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

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

February 19, 2025

Abstract

Agricultural extension programs often train a subset of farmers and rely on social networks for knowledge dissemination. We evaluate this approach through a two-stage experiment of an agronomy training program among Rwandan coffee farmers. The first stage randomized trainee concentration at the village level; the second randomly selected participants within villages. Training increased knowledge and self-reported adoption, with smaller effects on audited adoption. At first glance, the program appeared effective: trained farmers had 4.6% higher yields than non-trained applicants within the same village and stronger social ties with co-trainees. However, knowledge did not diffuse, and control farmers with more treatment friends reduced audited adoption and input use. Villages with high trainee concentrations showed suggestive evidence of negative spillovers, likely due to competition for inputs – mulch, fertilizer, and labor. Declines in control farmers' yields account for treatment-control differences, raising concerns both about this dissemination strategy and estimates that fail to consider potential negative spillovers.

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 thank Laura Derksen for helpful feedback. 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 for low-income economies, the successful adoption of yield-enhancing technologies is critical for well-being, particularly in sub-Saharan Africa, where agricultural productivity and technology use are increasingly behind global trends (Aker et al. 2022; Suri and Udry 2022; FAOSTAT 2022; Suri et al. 2024). Agricultural training could 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 approach to organizing agricultural training in developing countries is to train a select group of farmers in each community and rely on the organic spread of information through their social networks. This strategy raises a number of related questions. First, how well does this kind of information diffuse through social networks? Second, does treating some farmers and not others place untreated farmers at some disadvantage, through a market externality or some other form of crowding out? Third, does the group nature of the training alter the structure of social networks and thus potentially disrupt or amplify the impact of the intervention? If untrained farmers seek out training participants, the transmission of the intervention could accelerate (Comola and Prina 2021). Indeed, recent evidence shows that social networks can form strategically around new information sources (Derksen and Souza 2024). Conversely, if the intervention strengthens network ties among participants, it might disadvantage non-participants. For example, when inputs are limited, stronger ties between trained farmers could facilitate resource sharing within this group, potentially excluding others (Banerjee et al. 2021). Importantly, direct comparisons between trained and untrained farmers would then overstate the benefits of agricultural training by failing to account for these or other negative externalities.

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 is (Bridle et al. 2019), how profitable it is (Magnan et al. 2015), and the identity of the early adopters (Beaman et al. 2021).[1] Meanwhile, the potential for market and network externalities has not, to our knowledge, yet been investigated.

In this paper, we seek to fill this gap with an experiment conducted among coffee farmers in Rwanda. We designed a two-stage randomized controlled trial (RCT) and collected detailed social network data both pre- and post-intervention to: (i) assess the impact of an agronomy training program on the structure of social networks, (ii) test whether information shared during training spreads through existing networks, and (iii) estimate spillovers (positive or negative) to untrained farmers. We find evidence of endogenous network rewiring in response to the program, little evidence of knowledge diffusion within social and geographic networks, and consistent signs of negative spillovers to control farmers. These spillovers likely occurred through the crowding out of inputs in limited supply: soil inputs (fertilizers and mulch) and

[1] For recent reviews of this literature, see Suri and Udry (2022) and Suri et al. (2024).

hired labor. We conclude that a naive comparison of treatment and control farmers within a village would seriously overestimate the program's benefits.

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 follows the approach of 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 the villages, 25% of farmers who subscribed were selected for the treatment group, in another third, 50%, and in the final third, 75%. Finally, the 1,594 farmers who signed up for the program were randomly assigned to treatment and control groups according to the proportions assigned in each village.

The selected farmers received monthly instruction modules for the first year, followed by six refresher modules during the second year. These modules covered topics such as nutrition, pest and disease management, weed control, mulching, rejuvenation and pruning, shade management, soil and water conservation, and record keeping. The farmers were organized into groups based on geographical proximity and picked a lead farmer, whose plot was used for demonstrations of the agricultural practices.

We collected ten (including four post-treatment) rounds of surveys 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 the names of farmers to whom the head or spouse in each household talked about growing coffee, before and after the intervention. We use these data to build complete social network maps of all coffee-growing households in the subdistrict, covering more than 3,000 households in 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.

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

Our analysis yields five main findings. First, the agronomy training program fostered new social ties among treatment farmers, particularly in villages with 75% treatment farmers. Almost all these links came from connections with co-trainees within a training group, without significantly affecting interactions between treatment and control farmers. Of the 0.336 new friends gained by treatment farmers — a 28% increase over baseline — 0.319 were from the same training group. In contrast, the networks of control farmers appear to have been largely unaffected by the program.

Second, within villages, treatment farmers have greater knowledge (+1.24 SD) and self-reported adoption of practices (+0.32 SD) than control farmers, but differences in observable adoption outcomes are small. Tree audits indicate slightly greater adoption of the new practices among treatment farmers (+0.02 SD) and higher leaf health (+0.032 SD), with no significant differences in input use (measured by an index of labor days, mulch use, and amount of fertilizer applied). Yields of treatment farmers are 4.6% higher than those of control farmers though this effect is only significant at the 10% level. Importantly, these within-village comparisons could be either an overestimate or an underestimate of the actual treatment effects if the program had positive or negative spillovers on the control group. Either of those would violate the SUTVA assumption.

Third, we find no evidence of information diffusion to untrained farmers through pre-existing networks. We exploit exogenous variation in the number of farmers' friends who were assigned to the treatment group (conditional on the total number of friends at baseline) to examine whether the agronomic practices taught during training sessions diffuse to individuals with whom treatment farmers reported discussing coffee production at baseline. Using this variation, we find no evidence of spillovers to friends in the control group in terms of knowledge or adoption. Strikingly, we find that control farmers connected to more treatment farmers at baseline have significantly lower input use and lower adoption. Applying a similar strategy for geographic neighbors, we find no evidence of diffusion of the program's effects to neighbors either.

Fourth, we leverage the exogenous variation in village-level treatment shares created by our experimental design to test whether the concentration of treatment farmers in a village influences outcomes for control farmers. The results suggest large negative spillovers: control farmers' yields are 17% lower in villages with 75% treatment farmers than in 25% villages. The observed difference between treatment and control farmers on yields is actually negative in villages with few (25%) treatment farmers, and becomes more positive in villages with more treatment farmers. Due to the small number of villages, the p-values based on randomization inference are often larger than 5%, but overall, this suggests that the apparent positive impact of the program on yields from our within-village analysis is accounted for by lower yields among control farmers in high treatment intensity villages, rather than by an increase in the yields of treatment farmers relative to what they would have experienced without the program.

There is suggestive evidence that these apparent perverse effects on yields in the control group were caused by input crowd-out in villages with higher treatment concentration. For both soil inputs and hired labor, a consistent pattern emerges: very small and insignificant differences between treatment and control farmers in low intensity (25%) villages, while in high intensity (75%) villages a gap appears, driven by decreased input use among control farmers. Further evidence supporting the labor channel comes from wage impacts. Daily wage rates are higher in villages with 50% and 75% treatment concentration, with weeding wages, the most affected category, rising by 3% and 6%, respectively. While the limited number of villages in the study reduces precision, a regular pattern of results emerges, indicative of higher demand for labor.

Finally, we examine the supply side of the labor market. Treatment farmers work more days on other households' coffee farms, with this effect intensifying in villages with higher treatment densities; in villages with 75% treatment intensity, they work nearly five times as many days as control farmers. Combined with the program-induced social network changes and our findings on spillovers in labor and soil input use, all of which are particularly pronounced in

75% villages, this indicates that the inputs crowd-out mechanism may have involved reciprocal labor arrangements among treatment farmers, potentially related to the sourcing and application of inputs such as mulch and fertilizers.

The fact that the training program does not increase the yields of treatment farmers but decreases 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, consistent with other estimates (Zhang et al. 2017; Kurniawan et al. 2024), the production function has decreasing returns to fertilizer, in particular NPK, and labor, and no differences between treatment and control farmers. Input re-allocation is thus unlikely to have increased efficiency in aggregate coffee production in this sample.

Taken together, our findings offer a cautionary tale regarding the traditional model of social learning in agriculture. The agronomy training program shifted the allocation of scarce resources from control to treatment farmers, perhaps due to the increased concentration of social networks among treatment farmers. These results suggest that the canonical model, which assumes static networks and straightforward knowledge diffusion, may not fully capture the complexities of real-world agricultural settings. The fact that this may have been missed without a village-level randomization also calls for caution in the evaluation of agricultural extension programs. More broadly, we conclude that any strategy that involves helping some people and not others may backfire when markets are imperfect and poorly integrated.

Our study makes two main contributions. First, it is among the few to document endogenous network responses to a randomized intervention. Notable exceptions include Banerjee et al. (2021), Heß et al. (2021), and Derksen and Souza (2024). Banerjee et al. (2021) show that access to microfinance causes social networks to shrink, with reductions in informal lending and risk-sharing even among non-borrowers. Similarly, Heß et al. (2021) document fewer social links in villages that were randomly assigned to a community-driven development program in The Gambia, which they attribute to the unequal distribution of program benefits. In an experiment in Malawian schools, Derksen and Souza (2024) find that randomizing access to an information resource (Wikipedia) significantly reshaped social networks, with most of the new links formed between treatment and control pairs. In contrast, our findings reveal that the agronomy training program primarily increased social ties among treatment farmers, almost exclusively within their training groups, with no significant increase in connections between treatment and control farmers. While Derksen and Souza (2024) attribute the formation of new links to information-sharing, this explanation seems less relevant in our context. Instead, the strengthening of social ties among co-trainees likely facilitated resource sharing, while the control group was deprived of any benefits of the training program.

Second, to our knowledge, this study is the first to document (negative) market spillovers from an agricultural training intervention. Our finding that limited access to inputs may have hindered farmers from fully adopting the trained agronomic practices or experiencing positive spillovers aligns with the results of Jones et al. (2022), who show that labor constraints impeded the adoption of another productivity-enhancing technology, irrigation, in the context of Rwandan agriculture. In a different context, Crépon et al. (2013) demonstrate that job-training programs can increase employment for participants but displace non-participants by intensifying competition in the French labor market. Finally, Heß et al. (2021) show how the unequal distribution of benefits from a development program can reduce economic interactions within networks (likely driven by elite capture). Taken together, these findings and our own highlight that, in settings with imperfect markets or unequal access to resources, development interventions can inadvertently exacerbate inequalities or harm non-participants.

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 program focused on improving tree health and productivity through a series of labor-intensive practices. These included a balanced nutritional program using both organic and inorganic additives, such as homemade compost and a chemical fertilizer containing nitrogen, phosphorus, and potassium (NPK 22-12-6). It also involved integrated pest management and effective weed control through hand weeding and mulch application, with mulching additionally serving to maintain moisture and reduce soil erosion. No input subsidies were provided; instead, farmers were trained to apply these practices independently, without external support. Appendix A provides a more detailed description of the practices covered and the expected impacts as outlined by the NGO's agronomists.

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 advertised 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. Farmers who did not have coffee farms were not allowed to register for the program and only one person per household was allowed to register. In total, 1594 farmers registered interest in the program. Although the program was advertised 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 assigned to treatment, 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.

Attendance rates of each training session are reported in Appendix Table A6, categorized by village treatment concentration group. During the first year of training, farmers in the treatment group attended an average of around 8 out of the 11 meetings (an attendance rate of about 73%). Overall, attendance was slightly higher in villages with a greater proportion of farmers offered training: the average attendance rate was 69% in 25% concentration villages, 72% in 50% concentration villages, and 74% in 75% concentration villages.

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 pre-program census. As we describe in detail in Appendix B, 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. We collected data on these farmers to be able to map full social networks for the coffee farmers in the RCT sample. However, because the non-RCT sample is not a random subset of all the coffee farmers in these villages, we omit them from our analysis of the diffusion of results.

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. These audits only cover tree health and observable practices and were not designed to directly observe or measure fertilizer or labor use, so they cannot speak to these.

We use this data to construct the following indices of correlated outcomes:

1. **Knowledge**: this index is the simple average of 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 for optimal tree nutrition.

2. **Self-reported adoption**: this index is the mean of nine standardized measures of adoption of best agronomic practices. Since these are collected using survey questions, this index measures self-reported 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 farmer 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 not have a counterpart in the tree audits data.

3. **Adoption, audits**: this index is the mean of eight standardized measures of what the farmer adopts as per the observed tree audits. This index includes two measures of integrated pest management (whether old and dry berries are removed, whether the bark is smoothed or banded to control white borer), whether the tree canopy was mulched, whether the dripline was weeded, 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, audits**: 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).

5. **Inputs**: this index is the mean of five standardized measures of input use. These are total household labor days, total non-household (paid and unpaid) labor days, KGs of compost, KGs of NPK, and the share of coffee plots on which the household applied mulch. This is based entirely on farmer self-reports.

To measure social networks comprehensively, we collected the names of farmers to whom the head or spouse in each household talked about growing coffee - "coffee friends" - before and after the intervention. We use these data to build complete social network maps of all coffee-growing households in the sub-district, covering more than 3,000 households in 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. The resulting social networks include a mix of farmers in the treatment group, control group, and those who were not interested in participating in the program.

