Journal of Development Economics 182 (2026) 103733



Contents lists available at ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/devec



Regular article



Chat over coffee? Diffusion of agronomic practices and market spillovers in Rwanda☆

Esther Duflo [a], Daniel Keniston [b], Tavneet Suri [c], Céline Zipfel [d,*]

[a] Massachusetts Institute of Technology, United States of America
[b] Louisiana State University, United States of America
[c] Massachusetts Institute of Technology Sloan School of Management, United States of America
[d] Stockholm School of Economics, Sweden

ARTICLE INFO

JEL classification:
O12
Q13
Q16

Keywords:
Agricultural technology adoption
Networks
Spillovers

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. The program increased knowledge and self-reported adoption of the taught practices. It also reshaped farmers' social networks by creating new social ties primarily among co-trainees. At first glance, the intervention appeared modestly effective, with treated farmers exhibiting higher yields than control applicants within the same village, although this difference is not statistically significant. However, knowledge did not diffuse, and control farmers with more treatment friends at baseline reduced audited adoption and input use. Villages with high trainee concentrations showed suggestive evidence of negative spillovers, consistent with reallocation of labor and other shared inputs toward treated farmers, as formalized in a simple model in which social ties shape access to inputs. Relative declines in the yields of control farmers in these villages account for treatment-control differences, raising concerns both about this dissemination strategy and estimates that fail to consider potential negative spillovers.

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, 2021) 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, 2023). Indeed, recent evidence shows that social networks can form strategically around new information

☆ 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 and two anonymous Referees and the Editor for helpful feedback, as well as audiences at Stanford, the NOVAFRICA Conference in Development Economics, and the Swedish Conference in Economics. We gratefully acknowledge financial support from TechnoServe, United States of America (as part of their funding from the Gates Foundation) and the International Growth Centre.

* Corresponding author.
E-mail address: Celine.Zipfel@hhs.se (C. Zipfel).

https://doi.org/10.1016/j.jdeveco.2026.103733
Received 15 February 2025; Received in revised form 23 January 2026; Accepted 24 January 2026
Available online 1 April 2026
0304-3878/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

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 reallocation of inputs in limited supply toward trained farmers: soil inputs (fertilizers and mulch) and 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 1594 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.

To guide intuition on how changes in social networks may translate into program impacts, we first introduce a simple model in which the allocation of scarce inputs across farmers is shaped by social connections. The model is developed for labor, but its logic extends to other inputs that are commonly exchanged or shared among farmers, such as mulch and fertilizers. A key prediction of the model is that if training creates tighter social clusters among treated farmers, inputs are increasingly allocated within this group. As the village share of farmers assigned to the training rises, this reallocation tends to reduce access to inputs for control farmers, generating negative spillovers. Crucially, these effects arise even in the absence of learning spillovers, technology adoption, or productivity gains.

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

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 and coffee sales. 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 3000 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. Net revenues are constructed as gross revenues from coffee sales minus reported monetary expenditures on all inputs. 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 0.34 new friends gained by treatment farmers — a 21% increase over baseline — 0.32 were from the same training group. In contrast, the networks of control farmers appear to have been largely unaffected by the program. Thus, social networks became more tightly knit within the treatment group rather than across treatment and control farmers under 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 about 6% higher than those of control farmers, though this effect is not statistically significant. Estimated effects on net revenues are similar in magnitude and also insignificant. 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 likewise find no evidence of positive spillovers of the program for control farmers living near trained households or cultivating adjacent coffee plots.

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 19% 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 gains among treated farmers.

There is suggestive evidence that these apparent perverse effects on yields in the control group were caused by input reallocation 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 lower input use among control farmers. While the limited number of villages reduces precision, the regularity of this pattern is indicative of input reallocation from control to trained farmers as treatment intensity rises.

Further evidence supporting the labor channel comes from wage impacts. Daily wages increase with treatment intensity, but only for control group employers; wages paid by treated farmers remain largely unchanged. This asymmetric pattern is consistent with the program-induced tightening of social links within the treatment group reallocating hired labor toward treated farmers, rather than reflecting an increase in aggregate labor demand.

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. This pattern points to reciprocal labor arrangements concentrated among treatment farmers, potentially related to the sourcing and application of inputs such as mulch and fertilizers. Such arrangements align closely with the mechanism formalized in the model, in which training-induced network clustering reallocates labor and other shared inputs to treated farmers.

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 and labor, and no differences between treatment and control farmers. Input reallocation 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. Related, Emerick (2018) shows that farmer-to-farmer trading frictions limit the diffusion of a new rice seed variety in India and that some of these frictions (not all) can be attributed to the caste system. 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 4 kg 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 B 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. Theory

How might a reallocation of social network connections, such as the one induced by the program trainings, affect input markets and production? To guide intuition, we present a simple model of employer choice that captures some of the main features of the program. This model is, admittedly, abstract—our goal is to demonstrate how the mechanisms we propose might operate, rather than develop a detailed model for structural estimation. Formal results are shown in Appendix A.

Consider a village with M coffee employers (indexed j) and N potential workers (indexed i) on a given day. Within each group, a share π have received the training. Workers choose either coffee labor, earning wage w and utility $w + \nu_{ij}$ for each "job" performed, or an "outside option" with fixed utility $\ln V$.

We interpret the ν term to capture the unobserved "closeness" of worker and employer, both in terms of geographical proximity, but also social network connections. Specifically, we assume that if both worker i and employer j have been trained together, the mean of the ν distribution shifts upwards by $\ln s$ $(s > 1)$, capturing the new social network links.

Assuming an exogenous, village level wage w, and that $\nu_{ij} \sim EV1$, it is straightforward to show that untrained workers work for all employers with equal probability. However, trained workers will be particularly attracted to trained employers, and will work for them with probability $e^w s / (\pi M e^w s + (1-\pi) M e^w + V)$.

Let the relative attractiveness of the outside option be $k = V / (Me^w) > 0$. Labor supply to each farm owned by trained (L_t) and untrained (L_c) employers is then,

$$L_t = \frac{N}{M}\left[\frac{\pi s}{\pi s + 1 - \pi + k} + \frac{1-\pi}{1+k}\right],\ L_c = \frac{N}{M}\left[\frac{\pi}{\pi s + 1 - \pi + k} + \frac{1-\pi}{1+k}\right]$$

Clearly as s increases, trained workers are increasingly attracted to work for trained employers. More subtly, the share of the village trained π also affects relative labor supply to trained and control employers. This effect is clearest for the untrained employers, who lose employees both because the share of "neutral" untrained workers decreases, and because trained workers (some of whom still work for control employers) are more likely to find a trained employer with whom they have a strong connection. The effect is less unambiguous for treatment employers, but it can be shown that if the outside option is relatively attractive, compared to the social network effect of training, then L_t will increase with π.[2] This re-allocation of labor then implies a decrease in output for control farmers and a possible increase for treatment farmers as the share of trained coffee farmers increases.

