The Complexity of Multidimensional Learning in Agriculture

Rachid Laajaj, Universidad de Los Andes,

Karen Macours, Paris School of Economics and INRAE

Abstract:

Studies on agricultural technology adoption often focus on one input, practice or package, which is analytically useful, but may overlook the complexities involved with multidimensional learning needed for a lot of agricultural decisions. In Kenya, we study farmers' dynamic learning (from oneself and others) and adoption decisions over six seasons after randomly inviting them to participate in agronomic research trials, comparing different combinations of inputs during three consecutive seasons. As a response to the trials, adoption increases steadily despite profits being initially harmed by exposure to the trials. Know-how increases rapidly and faster for high skill farmers who experiment the most, at the cost of making new mistakes. The findings are consistent with a theoretical model with multidimensionality of input and practice decisions and differential learning from one's own experience by skills, where complementarities imply that adoption of an input requires finding how to re-optimize other dimensions, which adds to the cost of adoption.

Keywords: technology adoption, learning constraint, Sustainable agriculture, skill, multidimensional

Acknowledgements: Manuel Camargo, Matilda Chweya, Irene Clavijo, Katriel Friedman, Juan Miguel Jimenez, Freddy Felipe Parra Escobar, Jack Pfeiffer, Juan Restrepo Rivera and Sébastien Tamegnon Zinsou provided excellent research assistance at different stages of this project. We acknowledge the excellent field teams of IPA Kenya for key contributions during piloting and translation and for all efforts during the data collection. This paper builds on joint work with IITA and we gratefully acknowledge inputs from Cargele Masso, Edwin Mutegi, Violet 0menyo Moses Thuita, and Bernard Vanlauwe. This work was funded by INRAE, DFID-ESRC Growth Research Program Call 2 (JES-1362222), the LSMS team at the World Bank, the Standing Panel for Impact Assessment of the CGIAR under SIAC 1, the French National Research Agency (ANR), under Grant ANR-17-EURE-0001, and CEPREMAP. The paper benefitted from comments by Michael Carter, Pascaline Dupas, Kyle Emerick, Marcel Fafchamps, Lauren Falcao Bergquist, Eliana La Ferrara, Marco Gonzalez-Navarro, Kelsey Jack, Travis Lybbert, Craig McIntosh, Jonathan Robinson, Tavneet Suri, Leonard Wantchekon and seminar participants at Bergen, EUDN, FERDI, IFPRI, Namur, Paris Sorbonne, PSE, Stanford, Tufts, UC Berkeley, UC San Diego, UC Santa Barbara, UC Santa Cruz, U of Los Andes, U of San Francisco.

1. Introduction

Poverty reduction, food security and sustainable agriculture remain among the greatest challenges of the 21st century, as reflected in the first two Sustainable Development Goals. An abundant literature studies constraints to agricultural technology adoption, motivated by the belief that major progress towards these goals could be achieved if only smallholder farmers would adopt promising new technologies. Lack of information about the returns to a new technology or about the required know-how is often flagged as a major constraint to adoption (Foster and Rosenzweig, 1995; Bandiera and Rasul 2006, Conley and Udry 2010). Hefty increases in the price of mineral fertilizer, global targets for mitigating greenhouse gas emissions, and farmers' needs to shift technologies to adapt to changing climates, all provide impetus for a shift towards more sustainable agricultural practices (Lobell et al, 2008; Snapp et al 2023). Because such practices tend to be more knowledge-intensive, compared to the capital-intensive inputs that were the engine of the green revolution (Bationo et al. 2012, Vanlauwe et al. 2010), this transition is likely to raise the role of knowledge constraints. This is particularly the case as sustainable intensification often involves multidimensional adoption, with simultaneous decisions on various inputs and practices. Better understanding how farmers learn about multidimensional input combinations then becomes crucial for effective policy design.

This paper presents a theoretical and empirical analysis of the complexity of farmers' learning about multidimensional input combinations. We intensively exposed farmers to new information about the return and know-how for different input combinations, and study their dynamic decisions on inputs and practices, and the resulting profits. In randomly selected villages, we set up a series of agronomic 2x3 trials offering farmers a direct multi-season opportunity to learn from multiple side-by-side comparisons of various input combinations on their own land. While the trials demonstrated clear yield gains, and farmers started adopting inputs and practices tried on the trials, this did not increase profits initially, and if anything impacts on profits are negative. Remarkably, however, despite these negative effects on profits, farmers kept on further experimenting with the trial inputs and packages and impacts on profits gradually improved over the course of 6 seasons.

The paper proposes a version of the target-input model that can explain these and related results and empirically documents the dynamic learning and adjustments, building on an intensive data collection process, including a comprehensive baseline measure of skills, administrative data from the trials and six rounds of detailed measures of agricultural know-how, practices, input use, experimentation, beliefs about the returns to the new inputs, profits and social interactions. This allows to show how accounting for the multidimensionality of input decisions helps understand dynamic learning and farming decisions. We also show how this may vary by skills, with its implications on convergence of know-how and for farmers' decisions on learning-from-others.

The focus on multidimensionality and interdependence of decisions on various inputs and practices is informed by insights from agronomy. Examples of complementarities (and substitutabilities) between input and practice decisions abound. For instance, the optimal number of seeds per hole depends on whether it is an improved variety (which affects germination rates), planting in line is valuable when fertilizer is applied, and biological nitrogen fixation can be very cost effective but only if no chemical nitrogen is provided, while combining it with a P source increases its effectiveness. Similarly, there can be somewhat complex complementarities with contextual variables, where returns to fertilizer depend on the use of weed-resistant seeds, but only when those weeds are present. Such complementarities imply that changes in one input leads to the need to figure out the required adjustments in the other inputs, which we'll refer to as the need to re-optimize. Learning about these multidimensional adjustments is arguably underrepresented in most current learning models, motivating this paper.

The main empirical estimations are based on a sample of 10 farmers in each of the 96 villages, purposely sampled to represent both high and low skilled farmers. In 48 randomly selected villages, the 10 farmers were invited to participate to an agronomic research trial during 3 consecutive seasons on one of their parcels. The trials were composed of 6 subplots including a control subplot and 5 subplots where different combinations of modern inputs were tested. Farmers participated in all tasks, under the guidance of an agronomist from the International Institute for Tropical Agriculture (IITA). The inputs and practices followed the principles of Integrated Soil Fertility Management (ISFM), a sustainable intensification practice. In each treated village, 4 farmers were randomly selected to take part in a maize trial, 3 farmers were assigned to a soya trial and 3 farmers were assigned to an intercropping trial (combining maize and soya) with the randomization stratified by skills. The trials occupied a very small portion of each farmer's land, hence they have almost no direct economic impact. Instead, they represent an opportunity to learn about input and recommended practices.

The first set of results show that the 3-season-long exposure to the agricultural trials led to a progressive increase in the use of the trials' practices and inputs in the farmers' own plots, outside of the trials. The treatment initially reduces estimated profits for both low skill farmers (LSFs) and high skill farmers (HSF), but the effect grows over time and becomes positive for LSFs in the last season. We also evidence substantial and differential learning: despite starting at a higher level of knowledge, the treatment increased HSFs' know-how significantly more than the one of LSFs. These basic results point to two related puzzles. First why do input adoption grow over time despite an initial reduction in profit among treated farmers? This seems to go against a standard model of learning about the returns, in which adoption is driven by the high realizations of profits, thus increasing the perceived return to the input. By contrast, within the framework of a target input model[1] with farmers acquiring know-how, adoption can occur and increase despite a reduction in profit if farmers anticipate that, as they figure out the target, their future profits will increase. This then relates to the second puzzle however, as we find that the know-how between LSFs and HSFs diverge, while in the standard target-input model with one input we expect them to converge, since eventually all farmers should figure out the target and corresponding optimum allocation of the input. The two puzzles can be explained through a version of the target input model with multiple local maxima resulting from the multidimensionality of the input decisions.

The theoretical model is similar to the target input model (Foster and Rosenzweig 1995) and to Conley and Udry (2010), in the sense that Bayesian learning farmers update their beliefs about the shape of a profit function both through their own experience and from receiving signals from exogenous outside sources of information. They are also forward looking and incorporate the value of learning in their input decisions. To adapt the model to our focus and first set of results, we assume 1) a multi-dimensional and non-parametric profit function and 2) HSFs differ from LSFs in the precision of the signal that they extract from observation (as in Rosenzweig 1995). At each period, farmers have an initial prior, decide on a bundle of inputs and practices, and observe the realization of profits, which allows them to update their beliefs at that point of the profit function and its vicinity. The resolution of the problem follows a multi-armed bandit

[1] In a standard version of the target input model, the profit function takes the form $y = k - (\theta - x)^2$, where y is the profit, x is the input decision, k is a constant and θ is the uncertain target. Profits decrease with the distance between the input decision and the target, a loss that vanishes as farmers figure out the value of θ.

problem, where farmers' decisions balance the benefits from exploitation (immediate utility gain) with the ones from exploration (benefits from the acquisition of knowledge that can increase future gains). In a continuous space, possible strategies are the exploitation of a known local maximum, local exploration or alternatively "jumping into the wild", i.e. to explore mostly unknown areas of the profit function. However, because of risk aversion, the latter option is quite limited, unless an external source of information provides sufficient Information about this part of the profit function to reduce the related uncertainty.

The theory leads to the conclusion that farmers can be good at finding a local maximum but are limited in their capacity to explore further and to transit to a higher local maximum. The curse of multidimensionality occurs only in the presence of (unknown) complementarities and substituabilities between inputs, making it costly to re-optimize other dimensions. Exploration away from a known local maximum can, however, be triggered by external sources of information, such as the trial, or observation of input decisions and output realizations by farmers in one's network. For simplicity, both trials and network information are modelled as exogenous and unexpected. When exploring, farmers are willing to tolerate a reduction in profit in the short term if it is compensated by the expected gain from learning about the profit function, and the possibility to discover a higher local equilibrium. Because HSFs learn more with each observation, they tend to explore more and are willing to accept more profit reduction than LSFs, which is consistent with our first set of results. This also leads to the prediction that LSFs particularly value connections with HSFs because they can benefit from their experimentation.[2]

The third part of the paper then makes use of this framework to guide the analysis and interpretation of additional results on different outcomes and specifications. Several results confirm our central hypothesis of differential learning between LSFs and HSFs. The difference in learning is particularly pronounced in "subtle" learnings, such as adapting seed decisions to weed prevalence or deciding which fertilizer to use between two relatively good fertilizers tested in the trials. Additionally, consistent with HSFs learning from their own experimentation, we find that the HSFs' decisions to use the tested inputs is highly correlated with the yield response observed in their own trials (which is not the case for LSFs). We also find evidence of a possibly

[2] While the two terms refer to the same concept, we use "exploration [of the profit function]" in the context of the theoretical model to adopt the language used in the description of the arbitrage faced in armed bandit models, and "experimentation" in the empirical parts to reflect more standard vocabulary in technology adoption.

costly re-optimization as farmers adopt the new inputs. Indeed, the treatment makes farmers more likely to discover and use some complementary inputs over time, but they also become more likely to commit new mistakes due to “wrong combinations” of inputs and practices, and these patterns tend to be stronger for HSFs who experiment more. Exposure to new inputs may take farmers out of their comfort zone, at the cost of committing new mistakes and losing profits while they are exploring the profit function and identifying the right new input combinations.

A last set of findings enrich our understanding of learning from others conditional on skills. The treatment increases communication between farmers and results show they learn a lot from each other. For example, the farmers with the maize treatment significantly increased their knowledge about the inputs tested in the soya trials and vice versa. However, learning across crop treatments is limited when it comes to more subtle learning and actual adoption of the new inputs. The results show that LSFs tend to be more aware about the input use of HSFs than the one of LSFs. Finally, LSFs tend to adopt inputs and practices one to two seasons after HSFs do so. Together, the results point towards HSF being more able to learn from their own experience, acting as explorers, whose learning tend to spill over to other farmers.

This paper contributes to various strands of the literature. First, it broadly relates to papers testing interventions addressing farmers’ information constraints, such as demonstration plots and field days (Crane-Droesch 2018, Emerick and Dar, 2021), extension services or trainings (Kondylis et al. 2017, Beaman et al. 2021; Aker and Jack, 2023), or input subsidies (Carter et al. 2021, Gignoux et al. 2023) and to evidence that the difficulty to observe returns can lead to the persistence of lemon technologies (Bold et al, 2017). Rather than testing a specific intervention, this paper focuses on understanding the learning process about different inputs and practices. Other related works zoom into learning using experimental games or virtual environment (Barham et al 2018, Tjernström et al. 2021, Conlon et al. 2022). This paper also connects with Hanna et al. (2014)’s work on mistakes in learning from selective attention, and closely relates to Nourani (2019)’s model with farmers whose adoption and dis-adoption decision are affected by learning about both profitability and know-how. We add to this literature by causally documenting the intricacies of the dynamic learning process of farmers facing multidimensional input and practice decisions. For example, to the best of our knowledge, the observation of a growing adoption among treated farmers despite an effect on profit that is initially negative and then grows over time has not been evidenced and validates a prediction common to know-how oriented models with forward looking farmers. The limitations of LSFs to learn from their own experience and about more subtle learning, and the extent to which experimentation leads to new mistakes are also novel elements. The evidence stems from being able to closely track the dynamic learning over 6 seasons, which in addition tackles concerns about external validity resulting from stochastic processes (Rosenzweig and Udry, 2020).

This paper highlights that besides the know-how about a particular new technology itself, it can be the cost of re-optimizing other inputs and practices that affects adoption decisions. As such, it provides a common framework that helps interpret other recent papers finding that the success of an intervention depends on whether farmers succeed in making such adjustments (Emerick et al. 2016) or not (Ghosh 2018, Jones et al. 2022). In Emerick et al. (2016) a large share of the gains from a new flood-tolerant rice variety comes from crowding in other complementary investments, such as labor-intensive planting method, area cultivated, fertilizer usage, and credit utilization. By contrast, in Jones et al. (2022) many farmers fail to adopt irrigation because leveraging its benefits requires switching to horticulture, which is associated with increased use of labor, fertilizer, and seeds. Similarly, Ghosh (2018) finds that farmers in India lost profit because they failed to understand that, to reap the benefits of genetically modified Bt cotton, they needed to reduce pesticide use, a substitutable input. As our findings point to the potential gains of reducing this re-optimization cost for the farmers, this paper also connects to research exploring solutions to this complexity through decision support tools (Chandrasekhar et al 2022), or precision agriculture (Fabregas et al. 2019; Corral et al 2020).