4 Results

4.1 Program Impacts on Social Networks

This section examines the impact of the training on farmers' reported social network links to other farmers with whom they discuss coffee. We analyze the total number of "coffee friends" reported by each farmer, both across different geographical areas (within our outside the farmer's village) and for friends with different treatment statuses. The effects could *a priori* follow several directions; treatment farmers might have formed new friendships during the training sessions or found it more beneficial to interact with other treatment farmers, potentially reducing their interactions with control farmers. Such a shift could have limited the latter's opportunities for learning and risk-sharing. Conversely, control farmers might have sought out treatment farmers (Comola and Prina 2021) and created new friendships to gain information.

Table 1 displays the results. Panel A, which reports effects on total friends, shows that treatment households gain 0.336 friends who were also chosen for training (a 28% increase over baseline), far more than their (insignificant) increase in friendships with control group members. In other words, the program increases the share of the average treatment household's social network that is also part of the treatment group. By contrast, control group households gain no new friends from either the treatment or control groups. Both treatment and control households 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. Treatment 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 treatment friends can be attributed to additional friends from the same training group. Appendix Table A8 further reveals that the formation of these new links among treatment households is especially pronounced in villages with a higher treatment concentration, particularly those where 75% of the sample households were assigned to treatment.

The control group also increases treatment friends within a 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 treatment and control), a group likely more focused on coffee farming, strengthened within-village links among themselves over time. This 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 treatment households within a village (see Appendix Table A8), indicating that it is likely due to a time trend rather than a program effect.

In sum, the social impact of the training appears to have been primarily to create new friendships among farmers attending the same training meetings. There is no evidence that the control group actively formed new connections with treatment individuals to benefit from their increased coffee knowledge.

4.2 Within-Village Treatment Effects and Social Diffusion

Our experiment's central hypothesis is grounded in a canonical model of social learning, which hinges on two key elements: treatment effects and social diffusion. We test both elements in turn and present the results in Table 2. Panel A reports within-village estimates of treatment effects, defined as the mean treatment-control difference in outcomes within villages, obtained by running the following specification:

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

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 randomly assigned to the agronomy training program, γ_j are a set of village fixed effects and δ_t are survey round fixed effects.[2] We use OLS regressions for all indexed outcomes and a Poisson regression for the yields outcome, clustering standard errors at the household level. Throughout the paper, we use Poisson regressions for all outcomes with a non-negligible share of zeros and a long right tail, as this approach aligns with the distributional properties of such variables. As can be seen in Appendix Table A1, we find balance between treatment and control households across a wide variety of baseline outcomes

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

(the p-value on the joint F-test is 0.999). Appendix Table A7 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. To increase precision, we control for baseline outcomes selected by post-double selection LASSO (Belloni et al. 2014).

This specification compares treatment 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 invalid if information did indeed diffuse to the control group (in which case the treatment effect on knowledge would be underestimated), or if there were negative externalities on the control for some outcomes (in which case the treatment effect would be overestimated). We examine this assumption in detail below.

Panel A of Table 2 reports the treatment effects, β, for measures of knowledge and adoption of 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 practices: treatment farmers have a 1.24 standard deviation (henceforth SD) higher knowledge index and a 0.32 SD higher self-reported adoption index than control farmers within their village.

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 a small but significant within-village treatment effect on the adoption index constructed from the tree audits data (+0.02 SD). Column 4 of Table 2 suggests that the trees of treatment farmers are better nourished: the audits data reveal a 0.032 SD lower index of tree disease (yellow, curling, or rusting leaves) among treatment farmers.

Column 5 shows within-village treatment effects on the input quantities index. While treatment farmers report using slightly more inputs (labor, mulch, and fertilizer) on their coffee plots, with an increase of 0.034 SD, this effect is not statistically significant. Column 6 shows that yields are approximately 4.6% higher in the treatment group compared to the control group, although this estimate is only significant at the 10% level.

Taken together, the results in Panel A of Table 2 indicate that while the treatment led to significantly higher knowledge and self-reported adoption of agronomic practices, the differences in observable adoption outcomes are much smaller, with no significant differences in input use and a marginally significant effect on yields.[3]

In Panel B of Table 2, we analyze the diffusion of information through treatment farmers' networks as measured prior to the intervention. Our focus is on two types of networks: baseline "coffee friends" and neighbors. This section emphasizes baseline friends, but the results for diffusion through neighbors (household and plot) are broadly consistent and reported in Appendix Table A13. The identification strategy leverages exogenous variation in the number of treatment friends, controlling for the total number of RCT-sample friends (i.e. who entered the treatment lottery). We interact these network variables with the household's own treatment

[3] Appendix Table A9, which reports the estimates of the same regressions as in Panel A but without controlling for baseline covariates, shows that the coefficients are very similar for all outcomes, but not significant.

status to examine whether the effects differ between the control and treatment groups. Panel B therefore reports the results of the following specification:

$$
\begin{aligned}
y_{ijt} = \alpha &+ \beta_0 NumTreatFriends_{ij} + \delta_0 NumFriends_{ij} \\
&+ \beta_1 Treat_{ij} \times NumTreatFriends_{ij} + \delta_1 Treat_{ij} \times NumFriends_{ij} + \gamma_j + \delta_t + \epsilon_{ijt} \quad (2)
\end{aligned}
$$

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. As in Panel A, we cluster standard errors at the household level and we control for baseline outcomes selected by post-double selection LASSO; Appendix Table A9 shows the results of running the regression without controls.

The estimates in columns 1 and 2 report the diffusion effect of treatment friends on knowledge and self-reported adoption. The coefficient estimate in the second row of column 1 indicates that there is no diffusion of knowledge about agronomic practices to the control group. Consistent with this finding, column 2 (self-reported adoption) shows no spillover of actual practices to the control group, which is further corroborated by the absence of effects on leaf health and yields in columns 4 and 6, respectively. Additional results reported in Appendix Table A10 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 coefficients in the second row of columns 3 and 5 are surprising. These indicate that for an average farmer in the control group, having one more friend in the treatment group decreases the tree audit outcomes index by 0.057 SD and the inputs index by 0.042 SD.[4] Appendix Table A11, which disaggregates the audited adoption index by component, shows that control farmers with more treatment friends have fewer mulched and pruned trees and less weeded plots. Similarly, Appendix Table A12 indicates that these control farmers use less mulch (although this result is only significant at the 10% level), apply less chemical fertilizer (NPK), and hire less labor. The fact that control farmers apply fewer best practices (or apply them less intensively) and use fewer inputs if they have more links to treatment farmers at baseline suggests that the program may have had negative spillovers on control farmers. Our later analyses show further evidence of this mechanism at the village level.

In the fourth row of Panel B, we interact the number of treatment friends with household treatment status. The estimates suggest that any positive diffusion through baseline networks is confined to links within the treatment group itself. The coefficient on the interaction with treatment is positive and statistically significant for leaf health: column 4 shows that, for the average treatment farmer, each additional treatment friend leads to a 0.034 SD increase in the leaf health index. In contrast, the positive interaction between treatment status and the number of treatment friends merely offsets the negative effect of treatment friends on audited program

[4] Appendix Table A13, which shows similar results on neighbors of treatment farmers to Table 2, Panel B, also include negative effects on the audits adoption index.

adoption (column 3) and input use (column 5). Finally, we observe no effect of additional treatment friends on yields for the treatment group (column 6).

In summary, Table 2 highlights two key takeaways. First, we find smaller treatment effects on observed adoption compared to self-reported adoption in our basic within-village specification, with no significant effects on input use or yields at conventional levels. Second, we find no evidence of diffusion of the taught practices through baseline social networks in the control group. This lack of diffusion may stem from the effects not being large enough to spread, or alternatively, information sharing about the taught practices may have been more likely to occur within the newly formed network of co-trainees, as the program increased social links amongst them (Table 1) but not with control farmers. Finally, the negative network effects on inputs and audited adoption, which are present within the control group but much weaker among treatment farmers, suggest the potential for negative spillovers mediated by increased social network connections between trained individuals.[5] We now turn to this question.

4.3 Spillovers by Village 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 examine heterogeneity in program impacts by village treatment concentration. In Table 3, we regress the same outcomes as in Table 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. Throughout this section, we use OLS regressions for indexed and binary outcomes, and Poisson regressions for outcomes capturing quantities (yields in Table 3, fertilizer quantities in Table 4, and labor days in Table 5), clustering standard errors at the village level. Given the small number of villages, we also report p-values for the exact null hypothesis, using a randomization inference algorithm that takes the experimental design into account.[6] To address any baseline imbalances across village groups (Appendix Tables A2-A5), we include baseline covariates selected using post-double selection LASSO in all regressions.

Panel A of Table 3 shows the aggregate effects of treatment concentration, while Panel B shows the interactions of the village-level concentrations with treatment status. At the bottom of each panel, we report the results of several hypothesis tests, including comparisons of coefficients for the 50% and 75% treatment concentration indicators and their interactions with treatment, as well as tests of treatment effects within each village group. Generally, our estimates are not precise enough to statistically distinguish between the effects of 50% and 75% treatment concentration; hence, we highlight in the text only the p-values for cases where the null of equal impacts is rejected.

[5] To avoid repetition, we use "trained" and "treatment" interchangeably throughout the text. "Trained" refers to farmers selected for the training, regardless of attendance.

[6] Here, randomization inference involves 1,000 permutations of the sample. First, village-level treatment intensities are shuffled across villages, and then treatment-control status is assigned to households based on the re-assigned village proportions. Randomization inference is then conducted using the `ritest` Stata command (Heß 2017).

The effects reported in Panel A represent the net impacts of the program at the village level. We find little evidence of aggregate improvements in audited adoption or yields, despite increased knowledge and self-reported adoption of farming techniques in higher-intensity villages. The knowledge index increases significantly with the share of the village assigned to treatment, consistent with the findings in Table 2, Panel A, where treatment effects on knowledge are evident, and the higher share of treatment households in 75% villages naturally leads to a larger effect. Effects on self-reported adoption are positive for both 50% and 75% treatment shares, with the 75% concentration showing a larger magnitude (the randomization inference p-value of the test of equal effects is 0.19).

We do not see a monotonic relationship between effect sizes and village treatment shares for any of the other outcomes in Panel A. Importantly, the aggregate impact of treating 50% or 75% of farmers on yields is negative and insignificant (column 6), despite the larger number of farmers enrolled in the training in higher-intensity villages.

Panel B decomposes these net aggregate effects by household treatment status. The top row of columns 1 and 2 confirms the existence of "pure" treatment effects (i.e., in villages least confounded by spillovers) on knowledge and self-reported adoption. Column 1 also shows a larger gap in knowledge between treatment and control groups in villages with 75% treatment concentration villages, and columns 2 and 3 show that the treatment-control gaps in adoption (both self-reported and as measured by the audits index) are larger in 75% and 50% villages compared to 25% villages, although these differences are not significant. These results are consistent with the absence of positive knowledge spillovers to the control group.

While the average within-village treatment effects reported in columns 4–6 of Table 2, Panel A are all positive, the estimates in columns 4–6 of Panel B of Table 3 suggest that these results are driven by *negative* input and yield spillover effects on the control group in villages with higher treatment intensity. In these columns, we see a triple pattern of results that is consistent across virtually all subsequent outcomes related to input use, tree health, and yields. First, we find no positive treatment effects in 25% concentration villages in the top row. The yield estimate in column 6 is even negative, though it does not survive randomization inference (p=0.243). Second, the exact opposite is true in 75% treatment villages: in these areas, treatment households use significantly (p=0.02) more inputs and have significantly higher yields (p=0.08).[7] Third, there is no significant difference between treatment households in 75% villages and control households in 25% villages.[8] Thus the positive estimated effects in high-intensity villages seem more likely to be due to negative spillovers among the control group, rather than (for instance) increasing returns to training among treatment farmers. This is further supported by the estimates on the audited adoption index, which follow a similar pattern (column 3). These results are consistent with the individual-level spillover regressions, where we found negative (and significant) spillovers on program adoption and input use of control farmers who had more friends in the treatment group.

[7]The "p: Treatment + Treatment x 75% T=0" row reports the test of this hypothesis.

[8]The "p: Treatment + 75% T + Treatment x 75% T=0" row reports the test of this hypothesis.

Appendix Table A14, which disaggregates the estimates in Table 3 for the audits and leaf health indices, shows that spillovers are most pronounced for mulching, weeding, and yellowing of leaves. For example, in villages with 75% treatment intensity, the village indicator coefficients for mulching and weeding are negative for the control group, while the interaction terms with treatment are highly positive, with p-values of 0.010 for mulching and 0.034 for weeding, respectively. Additionally, yellowing of leaves, an objective measure of tree nutrition, is worse among the control group in 75% treatment villages (p=0.070), while the interaction term with treatment is highly positive (p=0.095). These outcomes reflect treatment-control gaps in the adoption of practices that require applying soil inputs (mulch and fertilizer) and which are labor-intensive.