While these results depend on a single village wage set outside the model, a natural extension is to consider a setting in which trained and untrained farmers endogenously set different wages. While a full comparative static analysis of this case is algebraically dense (see Appendix A), the results are intuitive. If farmers' production functions are concave in labor, trained farmers pay lower wages relative to control farmers as they face a lower elasticity of labor supply from their newly strengthened social networks. As the share of trained individuals π increases, for certain parameter values the control farmers may further raise wages and treatment farmers lower them.[3] This partially, though not fully, offsets the effect of the partial training on labor allocations and output.

In summary, a training program that creates social links within a fraction of a community may have real effects on labor allocation, wages and output if individuals prefer to work with others with whom they share social network ties. The model focuses on labor, but its results might apply to other inputs, such as compost and mulch, that are traded within networks of farmers who share social connections. Finally, we emphasize that these effects are completely independent of any real productivity effects of the training—which are absent in our proposed model.

4. Experimental design and data

4.1. 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 1594 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.[4]

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 A1, 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

[2] The technical condition that ensures a positive relationship is $\pi < (1+k)/\left(1+\sqrt{s}\right)$. Taking the model seriously, the share of workers employed outside the coffee sector is much greater than the social network effect on employment, implying $k > \sqrt{s}$. Thus this condition is likely to hold for any value of $\pi \in [0,1]$.

[3] In general, model predictions are ambiguous on the effect of the share trained on wages; these results hold unambiguously in the extreme case of a labor market segmented between treatment and control groups.

[4] Note that since the 1594 farmers who signed up for the training program covered only 30 villages, for power reasons, we were only able to create three treatment concentrations, and opted not to create a 0% treatment concentration arm.

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.

4.2. 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 C, 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 3000 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.

4.3. Empirical specifications

Our empirical analysis includes five main econometric specifications. First, we estimate the impact of the agronomy training program on farmers' social networks by applying the following OLS model to the baseline and endline rounds of social networks data:

$$Friends_{ijt} = \beta_1 Treat_{ij} \times Post_t + \beta_2 Control_{ij} \times Post_t + \beta_3 Treat_{ij} + \gamma_j + \epsilon_{ijt}, \quad (1)$$

where $Friends_{ijt}$ denotes the number of social connections who grow coffee ("coffee friends") reported by household i in village j at survey round t; $Treat_{ij}$ is an indicator equal to one if household i was randomly assigned to receive the agronomy training, $Control_{ij} = 1 - Treat_{ij}$, and $Post_t$ is an indicator for the endline round. β_1 captures the change in social connections for treatment households by the end of the program relative to their baseline and β_2 captures this corresponding change for control households. We cluster standard errors at the household level.

We run this specification for the following outcomes: (i) total number of coffee friends, (ii) friends within the village versus outside the village, (iii) friends by treatment status (treatment, control, or non-sample farmers), and (iv) friends belonging to the same training group.

Next, we use within-village random assignment of households to estimate mean treatment–control differences in outcomes within villages, using Eq. (2):

$$y_{ijt} = \beta Treat_{ij} + \gamma_j + \delta_t + \epsilon_{ijt} \quad (2)$$

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.[5] Here, too, we cluster standard errors at the household level.[6]

As can be seen in Appendix Table A2, we find balance between treatment and control households across a wide variety of baseline outcomes (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).

Specification (2) compares treatment and control farmers within the same village. For β to be interpreted as a treatment effect, we

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

[6] For brevity, we do not include randomization inference p-values in Tables 1 and 2 like we do for the results of regressions on village treatment concentrations (Tables 3–7) and only report standard errors clustered at the household level, i.e., the level of treatment in these within-village regressions.

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 with our analysis of social diffusion within villages in specification (3) and of the effects of treatment concentration at the village level in specifications (4) and (5).

To assess whether any program impacts diffuse through existing social networks, we leverage the exogenous variation generated by the experiment in the number of treatment friends, controlling for the total number of RCT-sample friends (i.e. who entered the treatment lottery). We examine two types of networks: baseline "coffee friends" and neighbors. The results presented in Section 5 focus on baseline friends, but the results for diffusion through neighbors (household and plot) are broadly consistent and reported in Appendix Table A14. We interact these network variables with the household's own treatment status to examine whether the effects differ between the control and treatment groups. The specification is as follows:

$$y_{ijt} = \alpha Treat_{ij} + \beta_0 NumTreatFriends_{ij} + \delta_0 NumFriends_{ij} + \beta_1 Treat_{ij} \times NumTreatFriends_{ij} + \delta_1 Treat_{ij} \times NumFriends_{ij} + \gamma_j + \delta_t + \epsilon_{ijt} \quad (3)$$

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 Eq. (2), 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.

Finally, we examine how program impacts vary with experimentally assigned village-level treatment concentration, using two additional specifications:

$$y_{ijt} = \beta_1 50\%T_j + \beta_2 75\%T_j + \delta_t + \epsilon_{ijt} \quad (4)$$

$$y_{ijt} = \alpha Treat_{ij} + \beta_1 50\%T_j + \beta_2 75\%T_j + \gamma_1 Treat_{ij} \times 50\%T_j + \gamma_2 Treat_{ij} \times 75\%T_j + \delta_t + \epsilon_{ijt} \quad (5)$$

Here, we cluster standard errors at the village level, since the experimental variation is the treatment intensity assigned at the village level.[7] Given the small number of villages, we report p-values for the exact null hypothesis, using a randomization inference algorithm that takes the experimental design into account.[8] To address any baseline imbalances across village groups (Appendix Tables A3– A6), we include baseline covariates selected using post-double selection LASSO in all regressions.

Throughout the paper, we estimate program impacts—both within-village treatment effects and effects of village treatment intensities—on yields and other non-negative, right-skewed outcomes (yields in Table 3, fertilizer quantities in Table 4, and labor days in Table 5) using Poisson pseudo-maximum likelihood, following Wooldridge (1999, 2010).[9] For binary outcomes, indices, and outcomes that can take negative values, such as net revenues, we instead use OLS.

[7] For brevity, we do not report these standard errors and only show randomization inference p-values in these tables.

[8] Randomization inference involves 1000 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).

[9] We show robustness to estimating these outcomes using OLS in Appendix Tables A10, A16 and A17.

5. Results

5.1. 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 or outside the farmer's village) and for friends with different treatment statuses.

Changes in the composition of these social networks can influence program spillovers in at least two ways. First, they may affect social learning, for instance if interactions between treatment and control farmers decline or if control farmers actively seek out treated peers. Second, social ties can shape economic interactions more directly, such as through preferential input sharing or labor exchange between connected farmers, as formalized in our model in Section 3.

Table 1 reports estimates from specification (1). Panel A, which reports effects on total friends, shows that treatment households gain 0.336 friends who were also chosen for training (a 21% 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. We find no evidence that control farmers actively formed new connections with treated individuals in order to benefit from their increased coffee knowledge. Taken together, these findings imply that the program increased the relative density of ties among treated farmers, even though the absolute number of links to control farmers did not fall. In relative terms, this pattern hints that any spillovers operating through social networks—whether via information or more direct economic interactions—may be more likely to operate within the treatment group.

