Learning Through Noticing: Theory and Experimental Evidence in Farming

PRELIMINARY AND INCOMPLETE

Rema Hanna, Harvard University

Sendhil Mullainathan, Harvard University

Joshua Schwartzstein, Dartmouth College

February 29, 2012

Abstract

We build and test a model of how people learn to use technologies. Existing learning models suggest people may fail to learn because they lack relevant data. Our model suggests a different barrier: they may fail to notice certain features in the data they do have. Such "learning through noticing" is particularly important given the large number of dimensions –most of them irrelevant – one could focus on when using a technology. The model has implications beyond the canonical models of learning by doing and learning about profitability. First, people may not learn important features of the production process even when there is readily available data showing their importance. Second, simply presenting a summary of this data back can induce learning: it can highlight features the agent had failed to notice. Finally, learning through noticing can lead to systematic disappointment upon adopting a technology: the presence of unnoticed features means the agent will not use the technology as effectively as expected and to experience lower returns. We test several of the model's predictions using data from a field experiment with seaweed farmers in Indonesia. We find evidence that farmers do not notice relevant features of the production function. We further find that, while farmers may not learn much from participating in experimental trials that exogenously vary inputs along these dimensions, they are more likely to learn when also presented with summaries of the trials' findings. More broadly, we illustrate how learning through noticing can change our understanding of the role of human capital, experience, and agricultural extension in promoting efficient use of technologies.

1 Introduction

Imagine going to a friend's house where you eat some delicious steaks for dinner. You decide to try your friend's recipe. You watched your friend cook, and the steps make sense. The ingredients are easy to get and the effort involved is low. The only equipment you don't have is a cast iron skillet, which you buy. How would our existing models of technological learning understand this skillet experience? One group of models emphasizes uncertainty over the benefits of a technology (Besley and Case 1993, 1994). Another emphasizes learning by doing, or uncertainty over the steps necessary to effectively use a technology (Jovanovic and Nyarko 1996; Foster and Rosenzweig 1995; Conley and Udry 2010). Suppose that in this example these learning barriers should be resolved. Tasting the steak should remove any uncertainty you may have about the value of using the recipe. Seeing your friend cook provides all the information you need to follow the recipe.[1]

We argue that another barrier still remains. Suppose that after seeing your friend make the steaks, you go home and repeat the steps exactly as you saw them. Sure enough, the first batch of steaks is delicious. The second, however, is not as good. The ones after that are even worse. On average, they're just not as good as ones you made before and you give up. What happened? There was an aspect to using the skillet that you did not notice. When washing the skillet, your friend did not use any soap. For skillets, this happens to be a central detail. Washing with soap removes the "seasoning" from the cookware, resulting in a sticky surface that cooks unevenly and produces less flavor.

But why did you fail to notice this central fact? Watching your friend cook and clean, you encountered a large array of details, from the time of day to the belt he was wearing. Which of these details you encoded or "noticed" was shaped by your knowledge of cooking. Not having unlimited attentional capacity, you needed to focus on some to the exclusion of others. The features you focused on were the ones you thought could affect the outcome, the quality of the steak. In

[1]Specifically, when made properly, the steaks are consistently delicious and worth making with the skillet. Further, everything necessary to make the steaks properly can be observed by watching your friend make them, and your friend's kitchen and skills are simliar enough to yours that this knowledge transfers well to you. Under the traditional learning models, you should then end up with delicious steaks.

your view (and experience), the precise details of washing the pan have never mattered, much like it does not matter which belt was worn. Interestingly, it is exactly because you have never used a cast iron skillet that these details escaped your attention.[2] In short, what you learn about a technology depends on what you notice; we call this learning through noticing.

In this paper, we adapt Schwartzstein's (2012) model of selective attention to build and test a model of learning through noticing. The model focuses on a technology that has many dimensions to manage. For example, there are many aspects to using fertilizer, from top coating to whether it should be applied before or after rains. The agent is a Bayesian who has beliefs about the importance of various dimensions. Attending to a dimension is costly or effortful. As a result, he only notices dimensions that he currently believes are important. The model has several implications. First, if the agent does not put sufficient weight on the importance of a dimension, he will fail to optimize along it even if he has access to readily available information that demonstrates the importance of the dimension and how it should be optimized. In other words, a failure to notice can mean a failure to learn, whether from data generated by his own actions, his neighbors', or participation in a trial. Second, when presented with a *summary* of the data on how different practices along a dimension perform, agents may learn. In other words, it can help to highlight features of the data the agent already had seen but failed to notice. In short, one cannot understand performance on a technology simply through understanding what data is available. One must also understand which features of the data the agent's priors would lead him to notice.[3]

We test the model's predictions by performing a field experiment with seaweed farmers. In seaweed farming, there are several dimensions that may affect productivity. Seaweed is farmed by attaching strands ("pods") of seaweed to lines submerged in the ocean. Distance between these lines, spacing between the pods and the size of the pods are all dimensions of the technology

[2] If your friend watched you attempt the recipe, he could point out your error, even when you thought you were doing exactly what he did. But it is not the case that explicit communication always bridges this gap. If your friend tells you exactly what to do, the problem may not be solved. He would need to know your tendency not to notice certain details and to tailor his instructions accordingly.

[3] While the model presented in this paper is specific to issues related to technology adoption and use, Schwartzstein (2012) presents a general model of belief formation when an agent is selectively attentive. For other more general approaches to modeling limited attention in economic settings, see Sims (2003), Bordalo et al. (2012), Koszegi and Szeidl (2011) and Gabaix (2011).

that could matter, but which agents may or may not notice. In the experiment, we first measure farmers' baseline practices and beliefs. These data already hint at a potential failure to notice. While most farmers know the distance between lines and spacing between pods, few farmers even know what their pod sizes are. Interestingly, this failure to notice size means that the farmers' own fields have random variation in pod sizing. After measuring farmers' practices and beliefs, we randomly selected a subset of the farmers to participate in an experimental trial. On the farmers' own plot we varied spacing between pods and pod size. We then measured productivity. On pod spacing–which farmers noticed–we find farmers were close to the optimum. In contrast, on pod size–which farmers did not notice–they were potentially very far from the optimum. Importantly, we find that simply participating in a trial had little effect on farmers' behavior on subsequent pod size. Though they participated in the experiment, they did not notice the impact of pod size differences. Instead, it was only when we presented the data on yield broken down by pod size that their behavior changed.