Why might higher treatment densities result in reduced input use, poorer leaf health, and lower yields for control farmers? One explanation is that increased demand for inputs from treatment farmers may have crowded out access for control farmers due to limited supply. To better understand the program's effects on input use, we then evaluate spillovers in soil inputs (mulch and fertilizers) and labor separately. The results of these additional analyses are presented in Tables 4 and 5.

Soil Inputs. Table 4 reports the results of the same specification used in Table 3, but focusing on outcomes related to the use of soil inputs: mulch, compost and NPK. Column 1 reports results on the index of standardized outcomes, while columns 2-7 report effects on each outcome separately.

Panel A indicates no aggregate program impacts on input use across varying village treatment intensities. In contrast, Panel B demonstrates the exact same pattern of results consistent with negative spillovers as in columns 3-6 of Table 3. Treatment effects are insignificant in 25% intensity villages, yet become larger and sometimes significant in villages with 75% intensity. Again, these results seem driven by decreased input use among control farmers in villages where many of their peers were trained.

Column 1 shows this pattern for the index of soil input use, where treatment effects in 75% villages are statistically significant (p=0.00). The treatment-control differences in 75% villages are also larger than in 50% villages (p=0.08 for the test of equal interactions with treatment). Breaking this down by components of the index, we find that results are strongest in mulch application (column 2) and the quantities of compost and NPK used (columns 5 and 7) in villages with higher treatment concentrations, particularly in the 75% group. For mulch, measured only at the extensive margin, we observe negative coefficients for the 50% and 75% treatment intensity indicators, alongside positive coefficients for their interactions with treatment (p=0.061 for the 75% interaction). This finding aligns with the audit index components results in Appendix Table A14, which show significantly more mulching under tree canopies in the treatment group compared to the control in 75% villages.

Columns 3-7 focus on compost and fertilizers. Columns 3, 4, and 6 present the results of OLS regressions on extensive margin outcomes of fertilizer use (columns 3 and 6) and an audit-based indicator of adherence to agronomy program standards (column 4). Columns 5 and

7 report the results of Poisson regressions on the total quantities of compost and NPK applied, respectively. For compost, there is no evidence of differential effects by village group on the first two outcomes (columns 3 and 4). However, column 5 reveals a pattern similar to that for mulch, with increasing gaps in intensive margin compost use between treatment and control groups as the treatment share rises, and the treatment effect becomes statistically significant in 75% villages (p=0.05).

Columns 6 and 7 examine NPK use. Similarly to compost, we find no evidence of differential effects for NPK at the extensive margin (column 6). However, for intensive margin use, column 7 shows a large treatment-control gap in total kilograms of NPK applied in villages with 75% treatment concentration. Here, the treatment effect in 75% villages is close to significant at conventional levels (p=0.06).

Labor Use. The next input we examine is labor. Many of the trained practices are labor-intensive, and since farmers grow coffee alongside other food crops, the demands on household members' time can quickly accumulate. Over two-thirds of the households in our RCT sample hired labor to assist with at least some coffee-related tasks during the study period. Evidence from Jones et al. (2022) shows that labor constraints significantly restrict the adoption of irrigation in Rwanda, highlighting labor market frictions as another potential source of negative spillovers on the control group in villages with higher treatment intensities. To examine whether labor is being crowded out in these villages, we run Poisson regressions on total labor days, first pooling across all tasks (column 1) and then separately for the most labor-intensive activities in our data - mulching, fertilizing, weeding, and harvesting - and all other tasks. Additionally, we distinguish between household and non-household labor, the latter encompassing both paid and unpaid, hired labor.[9]

The results are presented in Table 5. Panel A shows variation in the effects of 50% and 75% village shares across tasks and labor types. While some estimates are significant based on conventional standard errors, none remain so under the more conservative randomization inference. Still, there is suggestive evidence of increased hired labor days in 75% villages (column 2), particularly for mulching, fertilizing, and other tasks (columns 4, 6, and 12).

Once again, negative spillovers appear in Panel B, most prominently with the total days of non-household labor (column 2), showing the same pattern of results found in yields, leaf health, and soil inputs. There is essentially no treatment effect in 25% intensity villages, yet in 75% villages, trained farmers hire more labor than control farmers (though this is marginally significant at p=0.12). And again, this result seems driven by a negative effect on the control group. The opposite effect is observable for household labor (column 1), which control households appear to use more as the treatment intensity of the village increases.

Examining total days separately by task shows that the 75% concentration effects on hired labor seem to be driven by mulching (column 4), fertilizing (column 6) and harvesting (column 10), although the effects on mulching and fertilizing are not significant. The effects on

[9]Most non-household labor was compensated through wages or piece-rate payments; however, households also engaged unpaid labor for certain tasks, often providing non-monetary remuneration such as meals. In our endline survey rounds, 24.8% of non-household labor days in our RCT sample involved unpaid labor.

weeding labor are more mixed; here the impact on non-household hired labor use is largest in 25% concentration villages, though still positive in the 75% group. We find clear evidence of treatment-control gaps in the number of hired labor days for harvesting, the most labor-intensive task, in higher-intensity villages, with a highly significant and large effect in 75% villages.

Since harvesting labor has a nearly one-to-one relationship with harvested quantities, it is not surprising that the yield results observed in Table 3, column 6, reappear in our analysis of harvest labor. Within 25% concentration villages, treatment households use (insignificantly) less labor than controls. However, in 75% concentration villages, treatment households employ significantly more labor (p=0.01). This effect is driven by fewer workers hired by control households (p=0.047). In contrast, the effects on household labor are much smaller and even reversed in 75% villages (albeit insignificant), suggesting that the program effects operate through labor market transactions between households, rather than simply reducing labor demand.

Figure 1 illustrates the differences in the use of non-household and household labor between treatment and control groups in 75% villages, where the labor market effects are strongest. We observe that treatment farms employ more non-household labor than control farms (Panel B) but show either the opposite trend or no difference in household labor, with the exception of mulching (Panel A). The key takeaway is that treatment farms in higher-intensity treatment villages shift towards hiring more labor, potentially crowding out labor inputs for control households. These differences are partially compensated by offsetting changes in household labor, limiting the effects on the total amount of labor used.

Wages. A natural test of the input crowd-out hypothesis is to examine whether increased demand for inputs causes treatment farmers to bid up input prices. We investigate this in the next table, focusing on wages paid to hired labor, as price data for soil inputs are unavailable. Compost and mulch are primarily home-produced or collected by farmers, while NPK was initially provided free by the government.[10]

Table 6 presents the results of five OLS regressions on daily wage rates at the household-task level. In column 1, we stack observations for all tasks for which at least one household reported hiring external labor at a daily rate, trimming the top 1% to remove implausibly high values. In columns 2–4, the outcome is household-level daily wages paid per survey round for the four most labor-intensive tasks examined in Table 5. Column 1 includes controls for round and task fixed effects, while columns 2–4 control for round fixed effects.

The broad takeaway from this table is that the daily wage rates are higher in villages with 50% and 75% treatment concentration compared to 25% villages. Weeding, the task for which households hire the most external labor (column 4), and the stacked regression estimates (column 1) provide the most compelling evidence of higher daily wages in villages with greater treatment concentration.

[10]During our study, however, the government phased out its direct provision of NPK due to logistical challenges, transitioning to a system where coffee washing stations facilitated input provision. Under this new system, NPK was no longer free and was typically offered on credit, leading to a sharp decline in its use between 2011 and 2012 (from 298 reported users in our RCT sample in 2011 to 126 users in 2012).

Column 1 shows that daily wage rates are, on average, 22.19 RWF higher in 50% villages and 32.48 RWF higher in 75% villages compared to 25% villages. These correspond to a 3.4% and 4.9% increase relative to the mean daily rate in 25% villages, respectively. The 75% village effect is robust to randomization inference (p=0.034). Column 4 shows that daily wage rates for weeding are, on average, 19.41 RWF higher in 50% villages and 40.71 RWF higher in 75% villages (p=0.028), corresponding to a 3% and 6% increase relative to the mean in 25% villages.

The effects of treatment intensities on daily wage rates for mulching (column 2), fertilizing (column 3), and harvesting (column 5) are less conclusive, although most coefficient estimates are positive, except for the 50% intensity effect for harvesting.

These wage patterns corroborate the idea that the program increased demand for hired labor. Together with the results in Tables 4 and 5, they support the interpretation that the negative spillovers shown in Table 3 stem from input crowd-out in the control group in villages with higher treatment concentration. Even without statistical significance for some results, the consistent direction of effects across multiple outcomes — soil inputs, labor days, and wages — reinforces our confidence in this interpretation.

Labor Supply. Finally, we examine the supply side of the labor market using data on the number of days household members worked outside their own household. Table 7 presents village-level regression results for three outcomes: total days worked (column 1), days worked on other farmers' coffee plots (column 2), and days spent on non-coffee labor (column 3).

Panel A indicates that households in higher-intensity villages spend more time working outside their own households overall. This increase seems to be driven by non-coffee labor, as the estimates for coffee labor in column 2 are negative, while those for other types of labor in column 3 are positive. Breaking these results down by treatment status in Panel B reveals intriguing patterns in coffee labor: treatment farmers are more likely to work on other farmers' coffee plots in 25% villages (though this effect is not statistically significant), and this tendency becomes stronger with higher treatment intensity (column 2, Panel B, p=0.16). In contrast, the control group spends considerably less time working on other farmers' coffee plots (p=0.113) in higher-intensity villages, which likely explains the overall negative effects in Panel A.

These findings suggest that trained farmers may be more inclined to hire one another over members of the control group. This pattern could arise from program-induced changes in social networks (Table 1), particularly in 75% villages, where new connections between co-trainees are most prevalent (Appendix Table A8). The sharing of resources among treatment farmers may also contribute to this dynamic, which aligns with our finding of spillovers in input use (Table 4). For example, farmers might engage in reciprocal labor arrangements or assist one another with tasks such as sourcing inputs (e.g., mulch or compost) or performing labor-intensive activities. Although our labor supply data do not allow us to disaggregate the effects by specific tasks to further isolate these mechanisms, these results suggest that resource constraints might interact with social network effects to shape program outcomes.

5 Agricultural Production Function Estimation

The effects of village-level treatment concentrations in Tables 3-5 suggest a reallocation of inputs from control to treatment farmers in villages with higher treatment concentration. The impact of these transfers on aggregate output and efficiency is *a priori* ambiguous. If the treatment causes an increase in productivity among the treatment farmers, then such a reallocation would increase efficiency by assigning more inputs to the farmers who can use them most productively. Similarly, if the 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 training has no productivity effect, then concentrating more inputs among the treatment 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):

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

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 2, with sub-figure 2a showing the association between NPK and output, and sub-figure 2b showing the association between labor and output. In both cases, the relationship is concave. For NPK, this is consistent with the findings of agronomic experiments on the effects of fertilizer on coffee yields (Zhang et al. 2017; Kurniawan et al. 2024). 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 treatment/control 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. Rather, given the declining returns, it seems that it may have exacerbated misallocation.

6 Conclusions

This paper evaluates a common strategy to diffuse agricultural innovation and knowledge: a "cascade" model where some farmers get trained first, with the hope that the innovation will diffuse among their social contacts. Our results suggest that, at least in our context in Rwanda, this strategy is not particularly effective.

First, the hypothesized mechanism for information diffusion is undermined by the finding that the program increased social links only between co-trainees in the treatment group. While this does not imply that treatment farmers ceased interacting with the control group entirely, it is likely that they focused their "chats over coffee" — discussions about farming practices — with the other farmers with whom they attended the trainings. Second, although knowledge of best agronomic practices increased among treatment farmers, we find no evidence that treatment farmers shared their new knowledge with control households they were socially connected to or lived close to. Third, the consistent evidence that control households experienced *negative* spillovers when they were socially connected to more treatment households, as well as in high treatment concentration areas, suggests that much of the 4.6% higher yields of treatment farmers compared to the control group that we observe at endline is the result of these negative spillovers, rather than a net gain for the treatment 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 small effect we find from the tree audits on treatment farmers' own adoption suggests that they probably did not undertake enough of these new practices to see a meaningful difference.

Another possibility is that information provision alone was not sufficient to unleash the yield-boosting potential of these agronomic practices. The negative spillovers we find on control farmers appear to stem from input crowd-out in a context of limited supply, as seen in lower soil input use, reduced reliance on hired labor, and rising wages. These supply constraints may have also inhibited the treatment group farmers from taking full advantage of the techniques taught in the trainings. This finding complements those of Jones et al. (2022), who show that labor market failures constrained the adoption of irrigation in another part of rural Rwanda.

More broadly, 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). 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 yields 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 input markets 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, those distortions may give the misleading impression that the program is effective. This evidence suggests that this widely practiced strategy likely needs to be re-evaluated.

References

Aker, J., Bergquist, L. F., Jayne, T., and Magruder, J. R. (2022). Agricultural Technology in Africa.

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.*

Beaman, L., BenYishay, A., Magruder, J., and Mobarak, M. (2021). Can network theory-based targeting increase technology adoption? *American Economic Review*, 111(6):1918–1943.