5.2. Program impacts on key outcomes

Within-Village Treatment Effects and Social Diffusion through Baseline Networks. Panel A of Table 2 reports the within-village

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	1.567	1.316	1.446	4.328
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	1.356	1.116	0.939	3.410
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	0.211	0.200	0.507	0.918
Observations	3156	3156	3156	3156
	Panel D: Friends in same training group			
Treat × Post	0.319***			
	[0.0629]			
Baseline mean	0.755			
Observations	1678			

Notes: Standard errors clustered by household in brackets. For brevity, we do not report randomization-inference p-values in this table, unlike regressions using village-level treatment concentrations (Tables 3–7). 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 C). 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).

treatment effects estimated using specification (2) for measures of knowledge and adoption of agronomic practices, input use, yields, and net revenues. 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 (estimated using a Poisson model) are approximately 6% higher in the treatment group compared to the control, although this estimate is not statistically significant. Column 7 likewise shows a positive but statistically insignificant within-village treatment effect on net revenues, corresponding to an increase of about 6 RWF per tree (about 5% of the control mean).

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 or yields.[10]

Panel B reports diffusion effects through baseline social networks, estimated using specification (3). Columns 1 and 2 present these estimates for 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 (column 4). Additional results reported in Appendix Table A11 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.[11] Appendix Table A12, which disaggregates the

[10] 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.

[11] Appendix Table A14, which shows similar results on neighbors of treatment farmers to Table 2, Panel B, also include negative effects on the audits adoption index.

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	Input quantities index (labor + mulch + fertilizers)	Yield (kg/tree)	Net revenues (RWF/tree)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A						
Treatment	1.235	0.323	0.040	0.031	0.053	0.061	5.225
	[0.060]	[0.036]	[0.010]	[0.010]	[0.028]	[0.038]	[5.092]
Control mean	−0.000	0.000	0.000	0.000	−0.000	0.586	126.475
Observations	4666	4666	47 889	47 878	6252	6221	6219
	Panel B						
Treatment	1.266	0.330	0.057	−0.072	−0.132	0.088	9.692
	[0.091]	[0.051]	[0.015]	[0.016]	[0.051]	[0.064]	[8.053]
Number of treatment friends	−0.025	−0.008	−0.051	0.003	−0.037	−0.063	−8.949
	[0.029]	[0.024]	[0.007]	[0.007]	[0.022]	[0.027]	[3.847]
Number of sample friends	0.009	0.023	0.032	−0.019	0.012	0.070	9.684
	[0.017]	[0.014]	[0.005]	[0.005]	[0.015]	[0.017]	[2.510]
Treatment × num. treatment friends	0.040	−0.005	0.012	0.035	0.061	0.054	6.903
	[0.059]	[0.036]	[0.010]	[0.009]	[0.039]	[0.036]	[4.606]
Treatment × num. sample friends	−0.032	0.000	−0.011	0.012	0.031	−0.040	−5.296
	[0.042]	[0.025]	[0.007]	[0.007]	[0.026]	[0.026]	[3.223]
Control mean	−0.000	0.000	0.000	0.000	−0.000	0.586	126.475
Mean T friends	1.583	1.583	1.583	1.583	1.583	1.583	1.583
Mean tot. friends	2.896	2.896	2.896	2.896	2.896	2.896	2.896
p: Num. T friends + T × num. T friends = 0	0.651	0.724	0.005	0.000	0.305	0.607	0.711
Unit of observation	HH-round	HH-round	Tree-round	Tree-round	HH-round	HH-round	HH-round
Survey rounds	6, 8, 9	6, 8, 9	6, 8, 9	6, 8, 9	6, 7, 8, 9	6, 7, 8, 9	6, 7, 8, 9
Observations	4666	4666	47 889	47 878	6252	6221	6219

Notes: Standard errors clustered by household are in brackets. For brevity, we do not report randomization-inference p-values in this table, unlike regressions using village-level treatment concentrations (Tables 3–7). All specifications control for village and round fixed effects. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix C). 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. Column (7) reports the results of an OLS regression on net revenues per tree, calculated using self-reported coffee revenues from sales to the local washing station and private buyers net of reported expenditures on hired labor and on purchased inputs (fertilizer, pesticides, and mulch when applicable). All regressions control for baseline covariates selected via post-double selection LASSO.

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 A13 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. Columns 6 and 7 show that additional treatment friends also lead to lower yields and net revenues for the average farmer in the control group. In column 6, each additional treated friend reduces yields by about 6%, while column 7 shows a corresponding decline in net revenues, and both effects are statistically significant.

The fact that control farmers apply fewer best practices (or apply them less intensively), use fewer inputs, and have lower yields and net revenues 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. At the bottom of the table, we report the p-values for the test of the average effect of an additional treatment friend for treatment households, evaluated at the mean number of treatment friends in the sample.

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 trained friend leads to a 0.034 SD increase in the leaf health index relative to the effect for a control farmer. The linear combination test indicates that the overall effect of an additional trained friend for treatment farmers is significantly different from zero ($p = 0.000$).

In contrast, for audited adoption (column 3) and input use (column 5), while the positive interaction with treatment offsets much of the negative effect of additional trained friends for control farmers, the combined effect for treated farmers remains statistically indistinguishable from zero ($p = 0.342$ and $p = 0.285$). For yields (column 6), the interaction with treatment is positive, indicating that the negative effect of exposure to treated friends observed among control farmers does not extend to treated farmers. However, the linear combination test shows that the overall effect of an additional trained friend for treated farmers is statistically indistinguishable from zero ($p = 0.648$). A similar pattern is observed for net revenues (column 7).

The intervention may also have influenced farmers' marketing and sales of coffee. Interactions with trainers affiliated with the local coffee-processing wet mill, both directly and via social networks, may have encouraged farmers to sell their output to this mill. Furthermore, conversations with other farmers at training sessions may also have allowed attendees to identify sellers offering higher prices for coffee, either at the wet mill or elsewhere. We test the importance of both these hypotheses by examining the impact of the treatment on sales to the wet mill and average prices received.

The results, in Appendix Table A15, show that the training caused a 1–2 percentage point increase in the probability of a farmer selling to the wet mill, and a further 1.5–2.6 percentage point increase with each additional trained friend, with this social network effect highly statistically significant. However, these increased sales did not result in higher prices for treated farmers, consistent with statements from the wet mill management that their pricing strategy was simply to follow market prices. Overall, this channel does not appear to have substantially affected farmers' net revenues or other outcomes.

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 pattern is consistent either with effects that were too small to spread, or with information flows being more likely to occur within the newly formed network of co-trainees (Table 1). 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.[12] We now turn to this question.

Spillovers by Village Treatment Concentration. Panel A of Table 3 shows the aggregate effects of treatment concentration, estimated using specification (4) while Panel B shows the interactions of the village-level concentrations with treatment status (specification (5)). 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.

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 ($p = 0.216$, row [1]).

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 impacts of treating 50% or 75% of farmers on yields and net revenues are negative and insignificant (columns 6 and 7), 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. In both columns, the coefficients on the 50% and 75% village treatment share dummies are negative and statistically insignificant, consistent with the absence of positive spillovers to the control group on these outcomes in Table 2.