The paper further contributes to the large literature on information spillovers in agricultural technology adoption (Bandiera and Rasul, 2006; Conley and Udry 2010). Evidence shows that diffusion through social networks can be an effective strategy (BenYishay and Mobarak, 2019e-; Beaman et al, 2021), but diffusion about appropriate technology use and know-how through networks is often unequal (Beaman and Dillon 2018; Bandiera et al., 2023), slow (Chandrasekhar et al 2022) or even absent (Duflo et al 2023). We highlight that this heterogeneity can be related to the interaction between the complexity of the lessons to be shared and the farmers' abilities. While side-by-side comparisons of new and old inputs on demonstration plots can make farmers seek information from others (Kelley et al 2023), heterogeneity among farmers and their conditions are often seen as one of the reasons for slow diffusion of innovations (Munshi 2004, Magnan et al, 2015) and productivity dispersion (Gollin and Udry, 2021). By contrast, this paper shows evidence suggestive of LSFs benefitting from the presence and experimentation of HSFs, helping LSFs compensate for their limited ability to learn from their own experimentation compared to HSFs.

Finally, with the focus on farmers' skills, this paper speaks to the broader debate about the returns to human capital in the agricultural sector (Gollin et al., 2014; Hamory et al., 2021). The possibility of positive learning externalities from HSFs to LSFs implies that estimates based on individual returns may underestimate the aggregate return to skills in the agricultural sector. This also relates to a broader literature, according to which such spillovers cause productivity gains to be driven by the upper tail of the local skills distribution (Squicciarini and Voigtländer 2015). Since sustainable agriculture tends to be more knowledge intensive and context specific, if it is to become more prominent, this would raise the relevance of such models in the agricultural sector. At the same time, as our results highlight that learning from observation is very imperfect, suggesting that educational systems that teach children how to learn using scientific reasoning (as in Ashraf et al, 2021) could be particularly relevant even in agriculture, despite it often being considered the traditional sector.

The next section describes the context, the agricultural trials, and the study design. Section 3 provides a first set of results and highlights a set of paradoxes within them. Section 4 presents a theoretical model which addresses some of these puzzles and develops further predictions to take to the data in Section 5. Section 6 concludes and discusses policy implications.

2. Context, Agricultural Trials, and Study Design

2.1 Context

This project occurred in Siaya, Western Kenya, which had a poverty rate of 38.2% in 2016[3], while farming was the main economic activity and maize the main crop. In our sample, about a third of farmers is involved in commercial farming, and 77% of farmers used mineral fertilizer at baseline, mostly Calcium Ammonium Nitrate (CAN) and Diammonium Phosphate (DAP), which have long been used and are easily available, but can increase soil acidity over the long run (Uwiragiye et al. 2023). Research efforts in the region have focused on testing more context-specific and sustainable intensification packages, including combinations of mineral and

[3] Source: Kenya National Bureau of Statistics, https://kenya.opendataforafrica.org/urwhbig/poverty-estimates?region=1000260-siaya

organic fertilizer and new seeds and crops. The challenge that motivated this paper is about improving farmers' decisions about optimal combinations of inputs and agronomic practices.

We partnered with the International Institute for Tropical Agriculture (IITA), one of the CGIAR research centers, with a long track record in sustainable intensification research, including ISFM. ISFM aims at increasing efficiency while bringing environmental benefits through the right combinations of inputs and practices and local adaptation, but its full implementation is quite knowledge intensive (Bationo et al. 2012, Vanlauwe et al. 2010).[4]

2.2 The agronomic trials and farmers' learning opportunities

A set of agronomic trials was organized on the farmers' parcels, following a factorial design aimed at comparing different combinations of inputs with each other. Each trial is laid out in 12 by 16.5 meters (39 by 54 feet), representing only 1.3% of the median farmers' cultivated area. Treated farmers were assigned to either a maize trial, a soya trial or a trial with maize and soya intercropped. Using a 2 by 3 design, each trial was divided into 6 subplots where different input combinations were tested. The trials occurred during three consecutive seasons to reduce the influence of rainfall variability and to allow farmers to observe the benefits across seasons. The trials were part of IITA's agronomic research and would qualify as researcher-designed and farmer-managed trials (Franzel and Coe, 2002). This implies that agronomists had full control over the design of the trials, from the choice of inputs to the recommended management practices, while the farmers provided the land and labor, following their guidance.

The tested inputs were relatively new in the region and complied with the previously described logic of ISFM. They included biological nitrogen fixation (or biofertilizers, which provide a cheaper and sustainable substitute to the provision of nitrogen through mineral fertilizers), compost, seeds resistant to striga (the most common weed in the area), and mineral fertilizers that improve soils through micronutrients and a reduction of soil acidity (Njoroge et al. 2017). Despite our intent to include only inputs available in the local market, de facto, their availability at local retailers varied between the different inputs and by location. Availability for the trials themselves, however, was guaranteed for the 3 seasons. The trials were implemented to

[4] Vanlauwe et al. (2010) define ISFM as "a set of soil fertility management practices that necessarily include the use of fertilizer, organic inputs and improved germplasm, combined with the knowledge on how to adapt these practices to local conditions, aimed at maximizing agronomic use efficiency of the applied nutrients and improving crop productivity […]".

conduct standard agronomic research and were subsequently used for publications (Thuita et al. 2018, Laajaj et al 2020). On average, the results from the trials demonstrate substantial yield gains and profitability when comparing subplots with the input packages compared to the control subplot of each trial. Appendix A and Laajaj et al. (2020) describe the design of the trials, the inputs and the main agronomic findings in more detail.[5]

The trials provided farmers a 3-season-long opportunity to directly learn from the side-by-side comparisons of input combinations on their own land. Participating farmers were further exposed to information (signals) about inputs and practices through : i) about 3 visits of the agronomists per season, with the possibility to interact and ask questions; ii) a one-page description of the inputs provided to the farmer at the onset of the implementation; iii) signs placed in each subplot listing inputs to help farmers remember what is being tested in each one of them; and iv) a field day, including a collective visit to all trials in the village and a discussion session, organized towards the end of the third season. The salience of the trials is likely to have been further enhanced by weekly visits by a contact person (typically one of the ten farmers) to the trial plot to verify that the farmers fulfilled their responsibilities, and by the frequency of surveys, including seasonal comprehensive post-harvest surveys as well as targeted visits to the trial plots (treatment farmers only) to ask about input use and outputs on each of the subplots.

Note that this intensity of interactions was not designed to evaluate what an extension intervention with demonstration plots could look like hence the RCT was not designed to evaluate a scalable intervention. Instead, we use the intensity of learning opportunities that resulted from being exposed to these trials as an exogenous variation in information signals about various inputs and practices, to study farmers' learning about the input combinations being tested. The sources of learning are many: farmers can observe the yield differences between subplots, but even if the yield effects are not discernable, farmers may also consider the simple fact that certain practices are requested by the agronomists in the trials and that certain inputs were selected for the trials as signals on their own. The trials also provide an opportunity to discuss with the agronomists, with the contact person and with other farmers. Last but not least, as farmers learn from the trial, they may want to adjust what they do in their own parcels and learn from their own experimentation outside of the trials.

[5] Trial results show that the "best bet" input package increased yield by 950 kg/acre in the maize trials and by 283 kg/acre in the soya trials.

2.3 Sampling and Study Design

This section describes the sampling of villages and farmers within villages followed by the random selection of villages assigned to the treatment and random assignment of trial crop (maize, soya or intercrop) among the 10 farmers of each treated village. Figure 1 provides a visual representation of the study design. The study was conducted in 96 villages, located in five regions (Boro, Ugunja, Ukwala, Wagai and Yala) of Siaya County in western Kenya.

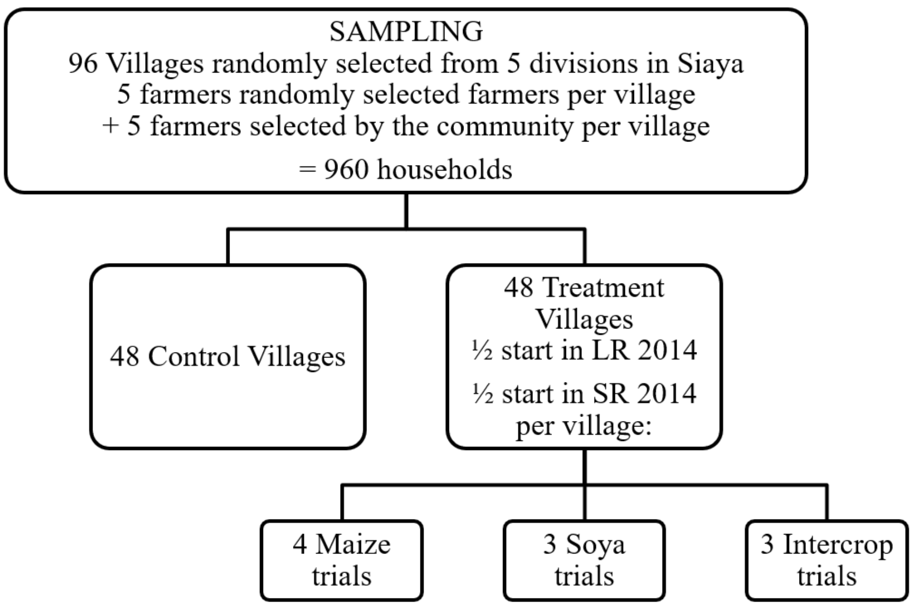

Figure 1: Graphical representation of the sampling and study design

In each village 10 farmers were selected to take part in the study, purposely sampled to represent different farmer types. First, to be able to randomly select 5 farmers, stratified by skill level, we first made lists of all farming households in the village with the help of each village's community health worker (CHW).[6] We asked the CHW 3 questions about every village member's 1) level of education 2) being active and motivated and 3) agricultural knowledge. We used the answers to the 3 questions (low, middle or high) to create an index, henceforth referred to as the ex-ante proxy measure of skill, to sort farmers into 5 skill quintiles. We then randomly

[6] The listing exercise was done with the CHW because they tended to be very knowledgeable about each household in their village since their work implies regular visits of all households. This started from a full mapping of the village, taking them part by part to minimize the risk of omitting households. In nearly half of the villages, CHWs already had a list of households, which was used as a starting point.

selected one farmer from each skill quintile, resulting in 5 farmers per village, henceforth referred to as the randomly selected farmers.

The 5 additional farmers per village were selected by the community after explaining the trials and study during a gathering in each village. Reflecting common practice in agronomic trials, village members were asked to nominate five farmers thought to be good farmers willing to participate in the trials, with the precise selection criteria left to the village members to decide (except that it should include at least 2 women). Qualitative observations of discussions during the gatherings clearly pointed at selecting farmers that the community thought would represent them well. We refer to these 5 additional farmers as the community-selected farmers.

In the 96 villages, before assigning treatments, all 10 selected farmers were visited to obtain consent for the trials and identify the potential trial parcel (chosen by the farmer, conditional on fitting with some criteria for suitability to the research trials). A small number of replacements occurred when farmers were not interested in participating or when they did not have any parcel that met the basic conditions (no extreme slopes and absence of obstacle, tree or shade). Replacement was done within farmer type, always keeping 5 farmers selected by the community and 5 random ones, balanced across skill quintiles.

After finalizing the selection of the farmers in all 96 villages, we randomized the selection of villages that would participate to the trials, using public lotteries organized in each sub-county, which implies that it was stratified at the sub-county level. We organized the trials in two waves: in 24 randomly selected villages trials started in the long rain of 2014 and in another 24 randomly selected villages trials started in the short rain of 2014. Each trial lasted for 3 seasons. In all treatment villages, 4 of the 10 selected farmers were randomly assigned to participate in the maize trial, 3 were randomly assigned to participate in the soybean trial and the remaining 3 farmers were assigned to the intercrop trial (of maize and soya together). This assignment was stratified by community versus randomly selected farmers and by quintile of skill level (based on the ex-ante proxy).

2.3 Data Collection, Timeline and the skills measure

The main data was collected through a baseline survey (with 2 visits), and five follow-up surveys. The timing of data collection is illustrated in Figure 2. Data collection was designed to capture the dynamics of the treatment effects over time, during the three seasons of treatment and up to two seasons after the treatment ends. Data was collected in cooperation with Innovation for Poverty Action (IPA), on farmer's skill measures, a plot-by-plot measure of characteristics, inputs, practices and production, a test of know-how related to the new inputs, input decisions, yield expectations under various scenarios and network information. Additionally, monitoring data from the trials were collected by the agronomists of IITA.

As the beginning of the trials was staggered (with 24 villages initiating in the short rain 2013 and another 24 villages in the long rain 2014), estimations tend to balance seasonal effects: for example when we look at the effects in the first season, half of the villages that are in their first season of treatment are in long rain and the other half is in short rain, and this is true for any number of seasons between 1 and 4. The effects of the treatment 5 seasons after the beginning of the treatment are only comprised of the first treated group once it reaches the short rain of 2016.

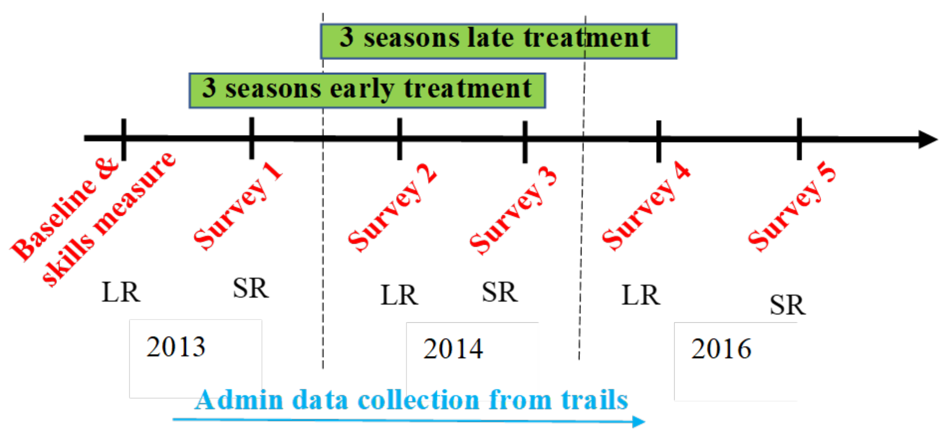

Figure 2 Timeline of interventions and data collection

Throughout the analysis, we separate the results by farmers' initial skills level. Given that skills measurement is prone to substantial measurement error, we tested, adapted and improved a set of skill measures to obtain one of the most extensive skills measures available for rural developing context (see appendix C for more details and Laajaj and Macours (2021) for evidence in reliability and validity of the measure). For each farmer, we administered twice a 2.5 hour survey instrument to calculate a non-cognitive index (using factor analysis), a cognitive index and an index of agricultural knowledge (using item response theory for the last two components).[7] Taking the non-weighted average of these three components provides us with a

[7] The design of the knowledge tests with our agronomist partners and its piloting among farmers qualitatively revealed the complexity of multidimensional decision-making as finding questions with unambiguous scoring of

comprehensive skill index. We then divide this index at the median, classifying half of the farmers in the sample as low skill farmers (LSFs) and the other half are high skill farmers (HSFs).[8] Among the community-selected group 62% are classified as HSF and among the randomly selected farmers 38% are (see Figure C2).