These data apply the model of learning through noticing to understand how agents use technologies. We further apply it to understand whether agents will use a technology–the issue of technology adoption. The key result here is that agents can experience disappointment. Farmers can observe data–perhaps in a demonstration–that leads them to believe that they can use a technology profitably. Yet when they try it themselves they will find low returns. This was explicit in the skillet example. We argue that some prior evidence can be interpreted in this light, prominently the fertilizer experiments by Duflo, Kremer and Robinson (2008) which study the impact of observing experimental trials on adoption. In traditional models, observing trials is sufficient to encourage adoption. In our model, the impact of this type of "extension" activity depends on whether the technologies are "filter congruent"; i.e., how well they line up with the dimensions farmers notice. When farmers adopt a technology they may inadvertently use it much differently than how it was used in the trial.

The paper proceeds as follows. Section 2 presents the model of learning through noticing given a choice of technology and develops our experimental hypotheses. Section 3 describes the experimental design and data. Section 4 presents the results. Section 5 extends the model to consider the technology adoption decision and further applications. Section 6 discusses broader issues, for example how learning through noticing changes our understanding of the role of human capital, experience, and education in promoting efficient use of technologies.

2 Basic Model

2.1 Setup

A farmer knows his production technology up to the choice of a single input, $s \in 1, 2, 3, \ldots, S$. Over the course of two periods, $t = 1, 2$, the farmer collects data which may inform his choice of this input. In the baseline model, at each time t, the farmer chooses which level of the input to use for each of N pieces of land. He can either measure to guarantee a particular level, which costs m, or choose not to measure. If the farmer does not measure, then the input level is random and uniformly drawn across possible levels $\mathcal{S} = \{1, 2, \ldots, S\}$. To illustrate, in our seaweed application, the farmer may not know the optimal pod size. In every period, he has the opportunity to choose particular sizes for each of his N pods or not to precisely measure.

Given input s, the yield from piece of land p at time t is given by:

$$y_{pt} = f(s_{pt}) + \varepsilon_{pt}, \tag{1}$$

where $\varepsilon_{pt} \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ is a random shock which is independent across time and pieces of land. Denote the true production process (the true value of $f(\cdot)$) by $f_0(\cdot)$, and the value of the input that maximizes this function by s_0 (if there are multiple such values, pick one). The farmer is initially uncertain whether and how the input affects the yield – that is, about $f_0(\cdot)$ – and updates his beliefs over time.

The farmer attaches prior weight π to a model, P, in which the input affects yield and weight

$1-\pi$ to a model, NP, in which the input does not affect yield. Under model P,

$$f(s) = \theta(s), \tag{2}$$

where each $\theta(s)$, $s \in \mathcal{S}$, is independently drawn from a $\mathcal{N}(\mu, \sigma^2)$ distribution. Under this S parameter model, yield almost surely varies in the input. On the other hand, under model NP,

$$f(s) = \theta, \tag{3}$$

where θ is drawn from a $\mathcal{N}(\mu, \sigma^2)$ distribution. Under this single parameter model, yield does not vary in the input. Combining Equations (1)-(3),

$$y_{pt} = \theta(s_{pt}) + \varepsilon_{pt} \tag{4}$$

under model P, and

$$y_{pt} = \theta + \varepsilon_{pt} \tag{5}$$

under model NP.

A fully attentive farmer updates his beliefs about whether and how the input affects yield using Bayes' rule, as applied to first period history:

$$h = (y_{p1}, s_{p1})_{p=1}^{N}. \tag{6}$$

It will greatly simplify matters to assume that the farmer updates his beliefs conditional on a model, but does not update the model weights. The benefit of assuming this deviation from full Bayesian updating is that it implies that the farmer's posterior expectations depend only on sample average yields. To illustrate, the standard updating formula for a normal learning model (e.g., DeGroot, 1970) tells us that, given h, a farmer's posterior expectation of yield for a given level of the input

s' is then given by:

$$
\begin{aligned}
E[f(s')|h] &= \pi E[f(s')|h,P] + (1-\pi)E[f(s')|h,NP] \\
&= \pi E[\theta(s')|h,P] + (1-\pi)E[\theta|h,NP] \\
&= \pi\left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2+n(s')\sigma^2}\mu + \frac{n(s')\sigma^2}{\sigma_\varepsilon^2+n(s')\sigma^2}\bar{y}_{s'}\right) + (1-\pi)\left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2+N\sigma^2}\mu + \frac{N\sigma^2}{\sigma_\varepsilon^2+N\sigma^2}\bar{y}\right),
\end{aligned} \tag{7}
$$

where $n(s')$ stands for the number of times the farmer used input level s' in the first period, $\bar{y}_{s'}$ is the average first-period yield across pieces of land with input level s', and $\bar{y}$ is the average first-period yield across all pieces of land.

We follow Schwartzstein (2012) and make the additional assumption that the farmer must expend effort to attend to the input level relative to yield, so he is only selectively attentive. The farmer faces a cost e of keeping track of the input level applied to a given piece of land if it was not measured, and will do so only if the expected benefits (from being able to use this information later on in choosing the input) exceed this cost. In the second period, he updates his beliefs not based on the full history h, but on what he noticed, i.e., based on

$$\hat{h} = (y_{p1}, \hat{s}_{p1})_{p=1}^{N}, \tag{8}$$

where

$$\hat{s}_{p1} = \begin{cases} s_{p1} & \text{if the farmer kept track of the input level applied to piece of land } p \\ \varnothing & \text{if the farmer did not keep track of the input level applied to piece of land } p. \end{cases} \tag{9}$$

2.1.1 Learning by Doing

The farmer is risk neutral and maximizes the expected undiscounted sum of yield across periods, net of costs of measuring and attending to the input.[4] For simplicity, we restrict the farmer to using strategies that are symmetric across pieces of land in a given period. For example, if he does not measure or attend to the input level applied to piece of land p' at $t = 1$, then he does not measure or attend to the level applied to piece of land p'' at $t = 1$.

Proposition 1. 1. *If π is sufficiently low, then the farmer does not measure or attend to the input level in the first period. In this case, the farmer's posterior expectation of $f(s)$ does not vary in s, and he chooses not to measure in the second period.*

2. *If π is sufficiently large and e and m are sufficiently low, then the farmer attends to the input level in the first period. In this case, with probability one, his posterior expectation of $f(s)$ will vary in s, and, with positive probability, he will measure the input in the second period. As $N \to \infty$, the farmer's posterior expectation of yield is almost surely maximized at s_0.*

Proof. All proofs in the Appendix. ■

Proposition 1 says that if the farmer does not place much initial weight on an input level affecting yield then he will not measure or attend to that input. Such a farmer will continue to be indifferent between input levels even though he has access to readily available information that is informative about the relationship between the level and yield. He will not change his farming practice with respect to that input over time. Note that this prediction is independent of the underlying production function $f_0(\cdot)$: the farmer can make arbitarily large mistakes by failing to attend to the input, as an incorrect belief that an input is unimportant is self-confirming given selective attention (Schwartzstein 2012).