While the average within-village treatment effects reported in columns 3–7 of Table 2, Panel A are all positive, the estimates in columns 3–7 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 yields and net revenues estimates in columns 6 and 7 are even negative, although they do not survive randomization inference ($p = 0.483$ and $p = 0.692$, respectively). Second, the exact opposite is true in 75% treatment villages: in these areas, treatment households use significantly ($p = 0.012$) more inputs and have higher yields ($p = 0.047$).[13] Third, there is no significant difference between treatment households in 75% villages and control households in 25% villages.[14] 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 (column 3; $p = 0.087$), leaf health (column 4), and net revenues (column 7), which follow a similar pattern, though the relevant tests of treatment effects in 75% villages for leaf health and net revenues are not significant ($p = 0.169$ for leaf health and $p = 0.205$ for net revenues).

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 baseline friends (or neighbors) in the treatment group. In both cases, greater exposure to treated farmers is associated with worse outcomes for control farmers, leading apparent treatment effects to emerge mechanically in within-village comparisons.

Appendix Table A18, which disaggregates the estimates in Table 3 for the audits and leaf health indices, shows that spillovers are most pronounced for practices that require applying soil inputs (mulch and fertilizer) and substantial labor inputs. In particular, the treatment–control gap differs significantly between 25% and 75% treatment villages for mulching and weeding (row [6] p-values of 0.016 and 0.017, respectively). Measures of tree nutrition also display a consistent directional pattern across outcomes—yellowing of leaves, curling of leaves, and leaf rust—pointing to worse conditions for control farmers in high-intensity villages, although these differences are not statistically significant.

Finally, we examine whether the program affected tree counts in ways that would parallel the patterns observed for inputs, audited adoption, yields, and net revenues.[15] Appendix Table A19 shows no evidence of reductions in either total or productive tree counts associated with treatment or village-level treatment intensity. If anything, point estimates suggest slightly higher tree counts for treatment households compared to control households in 75% villages. This pattern contrasts with our results for yields and net revenues, and supports the interpretation that the yield effects we document reflect changes in productivity per tree.

5.3. Inputs reallocation

Why might higher treatment densities result in reduced input use, poorer leaf health, and lower yields for control farmers? As formalized by our model in Section 3, the results of Tables 1–3 taken together hint at an inputs reallocation from treatment farmers to control farmers in villages with higher concentration of trained farmers, driven by farmers' tendency to share or source inputs within their social networks.

To better understand the program's effects on input use, we 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 2–6 report effects on each outcome separately.

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

[13] Row [6] p-values report the test of this hypothesis.

[14] Row [7] p-values report the test of this hypothesis.

[15] Note that over our study period, changes in yields per tree could only arise through plot transactions or through tree rejuvenation practices that temporarily render trees non-productive, since newly planted coffee trees take at least three years to reach maturity.

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)	Net revenues (RWF/tree)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A						
50% T in village	0.187	0.038	−0.008	0.056	−0.022	−0.102	−18.869
	(0.402)	(0.707)	(0.899)	(0.179)	(0.838)	(0.481)	(0.294)
75% T in village	0.687	0.163	−0.046	−0.017	0.003	−0.053	−10.759
	(0.000)	(0.098)	(0.554)	(0.708)	(0.982)	(0.741)	(0.598)
Sample mean, 25% villages	0.056	0.037	0.034	−0.000	0.037	0.644	142.101
[1] p-value: 50% = 75%	0.008	0.216	0.615	0.051	0.806	0.738	0.651
Observations	4666	4666	47 889	47 878	6252	6221	6219
	Panel B						
Treatment	1.133	0.289	−0.053	−0.049	−0.011	−0.063	−4.747
	(0.000)	(0.000)	(0.441)	(0.422)	(0.892)	(0.483)	(0.692)
50% T in village	−0.083	−0.037	−0.035	0.019	−0.037	−0.142	−22.281
	(0.747)	(0.748)	(0.668)	(0.713)	(0.769)	(0.327)	(0.223)
75% T in village	−0.108	−0.057	−0.141	−0.088	−0.162	−0.210	−23.852
	(0.706)	(0.649)	(0.095)	(0.168)	(0.160)	(0.186)	(0.262)
Treatment × 50% T in village	−0.003	0.011	0.080	0.100	0.036	0.100	9.154
	(0.983)	(0.913)	(0.352)	(0.209)	(0.770)	(0.383)	(0.546)
Treatment × 75% T in village	0.314	0.103	0.161	0.126	0.229	0.234	20.512
	(0.203)	(0.353)	(0.087)	(0.138)	(0.071)	(0.066)	(0.239)
Control mean, 25% villages	0.056	0.037	0.044	0.005	0.023	0.657	144.104
[2] p-value: 50% T = 75% T	0.926	0.881	0.254	0.086	0.291	0.682	0.940
[3] p-value: Treat × 50% T = Treat × 75% T	0.117	0.344	0.327	0.733	0.084	0.269	0.493
[4] p: Treatment + Treatment × 50% T = 0	0.000	0.000	0.620	0.296	0.753	0.637	0.659
[5] p: Treatment + 50% T + Treatment × 50% T = 0	0.000	0.009	0.920	0.163	0.909	0.469	0.333
[6] p: Treatment + Treatment × 75% T = 0	0.000	0.000	0.087	0.169	0.012	0.047	0.205
[7] p: Treatment + 75% T + Treatment × 75% T = 0	0.000	0.000	0.684	0.826	0.615	0.823	0.691
Unit of observation	HH-round	HH-round	Tree-round	Tree-round	HH-round	HH-round	HH-round
Survey rounds	6, 8, 9	6, 8, 9	6, 8, 9	6, 8, 9	6, 7, 8, 9	6, 7, 8, 9	6, 7, 8, 9
Observations	4666	4666	47 889	47 878	6249	6221	6219

Notes: Standard errors are clustered at the village level. For brevity, we do not report these and only show randomization inference p-values in parentheses. Each panel concludes with randomization inference p-values for various linear restriction tests. All specifications control for round fixed effects. All columns use data from endline rounds, i.e. rounds 6 through 9 (see Appendix C). 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. Column (7) reports the results of an OLS regression on net revenues per tree, calculated using self-reported coffee revenues from sales to the local washing station and private buyers net of reported expenditures on hired labor and on purchased inputs (fertilizer, pesticides, and mulch when applicable). All regressions control for baseline covariates selected via post-double selection LASSO.

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 ($p = 0.001$). The treatment-control differences in 75% villages are also larger than in 50% villages ($p = 0.057$ 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 4 and 6) 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 ($p = 0.052$ for the 75% interaction). The treatment effect is statistically significant in 75% villages ($p = 0.012$). This finding aligns with the audit index components results in Appendix Table A18, which show significantly more mulching under tree canopies in the treatment group compared to the control in 75% villages.

Columns 3–6 focus on compost and fertilizers. Columns 3 and 5 present the results of OLS regressions on extensive margin outcomes of fertilizer use. Columns 4 and 6 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 extensive margin (column 3). However, column 4 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 marginally significant in 75% villages ($p = 0.063$).

Columns 5 and 6 examine NPK use. Similarly to compost, we find no evidence of differential effects for NPK at the extensive margin (column 5). However, column 6 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 ($p = 0.057$).

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

Table 4
Impacts of village treatment concentration on soil inputs use.