2.4 Balance, Compliance and Attrition

Appendix Table B1 shows initial characteristics of the farmers in treated villages compared to other farmers in control villages. Out of the 14 variables, only one variable is significant, and the predetermined characteristics do not significantly predict the likelihood of being treated (p-value of the joint test=0.384).

There was high compliance with the trials, 95% of the farmers assigned to the treatment had the agronomic trial in their parcel at least one season, and no farmer in control villages participated to the IITA agronomic trials. [9] Additionally, since all treatment villages had multiple trials, treatment farmers that did not participate in a given season were still indirectly exposed through lessons from the trials among their neighbors. Survey attrition was also low: 2.9% in the control and 3.3% in the treatment, over all seasons, with a maximum of 4.6% in any given season. The difference in attrition between control and treatment group is not significant during any season (see Table B2).

3. Impact of Exposure to Trials on Agricultural know-how, Adoption and Profits

3.1 Estimation strategy

We estimate the treatment effects separately for low and high skill farmers. The following regression pools the treatment effect over the different follow-up rounds:

---

answers was far from straightforward, with very few input or practice decisions being optimal irrespective of other input and practices.

[8] Appendix table D1 displays the main results using the continuous skills measure instead, evidencing that the conclusions remain very similar. A binary index is used in the results presented in the main text to facilitate the reading and interpretation.

[9] Some treated farmers participated to the trials only in some seasons but not in others. In treatment villages, the first, second and third season of trials had participation rate of 91, 84 and 90% respectively. The main reasons for not participating were that they did not prepare their parcel for planting despite multiple warnings, or loss of interest.

$$Y_{it} = \beta_0 + \alpha T_{it} * LSF_i + \beta T_{it} * HSF_i + X_{it}\gamma + \varepsilon_{it} \quad (1)$$

Where Y_{it} is the outcome of interest, T_{it} is a dummy equal to one if farmer i is in a treatment village where the trials already started at period t, LSF_i and HSF_i are dummies for belonging to the group of low skill farmers or high skill farmers, respectively, X_{it} is a vector of controls that capture the baseline value of Y_{it} when available, dummies at the levels of stratification or heterogeneity (sublocations, ex-ante skill proxy groups and the LSF_i dummy) interacted with each round of data collection and ε_{it} is the error term, clustered at the village level. The coefficient α can thus be interpreted as the average effect of being offered to participate to the trials for LSFs, and β is interpreted as the treatment effect for HSFs. When the estimations are presented in tables, they encompass a test of the difference between α and β and of $\frac{\alpha+\beta}{2}$, the homogenous effect of the treatment, not separating by skill.

While stratification was based on the ex-ante skill quintiles, we use the more precise (ex-post) skill measure to split groups for the primary analysis investigating the heterogeneity of effects by skills. As robustness, we show in Appendix D that the results using the ex-ante proxy measure of skills lead to highly consistent results, with slightly noisier estimators. Appendix C displays a statistical analysis of the relationship between the ex-ante and more precise ex-post measure of skills. Additionally, Appendix Table C1 shows that the baseline skill level of farmers is correlated with many other characteristics, including age, gender, use of chemical fertilizer or soil conservation practices and benefits at baseline. While the empirical analysis will focus on differences by skill-level, in line with the theoretical model and the interest in learning, heterogeneity of impact by skill level could capture any of these observed, or indeed additional unobserved dimensions. Even if differences cannot be attributed to skill, it is arguably of interest to observe the different reactions to exposure to new information between the two types of farmers and many relevant policy implications do not require this attribution. Appendix Table D1 replicates all results without differentiating by skill level.

Besides the distinction between LSFs and HSFs treatment effects, we are particularly interested in observing the dynamics of the treatment effect over time. Hence our most common estimation separates the effects by number of seasons since the beginning of the treatment:

$$Y_{it} = \beta_0 + \sum_{s=1}^{s=5} \alpha_s T_{it}^s * LSF_i + \sum_{s=1}^{s=5} \beta_s T_{it}^s * HSF_i + X_{it}\gamma + \varepsilon_{it} \quad (2)$$

where T_{it}^{s} is a dummy = 1 if at time t, household i is in a village where the treatment started s seasons ago. Hence when s goes from 1 to 3, the evolutions of α_s and β_s reflect the increase in exposure to the trials, while when s goes from 3 to 5, the exposure does not increase anymore, and the coefficients α_s and β_s capture the dynamics after the end of the trials. Finally, as half of the treated farmers initiated their treatment in the short rain of 2013 and the other half initiated in the long rain of 2014, at any season s, half of the famers for which $T_{it}^{s} = 1$ are in a long rain season and the other half is in a short rain season. This design conveniently eliminates the seasonality from the interpretation of the dynamic effects.

For most of the outcomes of interest, Y_{it}, we use standardized indices aggregating information from a relatively large set of variables, which refer to the inputs, practices and outcomes measured on all the non-trial plots (based on plot-level data) of each farmer. See appendix C for definitions.

3.2 First set of results

Figure 3 illustrates both levels and treatment effects of being selected to participate in the trials, separated by LSF versus HSF and by number of seasons since the treatment started. The blue square presents the average of Y_{it} in the control group. The red diamond is the prediction of the average value of Y_{it} in the treatment group, obtained by simply adding the treatment effect to the average of the control group. The point estimates of the treatment effects are displayed next to each red diamond. The 90% confidence interval of the treatment group is obtained by adding the average of the control group to the confidence interval of the treatment effect. Hence a treatment effect is significant at the 90% confidence level if its confidence interval does not include the blue square. This presentation of the results facilitates the visual comparisons of both treatment effects and levels between LSFs and HSFs.

The upper left panel of Figure 3 shows the dynamic effects on know-how. As expected, even in the control group, HSFs have more know-how than LSFs. We observe a significant treatment effect on knowledge acquisition that grows over time between the first and third season of treatment (while the trials are implemented), and then stabilizes. The learning gains are large, reaching about half of a standard deviation after 3 seasons for the HSF. Strikingly, despite starting from a higher level of knowledge, HSFs learn faster and more than the LSFs.

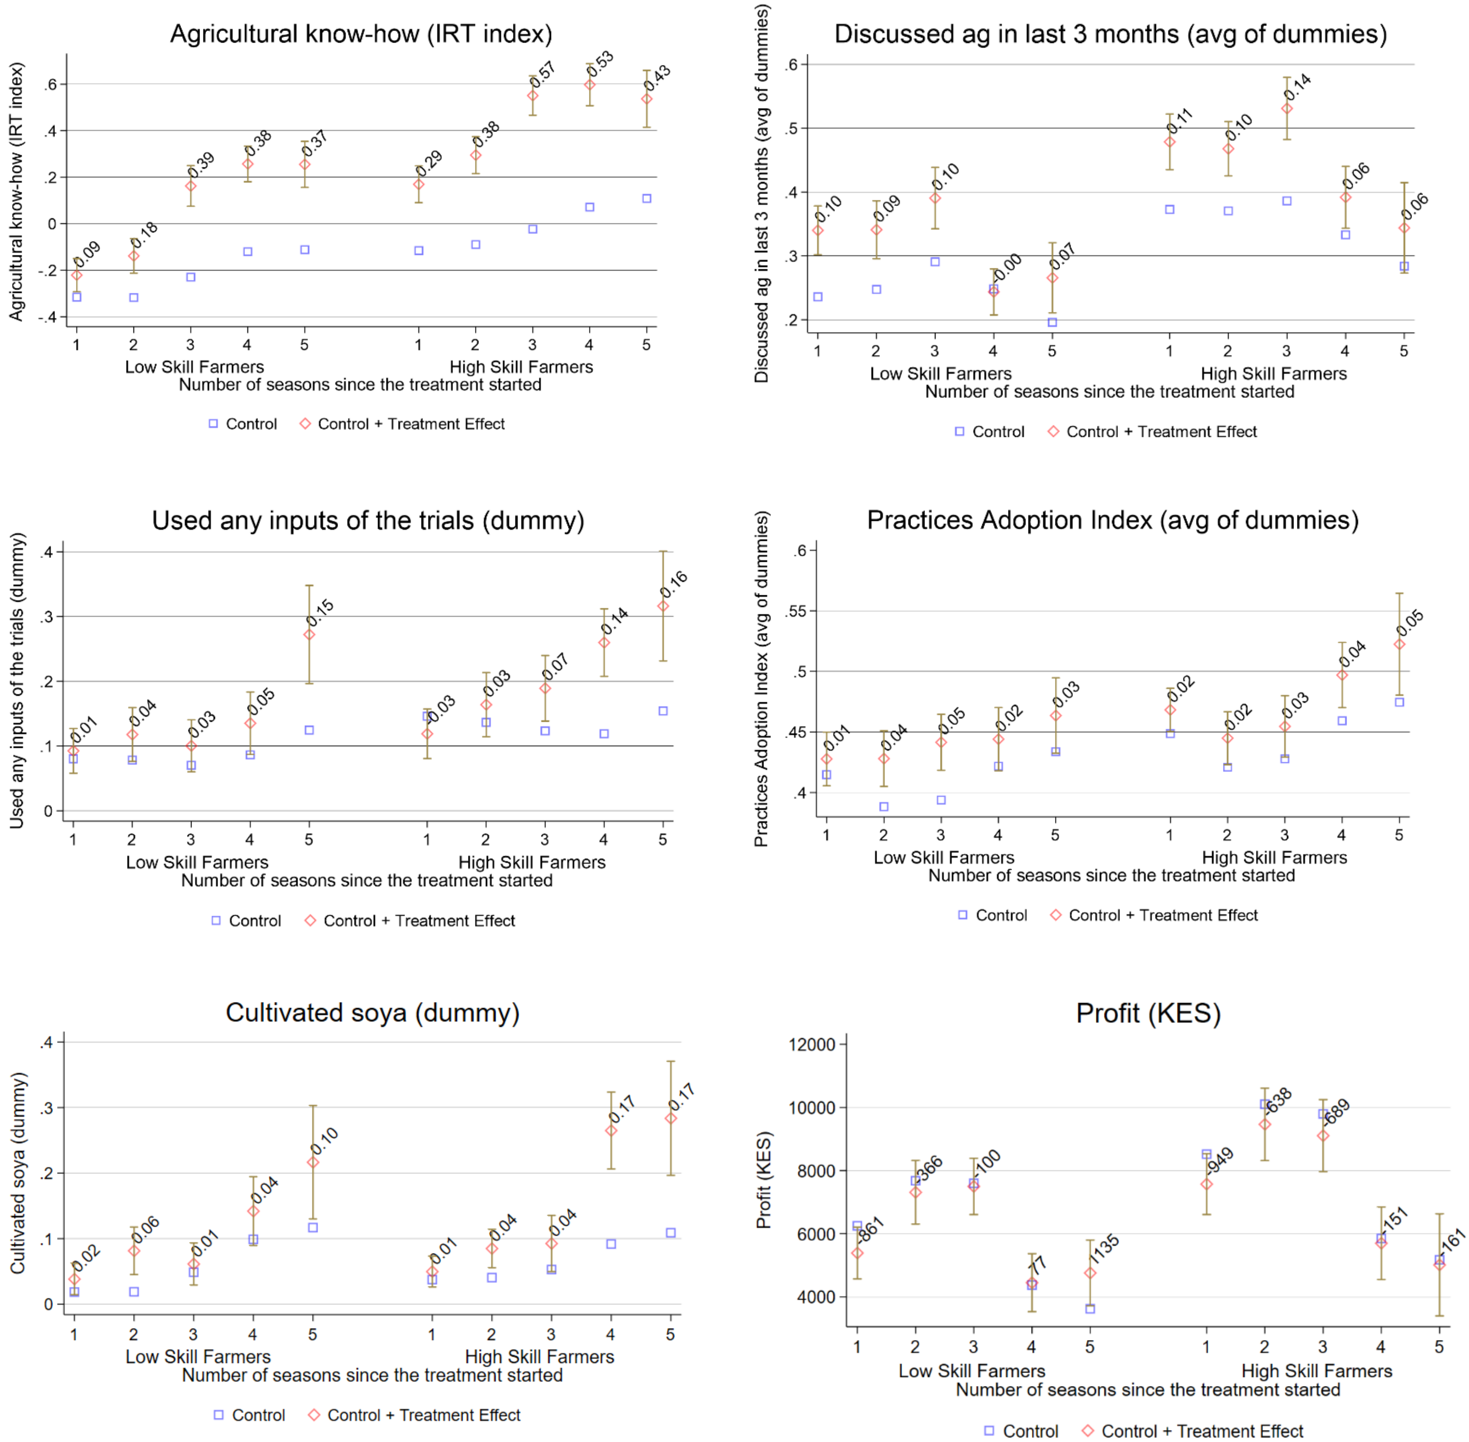

Figure 3 Treatment effects by skills and number of seasons since treated This figure displays the results from the estimation of equation (2). Blue squares represent the average among control farmers in the corresponding season and skill level. Red diamonds represent the expected value in the corresponding treatment group (average value in control group + treatment effect), the confidence bands around it represent the 90% confidence interval of the treatment effect (also adding the value in the control group to the lower and upper bound) and the value that appears next to each red diamond is the corresponding treatment effect.

The top-right panel displays the treatment effects on the likelihood of discussing agriculture with the other sample farmers in the same village. HSFs are substantially more likely to discuss agriculture at baseline, and there is a positive and significant treatment effect on the likelihood to discuss agriculture across almost all periods and skill levels. Discussion increased, in particular, during the trial seasons, but also, to a lower extent, after the trials are over. This points towards interest to learn at a time when there are more new technologies being experimented and more information to share, and suggests that learning may not only come from exposure to their own trials but also from discussing with others.

The middle panels of Figure 3 display the dynamic averages and treatment effects, by skill level, on adoption of inputs and practices. Key lessons from the middle panels are that the treatment effects are significant but initially modest in magnitude. Results for the binary variable of whether the farmer used any of the inputs tested in the trials in their own plots (excluding the agronomic trials), in the middle-left panel, show HSF tend to adopt before LSFs. By the last season, however, LSFs catch up in terms of input adoption. Practices (middle-right panel) tend to react faster than input use, but input adoption grows more steadily over time.