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.

Derksen, L. and Souza, P. C. L. (2024). Who Knows? Information Access and Endogenous Network Formation.

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.

Heß, S. (2017). Randomization inference with stata: A guide and software. *The Stata Journal*, 17(3):630–651.

Heß, S., Jaimovich, D., and Schündeln, M. (2021). Development projects and economic networks: Lessons from rural gambia. *The Review of Economic Studies*, 88(3):1347–1384.

Jones, M., Kondylis, F., Loeser, J., and Magruder, J. (2022). Factor market failures and the adoption of irrigation in rwanda. *American Economic Review*, 112(7):2316–2352.

Kurniawan, S., Nugroho, R. M. Y. A. P., Ustiatik, R., et al. (2024). Soil nitrogen dynamics affected by coffee (Coffea arabica) canopy and fertilizer management in coffee-based agroforestry. *Agroforestry Systems*, 98:1323–1341.

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.

Suri, T., Udry, C., Aker, J. C., Barrett, C. B., Bergquist, L. F., Carter, M., Casaburi, L., Osei, R. D., Gollin, D., Hoffmann, V., Jayne, T., Karachiwalla, N., Kazianga, H., Magruder, J., Michelson, H., Startz, M., and Tjernström, E. (2024). Agricultural technology in africa. *VoxDevLit*, 5(2).

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

Zhang, Z. X., Cai, Z. Q., Liu, G. Z., Wang, H., Huang, L., and Cai, C. T. (2017). Effects of fertilization on the growth, photosynthesis, and biomass accumulation in juvenile plants of three coffee (Coffea arabica l.) cultivars. *Photosynthetica*, 55(1):134–143.

Figure 1: Mean labor days by task and treatment status, in 75% villages.

Panel A.

Panel B.

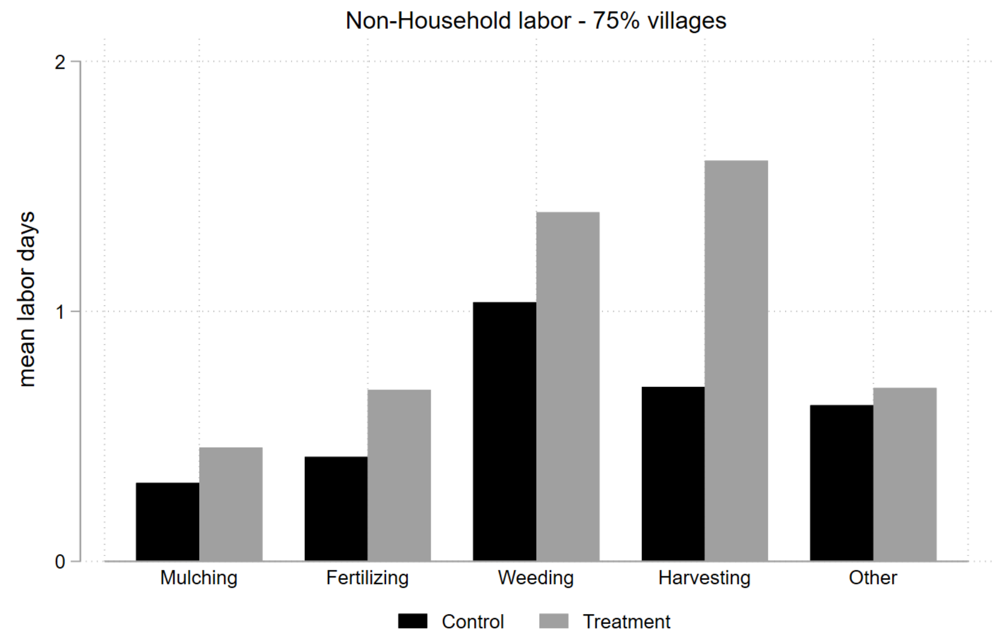

Notes: This figure shows the mean number of total labor days used by treatment and control groups in 75% treatment concentration villages, disaggregated by task and type of labor (household versus non-household, where the latter = unpaid + paid). We compute these means from the estimates in Table 5.

Figure 2: Coffee Cherry Production Function

(a) Relationship between Coffee Cherry Output and NPK Use

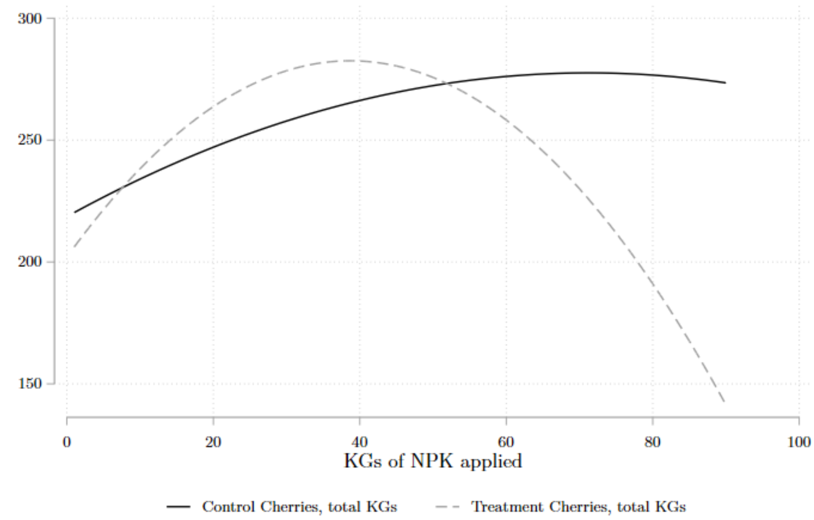

(b) Relationship between Coffee Cherry Output and Labor Use

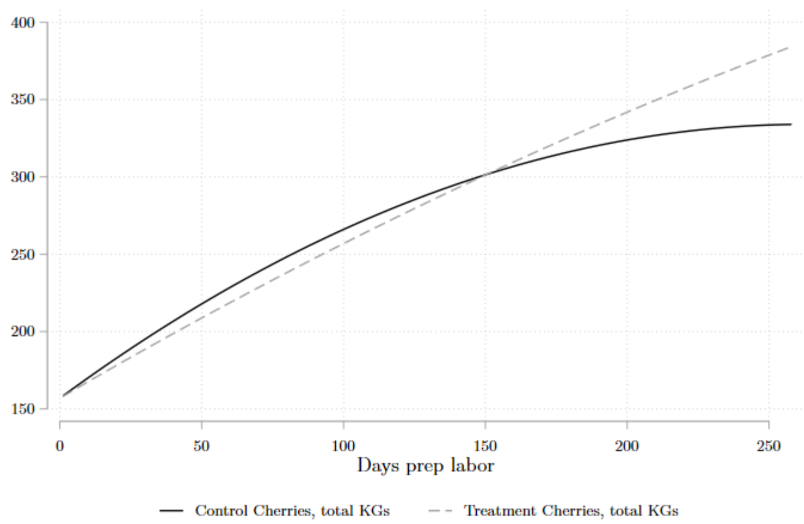

Figures show output of coffee cherries as a function of input quantity, using coefficients estimated from Equation 3. The relationship between output and each input quantity is plotted from 0 to the 99th percentile of that input intensity in the data, with other input quantities set to their mean values.

Table 1: Program Impacts on Social Networks.

	(1) Trained Friends	(2) Control Friends	(3) Non-sample Friends	(4) All Friends
	Panel A: Friends in All Villages			
Treat × Post	0.336	0.030	-0.164	0.201
	[0.057]	[0.052]	[0.051]	[0.119]
Control × Post	0.034	0.022	-0.150	-0.095
	[0.064]	[0.060]	[0.047]	[0.112]
T-C p. value	0.000	0.897	0.831	0.005
Baseline mean	1.196	1.004	1.567	3.766
Observations	3156	3156	3156	3156
	Panel B: Friends in Own Village			
Treat × Post	0.360	0.0870	-0.00119	0.446
	[0.0516]	[0.0466]	[0.0394]	[0.104]
Control × Post	0.0988	0.0893	0.0392	0.227
	[0.0542]	[0.0563]	[0.0428]	[0.110]
T-C p. value	0.000	0.970	0.408	0.025
Baseline mean	1.196	1.004	1.567	3.766
Observations	3156	3156	3156	3156
	Panel C: Friends Outside Own Village			
Treat × Post	-0.0238	-0.0572	-0.163	-0.244
	[0.0191]	[0.0155]	[0.0268]	[0.0408]
Control × Post	-0.0650	-0.0677	-0.189	-0.322
	[0.0216]	[0.0127]	[0.0284]	[0.0370]
T-C p. value	0.209	0.488	0.570	0.214
Baseline mean	1.196	1.004	1.567	3.766
Observations	3156	3156	3156	3156
	Panel D: Friends in Same Training Group			
Treat × Post	0.319			
	[0.0629]			
Baseline mean	0.412			
Observations	1678			

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 B). The outcome variable in column (1) is the count of a household's friends selected for training. The outcome variable in column (2) is the count of friends who applied for training but were not selected. The outcome variable in column (3) is the count of friends who did not apply for training. The outcome variable in column (4) is the sum of (1)+(2)+(3).

Table 2: Within-Village Treatment Effects and Diffusion through Baseline Social Networks.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Inputs index (labor + mulch + fertilizers)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A			
Treatment	1.243	0.321	0.020	0.032	0.034	0.045
	[0.060]	[0.036]	[0.010]	[0.010]	[0.023]	[0.027]
Control mean	-0.000	0.000	-0.000	0.000	0.001	2.724
Observations	4622	4622	47618	47618	6157	6154
			Panel B			
Treatment	1.279	0.331	0.039	-0.071	-0.115	0.081
	[0.092]	[0.051]	[0.015]	[0.016]	[0.043]	[0.047]
Number of treatment friends	-0.023	-0.006	-0.057	0.003	-0.042	-0.029
	[0.029]	[0.024]	[0.007]	[0.007]	[0.019]	[0.018]
Number of sample friends	0.007	0.021	0.034	-0.018	0.019	0.055
	[0.017]	[0.014]	[0.005]	[0.005]	[0.013]	[0.013]
Treatment X num. treatment friends	0.035	-0.004	0.029	0.034	0.060	-0.001
	[0.059]	[0.035]	[0.009]	[0.009]	[0.033]	[0.022]
Treatment X num. sample friends	-0.031	-0.002	-0.021	0.012	0.019	-0.009
	[0.042]	[0.025]	[0.007]	[0.007]	[0.020]	[0.016]
Control mean	0.682	0.179	0.001	0.014	0.039	2.724
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
Observations	4622	4622	47618	47618	6157	6154

Notes: Standard errors clustered by household are 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 B). 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) represents the average of five variables, each standardized using the Control group mean and standard deviation: total household labor days, total non-household labor days, the share of coffee plots where the household applied mulch, kilograms of compost, and kilograms of NPK. 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. All regressions control for baseline covariates selected via post-double selection LASSO.

Table 3: Impacts of Village Treatment Concentration.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Input index (labor + mulch + fertilizers)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A			
50% T in village	0.194	0.038	0.005	0.059	-0.016	-0.081
	[0.097]	[0.086]	[0.049]	[0.041]	[0.088]	[0.074]
	(0.410)	(0.715)	(0.950)	(0.167)	(0.867)	(0.497)
75% T in village	0.696	0.166	-0.046	-0.016	-0.008	-0.038
	[0.143]	[0.103]	[0.062]	[0.038]	[0.077]	[0.110]
	(0.000)	(0.102)	(0.501)	(0.744)	(0.946)	(0.752)
Sample mean, 25% villages	0.056	0.037	0.018	-0.002	0.027	2.967
p-value: 50%=75%	0.00 (0.01)	0.11 (0.19)	0.41 (0.47)	0.02 (0.06)	0.92 (0.92)	0.70 (0.71)
Observations	4622	4622	47618	47618	6157	6154
			Panel B			
Treatment	1.139	0.286	-0.082	-0.050	-0.022	-0.115
	[0.245]	[0.079]	[0.100]	[0.042]	[0.053]	[0.042]
	(0.000)	(0.000)	(0.206)	(0.416)	(0.771)	(0.243)
50% T in village	-0.081	-0.038	-0.017	0.021	-0.026	-0.126
	[0.046]	[0.100]	[0.061]	[0.046]	[0.085]	[0.067]
	(0.756)	(0.739)	(0.835)	(0.673)	(0.813)	(0.307)
75% T in village	-0.109	-0.053	-0.138	-0.087	-0.141	-0.192
	[0.064]	[0.115]	[0.092]	[0.043]	[0.067]	[0.131]
	(0.705)	(0.663)	(0.100)	(0.187)	(0.200)	(0.157)
Treatment X 50% T in village	-0.007	0.013	0.083	0.101	0.031	0.149
	[0.280]	[0.122]	[0.112]	[0.064]	[0.064]	[0.084]
	(0.967)	(0.900)	(0.296)	(0.204)	(0.755)	(0.218)
Treatment X 75% T in village	0.314	0.102	0.176	0.127	0.191	0.280
	[0.286]	[0.093]	[0.109]	[0.039]	[0.080]	[0.078]
	(0.197)	(0.371)	(0.049)	(0.145)	(0.071)	(0.040)
Control mean, 25% villages	0.056	0.037	0.037	0.004	0.020	2.967
p-value: 50% T= 75% T	0.70 (0.92)	0.88 (0.91)	0.14 (0.14)	0.01 (0.11)	0.13 (0.30)	0.59 (0.64)
p-value: Treat x 50% T = Treat x 75% T	0.11 (0.12)	0.40 (0.38)	0.12 (0.21)	0.63 (0.74)	0.02 (0.07)	0.18 (0.31)
p: Treatment + Treatment x 50% T=0	0.00 (0.00)	0.00 (0.00)	0.98 (0.99)	0.23 (0.25)	0.81 (0.89)	0.64 (0.66)
p: Treatment + 50% T + Treatment x 50% T=0	0.00 (0.00)	0.01 (0.01)	0.83 (0.82)	0.14 (0.15)	0.84 (0.84)	0.36 (0.48)
p: Treatment + Treatment x 75% T=0	0.00 (0.00)	0.00 (0.00)	0.02 (0.11)	0.00 (0.21)	0.00 (0.02)	0.01 (0.08)
p: Treatment + 75% T + Treatment x 75% T=0	0.00 (0.00)	0.00 (0.00)	0.52 (0.55)	0.81 (0.83)	0.73 (0.79)	0.81 (0.84)
Observations	4622	4622	47618	47618	6157	6154