	Soil inputs index	Used mulch	Used compost	Compost (Baskets/tree)	Used NPK	NPK (KG/tree)
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A					
50% T in village	−0.001	−0.029	0.046	−0.058	−0.006	0.158
	(0.995)	(0.731)	(0.128)	(0.711)	(0.704)	(0.471)
75% T in village	−0.036	−0.041	0.004	0.044	−0.012	0.052
	(0.676)	(0.634)	(0.906)	(0.772)	(0.494)	(0.812)
Sample mean, 25% villages	0.048	0.540	0.508	56.559	0.073	0.943
[1] p-value: 50% = 75%	0.630	0.869	0.155	0.481	0.745	0.605
Observations	6223	6214	6223	6220	6223	6220
	Panel B					
Treatment	0.010	−0.006	0.021	−0.174	0.015	0.282
	(0.877)	(0.818)	(0.510)	(0.378)	(0.245)	(0.457)
50% T in village	−0.019	−0.043	0.051	−0.105	−0.007	0.140
	(0.834)	(0.619)	(0.112)	(0.553)	(0.678)	(0.619)
75% T in village	−0.191	−0.101	−0.009	−0.280	−0.020	−0.440
	(0.047)	(0.207)	(0.829)	(0.175)	(0.312)	(0.240)
Treatment × 50% T in village	0.032	0.032	−0.021	0.175	−0.005	−0.078
	(0.725)	(0.419)	(0.620)	(0.515)	(0.779)	(0.858)
Treatment × 75% T in village	0.198	0.085	0.004	0.533	0.001	0.431
	(0.043)	(0.052)	(0.942)	(0.064)	(0.964)	(0.421)
Control mean, 25% villages	0.034	0.541	0.499	57.071	0.069	0.855
[2] p-value: 50% T = 75% T	0.075	0.467	0.115	0.382	0.514	0.136
[3] p-value: Treat × 50% T = Treat × 75% T	0.057	0.199	0.523	0.138	0.750	0.307
[4] p: Treatment + Treatment × 50% T = 0	0.454	0.313	0.981	0.994	0.347	0.497
[5] p: Treatment + 50% T + Treatment × 50% T = 0	0.821	0.817	0.138	0.577	0.887	0.213
[6] p: Treatment + Treatment × 75% T = 0	0.001	0.012	0.434	0.063	0.221	0.057
[7] p: Treatment + 75% T + Treatment × 75% T = 0	0.837	0.781	0.635	0.635	0.825	0.308
Observations	6223	6214	6223	6220	6223	6220

Notes: Standard errors clustered by village are in brackets. For brevity, we do not report these but show randomization inference p-values in parentheses. Each panel concludes with randomization inference p-values for various linear restriction tests. All specifications control for round fixed effects and use data from endline rounds, i.e. rounds 6 through 9 (see Appendix C). Columns 1, 2, 3, and 5 show the results of OLS regressions. Columns 4 and 6 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	0.006	−0.110	−0.071	−0.233	0.052	−0.098	0.073	0.118	−0.071	−0.222	−0.015	−0.216
	(0.953)	(0.534)	(0.534)	(0.536)	(0.696)	(0.903)	(0.441)	(0.335)	(0.508)	(0.545)	(0.904)	(0.344)
75% T in village	0.026	−0.000	−0.049	0.157	0.011	0.475	0.125	0.087	−0.036	−0.124	0.039	0.044
	(0.720)	(1.000)	(0.708)	(0.693)	(0.926)	(0.332)	(0.090)	(0.505)	(0.743)	(0.722)	(0.758)	(0.861)
Sample mean, 25% villages	34.237	4.212	3.795	0.356	6.262	0.393	7.113	1.208	9.348	1.594	7.719	0.660
[1] p-value: 50% = 75%	0.788	0.517	0.847	0.243	0.710	0.091	0.589	0.795	0.723	0.800	0.653	0.228
Observations	6252	6252	6252	6252	6252	6252	6252	6252	6252	6252	6252	6252
	Panel B											
Treatment	−0.114	0.213	−0.118	0.119	0.145	0.137	−0.099	0.597	−0.192	−0.220	−0.243	0.367
	(0.124)	(0.468)	(0.344)	(0.826)	(0.280)	(0.857)	(0.303)	(0.059)	(0.028)	(0.500)	(0.005)	(0.345)
50% T in village	−0.028	0.019	−0.107	−0.090	0.131	−0.024	0.060	0.385	−0.151	−0.181	−0.085	−0.015
	(0.736)	(0.929)	(0.365)	(0.826)	(0.365)	(0.968)	(0.569)	(0.033)	(0.161)	(0.652)	(0.509)	(0.957)
75% T in village	0.066	−0.271	−0.037	−0.087	0.202	0.120	0.185	0.063	−0.040	−0.856	0.023	0.078
	(0.476)	(0.336)	(0.819)	(0.881)	(0.231)	(0.839)	(0.086)	(0.792)	(0.747)	(0.052)	(0.863)	(0.835)
Treatment × 50% T in village	0.122	−0.355	0.129	−0.356	−0.220	−0.205	0.074	−0.740	0.257	0.025	0.264	−0.567
	(0.189)	(0.378)	(0.430)	(0.625)	(0.208)	(0.848)	(0.536)	(0.043)	(0.018)	(0.952)	(0.011)	(0.276)
Treatment × 75% T in village	0.018	0.213	0.060	0.233	−0.351	0.367	−0.019	−0.306	0.132	1.063	0.183	−0.262
	(0.860)	(0.581)	(0.740)	(0.799)	(0.072)	(0.719)	(0.875)	(0.463)	(0.284)	(0.041)	(0.106)	(0.620)
Control mean, 25% villages	34.690	3.863	3.854	0.345	5.865	0.363	7.161	0.973	9.710	1.612	8.100	0.570
[2] p-value: 50% T = 75% T	0.293	0.319	0.620	0.994	0.664	0.788	0.285	0.192	0.368	0.122	0.411	0.805
[3] p-value: Treat × 50% T = Treat × 75% T	0.267	0.128	0.668	0.413	0.453	0.497	0.420	0.201	0.262	0.003	0.403	0.569
[4] p: Treatment + Treatment × 50% T = 0	0.891	0.525	0.913	0.608	0.482	0.888	0.739	0.513	0.350	0.408	0.731	0.550
[5] p: Treatment + 50% T + Treatment × 50% T = 0	0.808	0.556	0.418	0.476	0.690	0.850	0.696	0.181	0.423	0.338	0.621	0.490
[6] p: Treatment + Treatment × 75% T = 0	0.178	0.123	0.644	0.526	0.119	0.368	0.186	0.270	0.505	0.005	0.419	0.779
[7] p: Treatment + 75% T + Treatment × 75% T = 0	0.682	0.453	0.433	0.540	0.978	0.200	0.510	0.028	0.365	0.970	0.777	0.512
Observations	6252	6252	6252	6252	6252	6252	6252	6252	6252	6252	6252	6252

Notes: Standard errors clustered by village. For brevity, we do not report those and only show randomization inference p-values in parentheses. Each panel concludes with randomization inference p-values for various linear restriction tests. All specifications control for round fixed effects and use data from endline rounds, i.e. rounds 6 through 9 (see Appendix C). All columns show the results of **Poisson** regressions and control for baseline covariates selected via post-double selection LASSO.