The bottom left panel analyzes effects on soya adoption, which is experimented in 6 out of the 10 trials per village (3 in soya trials and 3 in intercrop trials) and is new to most farmers at the time of baseline. We think of growing a new crop as a more drastic change, since it requires changing many practices and inputs at once. In this case, we also find significant and growing effects. By the 5th season after the treatment started, LSFs substantially increased their likelihood to grow soya, however, contrary to what we observed in the input adoption panel, they do not reach the level of HSFs who maintain treatment effects and adoption levels substantially higher than the LSFs. Interestingly, the most notable jump in input and soya adoption for HSFs occurs between seasons 3 and 4, just after the end of the trial, as if many farmers first waited to learn as much as possible from the trials before experimenting on their own. By contrast, LSFs have their largest jump in adoption during the 5th season, with a one-year lag compared to HSFs.

Finally, the bottom right panel shows the dynamic treatment effects on profits. While farmers' average profits remain positive, in the first period, for both LSFs and HSFs, the impact of exposure to trials on profits is negative and significant at the10% level.[10] However, the impact on profits progressively grows over time, becoming positive and significant at 10% level for LSFs in the 5th season, allowing the profit of LSFs to converge with the profit of HSFs. The positive trend of the treatment effect over time is significant (p=0.017).

Taking the results together we find two paradoxes. First, the estimated treatment effects on adoption of new inputs and the new crop increase over time even though profits are negatively affected at early stages. In a standard model of learning about the return to the new input (Conley and Udry 2010) adoption increases only if farmers are positively surprised when experimenting with the new inputs. Second, our results show a divergence in know-how, whereas a standard target-input model predicts a convergence in know-how between LSFs and HSFs, since ultimately, they all converge to the right target (Foster and Rosenzweig 1995, Bandiera and Rasul 2006). This could reveal a difference in how much information LSFs and HSFs acquire when exposed to a similar signal. The following section first lays out a model consistent with these findings, then uses it to derive additional predictions to take to the data, before analyzing additional outcomes to shed more light on the learning.

4. A model with multi-dimensional input decisions

This section lines out a model that provides a framing for the results and is consistent with most findings. We build on a learning model similar to the target input model, where forward looking farmers adjust their belief about a profit function when they observe new realizations of profit conditional on their input decisions. The most fundamental differences compared to the target input model are that the profit function is multidimensional and non-parametric, and that farmers' skills affect the precision with which the signals are perceived. We focus on the consequences of a costly learning process with a multi-dimensional profit function letting aside the numerous additional barriers to technology adoption and distinctions between LSFs and HSFs. We do not argue in any way that this is the only barrier to technology adoption nor the only driver of differences between LSFs and HSFs, but simply present a model that focuses on the aspects that we are most able to observe.

[10] The negative impact in profit is entirely driven by a reduction in the production value for LSFs while the increase in input costs is significantly higher for HSF and explains part of their reduction in profit (Appendix Table D1, columns 5 and 6).

In this theoretical model, farmers need to make a variety of input and practice decisions despite their limited and local information about the shape of the profit function. In this Bayesian optimization setting, each input decision generates a profit, which provides a signal at a given point of the profit function. HSFs and LSFs only differ in the precision of their signal: a HSF observes the signals with more precision and thus learns faster. The model helps characterize the initial stationary point, for a given farmer trying to maximize her long-term utility, before describing the expected reaction to an exogenous new signal about the profit function.

4.1.A Multi-Dimensional Profit Function

There exists a true profit function: $\pi(X) = f(X) + \varepsilon_{it}$, where $X = (x_1, x_2, \dots x_n)$ is defined over R_+^n and characterizes the n inputs and practices (with $n > 2$). Each continuous variable x_i can be the quantity of a commercial input (such as seeds, mineral fertilizer or bio-fertilizers), any practice decision (manure, spacing, dates of planting, weeding, soil conservation practices, etc.), or exogenous parameters (soil characteristics, climatic conditions, etc.).

The function $f(X)$ results from a Gaussian process. This stochastic process allows for flexibility in the shape of the function with convenient properties:

1) It is infinitely differentiable.
2) The covariance between two points X and X' is decreasing in the Euclidian distance between X and X'. This conveys the idea that closeness in the values of X increases the similarity in target values (Rasmussen and Williams 2006, chapter 4).
3) Every finite linear combination of the values at different points of X is normally distributed.

The second property guides the extent to which an observation at point X provides a signal that is informative about profits not only at this point but also in its vicinity in the form of a Kernel with a decreasing precision as the distance from X increases. The third property facilitates the use of Bayesian learning. Finally, the combination of the first and second property allows the farmers to also expect smoothness in the trend, hence if they find it to be increasing when moving in a certain direction it is more likely to keep being increasing when moving further in this direction (following a Bayesian optimization with gradients, as in Wu et al. 2017).

4.2 The farmer's learning and objective function

At any time t and any point X the farmer has the expected belief $f_t^b(X)$ about the profit function, with variance $V_t^b(X)$. The farmer's prior $f_0^b(X)$ has a large variance $V_0^b(X)$ in unexplored areas of the profit function. During a finite number of periods $t = 1, \dots, T$, at each period, the farmer decides all inputs $x_1, x_2 \dots x_n$ then observes the realization of the profit function $\pi(X_t)$ from which she derives an instantaneous utility $u(\pi(X_t))$ and updates her beliefs at and around X_t before moving to the next period. The instantaneous utility function $u(\pi)$ is increasing in π with Decreasing Absolute Risk Aversion (DARA).[11] For simplicity, we assume no savings nor any other way to smooth consumption over time.

At any given period, the farmer picks the vector of inputs X_t that maximizes her long-term utility during the current and future periods:

$$\max_{X_t} E\left[u(\pi(X_t)) + U_{t+1}\left(f_{t+1}^b(X_t), V_{t+1}^b(X_t)\right)\right] \quad (3)$$

where $U_{t+1}^b\left(f_{t+1}^b(X), V_{t+1}^b(X)\right) \equiv \; E[\sum_{s=t}^{T} \delta^{s-t} u(\pi(X_s)) | f_{t+1}^b(X), V_{t+1}^b(X)]$ represents the present value of utility at time t given that the farmer holds the beliefs $f_{t+1}^b(X)$ and $V_{t+1}^b(X)$, which are updated, after observing $\pi(X_s)$. And δ is the time discount.

Below we often refer to the utility gain from information, which we define as:

$$E[\sum_{s=t+1}^{T} \delta^{s} u(\pi(X_s)) | f_{t+1}^b(X), V_{t+1}^b(X)] - E[\sum_{s=t+1}^{T} \delta^{s} u(\pi(X_s)) | f_t^b(X), V_t^b(X)] \quad (4)$$

In words, it is the present value of the expected gain in utility that comes from the update of $f_t^b(X)$ and $V_t^b(X)$ when $\pi(X_t)$ occurs at X and its vicinity, given that the additional information will allow more informed input choices in the future. For simplicity, we first assume that the farmer makes decisions in isolation from other farmers, until we simulate an exogenous and unexpected signal and consider its consequences on the farmers' decisions.

4.3 LSFs and HSFs differ in the precision of their signals

Central to our analysis is the assumption that HSFs learn faster than LSFs because they observe the signals with a higher precision. Hence, in a situation where a HSF and a LSF start with the same $\left(f_{t+1}^b(X), V_{t+1}^b(X)\right)$ and experience the same $\pi(X_t)$, then for the HSF, at X and

[11] A utility function exhibits DARA if its absolute risk aversion is decreasing in wealth or consumption. This is a standard assumption and englobes most commonly used utility functions, including the Constant Relative Risk Aversion.

its vicinity, the expected value of $f_{t+1}^{b}(X)$ will be, in expectation, closer to the true $f(X)$ and the variance at $V_{t+1}^{b}(X)$ will be lower than the one of the LSF. Therefore, ceteris paribus, the expected utility gain from information is greater for the HSF than the LSF. In appendix E, we explain how the difference in the precision of the signal can result in a behavioral inattention that affects both farmers, but the HSFs have a larger attention span that allows them to observe more of the exogenous factors that affect $\pi(X)$ than the LSFs. The assumption is empirically supported by the difference in the treatment effects on know-how in Figure 3.

4.4 Resolution through Bayesian Optimization

Figure 4 illustrates the Bayesian optimization search, for a true profit function $f(X)$ with the confidence interval of the beliefs about its value, in a case where a farmer observed five realizations of $\pi(x)$. We use the figure, where X is a one-dimension vector, to provide a basic intuition of a Bayesian optimization, however it is important to remember that the actual optimization occurs in a multidimensional space X, substantially adding to the difficulty of the farmer's learning.

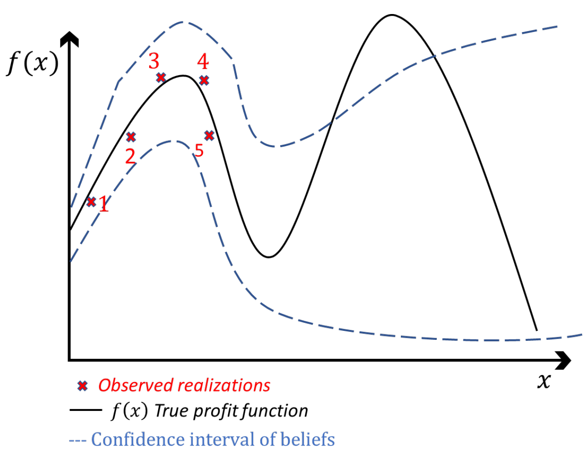

Figure 4: graphical representation of a Bayesian optimization search. The graphical representation is limited to a one-dimensional x to facilitate understanding, but the multidimensionality of the area to be explored is central to the model.

The Bayesian optimization consists in deciding, at each period, the vector X_t that offers the highest gain in expected utility from both the immediate utility derived from $\pi(X_t)$, and from the present value of the information gain from observing the realization of $\pi(X_t)$ and updating one's beliefs, as defined in equation (3). The optimal analytical solution is obtained by backward induction. [12] However, because it rapidly becomes too complex to be solved analytically, the programming literature for the optimization of the search of such black-box function, circumvents this complexity by using forward induction rules of thumb that approximate the optimal solution.[13] We borrow methods from this programming literature, but slightly modify it to reflect our distinct goal: while the programming literature aims at finding the global optimum, we aim to represent the behavior of a risk-averse farmer with a relatively small T, for whom each realization of π_t has large immediate consequences on the farmers' well-being, leaving her less room to explore and find the global maximum.

At each period, an acquisition function $AF_t(X)$ approximates the expected utility gain from exploring X. The Upper Confidence Bound (UCB) is one of the simplest and most commonly used acquisition functions.[14] The focus on the upper bound reflects the idea that areas of X with some possibility of a high realization have the highest exploration value. After incorporating risk aversion, we obtain:

$$AF_t(X) = \underbrace{E[u(\pi(X))]}_{Exploitation\ gain} + \underbrace{\alpha_{Ft}[f_t^b(X) + \lambda\, V_t^b(X)]}_{Exploration\ gain} \quad (5)$$

The first term of the right-hand side is the immediate expected utility gain of the risk averse farmer, which is increasing in $f_t^b(X)$ and decreasing in $V_t^b(X)$. The second term is the upper bound of the confidence interval, for example, if one choses $\lambda = 1.96$ then it is equal to the upper bound of the 95% confidence interval. α_t reflects the weight on exploitation relative to exploration, which is increasing in the farmers' patience and number of remaining periods. Also, to reflect the fact that a HSF acquires more information when exploring, and thus derives a higher utility gain from information, we set $\alpha_{HSF\ t} > \alpha_{LSF\ t}$, as the only difference in the model between LSFs and HSFs.

The farmer has 3 possible strategies. First, she can explore in the vicinity of the values of X where a realization was observed, making use of the reduction in the variance of the belief and of the information about the possible gradient around that point to explore in the direction that is

[12] See Lattimore and Szepesvári (2020) for a description of bandit problems, their resolution and standard trade-offs.
[13] For an overview of Bayesian optimization methods, see Garnett (2023).
[14] We chose the UCB for its analytical simplicity. Results are robust to any function with a marginal benefit from variance that is either linear or concave. Some of the more recent machine learning algorithms are shown to outperform the UCB method. Yet assuming UCB already probably overestimates the farmers ability to search close to optimally. We adapted the original UCB formula by including the concave utility function in the exploration part of the equation. See Contal et al. 2013 for evidence of the performance of Upper Confidence Bound when maximizing an unknown function for which evaluations are noisy and are acquired with a high cost.

most likely to be increasing. In Figure 4, this is consistent with the input decisions from 2 to 5. Second, the farmer can decide to stop exploring and exploit what appears to be the best option given the information available, applying minor adjustments as more information is revealed. In Figure 4, it is equivalent to selecting input decisions between X_3 and X_4 in the following periods. Finally, the farmer can decide on a third strategy which we refer to as a "jump into the wild" in which case the farmer explores an area where almost no information is available beyond the prior. This is equivalent to exploring the right side of Figure 4. In this model such "jump into the wild" can occur, either at very early stages, when the farmer only found a very unsatisfactory local equilibrium, or if an exogenous shock provides information regarding a possibly valuable new combination of inputs and practices in an unexplored area of X (e.g. a shock resulting from observing the results of the trials).

4.5 Summary of the propositions and their implications in our context

We derive 7 propositions that describe what to expect at equilibrium and then once a farmer is exposed to a new signal, such as the one of the trials. The 7 propositions are listed and discussed within the text below, emphasizing their implications in our specific context. Their formal statement, proof and intuitive understanding are developed in appendix F.

The results first describe the initial equilibrium, highlighting that, with enough time, farmers converge to a local maximum (proposition 1).[15] Hence farmers are good at finding a local maximum through finetuning and adjusting their decisions, but may be stuck in a local maximum that can be low compared to other maxima of the function if this requires making larger changes in their inputs or multiple changes at once. Because HSFs are better at acquiring information, they are expected to reach an equal or higher stationary point than LSFs (proposition 2). The curse of dimensionality results from the presence of complementarities and substituabilities in $f(X)$ (proposition 3). In practice if using a new seed or fertilizer, may require re-optimizing other inputs and practices, then this increases the likelihood of committing costly mistakes along the way. This will orient our empirical approach towards testing the risk of committing such mistakes.

[15] The fact that local optimization algorithms are often trapped in a local optimum that is highly dependent on the starting point is a common finding (Wu et al. 2017).