On the other hand, if the farmer places enough initial weight on the input affecting yield and the cost of measuring and keeping track of that input is sufficiently low, then he will attend to the

[4] Implicitly, we are making the simplifying assumption that the differential cost to the farmer of using different inputs is negligible.

input in the first period. In this case, he will come to expect some levels to achieve greater yields than others. As the number of pieces of land gets large, he will come to expect the optimal input to achieve the greatest yield. Such a farmer changes his farming practice with respect to the input with positive probability over time.[5]

Proposition 1 yields the following predictions:

Prediction 1. Farmers may not attend to all dimensions of the production function.

Prediction 2. When farmers do not attend to an input dimension, choices along that dimension may be far from optimal. When farmers attend, their choices are more likely to be optimal.

2.1.2 Learning from Others

In addition to learning through their own experience, the farmer could learn from observing the practices and outcomes of others. An extreme opportunity for such learning arises when the farmer can observe someone who actively experiments with different input choices. To match the main application, suppose the farmer learns from an experimental trial conducted by someone else, who we will label a researcher, on the farmer's own plot. Assuming that $N/S \equiv n$ is integer-valued, the researcher randomly assigns n pieces of land to input level $s = 1$, n to level $s = 2$, etc. The farmer chooses the input level in the second period. There are two regimes: one in which the researcher presents the farmer with results from the experimental trial prior to his choice of input (as fleshed out below), and one in which she does not. We consider the latter regime, labeled the "learning by observing regime," first.

Proposition 2. *If the farmer would not have attended to the input when learning by doing, then he also may not in the learning by observing regime. Such a farmer's posterior expectation of $f(s)$ does not vary in s, and he chooses not to measure the input in the second period.*

Proposition 2 says that when the farmer's failure to learn the relationship between the input level and yield is a consequence of selective attention, a demonstration that targets experimentation

[5] Note that all probabilities are measured with respect to the true underlying distribution; that is, given $f_0(\cdot)$.

alone may be ineffective. An inattentive farmer *already experiments*, albeit unintentionally, and does not learn the relationship between the input level and yield despite having access to freely available information which would aid in forming a belief about this relationship. Exogenously providing such a farmer with this same information in essentially the same way may not help.

Prediction 3. Farmers may not learn much from observing experimental trials that exogeneously vary inputs.

Adding Instruction A policy that makes it easier for the farmer to learn the empirical relationship between the input level and yield will be more effective. In particular, consider a regime (the “instruction regime") where, in addition to performing an experimental trial on a farmer’s plot, the researcher calculates and presents the farmer with summary statistics on results from the trial prior to his second period input choice. Namely, the farmer is told the level s^* that achieved the greatest sample average yield, as well as the corresponding sample average $\bar{y}_{s^*}$. Since the cost of keeping track of size relative to yield for a given piece of land is e, it is natural to assume that the cost of attending to $(\bar{y}_{s^*}, s^*)$ is also e.

It turns out that $(\bar{y}, \bar{y}_{s^*}, s^*)$ is a sufficient statistic relative to the farmer’s second-period decision (i.e., he will make the same choice whether he knows all of h or just $(\bar{y}, \bar{y}_{s^*}, s^*)$). This statement is formalized and proved in the Appendix. Thus, while the expected benefit of attending to the researcher’s recommendation is the same as that of keeping track of the input applied to *each* piece of land, the cost is lower by a factor of N (the number of pieces of land). This observation lies at the root of the following Proposition.

Proposition 3. *A farmer is more likely to attend to information that helps him estimate the relationship between the input level and yield when he is presented with summary statistics of results from the experiment: A farmer who would not keep track of the input level in the learning by observing regime may attend to $(s^*, \bar{y}_{s^*})$ in the instruction regime; a farmer who would not attend to $(s^*, \bar{y}_{s^*})$ in the instruction regime would also not keep track of the input level in the learning by observing regime. A farmer who attends to $(s^*, \bar{y}_{s^*})$ either uses s^* in the second period or does not* *measure.*

Proposition 3 shows that it can help to highlight features of the data the farmer had seen but failed to notice. The fact that it is enough to present the farmer with $(\bar{y}_{s^*}, s^*)$ for these purposes (i.e., that $(\bar{y}, \bar{y}_{s^*}, s^*)$ is a sufficient statistic) is rather specific to the way we specified the farmer's prior as well as the environment. More generally, the lesson is that a selectively attentive farmer is more likely to learn from some sufficient statistic relative to his second period-decision so long as it is less costly to process this statistic than to attend to the input level applied to every piece of land.

Prediction 4. Farmers are more likely to learn from summaries of experimental trial results.

2.2 Summary of Predictions

Summarizing, the model yields the following predictions, which we test in our main application:

1. Farmers may not attend to all dimensions of the production function.

2. When farmers do not attend to an input dimension, choices along that dimension may be far from optimal.

3. Farmers may not learn much from observing experimental trials that exogeneously vary inputs.

4. Farmers are more likely to learn from summaries of trial results.

3 Experimental Design and Data

3.1 Setting

This project was conducted with seaweed farmers in the Nusa Penida district in Bali Indonesia. The farmers in this area have been growing seaweed since it was introduced in the early 1980s.

While there are different methods to conduct farming, most farmers in this area use the bottom method: for their plot, farmers drive wooden stakes in the shallow bottom near the shore, and then attach lines across the stakes. They take raw seaweed from the last harvest and cut into pods; the pods are planted by attaching them at a given interval on the lines in the sea. The optimal size of the pod/seedling can be determined by a number of factors: for example, bigger seedlings may result in higher yields in still water, but may be more likely to break (or be lost completely) in a place in the ocean that faces a lot of waves. At low tide, the farmers walk out to tend the plots (remove debris, etc.). After about 35 to 40 days, the seaweed is harvested and dried, and then sold to local buyers. Most farmers initally grew a variety called spinosim, but many have moved to a different strain called cottoni due to intensive government and NGO extension programs. (Note that in August 2008, the price of spinosim was about a fifth of the price of cottoni).