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.[16]

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. None of these estimates are significant under randomization inference. 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 ($p = 0.123$). And again, this result seems driven by a negative effect on the control group.

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 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 ($p = 0.005$). This effect is driven by fewer workers being hired by control households ($p = 0.052$) and more workers being hired by treatment households ($p = 0.041$).

Fig. 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, in higher-intensity villages, some non-household labor is reallocated toward treatment farms, consistent with the model's prediction that the intervention reshapes within-village labor networks. 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 changes in the allocation of inputs caused differences in their prices across treatment groups. 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.[17]

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.

Column 1, Panel A shows that daily wage rates are, on average, about 16 RWF higher in 50% villages and 28 RWF higher in 75% villages compared to 25% villages. These correspond to a 2.5% and 4.3% increase relative to the mean daily rate in 25% villages, respectively. Column 4 shows that daily wage rates for weeding are, on average, about 12 RWF higher in 50% villages and 33 RWF higher in 75% villages ($p = 0.082$), corresponding to a 1.8% and 5% 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.

Panel B of Table 6 reveals a consistent pattern of heterogeneity in wage effects. Treated farmers pay similar, perhaps slightly higher, wages in 25% trained villages, but as the share of trained farmers increases the wage paid by control group farmers increases as well, with significant and substantial effects in several tasks in 75% trained villages. In contrast, there is no significant increase in wages paid by trained employers as the share trained grows. The Treatment X 75% interaction coefficient is negative for all tasks, and significant overall and for harvesting.

These wage patterns corroborate the idea that the program reallocated hired labor toward trained farmers through social networks. 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. Although wages in these village are likely determined by a wide range of factors outside the scope of this paper—tradition, dynamic contracts, asymmetric information, etc.—perhaps explaining the limited 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, and this tendency becomes stronger with higher treatment intensity ($p = 0.140$ for the treatment effect in 75% villages). In contrast, the control group spends considerably less time working on other farmers' coffee plots ($p = 0.117$) 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.[18] This pattern

[16] 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.

[17] 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).

[18] Appendix Table A20 presents within-village regressions of household labor supply (using the same three outcomes as in Table 7) on baseline networks. Consistent with the village-level results, labor supply outside the household is higher among treated farmers with more baseline treatment friends, particularly when both the worker and their network contacts are treated.

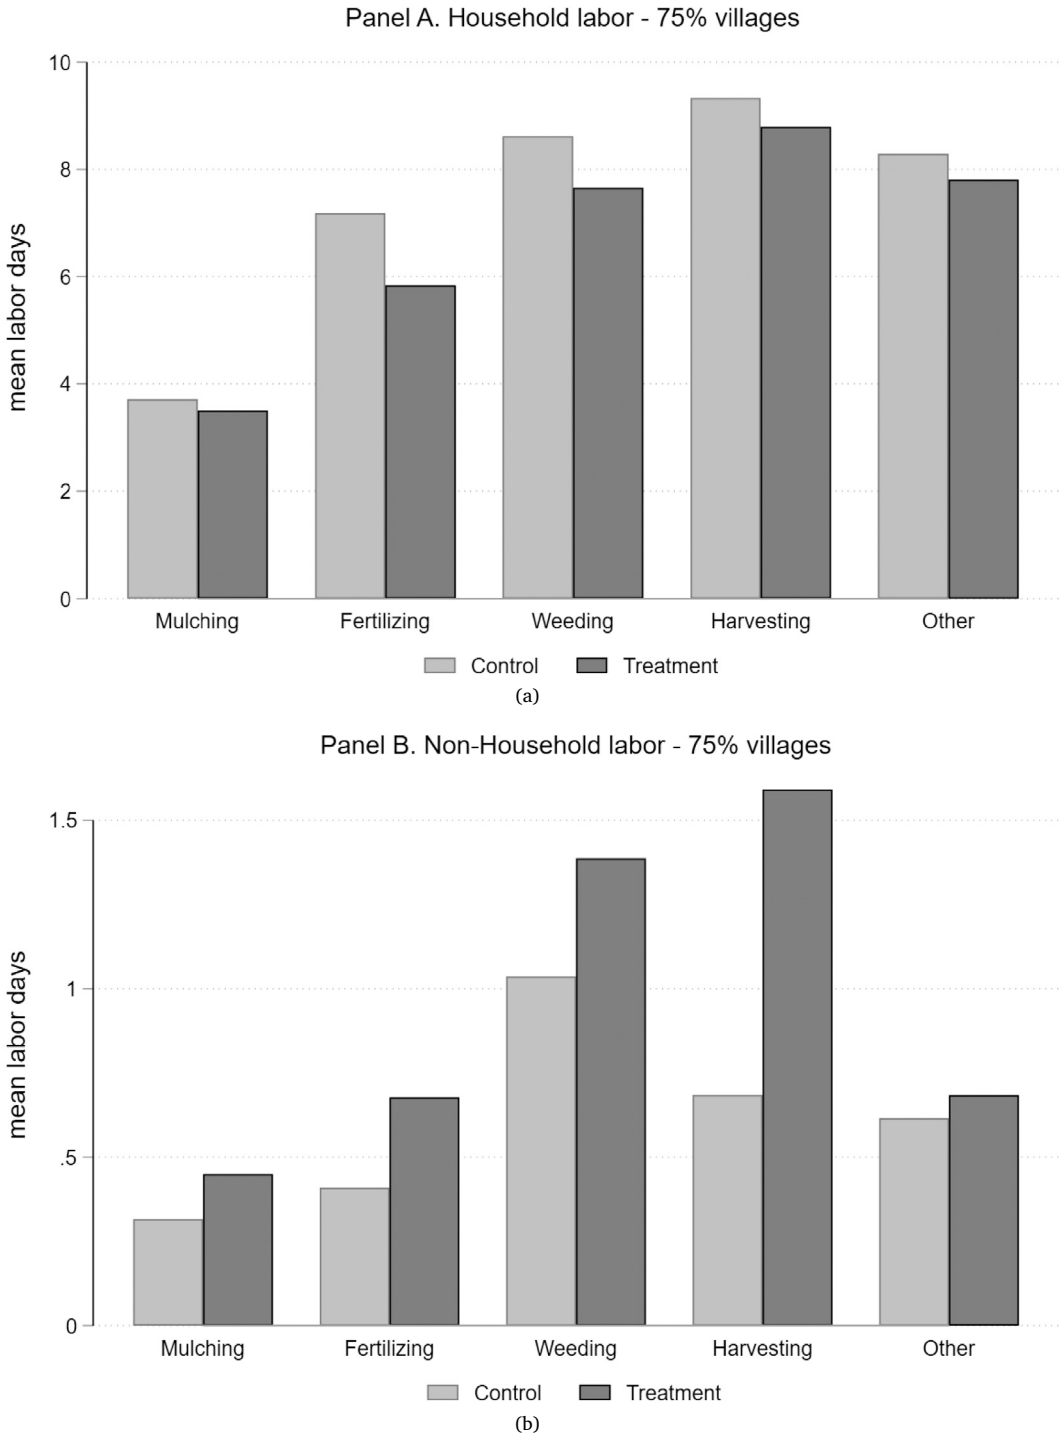


Fig. 1. Mean labor days by task and treatment status, in 75% villages.
Notes: This figure shows the mean number of total labor days used by treatment and control groups in 75% treatment concentration villages, disaggregated by task and type of labor (household versus non-household, where the latter = unpaid + paid). We compute these means from the estimates in Table 5.