We then explore the consequences of receiving new exogenous information about $f(X)$ at some unexplored area of X. The signal can trigger an exploration of the new area if its precision is strong enough and if the update reveals sufficiently high profits for the Acquisition Function at that point to exceed the one of the stationary points (proposition 4). The exploration can result in loss in expected utility in the short term, which comes from abandoning the farmers' "comfort zone", where the local optimum was identified, to investigate a new area, at the risk of making costly mistakes while trying to find a higher local maximum. HSFs can tolerate more immediate utility reduction than LSFs as their utility gain from knowledge acquisition is higher and they are willing to explore, as long as this gain exceeds the immediate reduction in expected utility (proposition 5).

Interestingly, we also find that the model makes it possible for subjective expectations about profit to move in a non-monotonic way. Assume that a farmer first performs poorly when trying a new input because of the lack of knowledge of complementary practices and inputs, but this productivity improves as these decisions are adjusted, then this can be reflected in the evolution of her subjective expectations (proposition 6). This prediction is not typical of models where the farmers underestimate the true (given) benefit of a technology nor in a pure learning-by-doing model. We find that it requires both elements (learning about the returns and increasing know-how) to generate the possible non-monotonicity of the expected evolution of subjective expectations over time. Finally, our model highlights that, when observing other farmers, homogamy can be valuable (if farmers with similar skills level tend to have similar practices). A LSF may also particularly value observing a HSF, however, given that she is more likely to explore and provide valuable knowledge in unknown areas, potentially pointing to a path towards a higher local maximum (proposition 7).

5. Additional results

The theoretical model was intentionally built to fit empirical patterns in section 3 and allows to provide a sensible explanation for the set of results. In this section we test the model's implications for additional outcomes, making use of the framework to guide the analysis and interpretation of the results.

5.1 Additional Outcomes: exploration and subjective expectations

Among the predictions of the model, proposition 4 foresees that if the trials provide sufficient information, they should trigger more exploration and proposition 6 opens the possibility of a non-monotonic change in subjective expectations about profit. Figure 5 shares the same format as Figure 3, displaying the results of equation (2) for a new set of outcome variables aimed at capturing those predictions.

The upper left panel shows the fraction of the farmers' input decisions that was modified compared to the last season, the upper right panel displays a similar index for changes in practices (a more detailed explanation is provided in Appendix C2). The means in the controls show there are substantial changes in practices (23 to 25%) each season, while changes in inputs are more limited (6 to 8%). Exposure to trials significantly increases changes in inputs and practices for both LSFs and HSFs. Hence, we conclude that the treatment triggers more experimentation. This increase in experimentation seems to be initiated by HSFs before LSFs and remains at least as large by the fifth season after the beginning of the treatment. According to the model, the fact that farmers are still exploring is consistent with the fact that profits have not yet reached their new local maximum and can help understand why the effects on HSFs' profits are still not positive by the 5^{th} season.

The bottom panels of Figure 5 show the treatment effects on subjective expectations about yields with the new inputs versus without them (see appendix on how this differential expectation was measured). Effects on expected returns are very limited during the first three seasons and remain low for the LSFs. Expected returns become somewhat larger for HSFs in the last two seasons (i.e. after completion of 3 seasons of the trials), though we only find a significant positive treatment effect on subjective expectations for HSFs with the soya input (right panel) package during the last two seasons. This is consistent with Laajaj et al. (2020)'s finding that the soya package is more responsive to the quality of management than the maize package. It is thus possible that only HSFs had a quality of management that was high enough to see good profits and adjust their beliefs upward.

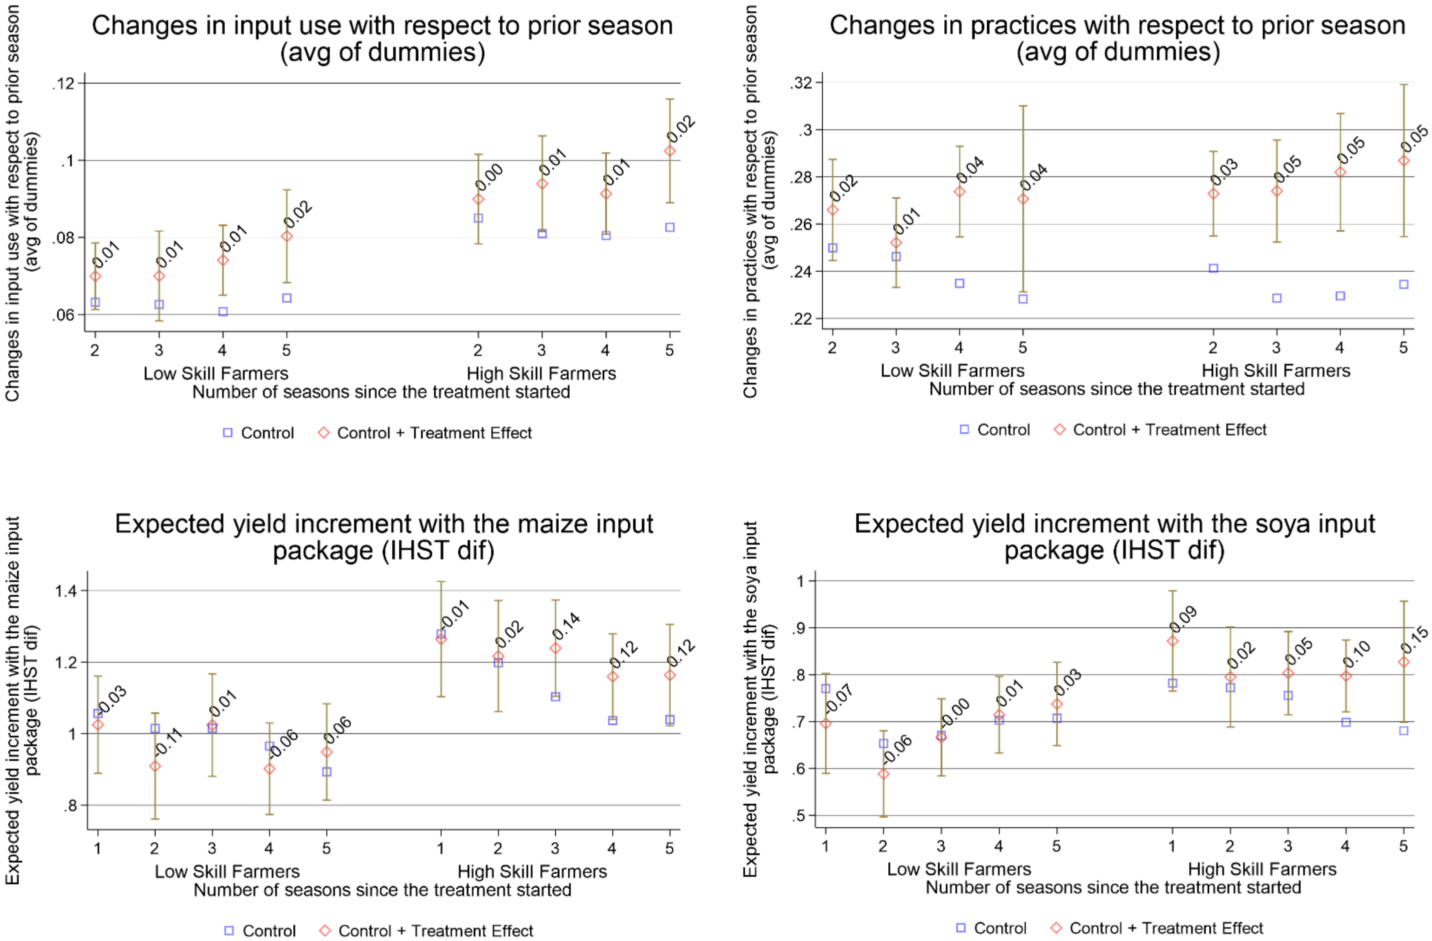

Figure 5 Treatment effects by skills and number of seasons since treated on additional outcomes. This figure displays the results from the estimation of equation (2). Blue squares represent the average among control farmers in the corresponding season and skill level. Red diamonds represent the expected value in the corresponding treatment group (average value in control group + treatment effect), and the value that appears above each red diamond is the corresponding treatment effect, the confidence bands around it represent the 90% confidence interval of the treatment effect (also adding the value in the control group to the lower and upper bound) and the value that appears next to each red diamond is the corresponding treatment effect.

Technology adoption models often assume a better technology implies an expected upward adjustment of the belief, as soon as they can observe a new technology (Conley and Udry 2010). In our context, adoption grows, however, despite initial effects on the expected beliefs about profits that are non-significant (and often negative). While these results need to be caveated by the possibility that noise in the expected yield measures is obscuring some effects, the findings are consistent with our model, which allows for more complex stories and movements than the standard model.

5.2 Subtle learning and complementarities

Proposition 3 highlights that the curse of dimensionality results from complementarities and substituabilities between inputs. In practical terms this reflects the fact that, as one changes a given input, the right decisions about other inputs and practices can become different, hence these need to be adjusted as well. Until farmers have found the right combinations, they are prone to making "mistakes" i.e. using combinations of inputs that lower profits compared to what farmers would do if they had access to better information (even though farmers are behaving optimally given their knowledge at time t and objective function). Based on what we know from the agronomists' expertise and from the trials results, we can dive into specific decisions and input combinations to see how the treatment affects the likelihood of such mistakes. This requires describing each variable before analyzing the corresponding results. For simplicity, Table 1 displays the results of estimating equation (1), pooling together the treatment effect over different seasons.

In the WTP section (see appendix for details), we ask what inputs the farmer would buy for maize production under two alternative scenarios: a plot with high level of striga and a plot with low striga infestation. Striga is a widespread weed in Western Kenya and very harmful to maize production. One of the seeds tested in the trial (the IR seed) is coated with a herbicide that makes it resistant to striga (there is therefore a complementarity between the presence of striga and the use of IR seeds). A farmer who understood the purpose of IR seeds should be more likely to select it under the scenario of high striga compared to the scenario of low striga. We therefore can take the difference in the use of IR seed under the two scenarios to test the farmers' understanding about the complementarity.[16] The results in column 1 of Table 1, show the effect of the treatment on adapting well the inputs to the conditions is significant and more than twice higher for a HSF (0.23) than for LSF (0.1). Hence it appears that this type of knowledge is acquired much faster by HSFs than LSFs.

Second, we draw on the fact that, in the soya trials, a fertilizer tested on some of the subplots (Sympal) strictly dominated the fertilizer tested on other subplots (Minjingu). Farmers who perceived this difference should select Sympal rather than Minjingu. Using the WTP

[16] While the non-incentivized WTP method is prone to experimenter demand effects, that should cancel out when we compare two different inputs tested on the trials, or in answers to questions asking WTP under two different conditions. As such it can tell us what farmers believe is the right choice, exempt from financial and input availability constraints.

questions for soya, column 2 shows that the treatment increases very significantly the likelihood of selecting Sympal rather than Minjingu for HSFs but not for LSFs. This arguably points to subtle learning, as it requires being able to compare the yields of two new inputs that each improve yields, which should have a lower signal-to-noise ratio than comparing yields with inputs to yields without inputs. We conclude that some subtle learning appears to be out of reach for LSFs and to be acquired only by HSFs.

Table 1: Subtle learning and the effect of exposure to new information on potential mistakes

Concepts	Adapting inputs to conditions	Picking the best of 2 inputs	Proper combinations of inputs and practices			Wrong combinations of inputs and practices	
VARIABLES	WTP IR seeds if high striga minus low striga	WTP Sympal vs Minjingu	Used hybrid and one seed per hole	Used both commercial and homemade fertilizer on at least one plot	WTP Biofertilizer together with Sympal (N & P sources)	Used hybrid maize seed from own production	WTP Sympal and Mijingu (2 P sources together)
Treatment * LSF	0.098***	0	0.04	-0.047*	0.007	0.026	0.032***
	(0.022)	(0.023)	(0.026)	(0.026)	(0.008)	-0.016	(0.008)
Treatment * HSF	0.229***	0.077***	0.065**	-0.041	0.070***	0.040**	0.080***
	(0.034)	(0.029)	(0.032)	(0.025)	(0.015)	(0.018)	(0.012)
Observations	3,742	2,893	4,431	4,552	2,893	4,431	2,893
Avg outcome in LSF control grp	0.0534	0.0202	0.245	0.345	0.007	0.107	0.00119
Avg outcome in HSF control grp	0.174	0.0249	0.442	0.49	0.012	0.142	0.00262
P-val of Treat. (low & high sk.)	0	0.0478	0.0119	0.0139	0.0002	0.00198	0
P-val Treat. * LSF = Treat. * HSF	0.0007	0.0302	0.559	0.868	0.0000	0.587	0.0013

Table 1 presents the results of estimating equation (1), which pools the treatment effects of all seasons together, but separates the treatment effects by LSFs and HSFs. Sections that use the WTP questions regarding maize (soya) inputs were only administrated to farmers who received the maize (soya) treatment, the intercrop treatment and a similar sized, random subset of the control group, hence the smaller number of observations on regressions with outcomes variables from this section. Standard errors, clustered at the village level, are in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

Columns 3 to 5 test whether exposure to the trials increased the likelihood of applying proper combinations of inputs and practices (according to the expert agronomists). First, farmers who use hybrid seeds are recommended to sow only one seed per hole, while with traditional seeds, more seeds per hole are recommended to compensate for their poor germination rate. Second, using chemical and non-chemical fertilizers together is one of the core principles of ISFM. Third, as done in the trials, combining a source of Nitrogen (like biofertilizer) and a source of Phosphorous (like Sympal or Minjingu) is the usual recommendation because of the expected complementarity between the two macronutrients. We find mixed effects of the treatment on these 3 recommended combinations. Treatment farmers appear more likely to adjust the seed and number of seed per hole together (only significant for HSF). On the other hand, they react to the treatment by reducing their joint use of chemical and commercial fertilizers and only HSFs significantly increased their combined use of a Nitrogen and Phosphorus sources. Hence both LSFs and HSFs catch the right combinations in some cases but seem to not follow the practices that agronomists believe to be best in other.

Finally, the last two columns look at how the treatment affects the likelihood of committing mistakes due to the lack of understanding of complementarities and substituabilities between the different inputs and practices. First, while it is fine to re-use one's seeds when a farmer uses traditional seeds, it can be harmful to the production (and profits) if a hybrid seed is re-used for planting. Yet the treatment significantly increased the likelihood of re-using hybrid seeds, likely because some farmers switched to hybrid seeds without realizing that they needed to modify the practice of re-using the seeds. Finally, the last column builds on the fact that, as shown in the soya trials, farmers only need to use one P-source (of phosphorus), such as Sympal or Minjingu, but combining 2 P-sources is not recommended for profit maximization. We see however that farmers in the treatment groups are 8 percentage points more likely (for HSFs) to want to combine Sympal and Minjingu. The results in the last 2 columns suggest that HSFs are most likely to commit these two mistakes. This is consistent with HSFs being the ones who explore the most and poit to a possible explanation for why HSFs experience a significant profit loss following the treatment. While these are only a small proportion of the many "mistakes" that can occur as a farmer gets out of her comfort zone to experiment a new package, they are illustrative that at least some mistakes may be more frequent during this exploration and possibly constrain technology adoption.