We became interested in seaweed farming for a variety of reasons. First, there had been a series of extension programs conducted in Indonesia, which provided advice to farmers on farming methodologies (including optimal pod size, harvest timeframe, distance between pods and lines, etc.), and few farmers had taken up the advice. Second, seaweed farming has a very short crop cycle and there are several crop cycles over the course of the year. This made it interesting as a crop for two reasons: first, as there are many crop cycles, there should be many opportunities for farmers to learn about optimal methods through their own experimentation. Second, the short crop cycle allowed us to conduct a trial with farmers and learn the results relatively quickly as compared to other types of crops. Finally, the seaweed farming process is similar to that of other types of crops, where lots of different decisions over inputs and processes add up to determine yields.

3.2 Experimental Design

From June 2007 to December 2007, we administered a survey to understand the demographic background, farming practices, and learning practices of seaweed farmers in Indonesia. From a census of about 2706 farmers located in seven villages (24 subvillages), commissioned by us in 2006, we drew a random sample of five hundred farmers for the baseline survey, stratified by subvillage.[6] Out of these, 489 were located and participated in the survey.

The baseline survey consisted of two parts: (1) a questionnaire that covered demographics, income, and farming methods, and (2) "show and tell" where the enumerators visited one of the farmer's plots to measure and document his actual farming methods (see the next subsection for an in depth description of the survey data). The questionnaire typically took about an hour and a half to two hours, with the enumerator reading each question and possible answers to the participant.

From the baseline survey list, we randomly selected 117 farmers to participate in an experimental trial on one of their plots to determine the optimal pod size for their given plot (stratified by sub-village). After the farmer had been surveyed in the baseline, we asked the farmer if he wanted to participate in the trial. The trials occurred between July 2007 – March 2008, shortly after the baseline was conducted for the farmer.[7] In each trial, the enumerators would vary the seaweed production methods across 10 lines of a plot, with the farmer's assistance. All of the farmers that we approached participated in the trials. Note that they were compensated for participating in the trials in several ways. First, we provided the necessary inputs for planting ten lines, and guaranteed a given income from each line, calculated from past experience so that the farmers at least break even if they participated. Second, we provided a small gift (worth $1) to each farmer to account for their time.

Farmers were randomly assigned into one of two trial subtreatments: *sort* (65 farmers) and *weight* (52 farmers). The sort subtreatment was designed to answer the following question: "Could farmers achieve a higher yield by reducing the variation in their production technologies?" The farmers were asked to cut pods as they usually would for the plot in question, and then they were sorted into small, medium and large groups. Working with the farmers, the enumerators attached the pods into the lines by groups (3 lines with small pods, 4 lines with medium pods, and 3 lines with large pods). The lines were then planted in the farmer's plot. The weight subtreatment was

[6]Once the initial 500 farmers were chosen, we showed the list to local village leaders. Those who had migrated prior to our survey or were no longer involved in seaweed farming were replaced.

[7]The exact timing of when the trial occurred after the baseline varied from farmer to farmer given the timing of when a farmer had a plot free, and also by the fact that our team could only manage a certain number of trials concurrently.

designed to answer the question: "What is the optimal planting technique?" There are three key variables that the farmers may choose when planting seaweed: the initial size of the seedling, the spacing between pods in a line, and the spacing between lines. We decided to focus on two of these: spacing between pods and initial weight. Average distance between pods is around 15cm, and in the past technical assistance programs have suggested larger spacing, so we tested 15cm and 20cm. To generate variation initial pod size, the pod weights were initially set at 60g to 140g (in intervals of 20g) for the first few trials, but in order to better reflect the ranges of weights used by farmers the weights were changed to 90g-180g for spinosim and 90g-210g for cottoni (both in intervals of 30g).[8] Note that the pods of different sizes were randomly distributed across the 10 lines, with the enumerator recording the placement of each pod (however, it is important to note that the farmers were present for the trials, and saw where the pods of each size were being planted on the lines).

For both trial groups, farmers were told to maintain the plots as they would normally maintain any other plot. The enumerators returned to reweigh the seedlings twice while in the sea: once around day 14, and around day 28 (exact dates depended up on the conditions of the sea). On around day 35, the seaweed was harvested and weighed for a last time (once again, the date of harvest varied to conditions in the sea, farmer availability, etc.). Note that in addition to being present for the planting, the farmers were present for all weighing, harvesting, and recording of results.

In April-May 2008, we conducted the first follow-up survey, which was designed to learn whether the farmers changed any of their methods after participating in a trial. The follow-up survey was conducted with a subsample of 232 farmers, which included all of the farmers who participated in the trials, as well as an additional set of farmers who were randomly selected from the baseline as a control group; 231 farmers completed the survey. The survey asked if the farmers had changed—in the previous four months or two cycles—the following seaweed growing practices: strain, distance between lines, distance between pods, or pod weight.

[8]Note that obtaining a seedling of the exact weight is quite difficult, so we ensured that the actual weights employed were within 10g of the target weight.

Shortly after the first follow-up (May to June 2008), the enumerators provided the results to each farmer. The enumerator gave each farmer a sheet of paper with a table that indicated the pod weight that had the best return on investment for their plot, and talked the farmer through the results.[9] Specifically, they were told the average weight of the pods that they typically used, and shown the difference between that pod weight and the optimal pod weight in a simple demonstration (the optimal pod weight for sorting trials was the average of the optimal size conditional on the optimum being small, medium or large). For the weight trials, they were also told whether the optimal distance between pods was 15cm or 20cm. Appendix Figure 3 presents examples of tables presented to the farmers.

Approximately two months after the results were given to a farmer (July –August 2008), we conducted a a second follow-up survey. Out of the original 232 farmers, 221 were found. The goal of this survey was to determine the effect of having received the results of their trials.

Appendix Figure 1 summarizes the experimental design and Appendix Figure 2 summarizes the sample design.

3.3 Data, Sample Statistics and Randomization Check

3.3.1 Data

We have four sources of data for the experiment: the baseline survey, the results from experimental trials with farmers, and both follow-up surveys. Details are as follows.