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.

Table 6
Impacts of village treatment concentration on wages.

	All stacked (1)	Mulching (2)	Fertilizing (3)	Weeding (4)	Harvesting (5)
	Panel A				
50% T in village	16.419	22.918	27.214	11.656	−20.684
	(0.368)	(0.440)	(0.108)	(0.559)	(0.434)
75% T in village	27.920	28.800	24.712	32.740	16.551
	(0.136)	(0.330)	(0.176)	(0.082)	(0.513)
Sample mean, 25% villages	654.004	680.680	693.041	647.184	643.251
[1] p-value: 50% = 75%	0.537	0.843	0.908	0.262	0.099
Observations	2514	267	544	661	409
	Panel B				
Treatment	10.863	32.318	−8.711	25.283	−9.304
	(0.206)	(0.132)	(0.561)	(0.012)	(0.535)
50% T in village	27.660	30.882	33.709	35.709	−25.523
	(0.108)	(0.334)	(0.062)	(0.041)	(0.317)
75% T in village	35.249	17.928	33.462	33.627	66.419
	(0.066)	(0.569)	(0.103)	(0.086)	(0.003)
Treatment × 50% T in village	−25.640	−33.044	−9.179	−54.354	12.262
	(0.000)	(0.003)	(0.042)	(0.000)	(0.940)
Treatment × 75% T in village	−15.479	−9.309	−5.677	−14.483	−52.004
	(0.194)	(0.978)	(0.243)	(0.902)	(0.002)
Control mean, 25% villages	652.574	670.876	697.248	639.908	647.719
[2] p-value: 50% T = 75% T	0.682	0.697	0.992	0.899	0.001
[3] p-value: Treat × 50% T = Treat × 75% T	0.887	0.142	0.613	0.324	0.000
[4] p: Treatment + Treatment × 50% T = 0	0.153	0.969	0.238	0.144	0.940
[5] p: Treatment + 50% T + Treatment × 50% T = 0	0.487	0.423	0.447	0.783	0.507
[6] p: Treatment + Treatment × 75% T = 0	0.852	0.562	0.350	0.817	0.001
[7] p: Treatment + 75% T + Treatment × 75% T = 0	0.211	0.346	0.384	0.122	0.847
Observations	2514	267	544	661	409

Notes: Standard errors clustered by village. For brevity, we do not report those and only show randomization inference p-values in parentheses. Here, the randomization inference permutations are applied to the restricted set of 1041 households out of 1594 in our sample who report hiring any external labor to work on their coffee farm in exchange for daily wages in endline rounds 6–9. All columns 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 2–5, the outcome is household-level average daily wages paid per survey round per task. Column 1 controls for round and task fixed effects. Columns 2–5 control for round fixed effects. All regressions control for baseline covariates selected via post-double selection LASSO.

Household Income. Finally, we investigate whether the negative spillovers on control farmers documented above were offset by increases in income earned outside farmers' own plots—particularly if treatment farmers worked more on one another's fields in high-intensity villages. Appendix Table A21 examines this by disaggregating total household income into coffee income earned off one's own plot, other agricultural income, and non-agricultural income.

In column (1), treatment households in higher-intensity villages report slightly higher (though insignificant) income from work on other farmers' coffee plots, consistent with the increases in off-farm coffee labor days documented in Table 7.[19] By contrast, columns (2) and (3) show no indication that control households compensated for reduced yields by increasing income from other agricultural or non-agricultural activities. Effects are small, imprecise, and lack a systematic pattern across village groups. Overall, although these estimates are noisy, they suggest that the spillovers we document for inputs, labor, and yields are not offset by clear income responses in other activities.

6. 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 the number of trees and four key inputs: labor, fertilizer, compost and mulch.[20] Farmers' productivity may also be a function of their household characteristics, so we augment the production function with measures of household head education, gender and the number of baseline coffee trees (which may proxy for tree age). 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.

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

[19] Because the income module relies on short recall windows and records relatively few paid labor transactions, these estimates should not be interpreted as a direct measure of the informal labor exchanges that may accompany program-induced links among trainees.

[20] Since we lack information on the exact quantities of mulch used, we include an indicator for whether the farmer uses any mulch at all.

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
	(0.128)	(0.225)	(0.104)
75% T in village	0.203	−0.382	0.218
	(0.062)	(0.459)	(0.053)
Sample mean, 25% villages	3.855	0.135	3.721
[1] p-value: 50% = 75%	0.659	0.651	0.688
Observations	4710	4710	4710
	Panel B		
Treatment	0.010	0.219	0.001
	(0.945)	(0.821)	(0.991)
50% T in village	0.092	−1.001	0.114
	(0.486)	(0.159)	(0.408)
75% T in village	0.110	−1.650	0.137
	(0.502)	(0.117)	(0.417)
Treatment × 50% T in village	0.119	0.546	0.118
	(0.568)	(0.685)	(0.594)
Treatment × 75% T in village	0.116	1.329	0.108
	(0.611)	(0.403)	(0.615)
Control mean, 25% villages	3.839	0.125	3.713
[2] p-value: 50% T = 75% T	0.920	0.532	0.899
[3] p-value: Treat × 50% T = Treat × 75% T	0.986	0.563	0.958
[4] p: Treatment + Treatment × 50% T = 0	0.270	0.349	0.316
[5] p: Treatment + 50% T + Treatment × 50% T = 0	0.069	0.727	0.068
[6] p: Treatment + Treatment × 75% T = 0	0.409	0.140	0.482
[7] p: Treatment + 75% T + Treatment × 75% T = 0	0.056	0.886	0.054
Observations	4710	4710	4710

Notes: Standard errors clustered by village. For brevity, we do not report those and only show randomization inference p-values in parentheses. All specifications control for round fixed effects. Each panel concludes with randomization inference p-values for various linear restriction tests. 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.

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.

Due to their flexibility the treatment/control production functions contain 80 coefficients each, so the formal expression of the production function is long and uninformative. In brief, the specification includes the main and squared effects of all household characteristics and coffee farm characteristics listed above all, plus all two-way interactions among them, plus all three-way interactions involving one household variable and two farm inputs. Rather than reporting all coefficients, we instead graph the marginal effect of the main inputs (labor, fertilizer and compost), holding constant all other variables at their respective sample means for the trained and untrained farmers.