Notes: Standard errors clustered by village are in brackets, and randomization inference p-values are in parentheses below these. Each panel concludes with p-values for various linear restriction tests, with those in parentheses corresponding to the randomization inference results. All specifications control for round fixed effects. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix B). 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) represents the average of five variables, each standardized using the Control group mean and standard deviation: total household labor days, total non-household labor days, the share of coffee plots where the household applied mulch, kilograms of compost, and kilograms of NPK. 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. All regressions control for baseline covariates selected via post-double selection LASSO.

Table 4: Impacts of Village Treatment Concentration on Soil Inputs Use.

	Soil inputs index	Used mulch	Used compost	Has a proper compost heap	Compost (Baskets)	Used NPK	NPK (KG)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				Panel A			
50% T in village	0.005	-0.034	0.046	-0.005	-0.060	-0.006	0.156
	[0.072]	[0.069]	[0.032]	[0.010]	[0.080]	[0.019]	[0.258]
	(0.942)	(0.699)	(0.134)	(0.789)	(0.696)	(0.773)	(0.486)
75% T in village	-0.028	-0.048	0.004	0.006	0.045	-0.011	0.049
	[0.073]	[0.057]	[0.034]	[0.019]	[0.135]	[0.016]	[0.234]
	(0.739)	(0.561)	(0.905)	(0.749)	(0.796)	(0.531)	(0.827)
Sample mean, 25% villages	0.035	0.543	0.502	0.171	56.891	0.073	0.950
p-value: 50%=75%	0.65 (0.65)	0.85 (0.85)	0.05 (0.15)	0.56 (0.54)	0.38 (0.49)	0.70 (0.74)	0.44 (0.61)
Observations	6157	6149	6157	4615	6157	6157	6157
				Panel B			
Treatment	0.040	-0.006	0.019	0.041	-0.156	0.016	0.332
	[0.070]	[0.024]	[0.033]	[0.026]	[0.107]	[0.014]	[0.392]
	(0.576)	(0.843)	(0.524)	(0.127)	(0.442)	(0.276)	(0.397)
50% T in village	-0.018	-0.048	0.051	-0.013	-0.101	-0.006	0.156
	[0.078]	[0.063]	[0.036]	[0.022]	[0.066]	[0.021]	[0.365]
	(0.840)	(0.560)	(0.139)	(0.542)	(0.597)	(0.725)	(0.597)
75% T in village	-0.180	-0.109	-0.009	-0.003	-0.270	-0.019	-0.425
	[0.076]	[0.042]	[0.035]	[0.026]	[0.105]	[0.018]	[0.338]
	(0.066)	(0.177)	(0.815)	(0.905)	(0.195)	(0.356)	(0.251)
Treatment X 50% T in village	0.027	0.030	-0.020	-0.004	0.154	-0.006	-0.139
	[0.080]	[0.036]	[0.044]	[0.039]	[0.118]	[0.016]	[0.465]
	(0.788)	(0.477)	(0.610)	(0.923)	(0.576)	(0.758)	(0.779)
Treatment X 75% T in village	0.176	0.085	0.005	-0.016	0.510	0.000	0.375
	[0.084]	[0.040]	[0.042]	[0.031]	[0.186]	[0.016]	[0.426]
	(0.075)	(0.061)	(0.901)	(0.657)	(0.061)	(0.991)	(0.484)
Control mean, 25% villages	0.035	0.543	0.502	0.171	57.234	0.069	0.856
p-value: 50% T= 75% T	0.04 (0.09)	0.31 (0.44)	0.01 (0.13)	0.74 (0.72)	0.08 (0.38)	0.48 (0.50)	0.02 (0.13)
p-value: Treat x 50% T = Treat x 75% T	0.01 (0.08)	0.18 (0.18)	0.53 (0.52)	0.72 (0.70)	0.03 (0.13)	0.56 (0.73)	0.09 (0.27)
p: Treatment + Treatment x 50% T=0	0.08 (0.23)	0.35 (0.33)	0.96 (0.95)	0.21 (0.06)	0.98 (0.99)	0.18 (0.35)	0.44 (0.48)
p: Treatment + 50% T + Treatment x 50% T=0	0.50 (0.59)	0.74 (0.79)	0.20 (0.14)	0.16 (0.24)	0.18 (0.54)	0.85 (0.85)	0.30 (0.21)
p: Treatment + Treatment x 75% T=0	0.00 (0.00)	0.01 (0.01)	0.38 (0.44)	0.13 (0.31)	0.02 (0.05)	0.03 (0.26)	0.00 (0.06)
p: Treatment + 75% T + Treatment x 75% T=0	0.64 (0.69)	0.61 (0.71)	0.70 (0.66)	0.28 (0.27)	0.57 (0.63)	0.85 (0.86)	0.36 (0.27)
Observations	6157	6149	6157	4615	6157	6157	6157

Notes: Standard errors clustered by village are in brackets, and randomization inference p-values are in parentheses below these. Each panel concludes with p-values for various linear restriction tests, with those in parentheses corresponding to the randomization inference results. All specifications control for round fixed effects and use data from endline rounds, i.e. rounds 6 through 9 (see Appendix B). Columns 1, 2, 3, 4 and 6 show the results of OLS regressions. Columns 5 and 7 show the results of **Poisson** regressions where the outcome is the amount of compost and NPK applied, respectively, by the household on their coffee farm. All regressions control for baseline covariates selected via post-double selection LASSO.

Table 5: Impacts of Village Treatment Concentration on Labor Days.

	Total labor days		Mulching		Fertilizing		Weeding		Harvesting		All other tasks	
	HH	non-HH	HH	non-HH	HH	non-HH	HH	non-HH	HH	non-HH	HH	non-HH
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A												
50% T in village	0.007	-0.080	-0.073	-0.206	0.052	-0.072	0.072	0.123	-0.072	-0.187	-0.015	-0.162
	[0.079]	[0.135]	[0.116]	[0.290]	[0.152]	[0.155]	[0.040]	[0.084]	[0.093]	[0.255]	[0.132]	[0.173]
	(0.929)	(0.681)	(0.499)	(0.557)	(0.702)	(0.920)	(0.429)	(0.336)	(0.491)	(0.645)	(0.906)	(0.535)
75% T in village	0.025	0.043	-0.051	0.217	0.014	0.520	0.124	0.097	-0.037	-0.080	0.040	0.115
	[0.065]	[0.130]	[0.092]	[0.337]	[0.150]	[0.248]	[0.073]	[0.093]	[0.075]	[0.283]	[0.096]	[0.204]
	(0.751)	(0.825)	(0.658)	(0.595)	(0.923)	(0.276)	(0.130)	(0.449)	(0.710)	(0.837)	(0.732)	(0.663)
Sample mean, 25% villages	34.620	4.082	3.842	0.339	6.324	0.381	7.192	1.198	9.466	1.545	7.795	0.619
p-value: 50%=75%	0.75 (0.80)	0.38 (0.49)	0.81 (0.84)	0.07 (0.22)	0.66 (0.76)	0.01 (0.10)	0.43 (0.59)	0.81 (0.81)	0.70 (0.74)	0.74 (0.74)	0.59 (0.65)	0.15 (0.23)
Observations	6157	6157	6157	6157	6157	6157	6157	6157	6157	6157	6157	6157
Panel B												
Treatment	-0.090	0.085	-0.100	-0.127	0.174	-0.000	-0.076	0.582	-0.173	-0.412	-0.232	0.190
	[0.078]	[0.121]	[0.122]	[0.404]	[0.173]	[0.288]	[0.066]	[0.211]	[0.101]	[0.197]	[0.074]	[0.179]
	(0.222)	(0.772)	(0.396)	(0.826)	(0.169)	(0.999)	(0.438)	(0.067)	(0.044)	(0.222)	(0.005)	(0.593)
50% T in village	-0.020	0.006	-0.103	-0.140	0.141	-0.044	0.067	0.375	-0.147	-0.193	-0.080	-0.020
	[0.082]	[0.168]	[0.116]	[0.351]	[0.143]	[0.212]	[0.047]	[0.138]	[0.109]	[0.328]	[0.145]	[0.222]
	(0.821)	(0.977)	(0.410)	(0.766)	(0.345)	(0.924)	(0.466)	(0.052)	(0.165)	(0.657)	(0.526)	(0.949)
75% T in village	0.083	-0.269	-0.022	-0.106	0.222	0.130	0.193	0.056	-0.030	-0.848	0.031	0.093
	[0.078]	[0.181]	[0.116]	[0.573]	[0.166]	[0.454]	[0.091]	[0.170]	[0.093]	[0.273]	[0.115]	[0.231]
	(0.403)	(0.367)	(0.880)	(0.863)	(0.206)	(0.867)	(0.078)	(0.830)	(0.807)	(0.047)	(0.792)	(0.799)
Treatment X 50% T in village	0.098	-0.217	0.108	-0.078	-0.252	-0.057	0.047	-0.717	0.239	0.227	0.251	-0.374
	[0.087]	[0.160]	[0.147]	[0.435]	[0.189]	[0.368]	[0.081]	[0.239]	[0.112]	[0.295]	[0.091]	[0.286]
	(0.264)	(0.564)	(0.478)	(0.916)	(0.132)	(0.946)	(0.672)	(0.061)	(0.027)	(0.613)	(0.012)	(0.440)
Treatment X 75% T in village	-0.021	0.342	0.026	0.499	-0.392	0.496	-0.046	-0.284	0.104	1.243	0.167	-0.085
	[0.098]	[0.236]	[0.165]	[0.546]	[0.202]	[0.405]	[0.085]	[0.241]	[0.118]	[0.301]	[0.101]	[0.347]
	(0.830)	(0.457)	(0.872)	(0.579)	(0.035)	(0.634)	(0.730)	(0.488)	(0.365)	(0.012)	(0.110)	(0.878)
Control mean, 25% villages	34.926	3.895	3.886	0.349	5.891	0.367	7.204	0.980	9.787	1.630	8.157	0.569
p-value: 50% T= 75% T	0.16 (0.28)	0.19 (0.36)	0.52 (0.58)	0.95 (0.95)	0.43 (0.63)	0.68 (0.75)	0.18 (0.25)	0.04 (0.23)	0.28 (0.32)	0.05 (0.15)	0.37 (0.39)	0.67 (0.80)
p-value: Treat x 50% T = Treat x 75% T	0.10 (0.18)	0.01 (0.13)	0.55 (0.58)	0.15 (0.45)	0.27 (0.41)	0.13 (0.51)	0.19 (0.43)	0.00 (0.21)	0.08 (0.18)	0.00 (0.01)	0.34 (0.38)	0.44 (0.57)
p: Treatment + Treatment x 50% T=0	0.84 (0.89)	0.21 (0.55)	0.92 (0.92)	0.21 (0.69)	0.31 (0.45)	0.80 (0.92)	0.54 (0.72)	0.18 (0.51)	0.17 (0.27)	0.40 (0.43)	0.72 (0.74)	0.41 (0.56)
p: Treatment + 50% T + Treatment x 50% T=0	0.89 (0.90)	0.30 (0.57)	0.43 (0.41)	0.32 (0.44)	0.71 (0.67)	0.71 (0.86)	0.32 (0.71)	0.08 (0.19)	0.40 (0.46)	0.14 (0.36)	0.65 (0.60)	0.26 (0.52)
p: Treatment + Treatment x 75% T=0	0.07 (0.10)	0.04 (0.12)	0.51 (0.50)	0.31 (0.50)	0.03 (0.10)	0.08 (0.40)	0.02 (0.17)	0.01 (0.26)	0.26 (0.38)	0.00 (0.01)	0.34 (0.38)	0.73 (0.77)
p: Treatment + 75% T + Treatment x 75% T=0	0.69 (0.74)	0.28 (0.46)	0.30 (0.43)	0.46 (0.56)	0.98 (0.98)	0.02 (0.23)	0.27 (0.44)	0.01 (0.03)	0.25 (0.36)	0.96 (0.96)	0.75 (0.78)	0.42 (0.49)
Observations	6157	6157	6157	6157	6157	6157	6157	6157	6157	6157	6157	6157

Notes: Standard errors clustered by village are in brackets, and randomization inference p-values are in parentheses below these. Each panel concludes with p-values for various linear restriction tests, with those in parentheses corresponding to the randomization inference results. All specifications control for round fixed effects and use data from endline rounds, i.e. rounds 6 through 9 (see Appendix B). All columns show the results of **Poisson** regressions and control for baseline covariates selected via post-double selection LASSO.