5.3 Adoption in the intensive margin

The results on adoption presented so far use dummies (or the average between dummies), reflecting the fact that in real life, the fixed cost of adoption of a new input (or practice) contributes to the challenge of adopting any positive quantity. As the theoretical model describes an adoption and learning process driven by variation in the intensive margin, we also estimate impacts on the total amount spent on trial inputs and the area dedicated to soya cultivation.

Figure 6 displays a visual representation of households' "trajectories" over the five seasons among adopters in the treatment group in terms of the value of trial inputs (with the parallel figure for area dedicated to soya in Figure D4). Trajectories are diverse, however an upward trend over time appears to stand out for most farmers. This could indicate farmers progressively getting closer to optimum as they explore the profit function. The patterns also suggest farmers tend to experiment on a small scale before expanding, possibly as they become more confident about the potential benefits of the new technology and how to use it.

Trajectories in amount spent on trial inputs among adopters in treatment

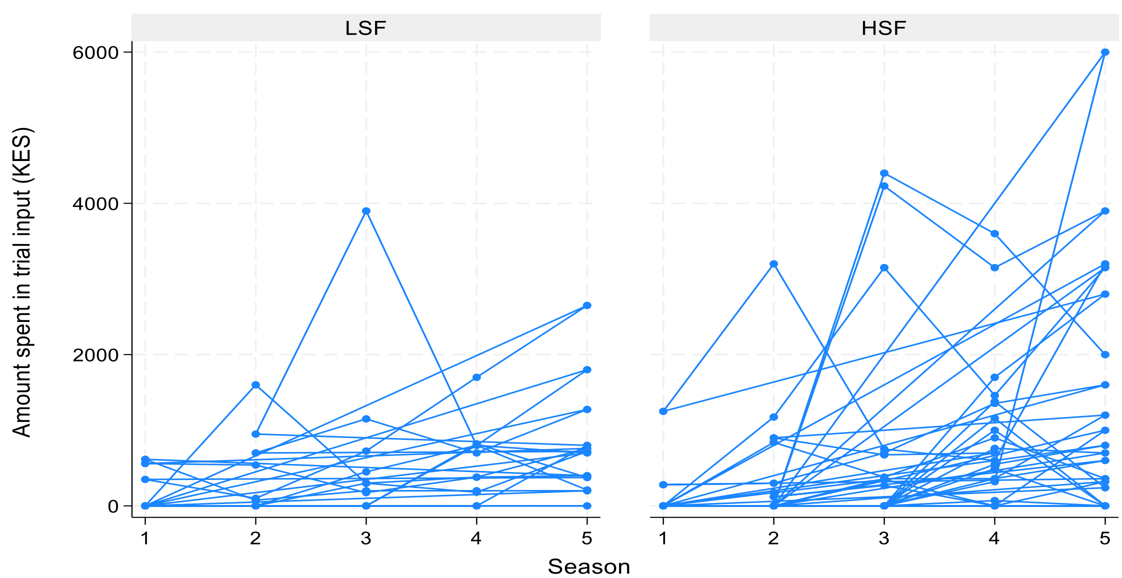

Figure 6: Each blue line connects a given household across the 5 seasons of the study. Only treated farmers are included (and only the ones with a strictly positive value in the amount spent on trial inputs in the last 2 seasons.

Appendix Figure D2 shows that the effects by season and skill level of these continuous variables are very similar to the adoption results in Figure 3. Additionally, the histograms of Figure D3 show that at the intensive margin, the two variables tend to increase over time.

5.4 Learning from oneself and from others

5.4.1 Learning from one's own experience in the trials

At the core of the theoretical model is the hypothesis that HSFs are better than LSFs at extracting valuable information from their own observations of profit realizations. Figure 3 provided evidence that HSFs tend to acquire know-how faster than LSFs. To complement we analyze whether the input use decisions of the farmers are consistent with the results obtained in their own agronomic trials. First, we estimated the yield increment observed in the agronomic trial of each farmer by comparing the yield in the most promising sublots to the control subplot (see details on the calculation of yield increment in appendix A). While the yield increments that a farmer observes from these small subplots can be noisy proxies for what actual yields would be at a larger scale, it provides a direct measure of the signal that farmers get and are expected to respond to. We then regress the decision to purchase the corresponding new inputs on this yield increment (averaged over all prior seasons for each period) with interactions that separate the effect for LSFs and the one for HSFs. These regressions are not as cleanly identified as the previous estimations, as they no longer only use the random treatment assignment, but the actual realizations in the trials. Yet they are suggestive of whether decisions are consistent with what would occur if farmers were extracting a signal from the profit realizations in their own trial.

Table 2: Input decisions and yield increment from one's own trials

VARIABLES	used any maize inputs of the trials (dummy)	used any soya inputs of the trials (dummy)
Maize yield increment in trial* LSF	-0.0145	
	(0.0575)	
Maize yield increment in trial* HSF	0.196***	
	(0.0485)	
Soya yield increment in trial* LSF		-0.00551
		(0.0375)
Soya yield increment in trial* HSF		0.0707**
		(0.0320)
Observations	1,412	1,141
Dep. var. mean	0.0880	0.0111
p-val LSF effect = HSF effect	0.0156	0.137

This table estimates whether the decision to use a maize (soya) input that was tested in the trial is affected by the yield increment observed in the agronomic trial. Yield increment refers to the difference between treatment subplots and the control subplots in the agronomic trials (a proxy for the performance of the tested inputs in this context). In each season we estimate the average yield increment over all preceding seasons to consider all prior realizations observed by the farmer.

Table 2 shows a striking difference between LSFs and HSFs. For both soya and maize, we find that HSFs who obtained a higher yield increment in trial sublots where the new inputs were used (compared to the control subplot) were significantly more likely to use at least one of these inputs in the subsequent seasons. By contrast for LSFs, we find no relationship between the yield increment in the trials and use of the corresponding inputs.[17] The results are consistent with the HSFs (but not LSFs) extracting information from their own trials and acting accordingly. Recall also that LSFs do adopt inputs, although with some delay compared to HSFs (Figure 3). The results in Table 2 provide further evidence that LSFs tend to imitate rather than learn from the observation of their own trials to guide their decisions. Again, these findings are only suggestive, since we cannot fully exclude the presence of unobservable characteristics that are correlated with higher yields in the trials and with the purchase of new inputs (though it would need to be the case only for HSFs but not for LSFs to generate these heterogenous results).

5.4.2 What is learned from others (and what is not)?

We now make use of the fact that, within a treated village, different farmers participated to different trials, and hence were expected to learn about different inputs and practices directly from their trials. To do this, we estimate the effect by crop treatment, using the following specification:

$$Y_{it} = \beta_0 + \sum_{(c=m,s,i)} \alpha_c \, T_{it}^c * LSF_i \; + \sum_{(c=m,s,i)} \beta_c \, T_{it}^c * HSF_i + X_{it} \, \gamma + \varepsilon_{it} \quad (3)$$

T_{it}^c is the crop treatment where c could be either maize, soya or intercropping (of maize and soya) and specifies the crop that was tested in the trials assigned to farmer i. Hence α_c and $\beta_c$are the treatment effects for a LSFs and HSFs respectively, to whom a trial with crop c was assigned. Y_{it} is the outcome, which tends to be crop specific in this set of regressions.

Given that farmers interact, and in fact significantly increased their interactions as a reaction to the treatment (as shown in Figure 3), a farmer who receives treatment T^c could also learn about optimal practices and inputs for a different crop than c through her interaction with the farmers in the same village with trials for the other crops. The main purpose of this estimation is to tease out this indirect knowledge acquisition from direct knowledge acquisition, based on the assumption that maize (soya) specific knowledge is directly acquired only by farmers assigned to a maize (soya) trial or intercrop trial. Figure 7 displays the results of

[17] The distributions of yield increment for LSFs and HSFs are displayed in figure A2. It shows somewhat similar averages and spread among LSFs and HSFs, hence the difference in adoption levels of the trial inputs and predictive power of yield increment on adoption cannot be explained by differences in their distribution.

equation (3), following a format similar to the ones of Figures 3 and 5, though instead of separating treatment effects by the number of seasons since the beginning of the treatment, it separates them by crop treatment.

The top and middle panels report treatment effects on a maize specific know-how index and a soya specific know-how index (each one using the relevant crop-specific subset of items from the aggregate know-how index). The color scheme aims to facilitate the reading: if the crop treatment is written in green (orange), the effect is expected to be direct (indirect).[18] For example if farmers who were assigned to the maize trials improved their maize cultivation know-how, we can expect this learning to come from their own exposure to the trials (and perhaps also from the experience of others) but if we observe an increase in maize know-how among farmers assigned to the soya trials then this points to a spillover from other farmers in the village who had the maize or intercrop treatment. The top panel shows a high level of spillover, as treatment effects on know-how that are not significantly different among directly treated farmers and those who could only have acquired the corresponding knowledge from the experience of their peers. This conclusion stands for LSFs as well as HSFs and for maize as well as soya specific know-how.

To investigate whether spillover in knowledge translates in spillover in input adoption, the middle panel displays the treatment effects on the adoption of the new inputs separating the inputs tested in maize trials from the ones tested in soya trials. The results are strikingly different from the top panel. The treatment effects on maize (soya) trial inputs are only significant among farmers assigned to the maize (soya) treatment and, to some extent, the intercrop treatment. In other words, in this context, farmers learned a lot of know-how from each other, but when it came to buying and using the new inputs, direct exposure to the trials was necessary.

If farmers acquired know-how both through direct and indirect exposure, then why was this insufficient for indirectly exposed farmers to adopt the corresponding inputs? To shed light on this question, in the bottom panel, we inspect whether farmers with direct and indirect exposure to a crop were also able to acquire the more subtle knowledge specific to that each, looking at crop-specific outcomes of Table 1. The bottom left looks at the crop treatment effects on the choice of IR maize seeds in the presence of striga minus without striga (see explanation in section 5.2). This knowledge should be acquired in maize trials, and we see indeed very

[18] No color is assigned to the intercrop given that it can be somewhat in between with partial learning made possible for either maize specific or soya specific knowledge.

significant effects of the maize treatments on this outcome variable. By contrast the soya treatment effect (associated to indirect learning) is not significant among LSFs and significant but substantially lower among HSFs. Similarly, we look at whether farmers chose the Sympal P-source rather than Minjingu and we find that only HSFs who received the soya treatment significantly increased their likelihood to pick the higher-performing P-source (Sympal).

We conclude that more complex lessons are far more likely to be acquired only with direct exposure (and only by HSFs for the most difficult ones). This result provides insights on multiple puzzles that our initial results seemed to raise. First, the fact that some kinds of knowledge require both direct exposure and high skills can explain the divergence in know-how over time. Second, it becomes more understandable that farmers who did not have this direct exposure may not feel ready to use the new inputs. This is consistent with proposition 6 of the model, emphasizing that the precision of the new information needs to be sufficient to trigger a new exploration, which may not be the case when farmers are indirectly exposed. This points to a potentially important limitation for diffusion of complex technology bundles through networks.

Within this sample the lack of adoption of maize (soya) inputs among farmers who received the soya (maize) trial treatment could be because farmers prioritized learning about their assigned crop treatment. Table D2 therefore displays the spillover effects on 5 additional randomly selected farmers in each village (outside of the 10 farmers that participated to the study) surveyed during the last 2 seasons. Results show modest spillovers in learning outcomes but no spillover in the adoption of trial inputs, allowing us to reject this alternative explanation.

5.4.3 Learning from whom?

Proposition 7 points to the possibility that LSFs may value learning from HSFs over homogamy. While lessons from agents who are more similar are more likely to readily apply (since they have more similar $X_i s$), the LSFs may still value learning from HSFs more if they particularly value information on parts of the profit function that they are less familiar with, and that HSFs are more likely to explore these areas. To investigate this, we asked each farmer about her connections and interactions with each of the nine other sampled farmers in the village. We use this data to learn about how much LSFs and HSFs are aware of what other farmers are doing and who they pay attention to the most.

