: NAME:
LOCATION:

HOW MANY KILOGRAMS DID YOU HARVEST TODAY?

1. Hang the scale to a fix and stable place ;
2. Hang the bag (with the cherries in it) to the scale ;
3. Read the number on the scale.

	MAY (05)	DAY	Write here the coffee harvest that you have just weighed	Write here the coffee harvest that you are going to sell and indicate the type of coffee (cherries, wet or dry parch)
1	1-MAY-2012	TUESDAY	Kg	Kg
2	2-MAY-2012	WEDNESDAY	Kg	Kg
3	3-MAY-2012	THURSDAY	Kg	Kg
4	4-MAY-2012	FRIDAY	Kg	Kg
5	5-MAY-2012	SATURDAY	Kg	Kg
6	6-MAY-2012	SUNDAY	Kg	Kg
7	7-MAY-2012	MONDAY	Kg	Kg
8	8-MAY-2012	TUESDAY	Kg	Kg
9	9-MAY-2012	WEDNESDAY	Kg	Kg
10	10-MAY-2012	THURSDAY	Kg	Kg
11	11-MAY-2012	FRIDAY	Kg	Kg
12	12-MAY-2012	SATURDAY	Kg	Kg
13	13-MAY-2012	SUNDAY	Kg	Kg
14	14-MAY-2012	MONDAY	Kg	Kg
15	15-MAY-2012	TUESDAY	Kg	Kg
16	16-MAY-2012	WEDNESDAY	Kg	Kg
17	17-MAY-2012	THURSDAY	Kg	Kg
18	18-MAY-2012	FRIDAY	Kg	Kg
19	19-MAY-2012	SATURDAY	Kg	Kg
20	20-MAY-2012	SUNDAY	Kg	Kg
21	21-MAY-2012	MONDAY	Kg	Kg
22	22-MAY-2012	TUESDAY	Kg	Kg
23	23-MAY-2012	WEDNESDAY	Kg	Kg
24	24-MAY-2012	THURSDAY	Kg	Kg
25	25-MAY-2012	FRIDAY	Kg	Kg
26	26-MAY-2012	SATURDAY	Kg	Kg
27	27-MAY-2012	SUNDAY	Kg	Kg
28	28-MAY-2012	MONDAY	Kg	Kg
29	29-MAY-2012	TUESDAY	Kg	Kg
30	30-MAY-2012	WEDNESDAY	Kg	Kg
31	31-MAY-2012	THURSDAY	Kg	Kg