Table 6: Impacts of Village Treatment Concentration on Wages.

	All stacked (1)	Mulching (2)	Fertilizing (3)	Weeding (4)	Harvesting (5)
50% T in village	22.191	23.218	27.969	19.410	-20.684
	[9.113]	[31.475]	[11.447]	[14.347]	[15.502]
	(0.177)	(0.426)	(0.108)	(0.335)	(0.368)
75% T in village	32.480	26.678	25.031	40.713	16.551
	[14.097]	[31.045]	[16.051]	[14.422]	[16.937]
	(0.034)	(0.367)	(0.174)	(0.028)	(0.496)
Sample mean, 25% villages	656.337	680.680	693.041	647.184	643.251
p-value: 50%=75%	0.47 (0.62)	0.88 (0.91)	0.85 (0.89)	0.20 (0.28)	0.05 (0.10)
Observations	2459	267	544	665	409

Notes: Standard errors clustered by village are in brackets, and randomization inference p-values are in parentheses below these. The table concludes with p-values for the linear restriction test of equal effects of the 50% and 75% village treatment intensities, with those in parentheses corresponding to the randomization inference results. Here, the randomization inference permutations are applied to the restricted set of 1,041 households out of 1,594 in our sample who report hiring any external labor to work on their coffee farm in exchange for daily wages. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix B) and report the results of OLS regressions. Piece rate payments are excluded from the sample. In column (1), all household-task-level average daily wage values are included except the top 1%, to remove implausibly high values. In columns 2-4, the outcome is household-level average daily wages paid per survey round per task. Column 1 controls for round and task fixed effects. Controls 2-4 control for round fixed effects. All regressions control for baseline covariates selected via post-double selection LASSO.

Table 7: Impacts of Village Treatment Concentration on Labor Supply.

	Total labor days outside of HH	Labor days, coffee	Labor days, other
	(1)	(2)	(3)
		Panel A	
50% T in village	0.156	-0.611	0.175
	[0.089]	[0.371]	[0.093]
	(0.143)	(0.212)	(0.115)
75% T in village	0.203	-0.382	0.218
	[0.081]	[0.442]	[0.085]
	(0.056)	(0.461)	(0.045)
Sample mean, 25% villages	3.855	0.135	3.721
p-value: 50%=75%	0.56 (0.67)	0.59 (0.60)	0.60 (0.71)
Observations	4710	4710	4710
		Panel B	
Treatment	0.010	0.219	0.001
	[0.178]	[0.618]	[0.191]
	(0.950)	(0.828)	(0.997)
50% T in village	0.092	-1.001	0.114
	[0.098]	[0.567]	[0.103]
	(0.505)	(0.159)	(0.410)
75% T in village	0.110	-1.650	0.137
	[0.130]	[0.668]	[0.133]
	(0.531)	(0.113)	(0.431)
Treatment X 50% T in village	0.119	0.546	0.118
	[0.208]	[0.920]	[0.223]
	(0.575)	(0.681)	(0.581)
Treatment X 75% T in village	0.116	1.329	0.108
	[0.213]	[0.778]	[0.227]
	(0.621)	(0.406)	(0.645)
Control mean, 25% villages	3.839	0.125	3.713
p-value: 50% T= 75% T	0.89 (0.92)	0.38 (0.50)	0.86 (0.89)
p-value: Treat x 50% T = Treat x 75% T	0.98 (0.99)	0.35 (0.56)	0.95 (0.96)
p: Treatment + Treatment x 50% T=0	0.23 (0.31)	0.26 (0.31)	0.30 (0.35)
p: Treatment + 50% T + Treatment x 50% T=0	0.05 (0.11)	0.64 (0.70)	0.05 (0.10)
p: Treatment + Treatment x 75% T=0	0.28 (0.39)	0.00 (0.16)	0.37 (0.47)
p: Treatment + 75% T + Treatment x 75% T=0	0.01 (0.05)	0.83 (0.87)	0.01 (0.05)
Observations	4710	4710	4710

Notes: Standard errors clustered by village are in brackets, and randomization inference p-values are in parentheses below these. All specifications control for round fixed effects. Each panel concludes with p-values for various linear restriction tests, with those in parentheses corresponding to the randomization inference results. All regressions use data from rounds 6, 7, and 8 (we did not collect this module in the final endline survey, round 9). All columns show the results of **Poisson** regressions and control for baseline covariates selected via post-double selection LASSO.

APPENDIX

A 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.

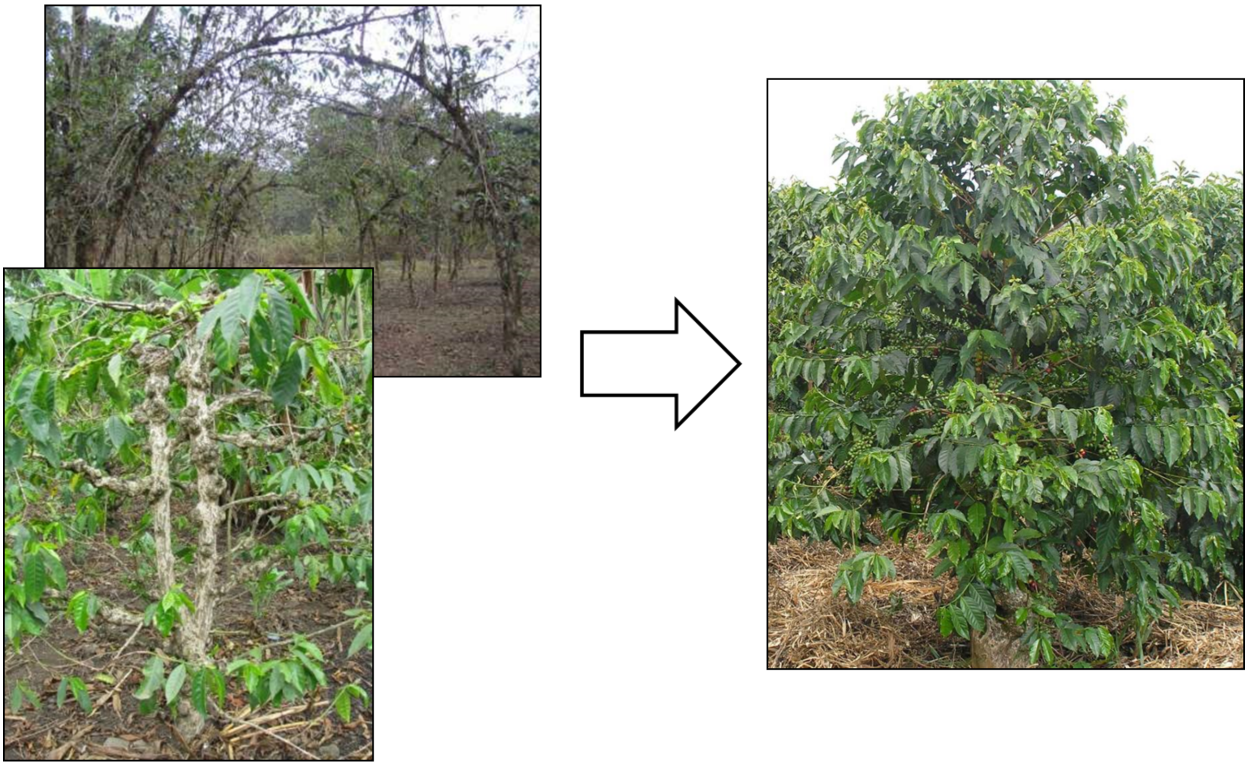

- Nutrition: a balanced nutritional program based on organic and inorganic additives, with the exact requirements determined by soil analysis. In the sub-district where our program evaluated took place, this included a combination of homemade compost and NPK (nitrogen, phosphorus, and potassium) as the main recommended chemical fertilizer. The specific type of NPK recommended for coffee in the area was 22-12-6.

- 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.

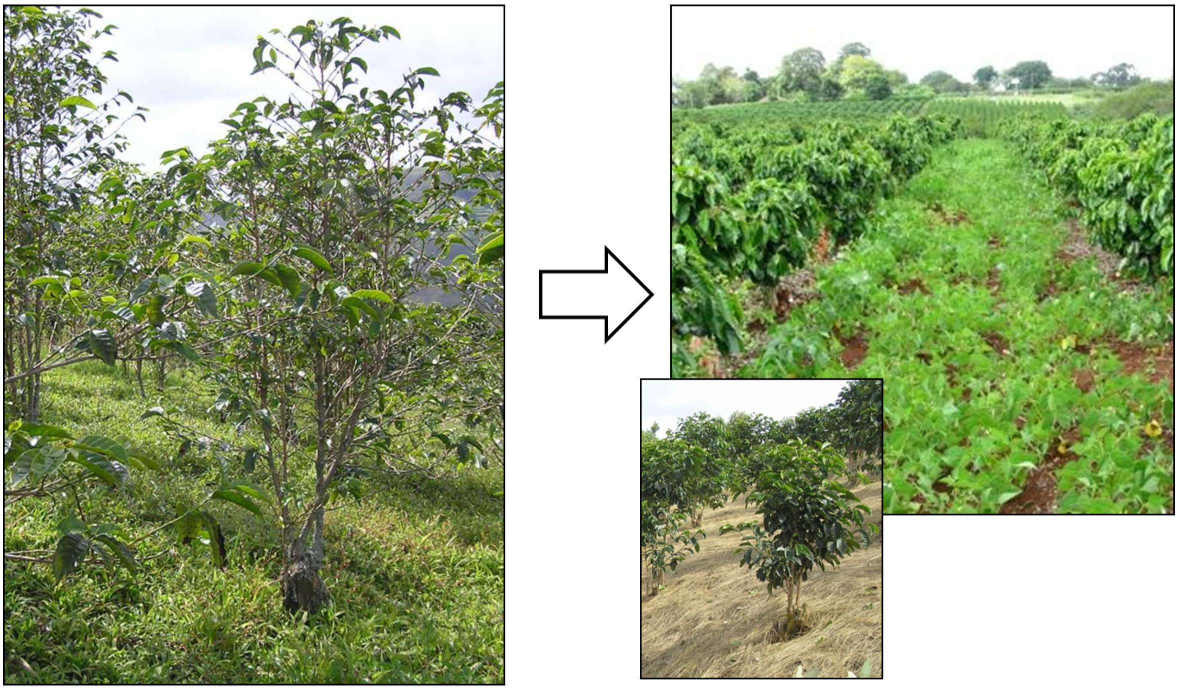

- Mulching: techniques to conserve moisture, add organic matter, and control soil erosion.

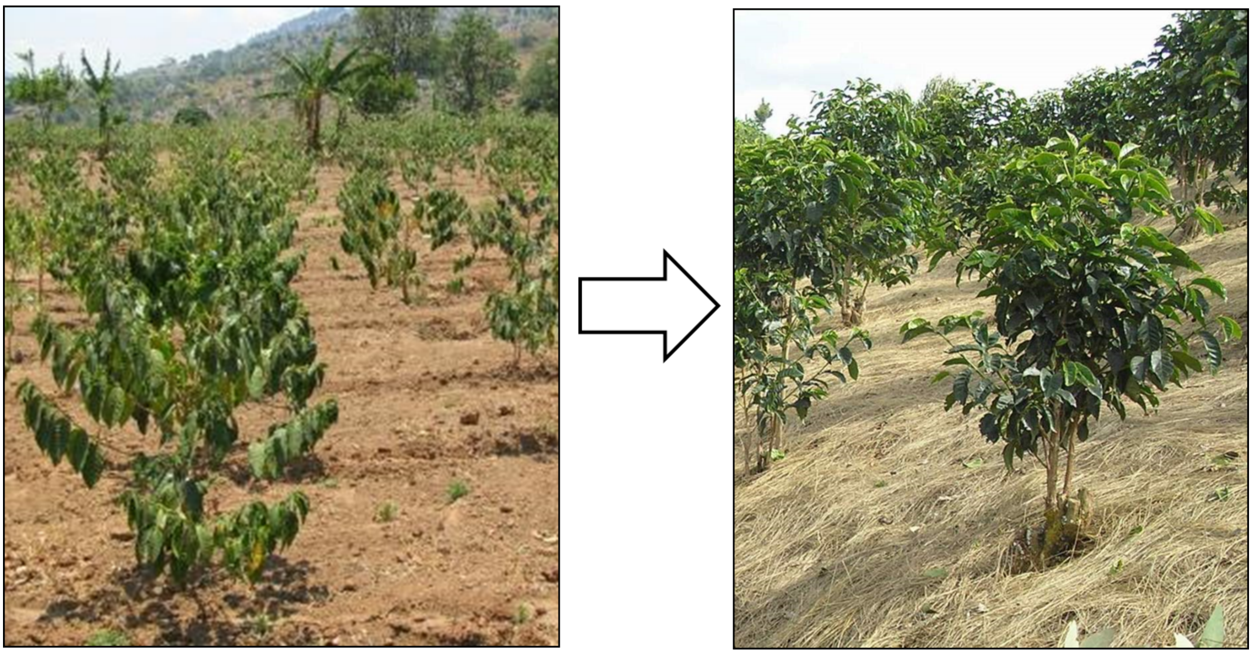

- Soil and water conservation: use of a number of techniques such as mulching, terracing and water traps to control soil erosion and maintain soil fertility. Encourage the management of water resources through conservation zones.