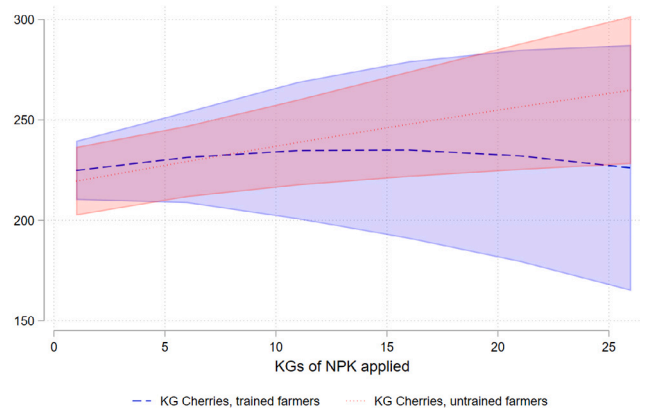
Results of the production function estimation are displayed in Fig. 2, with Fig. 2(a) showing the association between NPK and output, Fig. 2(b) showing the association between labor and output, and Fig. 2(c) showing the association between compost use and output. In all three cases, we cannot reject that the relationship is linear, and the only case in which the production function appears even possibly convex is for the *untrained* farmers' labor use. Concavity in the effects of NPK 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 compost relationship exhibits a similar pattern of diminishing returns, with no evidence of systematically higher productivity among treated farmers at any level of compost use. Across all inputs, the magnitude of predicted output is very similar between trained and untrained groups, suggesting that treatment 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.

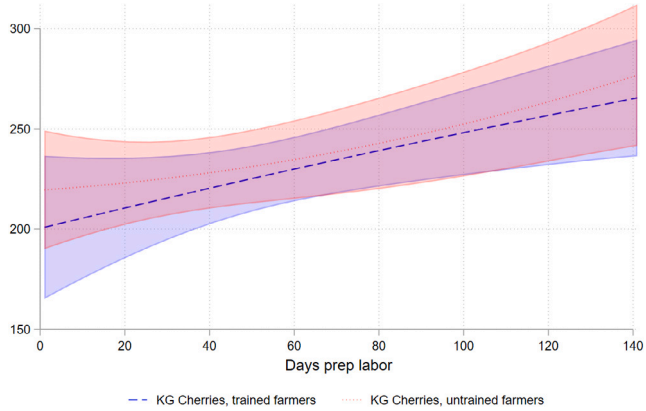
7. Conclusions

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.

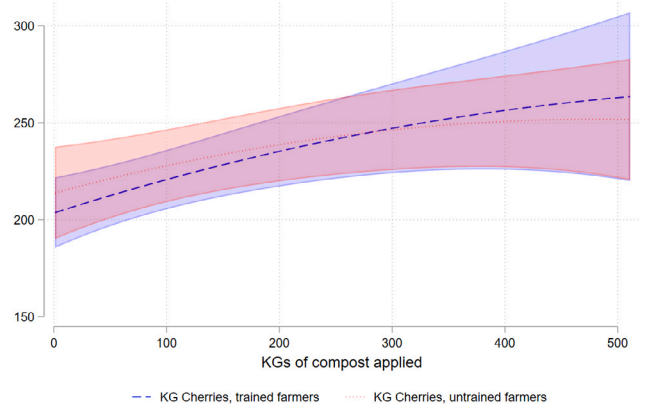
First, the hypothesized mechanism for information diffusion is undermined by the finding that the program increased social links only between co-trainees in the treatment group. While this does not imply that treatment farmers ceased interacting with the control group entirely, it is likely that they 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 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.



(a) Relationship between coffee cherry output and NPK use



(b) Relationship between coffee cherry output and labor use



(c) Relationship between coffee cherry output and compost use

Fig. 2. Coffee cherry production function.
Figures plot predicted coffee cherry output as a function of input quantities, based on the production function estimated in Section 6. 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.

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 net revenues (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 training only a subset of farmers can reshape social connections within villages in ways that reallocate scarce labor and other inputs, even in the absence of learning or productivity gains. 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.

CRediT authorship contribution statement

Esther Duflo: Writing – review & editing, Writing – original draft, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Daniel Keniston:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Methodology, Investigation, Formal analysis, Data curation. **Tavneet Suri:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Céline Zipfel:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Investigation, Formal analysis, Data curation.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: I hereby declare that I have no relevant or material financial interests that relate to the research described in this paper. Research funding was received from the International Growth Center and Technoserve. No party had the right to review the paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jdeveco.2026.103733.

Data availability

Data will be made available on request.

References

Aker, J., Bergquist, L.F., Jayne, T., Magruder, J.R., 2022. Agricultural technology in Africa.

Ashraf, N., Giné, X., Karlan, D., 2009. Finding missing markets (and a disturbing epilogue): Evidence from an export crop adoption and marketing intervention in Kenya. Am. J. Agric. Econ. 91 (4), 973–990.

Banerjee, A.V., Chandrasekhar, A.G., Duflo, E., 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., Mobarak, M., 2021. Can network theory-based targeting increase technology adoption? Am. Econ. Rev. 111 (6), 1918–1943.

Belloni, A., Chernozhukov, V., Hansen, C., 2014. Inference on treatment effects after selection among high-dimensional controls. Rev. Econ. Stud. 81 (2 (287)), 608–650, (Accessed 2023-05-01).

Bridle, J., McIntosh, C., Suri, T., 2019. Experimental insights on the constraints to agricultural technology adoption.

Chandrasekhar, A.G., Duflo, E., Kremer, M., F. Pugliese, J., Robinson, J., Schilbach, F., 2022. Blue Spoons: Sparking Communication About Appropriate Technology Use. Working Paper 30423, National Bureau of Economic Research.

Cole, S., Fernando, A., 2021. 'Mobile'izing agricultural advice technology adoption diffusion and sustainability. The Economic Journal.

Comola, M., Prina, S., 2023. Treatment effect accounting for network changes. Rev. Econ. Stat. 105 (3), 684–700.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., Zamora, P., 2013. Do labor market policies have displacement effects? Evidence from a clustered randomized experiment. Q. J. Econ. 128 (2), 531–580.

Derksen, L., Souza, P.C.L., 2024. Who knows? Information access and endogenous network formation.

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

Emerick, K., 2018. Trading frictions in Indian village economies. J. Dev. Econ. 132, 32–56.

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

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

Heß, S., Jaimovich, D., Schündeln, M., 2021. Development projects and economic networks: Lessons from rural Gambia. Rev. Econ. Stud. 88 (3), 1347–1384.

Jones, M., Kondylis, F., Loeser, J., Magruder, J., 2022. Factor market failures and the adoption of irrigation in Rwanda. Am. Econ. Rev. 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. Agrofor. Syst. 98, 1323–1341.

Magnan, N., Spielman, D.J., Lybbert, T.J., Gulati, K., 2015. Leveling with friends: Social networks and Indian farmers' demand for a technology with heterogeneous benefits. J. Dev. Econ. 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. Rev. Econ. Stat. 103 (4), 666–679.

Suri, T., Udry, C., 2022. Agricultural technology in Africa. J. Econ. Perspect. 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., Tjernström, E., 2024. Agricultural technology in Africa. VoxDevLit 5 (2).

Wooldridge, J.M., 1999. Distribution-free estimation of some nonlinear panel data models. J. Econometrics 90 (1), 77–97.

Wooldridge, J.M., 2010. Econometric Analysis of Cross Section and Panel Data, second ed. MIT Press, Cambridge, MA.

Zhang, Z.X., Cai, Z.Q., Liu, G.Z., Wang, H., Huang, L., 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.