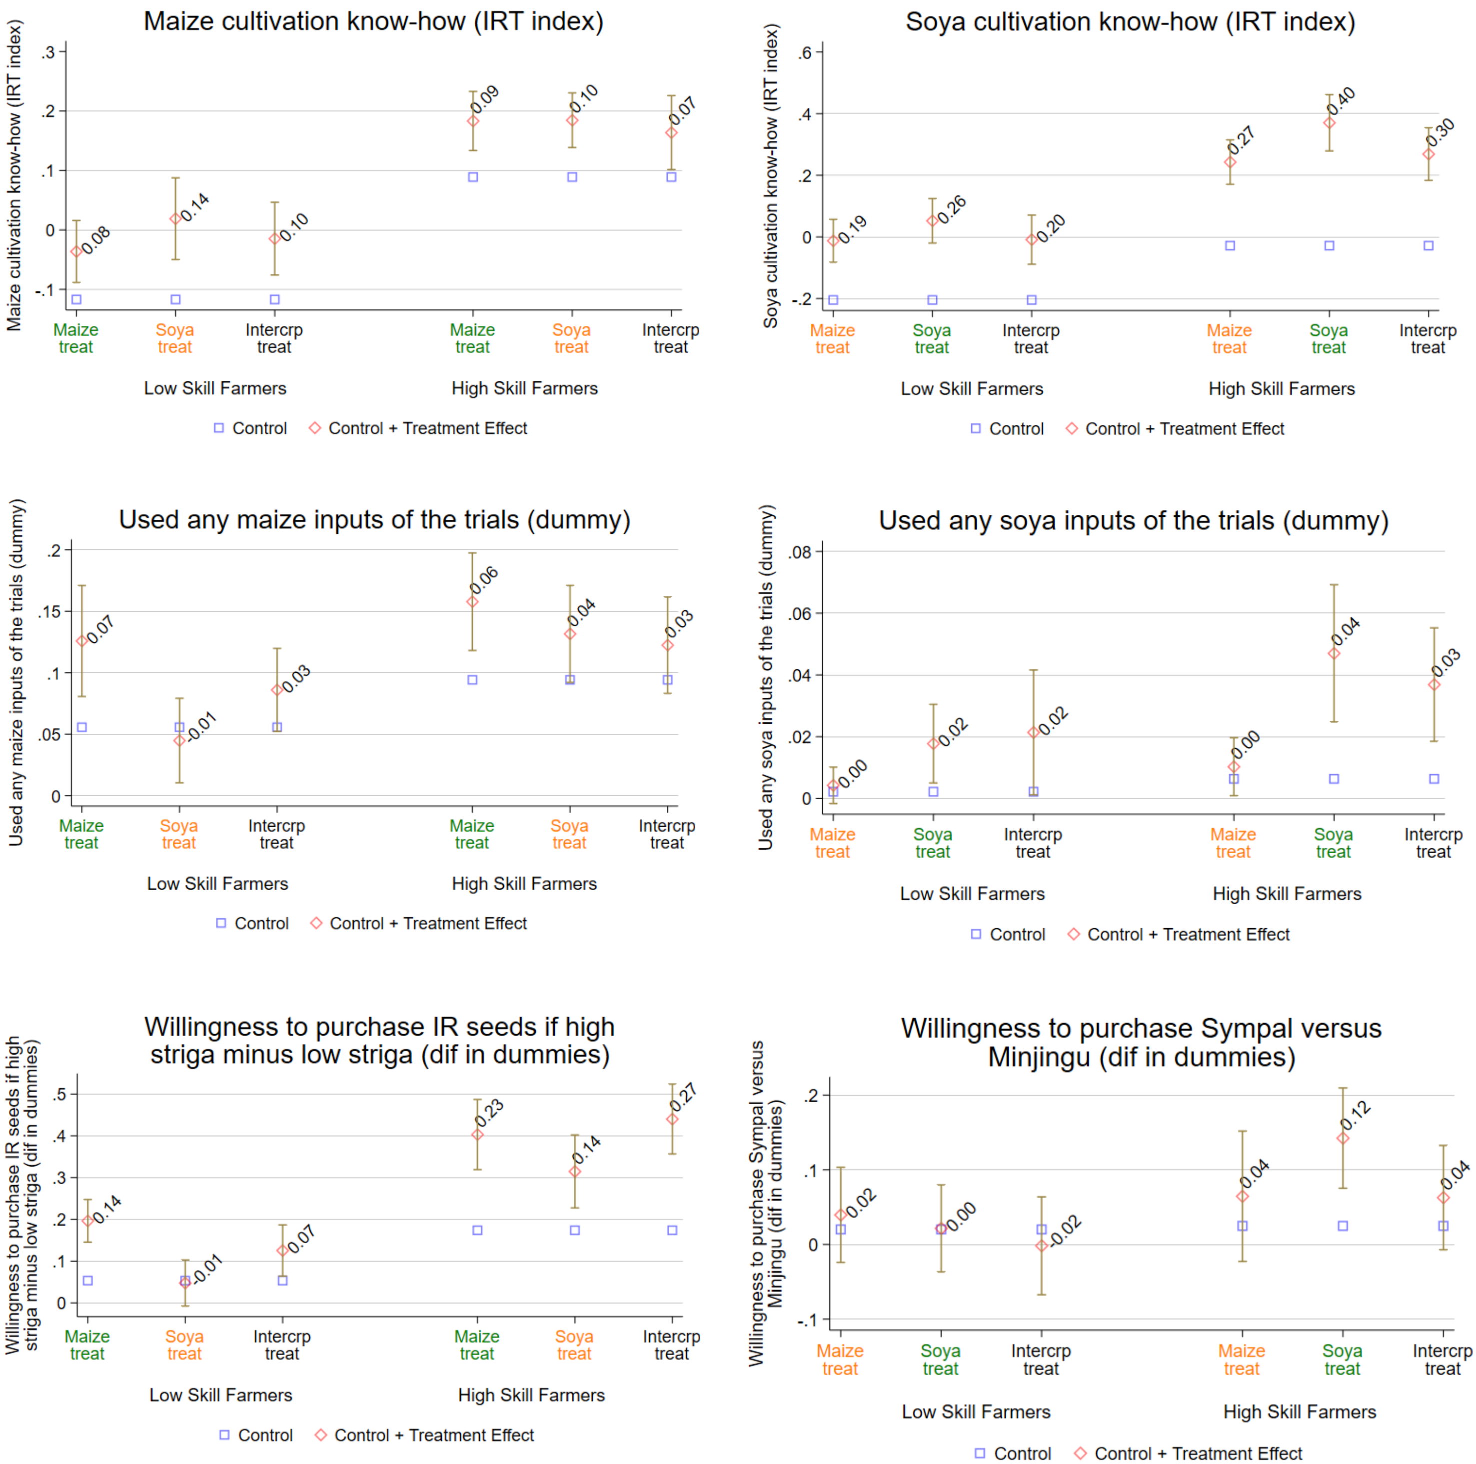

Figure 7 Treatment effects by skills and crop treatments on crop specific input use and learning. The figure displays the results from estimating equation (3), where each crop treatment (maize, soya, intercrop, each randomly assigned) is separated. Blue squares represent the average among control farmers by skill level. Red diamonds represent the expected value in the corresponding treatment group (average value in control group + treatment effect), and the value that appears above each red diamond is the corresponding treatment effect. When the crop is written in green, it is a direct effect (for example how the maize trial treatment affects maize know-how). When it is orange, it should be interpreted as a cross learning.

The upper left panel of Figure 8 shows the average skill level of LSFs, HSFs and of the person (among the other 9 farmers) they indicate has the most influence on their agricultural decisions. The upper right panel uses data from a question on the likelihood of talking about agriculture during the last three months with each of the 9 other farmers. We consider respondents that are LSF or HSF and whether they talked to a LSF or HSF in the same village. We also separately show whether they talked to the people they ranked in their top three of the most influential persons (among the 9 sampled farmers). The bottom panels share the same format, but the variables displayed assess the respondent's knowledge about the other 9 farmers' input use and crop tested in the trials (by comparing their answers about another farmer's input use and trial crops with the input use measured in the survey when asking that other farmer and their trial assignment). Such correct assessments are arguably a necessary (though not sufficient) condition to acquire valuable knowledge from other farmers in the village.[19]

The first panel of Figure 8 shows that the average skill index of the influencers of both LSFs and HSFs is very close to the average skill index of HSFs. Hence, based on their own statement, even LSFs tend to follow the advice of farmers who on average have a skill level close to the one of the HSFs. The other 3 panels of Figure 8 show that: 1) the influencers always have the highest values, which validates the farmers' claim about who they follow the most and thus supports the conclusions from the first panel; 2) HSFs talk significantly more to other farmers and know significantly more about the inputs that others use, which points towards HSFs being better not only at learning from their own experience, but also at learning from others; and 3) LSFs appear to discuss with and know about the treatment status of the HSFs in their village at least as much as about the LSFs, and they know significantly more about the input use of the HSFs compared to the one of the LSFs in their village. The latter is consistent with proposition 7, according to which LSFs may prefer observing HSFs because they tend to experiment more. Finally, the low levels of awareness are also quite interesting: on average respondents can accurately guess at least one input for only about 37% of the other farmers, which indicates that learning from one's neighbors faces important frictions.

[19] See Wolitzky (2018) for a theoretical discussion on the inefficiencies for learning when farmers can only observe the outcomes of other agents' decisions.

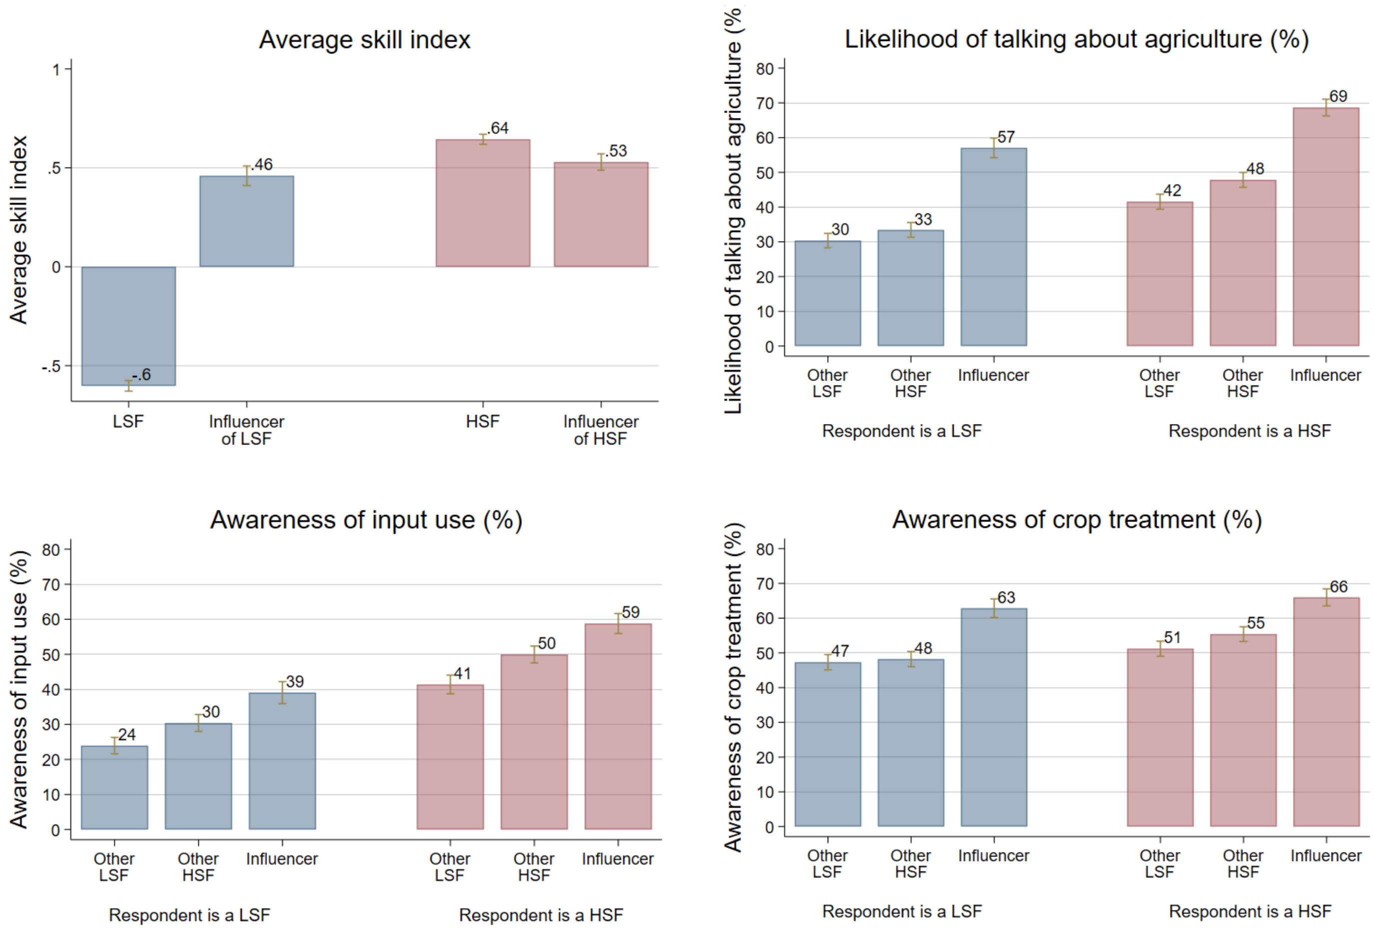

Figure 8: Descriptive statistics about the farmers and their influencers' skills and communication. The figure presents simple averages and their 95% confidence intervals for different groups and dyadic relationships. LSF and HSF refer to low-skill farmer and high-skill farmer, respectively. The "influencer" refers to the farmers response when asked who they follow the most for their agricultural decisions. "Influential contacts" refer to up to 3 most influential farmers, according to the respondent.

5.4.4 Potential Caveats and Related Robustness

The study has various caveats. First (as discussed in section 2.4), the theoretical model attributes the differences between LSFs and HSFs to the skill itself and its consequences on the ability to learn. We cannot directly test that the differences are due to skills rather than other systematic differences between LSFs and HSFs. For this purpose, an ideal experiment would require not only an exogenous variation in exposure to new information (like the trials), but also an exogenous (and sizeable) variation in a comprehensive set of skills, which would be difficult to achieve. We thus interpret our empirical results as different treatment effects between LSFs and HSFs, without assigning it to the skill difference. Moreover, skills may affect investment and management decisions (as a response to an information shock) not only through the ability to learn, but also through many other channels as they can capture being more risk taking, patient, pro-active, open to change or hardworking for example. Various results, such as different speed of learning between LSFs and HSFs, or negative effects in expectation on profit associated with increasing adoption are consistent with the predictions of our theoretical model, and less with heterogeneity driven by such other differences. However, we fully acknowledge that the design of the RCT does not allow us to isolate the effect of learning from other possible channels and drivers of heterogeneity. A related but distinct concern is the relatively high correlation between the selection process (community selected or not) and the skill measure, raising the possibility that the different effects for HSFs versus LSFs are driven by heterogeneity of impact due to the selection process. This could occur if being selected by the community (rather than randomly) affects farmer's involvement and behavior as a response to the treatment. To rule out this interpretation appendix Figure D1 shows that controlling for an interaction of the treatment with being community selected does not affect the main results.

Second the model ignores other constraints to technology adoption, including behavioral or internal ones. We fully acknowledge the possibility of bounded rationality in learning processes and connect to this literature with the assumption that limited attention and heterogenous attention capacity drive the difference in learning between LSFs and HSFs (appendix E). Still, the Bayesian optimization puts an upper bound on how well a farmer can manage the uncertainty, highlighting that even under such rational optimization, learning in agriculture is particularly complex and the challenges would only be accentuated once we incorporate human limitations in learning. Furthermore, farmers also often face multiple outside constraints, not incorporated in the model, such as credit constraint, access to input markets, uncertainty about output prices and access to output markets. While all these constraints are fully relevant, this article aims to focus on the complexity that arises from the multiplicity of input decisions, which is likely to further increase when the various other limitations are considered.

Third, one can wonder whether the negative profits after multiple seasons simply mean that the inputs tested in the trials are not profitable in the farmers' conditions. We believe this is unlikely for two reasons. First, it is inconsistent with the trial results themselves (which showed positive benefits of the input package compared to the control plots). Second, it is not consistent with the fact that the treatment appears to first lower profit, but the treatment effect on profit grows significantly over time, as adoption increases and becomes positive (though only significantly for LSFs). If it was simply due to bad inputs the reduction in profits should increase proportionally with adoption. By contrast, the dynamics that we observe are much more in line with increasing know-how as farmers make mistakes while figuring out how to combine the different inputs and practices. Additionally, even if the inputs and practices tested in the trials were simply non-profitable (even when the related know-how is acquired), results would not be incompatible with the model, where unsuccessful exploration is a possible scenario. Under this scenario, the time spent by farmers trying new inputs before realizing that the inputs are not worth adopting would strengthen our argument that exploration can be costly when compared to the expected benefits of an uncertain reward.