- 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

B 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 subdistrict 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.

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

C Additional Tables and Figures

Table A1: Balance Checks: Treatment vs. Control Households.

	Control Mean	Treatment Coeff.	Std Error	P-value
Head, Years of Schooling	3.615	-.061	.146	.678
Female Headed Household	.32	.025	.013	.075
Household Size	5.035	-.037	.141	.794
Average Schooling of Household	3.211	.103	.09	.264
Yield, total KGs per tree	.783	.017	.04	.682
Total Trees	240.385	1.294	11.616	.912
Fraction Unproductive Trees	.305	-.012	.014	.398
Cut Stems	.102	-.009	.014	.546
Book Keeping Done	.028	-.007	.007	.356
Removed Dead Branches	.771	-.043	.018	.023
Removed Suckers	.913	-.018	.011	.131
Removed Weeds	.992	-.005	.004	.256
Applied Compost	.716	.017	.024	.485
Applied NPK	.178	0	.015	.997
Applied Lime	.019	.01	.01	.336
Applied Pesticides	.767	-.022	.019	.264
Applied Mulch	.878	-.022	.016	.179
p-value of joint F-test				.9998

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

Table A2: Inter-village Balance Checks: Treatment Concentration Groups.

	25% Mean	50% coeff.	50% S.E.	p-val	75% coeff.	75% S.E.	p-val	p:50%=75%
Head, Years of Schooling	3.75	-.229	.239	.348	-.35	.211	.109	.62
Female Headed Household	.35	-.041	.031	.195	-.018	.03	.567	.511
Household Size	5.01	.01	.196	.96	.002	.15	.991	.959
Average Schooling of Household	3.36	-.226	.2	.27	-.045	.233	.848	.319
Yield, total KGs per tree	.83	-.031	.065	.643	-.129	.07	.077	.253
Total Trees	228.29	3.304	31.212	.917	53.676	40.582	.197	.282
Fraction Unproductive Trees	.27	.037	.036	.322	.077	.021	.001	.264
Cut Stems	.09	.015	.023	.518	.021	.025	.405	.776
Book Keeping Done	.03	0	.007	.996	-.005	.006	.451	.424
Removed Dead Branches	.77	-.036	.028	.199	-.026	.036	.466	.806
Removed Suckers	.92	-.03	.028	.294	-.019	.022	.392	.689
Removed Weeds	.99	-.001	.008	.902	.002	.007	.739	.674
Applied Compost	.76	-.085	.034	.018	-.008	.031	.805	.075
Applied NPK	.19	-.056	.057	.342	.029	.055	.601	.164
Applied Lime	.02	.004	.015	.811	.004	.011	.742	.996
Applied Pesticides	.77	-.05	.052	.341	.017	.056	.759	.276
Applied Mulch	.89	-.059	.033	.088	-.013	.026	.628	.142
p-value of joint F-test				.2522			.0556	

Notes: Standard errors are clustered by village. This table compares mean outcomes across different village groups without distinguishing between treatment and control status.

Table A3: Intra-village Balance Checks: Treatment vs. Control within 25% villages.

	Control Mean	Treatment Coeff.	Std Error	P-value
Head, Years of Schooling	3.875	-.466	.268	.126
Female Headed Household	.344	.034	.027	.256
Household Size	4.993	.079	.33	.819
Average Schooling of Household	3.333	.109	.237	.659
Yield, total KGs per tree	.808	.082	.142	.581
Total Trees	217.404	42.206	20.726	.081
Fraction Unproductive Trees	.268	.005	.04	.896
Cut Stems	.092	-.021	.018	.295
Book Keeping Done	.035	-.035	.007	.001
Removed Dead Branches	.784	-.049	.043	.291
Removed Suckers	.929	-.031	.026	.273
Removed Weeds	.989	0	.014	.976
Applied Compost	.738	.089	.039	.058
Applied NPK	.202	-.029	.035	.445
Applied Lime	.014	.027	.026	.335
Applied Pesticides	.773	.002	.056	.966
Applied Mulch	.89	.018	.03	.561
p-value of joint F-test				.0049

Notes: Standard errors are clustered by village.

Table A4: Intra-village Balance Checks: Treatment vs. Control within 50% villages.

	Control Mean	Treatment Coeff.	Std Error	P-value
Head, Years of Schooling	3.556	-.061	.236	.803
Female Headed Household	.31	.002	.018	.908
Household Size	4.98	.085	.204	.686
Average Schooling of Household	3.086	.097	.129	.471
Yield, total KGs per tree	.796	.005	.054	.922
Total Trees	235.002	-6.819	10.615	.537
Fraction Unproductive Trees	.313	-.015	.017	.395
Cut Stems	.112	-.02	.025	.437
Book Keeping Done	.03	-.007	.011	.549
Removed Dead Branches	.756	-.042	.024	.113
Removed Suckers	.894	-.006	.02	.757
Removed Weeds	.99	-.003	.006	.595
Applied Compost	.67	.011	.043	.808
Applied NPK	.147	-.015	.021	.502
Applied Lime	.013	.023	.012	.093
Applied Pesticides	.726	-.006	.024	.817
Applied Mulch	.852	-.032	.024	.21
p-value of joint F-test				.6889

Notes: Standard errors are clustered by village.

Table A5: Intra-village Balance Checks: Treatment vs. Control within 75% villages.

	Control Mean	Treatment Coeff.	Std Error	P-value
Head, Years of Schooling	3.251	.205	.173	.269
Female Headed Household	.296	.052	.019	.023
Household Size	5.224	-.279	.247	.291
Average Schooling of Household	3.235	.107	.164	.531
Yield, total KGs per tree	.71	-.012	.047	.801
Total Trees	293.822	-15.845	26.538	.567
Fraction Unproductive Trees	.359	-.018	.024	.475
Cut Stems	.099	.012	.024	.625
Book Keeping Done	.013	.011	.011	.319
Removed Dead Branches	.776	-.042	.036	.269
Removed Suckers	.921	-.025	.013	.092
Removed Weeds	1	-.011	.006	.101
Applied Compost	.77	-.023	.031	.493
Applied NPK	.197	.035	.017	.076
Applied Lime	.039	-.02	.016	.261
Applied Pesticides	.836	-.06	.024	.038
Applied Mulch	.908	-.034	.028	.26
p-value of joint F-test				.0042

Notes: Standard errors are clustered by village.

Table A6: Attendance Rates, by Training Session and Village Treatment Concentration.

	Mean attendance rate		
	25% T	50% T	75% T
Integrated Pest Management	.694	.743	.761
Nutrition	.663	.744	.737
Harvesting	.724	.727	.759
Weeding	.663	.659	.723
Mulching	.694	.705	.741
Pruning and Rejuvenation	.612	.702	.737
Pesticide use	.694	.708	.715
Composting	.643	.685	.699
Erosion Control	.724	.679	.717
Shade	.653	.672	.686
Nutrition Review	.704	.774	.768
Harvesting Review	.735	.725	.701
Sustainability	.704	.748	.765
Composting Review	.765	.754	.772
Pruning and Rejuvenation Review	.735	.741	.761

Notes: Session-specific average attendance rates for each village treatment concentration group. In total there were 38 training groups attended by the treatment group (from 27 villages). All sessions were taught on (one of) the coffee plot(s) of the training group's assigned focal farmer by a TechnoServe farmer trainer. Sessions are listed in chronological order.

Table A7: 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 A8: Program Impacts on Social Networks by Village Treatment Concentration.

	(1) Treatment friends, same village	(2) Same village friends	(3) Different village friends
Treatment	0.0644	0.205*	-0.0366
	[0.0727]	[0.100]	[0.0685]
Treatment × post	0.121	0.222	-0.221*
	[0.117]	[0.239]	[0.108]
Control × post	0.0255	0.246	-0.352***
	[0.0579]	[0.144]	[0.0574]
50% T × Treatment × post	0.228	0.163	-0.0827
	[0.146]	[0.295]	[0.118]
75% T × Treatment × post	0.300**	0.316	0.0113
	[0.142]	[0.291]	[0.111]
50% T × Control × post	0.142	0.0797	0.0583
	[0.105]	[0.261]	[0.0670]
75% T × Control × post	0.0714	-0.249	0.0307
	[0.167]	[0.266]	[0.0690]
25% T village mean	0.651	3.399	0.857
p-value: Treatment x 50% T X post = Treatment x 75% T X post	0.577	0.544	0.164
p-value: Control x 50% T X post = Control x 75% T X post	0.697	0.280	0.671
Observations	3156	3156	3156

Notes: Standard errors clustered by household in brackets. All specifications control for village fixed effects and use household-level data from the baseline and round 9 social network surveys (see Appendix B). Column (1) is the number of a household's friends residing in their village who were selected for training (Table 1, Panel B, column 1). Column (2) is the total count of friends from their village (Table 1, Panel B, column 4). Column 3 is the total number of friends residing outside their village (Table 1, Panel C, column 4).

Table A9: Within-village Treatment Effects and Diffusion through Baseline Social Networks, without controls.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Inputs index (labor + mulch + fertilizers)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A			
Treatment	1.239	0.329	0.019	0.043	0.052	0.041
	[0.075]	[0.041]	[0.030]	[0.028]	[0.039]	[0.047]
Control mean	-0.000	0.000	-0.000	0.000	0.001	2.724
Observations	4622	4622	47618	47618	6157	6154
			Panel B			
Treatment	1.259	0.342	0.031	-0.053	-0.140	0.056
	[0.114]	[0.060]	[0.044]	[0.040]	[0.064]	[0.076]
Number of treatment friends	-0.013	0.002	-0.058	0.000	-0.048	-0.012
	[0.033]	[0.025]	[0.021]	[0.019]	[0.033]	[0.033]
Number of sample friends	0.024	0.044	0.041	-0.013	0.077	0.040
	[0.018]	[0.015]	[0.013]	[0.012]	[0.021]	[0.021]
Treatment X num. treatment friends	0.018	-0.018	0.026	0.042	0.030	-0.005
	[0.069]	[0.040]	[0.028]	[0.026]	[0.056]	[0.039]
Treatment X num. sample friends	-0.017	0.004	-0.017	0.005	0.048	-0.004
	[0.049]	[0.030]	[0.020]	[0.017]	[0.033]	[0.026]
Control mean	0.682	0.179	0.001	0.014	0.039	2.724
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
Observations	4622	4622	47618	47618	6157	6154

Notes: Standard errors clustered by household are in brackets. All specifications control for village and round fixed effects and use data from endline rounds, i.e. rounds 6 through 9 (see Appendix B). 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) represents the average of five variables, each standardized using the Control group mean and standard deviation: total household labor days, total non-household labor days, the share of coffee plots where the household applied mulch, kilograms of compost, and kilograms of NPK. 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.

Table A10: 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)
Num. 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]
Num. 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. All specifications control for 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 A11: Within-village Treatment Effects and Social Diffusion:
Audits Adoption and Leaf Health Index Components.