Baseline Survey The baseline survey was designed to learn about current farmring practices as well as beliefs on farming methodologies. It consisted of two parts: (1) a questionnaire and (2) "show and tell" where the enumerators visited one of the farmer's plots to measure and document his actual farming methods. The questionaire was broken into several sections. In the first section, we collected basic demographic information on all farmers (household size, education background,

[9]We had worked with a local NGO, who provided extension services to seaweed farmers and worked with us in conjunction for this survey, to design a simple, easy to understand table to summarize the trial results.

etc). In the second section, we asked for information on how the farmer became involved in seaweed farming and how he obtained his plots of land, while in the third section, we tested the farmer for his knowledge on the "best practices" in seaweed farming. In the fourth section, we gathered data that would allow us to estimate each farmer's production function: labor costs, capital inputs, technologies employed, difference in methods based on seasonality and plot location, crop yields, etc. In the next section, we collected data on both learning and experimentation. The "learning" questions focus on external learning, such as: Where did the farmer gain his knowledge on his production methods? Where does he go to seek new information, or to learn new techniques? The "experimentation" questions focus on internal learning/learning by doing, such as: Has the farmer experimented with different techniques? If yes, which types of techniques has he varied? Did he change all of his production at once, or was it a step-wise process? The final section of the survey focused on the income/expenditures, assets and loans to provide a picture of income status of each farmer and level of indebtedness.

After the questionnaire was complete, the enumerator conducted the "show and tell."[10] This typically took about 15 minutes to two hours, depending on how hard it was to access the farmer's plot. In the show and tell, we documented each farmer's actual production methods. For example, we collected information on type of line used, sizes for a random sample of pods, distance between seedlings, distance between seaweed lines, and so forth.

Experimental Trial Results We collected all data from the experimental trials. Thus, we have data on where the plot was located, where each pod was planted, the weight of each pod during each of the three weighings, and the average return per pod size for each weight.

Follow-up Surveys From both follow-up surveys, we have data on self-reported changes in farming techniques. In addition, we conducted the "show and tell" module again for each follow-up to measure actual changes in production techniques. As in the baseline, we collected information

[10]For some farmers, the show and tell was conducted on another day, as the main constraint to the show and tell was that it had to be conducted during the low tide periods to access the plot.

on the type of line used, sizes for a random sample of pods, distance between seedlings, distance between seaweed lines, and so forth.

3.3.2 Sample Statistics and Randomization Check

As Table 1 illustrates, most farmers (83 percent) are literate. Most farmers have been farming for about 18 years. About half learned how to farm from their parents. As is typical in the agricultural sectors in developing countries, many farmers (about a quarter) have a loan from the buyer to whom they sell their crop to.

In Appendix Table 1, we provide a randomization check across the control and both subtreatment groups. We test for differences at the baseline across the groups on 10 variables, both demographic information and key farming characteristics. Columns 1 - 3 provide the means for the control, sort and weight groups respectively. Columns 4 - 6 provide estimates of the differences in means across groups; Columns 7 - 9 provide estimates of the difference in these means, conditional on subvillage fixed effects. As illustrated in Columns 4 through 6, out of 30 comparisons, 3 are significant at the 10 percent level (consistent with what you would expect by random chance). Controlling for subvillage fixed effect, we again find few significant differences between the groups (two coefficients are significant at the 10 perent level and one is significant at the 5 percent level). We additionally compute the p-value from a test of joint significance across all 10 variables; we fail to reject the null that all the differences are equal to zero for each of the three groups.

4 Results

4.1 Learning by Doing

A basic implication of learning through noticing is that some farmers will not keep track of input dimensions that influence returns. Table 2 presents sample statistics on seaweed farming practices and beliefs for the 489 farmers included in the baseline survey. Specifically, we examine data on existing farming practices (Panel A), self-reported beliefs on farming (Panel B), percent of those who have no opinion on different farming techniques (Panel C), reasons why farmers might experiment with pod size (Panel D) or new methods (Panel E), percent of farmers who notice differences in farming methods (Panel F), and percent of those who reported making changes in their methods in the past (Panel G).

The data reveal that a vast majority of farmers were inattentive to some input dimensions, such as pod size. Eighty six percent did not have an opinion on their current pod size in the baseline, while 87 percent of farmers did not even want to hazard a guess about what the optimal pod size should be (Panel C). Seventy-nine percent of the farmers stated that they do not notice when the size of a pod is significantly different from others in a given cycle (Panel F). Since many farmers did not notice key facts about their pod sizes, it is perhaps unsurprising that a broad range of pod sizes is observed (see Figure 1) and that about 88 percent of the farmers indicated they had not tried different sizes in the past (Panel G).

On the other hand, farmers appear to have been attentive to other input dimensions; in particular, they appeared attentive to both the distance between knots that secure the pods to a line and the distance between lines. Very few farmers (about 1 to 2 percent) did not have an opinion on the distance between the lines, and additionally most farmers had an opinion about the optimal distance between both the knots and lines. Furthermore, using the means and standard deviations in Panel A, we can compute the coefficient of variation to compare the relative variances across inputs: the fact that the coefficient of variation for the distance between the lines (0.13) and the pods (0.10) is much smaller relative to that of pod size (0.27) indicates that the practices for these inputs are relatively less variable across farmers than that of pod size. This is consistent with a story where farmers notice (and perhaps even measure) the distance inputs more often than that of pod size. However, only 21 percent indicate that they would notice if the distance between two knots is significantly different from the rest in a given cycle, which is the same percentage that indicates that they notice when pod sizes differ (Panel F). An extremely small percentage of farmers reported trying out different distances between lines and pods in the past (Panel G), suggesting that that social and public learning may contribute to farmers' beliefs about the optimal distance between pods and lines.

A second implication of the model is that when farmers do not attend to some input dimension, then their choices along that dimension can be far from optimal. For the 117 farmers that participated in the experimental trials, we have data on optimal pod size. The outcomes of the trials are summarized in Figure 2, Figure 3 and Table 3. Figure 2 provides information on the percentage of farmers whose measured baseline size was close to the experimental optimal, as well as the percentage of farmers that could gain from increasing or decreasing their pod size. This information is presented for the full trial group, as well as separately for the sort and weight sub-treatment groups. In Figure 3, we provide the average recommended change (for those who were larger or smaller than the optimal); once again, this information is provided for the total trial treatment group, as well as for each sub-treatment. Finally, in Table 3, for each sub-treatment, we compute the predicted change in income from changing behavior.[11]

On net, the trial data indicate that farmers can benefit from changing their pod size. As Figure 2 illustrates, very few farmers (about 3 percent) were at the optimum. Fifty-three percent of the farmers were advised to go larger, while the remaining 44 percent were told to reduce their pod sizes. The distribution of those who were told to increase or decrease their pod size were similar across both sub-treatments. For those who were told to go larger, the average estimated increase was 45 grams (or 43 percent larger than the mean pod size in the baseline); those who were told to go smaller were told to reduce their pod size by 39 grams. Unsurprisingly, as a larger range of weights were considered in the weight sub-treatment than the sort sub-treatment, the weight treatment resulted in a relatively larger average change than the sort treatment, conditional on being recommended to change. The difference from the optimum translated into real losses of income, as farmers could increase their predicted income by 20-30 percent by moving to the recommended size (Table 3). Interestingly, the results from the sort condition indicate that the farmers would

[11] Note that to compute these predicted changes to income, we make several strong assumptions. First, we assume that prices for seaweed in the past are consistent with prices for seaweed in the future; this may be unrealistic as the price may fall if all farmers increase their yields. Second, we assume that farmers do not change other methods (have fewer cycles, harvest earlier, etc.) if their yields change. Thus, this evidence should be viewed more as suggestive, rather than as the counterfactual of the change in income if farmers changed their methods.

even do much better if they changed their practice to systematically use a specific size within the range of sizes *they already use*: moving from the pod size with the lowest associated return among sizes used to that with the highest associated return would raise the farmers' income by roughly 20 percent on average. This implies that inattention to pod size may have had real effects on household income and well-being.