Finally, one can wonder whether the sustainable inputs can take various seasons to fully pay off, thus explaining initial losses and continued take-up? While this is technically possible, we note that in the trials, the benefits of the tested inputs materialized within the 3 seasons. It would also not explain most other empirical results, such as the increase in experimentation (frequency of changes in input and practice decisions), the increase in the likelihood to make new mistakes and to some extent, the increased discussions about agriculture and following of HSFs. Appendix D4 further explores this possible explanation by testing whether the treatment significantly increased soil quality and shows the evidence does not support this hypothesis.

6. Conclusion and Policy Implications

This paper uses farmers' intensive exposure to new information through their participation in multi-season agronomic trials to study the complexity of farmers' learning when they are given information that allows them to re-optimize the use of multiple inputs with possible complementarities and substitutability. The empirical findings provide insights in this complex multidimensional learning process over a 6-season period but also raise multiple questions. Why does providing information appear to make farmers worse-off in the short term? Why do farmers increase their use of the new inputs over time despite their initial reduction in profit? Why does farmers' knowledge diverge over time? If homogeneity makes lessons more applicable, why do LSFs know more about what HSFs do? The predictions of the theoretical model help provide answers to those questions, attributing the loss in profit and mistakes to the costs of taking farmers away from their comfort zones to explore new input combinations. Since HSFs are more able to learn and willing to explore, they end up bearing a higher share of the exploration cost and provide valuable information to LSFs.

Exposure to the trials not only increased know-how, but also increased farmer's own exploration and a sustained increase in the adoption of new inputs and practices despite a loss in profits in the first seasons. Farmers' know-how increased immediately after initial exposure and diverged between LSFs and HSFs, consistent with the latter also being significantly more able to learn subtle and complex lessons, including choosing between two inputs, adapting an input use decision to local conditions and learning about proper combinations of inputs. Exposure to new information through trials also increased the frequency at which farmers modify their input decisions on other plots, which came at the cost of an increased likelihood of committing new mistakes. And while there are spillovers in simple know-how, that is less the case for more complex learning or for the adoption of trial inputs. The model and the results together point to a possibly underappreciated constraint to technology adoption: if using a new input requires changes in other inputs or practices, farmers are likely to make mistakes along the way. Finding the right combinations may take time and come at a substantial cost that many may not be able to afford.

The findings are humbling in the sense that they highlight the complex challenges faced by small scale farmers in developing countries as they navigate the need to shift to more sustainable practices, often involving combinations of inputs and practices rather than marginal changes in one input. They highlight that top-down information interventions may be a complement rather than a substitute for local skills and exploration. The findings call for strategies that properly combine top-down information about what combinations of inputs are worth exploring with attempts to lower the cost of the exploration and adjustments required for the farmers. First, when some of the adjustments are homogenous enough that they can be anticipated, extension interventions can incorporate advice about what input adjustments farmers should pay attention to (Hanna et al, 2014). Second, the design and targeting of field days (Kelley et al. 2023), extension through on-farm experimentation (Lacoste et al. 2022) and citizen science approaches (van Etten et al. 2019; Ebitu, et al. 2021) or attempts to help farmers learn how to experiment and learn (Ashraf et al. 2021) could specifically aim to minimize the cost of local exploration. The ability to locally discover and share knowledge about effective combinations of inputs and practices may very well be a determining factor in the possibility to transition from an input intensive agriculture to a knowledge intensive and sustainable one.

REFERENCES

Aker, J and K. Jack, 2023. Harvesting the rain: The adoption of environmental technologies in the Sahel, *Rev. Econ. Stat.*, forthcoming.

Ashraf, N., Banerjee, A. and Nourani, V., 2021. Learning to teach by learning to learn. mimeo, University of Chicago and Makerere University.

Bandiera, O. and Rasul, I., 2006. Social networks and technology adoption in northern Mozambique. *Econ. J.*, *116*(514):869-902.

Bandiera, O., Burgess, R., Deserranno, E., Morel, R., Sulaiman, M. and Rasul, I., 2023. Social incentives, delivery agents, and the effectiveness of development interventions. *J. Pol. Econ. Micro.*, 1(1):162-224.

Barham, B.L., J.P. Chavas, D. Fitz, and L. Schechter. 2018. Receptiveness to Advice, Cognitive Ability and Technology Adoption. *J. Econ. Behav. Organ.* 149: 239–268.

Bationo, A., Fairhurst, T., Giller, K., Kelly, V., Lunduka, R., Mando, A., Mapfumo, P., Oduor, G., Romney, D., Vanlauwe, B. and Wairegi, L., 2012. *Handbook for integrated soil fertility management*. Africa Soil Health Consortium.

Beaman, L. and A. Dillon, 2018. The diffusion of agricultural technologies within social networks: Evidence from composting in Mali, *J. Dev. Econ.*, 133:147-161.

Beaman, L., BenYishay, A., Magruder, J. and Mobarak, A.M., 2021. Can network theory-based targeting increase technology adoption?. *Am. Econ. Rev.*, *111*(6):1918-43.

BenYishay, A. and Mobarak, A.M., 2019. Social learning and incentives for experimentation and communication. *Rev. Econ. Stud.*, *86*(3):976-1009

Bold, T., Kaizzi, K.C., Svensson, J. and Yanagizawa-Drott, D., 2017. Lemon technologies and adoption: measurement, theory and evidence from agricultural markets in Uganda. *Q. J. Econ*, *132*(3):1055-1100

Carter, M., Laajaj, R. and Yang, D., 2021. Subsidies and the green revolution: Direct effects and social network spillovers of randomized input subsidies in Mozambique. *Am. Econ. J: Appl. Econ.*, *13*(2):206-229.

Chandrasekhar, A. G., Duflo, E., Kremer, M., F. Pugliese, J., Robinson, J., and Schilbach, F., 2022. Blue spoons: Sparking communication about appropriate technology use. *NBER Working Paper* 30423.

Conley, T., & Udry, C. 2010. Learning about a New Technology: Pineapple in Ghana. The

Am. Econ. Rev., 100(1), 35-69.

Conlon, J.J., Mani, M., Rao, G., Ridley, M.W. and Schilbach, F., 2022. Not Learning from Others *NBER working paper. w30378*.

Contal, E., Buffoni, D., Robicquet, A. and Vayatis, N., 2013. Parallel Gaussian process optimization with upper confidence bound and pure exploration. In *Machine Learning and Knowledge Discovery in Databases: European Conference, Proceedings, Part I 13*:225-240.

Corral, C., Giné, X., Mahajan, A. and Seira, E., 2020. Autonomy and specificity in agricultural technology adoption: evidence from Mexico *NBER working paper* w27681.

Crane-Droesch, A., 2018. Technology diffusion, outcome variability, and social learning: evidence from a field experiment in Kenya. *Am. J. Agric. Econ.*, *100*(3):955-974.

Duflo, E., D. Keniston, T Suri and C. Zipfel, 2023. "Chat over Coffee? Diffusion of Agronomic Practices and Market Spillovers in Rwanda", *NBER working paper* 31368.

Ebitu, L., Avery, H., Mourad, K.A. and Enyetu, J., 2021. Citizen science for sustainable agriculture–A systematic literature review. *Land Use Policy*, *103*, p.105326.

Emerick, K. and Dar, M.H., 2021. Farmer field days and demonstrator selection for increasing technology adoption. *Rev. Econ. Stat.*, *103*(4):680-693.

Emerick, K., de Janvry, A., Sadoulet, E. and Dar, M.H., 2016. Technological innovations, downside risk, and the modernization of agriculture. *Am. Econ. Rev.*, *106*(6):1537-1561.

Fabregas, R., Kremer, M. and Schilbach, F., 2019. Realizing the potential of digital development: The case of agricultural advice. *Science*, *366*(6471), p.eaay3038.

Foster, A., & Rosenzweig, M. 1995. Learning by doing and learning from others: Human capital and technical change in agriculture. *J. Pol. Econ., 103*(6), 1176-1209.

Franzel, S. & Coe, R. 2002. Participatory on-farm technology testing: the suitability of different types of trials for different objectives. In Bellon, M.R., and J. Reeves (eds.). *Quantitative Analysis of Data from Participatory Methods in Plant Breeding*. Mexico: CIMMYT

Garnett, R., 2023. *Bayesian optimization*. Cambridge University Press.

Ghosh, S., 2018. Selective and Multidimensional Learning: A case study of bt cotton farmers in India, unpublished

Gignoux, J., Macours, K., Stein, D. and Wright, K., 2023. Input subsidies, credit constraints, and expectations of future transfers: Evidence from Haiti. *Am. J. Agric. Econ.*, 105 (3), 809-835.

Gollin, D., Lagakos, D. and Waugh, M.E., 2014. The agricultural productivity gap. *Q. J. Econ.*, *129*(2):939-993.

Gollin, D. and Udry, C., 2021. Heterogeneity, measurement error, and misallocation: Evidence from African agriculture. *J. Pol. Econ.*, *129*(1): 1-80.

Hamory, J., Kleemans, M., Li, N.Y. and Miguel, E., 2021. Reevaluating agricultural productivity gaps with longitudinal microdata. *J. Europ. Econ. Assoc.*, *19*(3):1522-1555.

Hanna, R., Mullainathan, S., & Schwartzstein, J. (2014). Learning through noticing: Theory and evidence from a field experiment. *Q. J. Econ*. 129(3), 1311–1353.

Jones, M., Kondylis, F., Loeser, J. and Magruder, J., 2022. Factor market failures and the adoption of irrigation in Rwanda. *Am. Econ. Rev.*, *112*(7):2316-52.

Kelley, E., M. Dar, A. de Janvry, K. Emerick and E. Sadoulet, 2023. Casting a Wider Net: Sharing Information Beyond Social Networks, mimeo, UC Berkeley and Tufts University.

Kondylis, F., Mueller, V. and Zhu, J., 2017. Seeing is believing? Evidence from an extension network experiment. *J. Dev. Econ.*, *125*:1-20.

Laajaj, R. and Macours, K., 2021. Measuring skills in developing countries. *J. Hum. Resour.*, *56*(4):1254-1295.

Laajaj, R., Macours, K., Masso, C., Thuita, M., Vanlauwe, B., 2020. Reconciling yield gains in agronomic trials with returns under African smallholder conditions. *Sci. Rep.*, *10*(1):14286.

Lacoste, M., Cook, S., McNee, M., Gale, D., Ingram, J., Bellon-Maurel, V., MacMillan, T., Sylvester-Bradley, R., Kindred, D., Bramley, R. and Tremblay, N., 2022. On-Farm Experimentation to transform global agriculture. *Nat. Food*, *3*(1):11-18.

Lattimore, T. and Szepesvári, C., 2020. Bandit algorithms. *Cambridge University Press.*

Lobell, D., Burke, M., Tebaldi C., Mastrandrea, M., Falcon, W., and Naylor, R., 2008. Prioritizing Climate Change Adaptation Needs for Food Security in 2030. *Science* **319**, 607-610.

Magnan, N., Spielman, D. J., Lybbert, T. J., and Gulati, K. 2015. Leveling with friends: Social networks and Indian farmers' demand for a technology with heterogeneous benefits. *J. Dev. Econ.*, 116:223–251.

Munshi, K., 2004. Social learning in a heterogeneous population: technology diffusion in the Indian Green Revolution. *J. Dev. Econ.*, *73*(1):185-213.

Njoroge S, Schut AGT, Giller KE, Zingore S. 2017. Strong spatial-temporal patterns in maize yield response to nutrient additions in African smallholder farms. *Field Crop Res.*;214:321–330.

Nourani, V. 2019. Multi-object Social Learning and Technology Adoption in Ghana: Learning from Friends and Reacting to Acquaintances. MIT working paper.

Rasmussen, C.E. and Williams, C.K., 2006. *Gaussian processes for machine learning* (Vol. 1, p. 159). Cambridge, MA: MIT press.

Rosenzweig, 1995. Why Are There Returns to Schooling?, *Am. Econ. Rev.*, 85(2): 153-158.

Rosenzweig, M.R. and Udry, C., 2020. External validity in a stochastic world: Evidence from low-income countries. *Rev. Econ. Studies*, *87*(1): 343-381.

Snapp, S., Sapkota, T.B., Chamberlin, J. *et al.* 2023. Spatially differentiated nitrogen supply is key in a global food–fertilizer price crisis. *Nat. Sustain.*, **6**:1268–1278

Squicciarini, M.P. and Voigtländer, N., 2015. Human capital and industrialization: Evidence from the age of enlightenment. *Q. J. Econ.*, *130*(4):1825-1883.

Tjernström, E., T.J. Lybbert, R. Frattarola Hernández, and J.C Correa, 2021. Learning by (virtually) doing: Experimentation and belief updating in smallholder agriculture. *J. Econ. Behav. Organ*., 189: 25-50.

Thuita, M., Vanlauwe, B., Mutegi, E. and Masso, C., 2018. Reducing spatial variability of soybean response to rhizobia inoculants in farms of variable soil fertility in Siaya County of western Kenya. *Agr. Ecosyst. Environ.*, *261*:153-160.

Uwiragiye, Y., Ngaba, M.J.Y, Yang, M., Elrys, A.S., Chen, Z. and Zhou, J. 2023. "Spatially Explicit Soil Acidification under Optimized Fertilizer Use in Sub-Saharan Africa." *Agronomy* 13(3):632.

Van Etten, J., de Sousa, K., Aguilar, A., Barrios, M., Coto, A., Dell'Acqua, M., Fadda, C., Gebrehawaryat, Y., van de Gevel, J., Gupta, A. and Kiros, A.Y., 2019. Crop variety management for climate adaptation supported by citizen science. *PNAS. 116*(10):4194-4199.

Vanlauwe, B., Bationo, A., Chianu, J., Giller, K.E., Merckx, R., Mokwunye, U., Ohiokpehai, O., Pypers, P., Tabo, R., Shepherd, K.D. and Smaling, E.M.A., 2010. Integrated soil fertility management: operational definition and consequences for implementation and dissemination. *Outlook Agric.*, *39*(1):17-24.

Wolitzky, A. 2018. "Learning from Others' Outcomes." *Am. Econ. Rev*, 108 (10): 2763-2801.

Wu, J., Poloczek, M., Wilson, A.G. and Frazier, P., 2017. Bayesian optimization with gradients. *Advances in neural information processing systems*, *30*.