	Tree canopy has mulch	Dripline is weeded	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)
						Panel A					
Treatment	0.035	0.030	0.002	-0.004	-0.004	0.001	-0.018	0.001	0.017	0.003	0.007
	[0.005]	[0.005]	[0.005]	[0.003]	[0.005]	[0.005]	[0.005]	[0.002]	[0.004]	[0.004]	[0.004]
Control mean	0.549	0.583	0.330	0.911	0.434	0.437	0.408	0.041	0.273	0.194	0.152
Observations	47598	47453	47617	47616	47616	47617	47618	47613	47595	47603	47593
						Panel B					
Treatment	0.034	0.050	-0.004	0.012	-0.005	-0.012	-0.015	-0.003	-0.012	-0.033	-0.016
	[0.008]	[0.008]	[0.007]	[0.004]	[0.008]	[0.008]	[0.008]	[0.003]	[0.007]	[0.006]	[0.006]
Number of treatment friends	-0.022	-0.013	-0.010	0.004	-0.014	-0.030	-0.004	-0.013	0.002	0.001	0.002
	[0.004]	[0.004]	[0.003]	[0.002]	[0.004]	[0.004]	[0.004]	[0.002]	[0.003]	[0.003]	[0.003]
Number of sample friends	0.016	0.015	0.004	0.000	0.006	0.016	0.001	0.004	-0.009	-0.005	-0.003
	[0.002]	[0.002]	[0.002]	[0.001]	[0.002]	[0.002]	[0.002]	[0.001]	[0.002]	[0.002]	[0.002]
Treatment X num. treatment friends	0.020	0.020	0.001	-0.001	0.005	0.007	0.004	-0.000	0.018	0.000	0.011
	[0.004]	[0.005]	[0.004]	[0.003]	[0.005]	[0.005]	[0.005]	[0.002]	[0.004]	[0.004]	[0.003]
Treatment X num. sample friends	-0.011	-0.017	0.001	-0.005	-0.002	-0.000	-0.003	0.002	-0.001	0.011	0.001
	[0.003]	[0.003]	[0.003]	[0.002]	[0.003]	[0.003]	[0.003]	[0.001]	[0.003]	[0.003]	[0.002]
Control mean	0.564	0.592	0.333	0.908	0.429	0.434	0.400	0.041	0.283	0.193	0.153
Mean T friends	1.596	1.596	1.596	1.596	1.596	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	2.925	2.925	2.925	2.925	2.925
Observations	47598	47453	47617	47616	47616	47617	47618	47613	47595	47603	47593

Notes: Standard errors clustered by household in brackets. 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 Audits Adoption index in column 3 of Table 2. Columns 9-11 are the components of the Leaf Health index in column 4 of Table 2. All regressions control for baseline covariates selected via post-double selection LASSO.

Table A12: Within-village Treatment Effects and Social Diffusion: Inputs Index Components.

	Used mulch	Compost (Baskets)	NPK (KG)	Total HH labor days	Total non-HH labor days
	(1)	(2)	(3)	(4)	(5)
			Panel A		
Treatment	0.035	0.053	0.434	-0.061	-0.007
	[0.012]	[0.050]	[0.104]	[0.031]	[0.098]
Control mean	0.502	55.195	0.921	36.438	3.915
Observations	6149	6157	6157	6157	6157
			Panel B		
Treatment	0.043	0.164	-0.057	-0.066	-0.004
	[0.018]	[0.085]	[0.158]	[0.046]	[0.141]
Number of treatment friends	-0.016	-0.071	-0.244	-0.022	-0.240
	[0.009]	[0.053]	[0.083]	[0.026]	[0.090]
Number of sample friends	0.011	0.061	0.097	0.013	0.202
	[0.005]	[0.032]	[0.052]	[0.013]	[0.058]
Treatment X num. treatment friends	0.006	0.067	0.178	0.058	0.322
	[0.011]	[0.065]	[0.126]	[0.030]	[0.135]
Treatment X num. sample friends	-0.006	-0.052	0.027	-0.030	-0.178
	[0.008]	[0.045]	[0.086]	[0.018]	[0.080]
Control mean	0.510	59.592	1.110	35.757	4.307
Mean T friends	1.596	1.596	1.596	1.596	1.596
Mean tot. friends	2.925	2.925	2.925	2.925	2.925
Observations	6149	6157	6157	6157	6157

Notes: Standard errors clustered by household in brackets. 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. Each column is a component of the Inputs index in column 5 of Table 2. All regressions control for baseline covariates selected via post-double selection LASSO. Columns 2-5 show the results of **Poisson** regressions.

Table A13: Diffusion via Geographic Neighbors.

	Knowledge index	Self-reported adoption index	Adoption index, audits	Leaf health index, audits	Input quantities index (labor + fertilizers)	Yield (kg/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Treatment neighbors					
Num. treatment HH neighbors	0.019	0.016	-0.055	-0.007	0.001	-0.001
	[0.025]	[0.029]	[0.008]	[0.008]	[0.024]	[0.022]
Num. sample HH neighbors	-0.019	-0.005	0.020	-0.002	0.009	0.005
	[0.015]	[0.016]	[0.005]	[0.005]	[0.013]	[0.015]
Outcome mean	-0.004	-0.002	0.000	0.000	0.002	2.739
Observations	2122	2122	21390	21390	2826	2825
	Panel B: Treatment plot neighbors					
Num. treatment plot neighbors	0.024	0.028	-0.008	0.006	-0.004	0.039
	[0.018]	[0.016]	[0.005]	[0.004]	[0.014]	[0.012]
Num. sample plot neighbors	-0.015	-0.019	0.000	-0.006	0.009	-0.021
	[0.011]	[0.010]	[0.003]	[0.003]	[0.009]	[0.007]
Outcome mean	-0.001	0.001	-0.001	-0.001	0.003	2.734
Observations	2131	2131	21450	21450	2838	2837

Notes: Standard errors clustered by household are 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 B). 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) represents the average of five variables, each standardized using the Control group mean and standard deviation: total household labor days, total non-household labor days, the share of coffee plots where the household applied mulch, kilograms of compost, and kilograms of NPK. 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. All regressions control for baseline covariates selected via post-double selection LASSO.

Table A14: Impacts of Village Treatment Concentration on Audits Adoption and Leaf Health Index Components.

	Tree canopy has mulch	Dripline is completely weeded	Signs that tree was re-juvenated in past 6-7 years	Dead branches have been removed	Branches touching the ground have been removed	Centers opened (crossing branches removed)	Unwanted Suckers removed	Old and dry berries are removed	Berry borer traps used	Bark is smoothed or banded to control white borer	Few signs of yellowing on leaves	Few signs of curling on leaves	Few signs of leaf rust
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Panel A													
50% T in village	-0.055	-0.005	-0.002	0.044	-0.016	0.032	0.011	0.017	-0.001	-0.007	0.002	0.008	0.028
	[0.071]	[0.037]	[0.015]	[0.019]	[0.014]	[0.023]	[0.026]	[0.019]	[0.009]	[0.013]	[0.014]	[0.021]	[0.041]
	(0.526)	(0.916)	(0.865)	(0.050)	(0.545)	(0.263)	(0.660)	(0.500)	(0.908)	(0.613)	(0.889)	(0.696)	(0.567)
75% T in village	-0.017	-0.018	-0.007	0.020	-0.011	-0.018	-0.011	-0.001	-0.004	-0.014	-0.017	-0.015	0.008
	[0.064]	[0.043]	[0.014]	[0.023]	[0.022]	[0.027]	[0.029]	[0.023]	[0.011]	[0.011]	[0.013]	[0.018]	[0.032]
	(0.845)	(0.708)	(0.629)	(0.394)	(0.715)	(0.543)	(0.670)	(0.979)	(0.691)	(0.336)	(0.283)	(0.494)	(0.897)
Sample mean, 25% villages	0.602	0.607	0.851	0.311	0.918	0.425	0.437	0.397	0.022	0.049	0.165	0.197	0.268
p-value: 50%=75%	0.60 (0.65)	0.76 (0.79)	0.69 (0.71)	0.22 (0.27)	0.82 (0.88)	0.01 (0.06)	0.24 (0.38)	0.44 (0.46)	0.74 (0.77)	0.62 (0.61)	0.12 (0.19)	0.18 (0.25)	0.63 (0.69)
Observations	47598	47453	47603	47617	47616	47616	47617	47618	31853	47613	47593	47603	47595
Panel B													
Treated HH	-0.032	-0.018	-0.015	-0.011	-0.003	-0.016	-0.021	-0.032	-0.014	-0.000	-0.026	-0.011	-0.010
	[0.023]	[0.033]	[0.040]	[0.025]	[0.009]	[0.024]	[0.021]	[0.033]	[0.013]	[0.019]	[0.014]	[0.016]	[0.023]
	(0.306)	(0.520)	(0.457)	(0.673)	(0.790)	(0.558)	(0.388)	(0.172)	(0.105)	(1.000)	(0.159)	(0.586)	(0.684)
50% T in village	-0.079	-0.015	-0.005	0.042	-0.013	0.025	0.016	0.017	-0.014	-0.008	-0.011	0.005	0.012
	[0.064]	[0.037]	[0.019]	[0.018]	[0.017]	[0.022]	[0.030]	[0.024]	[0.011]	[0.014]	[0.017]	[0.024]	[0.041]
	(0.344)	(0.751)	(0.768)	(0.090)	(0.635)	(0.406)	(0.561)	(0.538)	(0.157)	(0.605)	(0.508)	(0.820)	(0.822)
75% T in village	-0.083	-0.070	-0.001	0.004	-0.010	-0.024	-0.048	-0.003	-0.012	-0.022	-0.037	-0.032	-0.013
	[0.060]	[0.043]	[0.020]	[0.019]	[0.021]	[0.034]	[0.043]	[0.027]	[0.011]	[0.010]	[0.015]	[0.021]	[0.038]
	(0.372)	(0.153)	(0.956)	(0.875)	(0.756)	(0.505)	(0.123)	(0.930)	(0.262)	(0.213)	(0.070)	(0.227)	(0.807)
Treatment X 50% T in village	0.064	0.029	0.013	0.010	-0.005	0.022	-0.002	0.015	0.032	0.003	0.041	0.013	0.038
	[0.033]	[0.042]	[0.045]	[0.032]	[0.014]	[0.026]	[0.028]	[0.043]	[0.014]	[0.021]	[0.018]	[0.025]	[0.027]
	(0.104)	(0.432)	(0.610)	(0.788)	(0.778)	(0.541)	(0.961)	(0.610)	(0.004)	(0.871)	(0.086)	(0.642)	(0.171)
Treatment X 75% T in village	0.108	0.082	0.002	0.028	0.000	0.018	0.063	0.024	0.020	0.011	0.044	0.030	0.035
	[0.026]	[0.036]	[0.045]	[0.038]	[0.013]	[0.033]	[0.041]	[0.040]	[0.016]	[0.020]	[0.016]	[0.017]	[0.026]
	(0.010)	(0.034)	(0.960)	(0.448)	(0.982)	(0.618)	(0.077)	(0.485)	(0.099)	(0.550)	(0.095)	(0.299)	(0.287)
Control mean, 25% villages	0.602	0.607	0.851	0.311	0.918	0.425	0.437	0.397	0.022	0.049	0.165	0.197	0.268
p-value: 50% T= 75% T	0.95 (0.97)	0.08 (0.27)	0.88 (0.86)	0.07 (0.23)	0.88 (0.91)	0.09 (0.17)	0.09 (0.04)	0.53 (0.52)	0.65 (0.89)	0.28 (0.41)	0.07 (0.25)	0.11 (0.16)	0.58 (0.62)
p-value: Treat x 50% T = Treat x 75% T	0.11 (0.26)	0.08 (0.12)	0.72 (0.64)	0.64 (0.60)	0.75 (0.75)	0.89 (0.93)	0.10 (0.05)	0.84 (0.79)	0.18 (0.29)	0.42 (0.60)	0.83 (0.90)	0.43 (0.48)	0.86 (0.90)
p: Treatment + Treatment x 50% T=0	0.14 (0.19)	0.62 (0.57)	0.92 (0.88)	0.95 (0.94)	0.52 (0.45)	0.56 (0.77)	0.17 (0.26)	0.54 (0.36)	0.00 (0.01)	0.74 (0.81)	0.14 (0.31)	0.95 (0.95)	0.01 (0.11)
p: Treatment + 50% T + Treatment x 50% T=0	0.53 (0.60)	0.94 (0.96)	0.64 (0.68)	0.05 (0.11)	0.12 (0.43)	0.15 (0.33)	0.81 (0.78)	0.99 (0.99)	0.70 (0.64)	0.66 (0.70)	0.83 (0.84)	0.80 (0.80)	0.33 (0.38)
p: Treatment + Treatment x 75% T=0	0.00 (0.02)	0.00 (0.02)	0.55 (0.49)	0.59 (0.51)	0.78 (0.81)	0.91 (0.93)	0.24 (0.10)	0.74 (0.71)	0.39 (0.43)	0.10 (0.35)	0.06 (0.32)	0.13 (0.37)	0.11 (0.24)
p: Treatment + 75% T + Treatment x 75% T=0	0.92 (0.94)	0.90 (0.90)	0.28 (0.37)	0.41 (0.39)	0.59 (0.67)	0.38 (0.51)	0.85 (0.84)	0.67 (0.68)	0.67 (0.54)	0.38 (0.48)	0.22 (0.27)	0.50 (0.55)	0.72 (0.84)
Observations	47598	47453	47603	47617	47616	47616	47617	47618	31853	47613	47593	47603	47595

Notes: Standard errors clustered by household in brackets, with randomization inference p-values in parentheses below those. Each panel concludes with p-values for various linear restriction tests, with those in parentheses corresponding to the randomization inference results. All specifications control for 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 Audits Adoption index in column 3 of Table 3. Columns 9-11 are the components of the Leaf Health index in column 4 of Table 3. All regressions control for baseline covariates selected via post-double selection LASSO.