Taken together, the findings suggest that farmers failed to notice pod size as part of the production process and that many farmers are not optimizing pod size. The results provide *suggestive evidence* that it is precisely this inattention to pod size that contributes both to a failure to experiment along the size dimension and to a failure to optimize pod size. As such, inattention may hinder learning by doing. We now turn to analyzing the evidence on the farmers' response to the experimental trial, which will provide a stronger test for whether inattention actually *leads* to suboptimal production decisions.

4.2 Learning from Others

The model makes predictions about how farmers should respond to the experimental trial: we would expect that participating in the trial may not have large effects on future behavior if farmers tend not to notice (or care) about specific parts of the production function. However, farmers should be more likely to respond when they are presented with a summary of the trial findings, as the summary is easier to process. When farmers respond, they should bring their practice more in line with what performed well at the trial. If farmers are further from the optimum for some inputs (such as pod size) than others (such as distance between pods), we would expect that learning the trial results should have larger effects on the inputs that are further from the optimum.

We begin by exploring the effect of participating on a trial on production outcomes. Specifically, for each farmer i in subvillage v, we estimate the following model:

$$Y_{ivt} = \beta_0 + \beta_1 P1_t + \beta_2 P2_t + \beta_3 Treat_{iv} + \beta_4 Treat_{iv} * P1_t + \beta_5 Treat_{iv} * P2_t + \alpha_v + \mu_{ivt} \quad (10)$$

where Y_{ivt} is the production choice at time t, $P1_t$ is an indicator variable that denotes that the outcome variable comes from the first follow-up after the experimental trial, $P2_t$ is an indicator variable that denotes that the outcome variable comes from the second follow-up, and $Treat_{iv}$ is an indicator variable that denotes that the farmer participated in the trial. We also include a subvillage fixed effect, α_v, as the randomization was stratified along this dimension.[12] There are two key parameters of interest: β_4 provides the effect on production outcomes after farmers participate in the trial, but before they obtain a summary of the results, while β_5 provides the effect after the results of the trial are provided.

Table 4 presents these results. In Columns 1 and 2, the outcome of interest is the self-reported measure of whether the farmer has made any changes in their production techniques (except for a change in strain). In Column 1, we report the coefficient estimates from Equation (10); in Column 2, we report the estimates from a model that additionally includes farmer fixed effects. Columns 3 and 4 replicate the analysis in the first two columns, but with the enumerator measured pod size as the outcome variable. Note that all standard errors are clustered by farmer.

Consistent with the model's predictions, just participating in the trial had no effect on the subsequent production outcomes, but learning the results led to large changes in production methods.[13] In the follow-up after the trial, we do not observe a significant change in the number of farmers reporting that they changed a farming technique (Column 1). However, about 16 percent more farmers reported making any change in technique after receiving the results, which is about double the mean of the dependant variable (Column 2). Adding farmer fixed effects does not significantly alter the estimated coefficient (Column 2). It is possible that some of the results from this self-reported measure may be driven by farmers wanting to please the enumerators after participating in the trial, though this is unlikely as the control group had also been receiving visits from enumerators as well to both survey and measure their farming practices. Turning to the enumera-

[12]Note that this inclusion of a subvillage fixed effect does not significantly influence the results.

[13]We only varied distance between pods in the weight treatment, as in the sort treatment we tried to keep all aspects of the farmer's production methods constant, except to sort the pods by their natural variation. As a result, we do not estimate the effect of the trial on distance in Table 4, but rather in Table 5 where we disaggregate results by the subtreatments.

tor measured results, we find that after farmers participated in the trial, pod size does not change (Columns 3 and 4). After receiving the summary of the results, pod size increases by about 7 grams (on average) for those those in the treatment group.[14] This is significant at the 10 percent level in the basic specification (Column 3). However, while the coefficient estimates remain roughly the same (7.3 grams) when including farmer fixed effects, the significance level falls below conventional levels of significance (p-value = 0.14) due to an increase in the standard error (Column 4). Thus, while farmers did not appear to consider pod size to be an important part of the production process prior to the trials, providing summary information on their optimal pod size appeared to change their use of this production input.[15]

We next explore the sort and weight subtreatments. Here, we provide coefficient estimates when we include dummy variables for the sort and weight treatments as controls (S_{iv} and W_{iv}), as well as the interaction of these variables with the indicators for follow-up status:

$$\begin{aligned} Y_{ivt} = \beta_0 + \beta_1 P1_t + \beta_2 P2_t + \beta_3 S_{iv} + \beta_4 W_{iv} + \beta_5 S_{iv} * P1_t + \beta_6 S_{iv} * P2_t \\ + \beta_7 W_{iv} * P1_t + \beta_8 W_{iv} * P2_t + \alpha_v + \mu_{ivt}. \end{aligned} \tag{11}$$

β_5 through β_8 are the coefficients of interest providing the effect of the respective trial in the respective time period. Once again, the first two columns report the coefficient estimates where

[14]At first it may seem surprising that presenting farmers with results from the trial appears to have led them to increase their size on average since roughly 44 percent of farmers were advised to *decrease* their size. However, *on average*, farmers were advised to increase their size by roughly 6.69 grams, which matches the sign and magnitude of $\hat{\beta}_5$ in the third and fourth columns of Table 4.

[15]In Appendix Table 3, we disaggregate the results by whether farmers were told to increase or decrease pod size on their actual pod size in the follow-ups. To do so, we interact the interaction of the treatment and post variables with an indicator variable that denotes whether farmer should increase pod size (Increase), controlling for the main effect of the recomendation (as well as all double interactions between the recommendation and the treatment variable and the recommendation and the post variable). Once again, we present results with (Column 1) and without (Column 2) individual fixed effects; all standard errors are clustered at the farmer level. We observe that farmers who were told to increase pod size did so both after the first and second followup. However, the recommendation itself is an endogenous variable; thus, it is possible that this is simply capturing the fact that if farmers are randomizing with respect to pod size, those who "should go bigger" are those who had abnormally low draws of pod size at baseline. Thus, in expectation, those farmers would in fact go bigger the next period even if they continue to randomize. Therefore, these results provide little real content.

self-reported change is the outcome of interest, while the next two columns report on specifications where enumerator pod size is the outcome. We only varied distance between pods in the weight treatment; this is the outcome variable in Columns 5 and 6. Once again, we present results with (odd columns) and without (even columns) individual fixed effects; all standard errors are clustered at the farmer level.

The findings appear consistent with the predictions of the model. Farmers do not appear to make significant changes in production methods after participating in either type of trial (Column 1 and Column 2); however, in both types of trials, they report making large and significant changes in their production techniques upon being presented with the results. We once again find no effect on pod size after particpating in the trial (Columns 3 and 4); we observe positive effects on both the sort and weight subtreatments on pod size, but only the sort treatment is significant. Finally, we find no effect of the weight treatment on the enumerator measured distance between pods in either the first or second followup. This is consistent with the model, since, unlike pod size, farmers appeared to notice distance in the past, and appeared to already have strong beliefs on the optimal distance. As a result, we should not expect large changes in distance as a result of either participating in the trial or receiving its results. However, it is important to note that while this result is consistent with the model, it is not fully conclusive as the insignificant result on distance could be driven by the smaller sample size as we only varied distance in the trials for farmers in the weight treatment.

5 Model Extensions and Further Applications

This section presents extensions to the basic model that are important in settings outside of our main seaweed application, and discusses further applications.

5.1 Adopting a Technology

So far, we have taken the technology adoption decision as given. Now suppose that in addition to deciding how to use a technology, the farmer also decides *whether* to use the technology. Extend the learning by doing model so that, each period, the farmer faces an additional decision of whether to farm at all. If he does not farm, he receives payoff $\bar{u}$, normalized to equal 0.

To limit the number of cases considered, suppose that $\mu > 0$, which implies that the farmer will farm in the first period. Recalling that s_0 denotes the input level that maximizes the true production function, $f_0(s)$, further suppose that $\max\left\{f_0(s_0) - m, \frac{1}{S}\sum_{s'} f_0(s')\right\} > 0$, so a best-practices farmer is better off farming than not given the underlying technology.

Whether a selectively attentive farmer settles on adopting a profitable technology depends on whether he attends to the proper inputs. We will say that, given what a farmer selectively attends to, a technology is *filter incongruent* if the technology would no longer be profitable if the farmer only optimizes with respect to attended to inputs and randomizes with respect to the other inputs. Formally:

Definition 5.1. Fix whether the farmer attends to the input level in the first period. The technology is *filter incongruent* if the farmer does not attend and $\frac{1}{S}\sum_{s'} f_0(s') < 0$.

Intuitively, the farmer is less likely to settle on adopting a filter incongruent technology. While these technologies would be profitable if adopted correctly, the farmer will not learn to optimize these technologies because he fails to attend to important input dimensions, resulting in a belief that the technology is unprofitable. To formalize this intuition, we will simplify by assuming that the number of pieces of land N is "large enough" that with high probability the farmer accurately estimates returns conditional on his attentional filter.

Proposition 4. *Suppose $N = \infty$. Then with probability 1:*

1. *The farmer does not continue to farm in the second period if the technology is filter incongruent.*

2. *The farmer continues to farm in the second period if the technology is filter congruent and π is sufficiently large.*

Proposition 4 indicates that farmers are less likely to adopt filter incongruent technologies. Assuming their priors come from experience with other technologies, this suggests that they are less likely to adopt technologies when the relevant inputs differ from those of technologies they have experience with. This provides one reason why farmers may be quicker to adopt technologies that are "compatible" or "congruent" with current practices (Rogers and Shoemaker 1971, Brandner and Strauss 1959, Tornatzky and Klein 1982, Fliegel and Kivlin 1966). For example, Brandner and Strauss (1959) argue that farmers who previously adopted hybrid corn were more likely to subsquently adopt hybrid sorghum, which could result from the common practices involved in optimally planting hybrid varieties.

5.2 Learning from a Best-Practice Demonstrator

Suppose now that in the first period the farmer does not farm his own plot, but rather observes a demonstration by an individual with knowledge of the underlying production function. We will call this individual a *best-practice demonstrator*. The best practice demonstrator optimally farms the farmer's plot. Since the farmer is selectively attentive, however, he may not attend to the demonstrator's choice of input s. Since the farmer may not attend to the demonstrator's choice, he must hold some beliefs over what the demonstrator does. We assume that, whatever these beliefs are, they are symmetric across pieces of land and input levels. That is, the farmer's prior reflects ignorance: he has no reason to suspect that any one piece of land is farmed differently than another, or that any one input level is more likely to be selected than another. This assumption is likely to be appopriate so long as the demonstrator does not communicate his input choice to the farmer. (More on this below.)

Since the technology when used optimally is profitable, the farmer will decide to adopt the technology himself in the second period with high probability so long as N is large. Whether the farmer profitably uses the technology, however, depends on whether the technology is filter congruent. Formally:

Proposition 5. *Consider a farmer who learns from a best-practice demonstrator in the first period and then decides whether and how to farm in the second. If $N = \infty$, the farmer chooses to farm in the second period with probability* 1*. Moreover, with probability* 1*:*

1. *If the technology is filter incongruent, then the farmer is not profitable in the second period.*

2. *If the technology is filter congruent, then the farmer is profitable in the second period so long as $f_0(s_0) > \mu$, and m is sufficiently small.*

Proposition 5 indicates that agents can experience disappointment. Farmers can observe data–perhaps in a demonstration–that leads them to believe that they can use a technology profitably. Yet when they try it themselves they will find low returns. Some prior evidence can be interpreted in this light. Prominently, the fertilizer experiments by Duflo, Kremer and Robinson (2008) study the impact of being invited to observe a trial on a farmer's plot and of having a trial performed on one's own plot. In both cases, they find that a modest percentage of farmers adopt fertilizer the next year. Most interesting for our purposes, though, is what happens in the following years. The rates of adoption in subsequent years declines. It appears as if farmers were disappointed. In traditional models, participation in trials or observing trials is sufficient to encourage adoption. In our model, the impact of this type of "extension" activity depends on whether the technologies are "filter congruent"; i.e., how well they line up with the practices farmers notice. When farmers adopt a technology they may inadvertently use it much differently than how it was used in the trial.

5.2.1 Adding Communication

Note that an additional form of communication can help when the technology is filter incongruent: the best-practice demonstrator can point out his input choice, s_0, to the farmer either before or after the demonstration. Since s_0 is a single value, the farmer would then effectively have access to the full history h when updating his beliefs (so long as he processes the communication). As a result, this farmer is more likely to profitably adopt the technology.

In practice, the demonstrator makes many input choices and communicating these choices to the farmer is costly (Niehaus 2011). He only wants to communicate choices along dimensions the farmer is predisposed not to attend to. The effectiveness of communication is then increasing in the degree to which the demonstrator has knowledge of the farmer's mental model; i.e., of what he does and does not attend to.

This discussion speaks to the difficulty of agricultural extension (Evenson 2001, Gautam 2000, Anderson and Feder 2004). Simply sending an extension worker will not guarantee subsequent adoption and proper use of a profitable technology, even if the worker demonstrates how to use it. For certain technologies, effective communication is key.

6 Broader Issues

Learning through noticing modifies how we think about a range of issues in technology adoption and use.

Human Capital

The model and experimental results highlight two previously overlooked aspects of human capital, one specific and one general. First, the analysis suggests the importance of knowing which features of a given technology to attend to (the level of π in the formal model). Absent such capital, agents will fail to optimize along important dimensions and may fail to adopt a profitable technology. Second, it suggests a more general form of human capital: the ability to attend to and encode information along many dimensions (low e in the formal model). Such capital can have large returns across applications.

This more nuanced view of human capital forces us to re-think the role of experience and education in promoting efficient use of technologies.

The Role of Experience

The standard intuition is that there is mainly a significant role for learning to improve on outcomes when there is a new technology (Nelson and Phelps 1966, Schultz 1975, Foster and Rosenzweig 2010; for a recent empirical paper relying on this assumption, see Suri 2011). Underlying this intuition is that experience using a technology guarantees effective learning. The model and experimental results challenge this intuition: people may be persistently away from the production possibility frontier when they fail to notice key dimensions of the production process. While experience will lead people to optimize along dimensions they attend to, they may be far from optimizing along dimensions they fail to notice. In the context of our field experiment, we saw that farmers were much more likely to know their distance between lines than their pod size and were closer to optimizing along this dimension.[16]

This insight may help us understand persistent differences in practices across agents, for example farmers. A large literature posits that persistent differences must reflect differences in endowments or constraints (Griliches 1957, Schultz 1963, Gerhart 1975, Croppenstedt et al. 2003). A prominent recent example of this is Suri (2011) which argues that heterogeneity in adoption of hybrid maize in Kenya can be explained by heterogeneity in benefits and costs of adoption, under the identifying assumption that experience has supplied farmers with sufficient knowledge of these benefits and costs. For example, while the estimated gross returns of hybrid over nonhybrid maize of a subsample of non-adopters is quite large, these farmers also face plausibly higher costs to acquiring fertilizer, which is a complementary input to hybrid production.

Learning through noticing provides a different explanation for persistent heterogeneity in farming practices: farmers may act differently not because of heterogeneity in the inherent profitability of different practices, but because of differences in what they notice. For example, our theoretical results on technology adoption imply that farmers are more likely to adopt a profitable technology when they attend to the proper inputs. A measured comparative advantage of some farmers over others in profitably using a technology may stem from differences in what they notice, rather than from fundamental differences in endowments or constraints.

To test this hypothesis, it would be necessary to collect data on what people notice, extending our approach from the context of seaweed farming. For example, in the context of maize production in Kenya, it would be useful to understand whether persistent hybrid adopters are more likely to

[16]While bandit models (Gittins 1979) and "local" models of learning by doing (Conley and Udry 2010) also allow for farmers to persitently fail to optimize, the predictions of such models are very different. When, as in such models, the failure to learn comes from a failure of experimentation, the binding constraint is a lack of data. To diagnose problems, we should look for indications that people do not have sufficient data available, for example that they insufficiently experiment with different inputs; to solve these problems, we merely need to provide the relevant data, or to subsidize their collection. When the binding constraint is instead a failure to notice, knowledge of what data people have available is insufficient to diagnose problems; providing this data is not enough to solve them.

be able to answer questions about what they do along dimensions important when using hybrids, e.g., about when and how they apply fertilizer or exactly how much fertilizer they use (Duflo et al. 2008). Importantly in this context, even nonadopters are likely to have used hybrid maize at some point in the past (Suri 2011).

The model also suggests that not all experience is necessarily helpful. In particular, experience with one technology may actually make it difficult to profitably use a new technology if the relevant inputs are dissimilar, leading the new technology to be filter incongruent. This indicates that knowledge of the technologies farmers have used in the past may predict their propensity to adopt a new technology. The example of farmers being more likely to adopt hybrid sorghum when they previously adopted hybrid corn is suggestive (Brandner and Strauss 1959).

The Role of Education

Learning through noticing suggests that education, for example through agricultural extension, can be effective even when people are using old technologies they have much experience with. To understand when a failure to notice results in a failure to learn, we should look at indicators of insufficient attention to key input dimensions, for example when people cannot answer important questions about what they do as was the case with pod size in our field experiment. To promote better decisions, it is necessary not just to draw attention to these dimensions and to provide relevant information, but to provide this information in a way that is easy to process.

This last point speaks to the difficulty of agricultural extension (Evenson 1997, Gautam 1999, Anderson and Feder 2004). As our experimental results indicate, simply sending an extension worker will not guarantee subsequent adoption and proper use of a profitable technology, even if the worker demonstrates how to use it. Farmers did not change their practices with respect to pod size in response to the experimental trial, but did so after being presented with a summary of the trial results. Even when it looks like agricultural extension is ineffective, there may still be a role for extension to improve on outcomes. For certain technologies, effective communication is key.

Education can also be beneficial through increasing agents' ability to process information along multiple dimensions (through lowering e). In fact, one major role of education is to supply this general form of human capital. This provides a microfoundation for the widely held belief that schooling can increase the speed of learning and the ability to adapt to new circumstances (Nelson and Phelps 1966, Schultz 1975, Rosenzweig 1995, Foster and Rosenzweig 2010). The model suggests that general schooling rather than specific training is particularly valuable because it provides human capital that is complementary to learning novel technologies.

References

Anderson, J.R. and G. Feder, "Agricultural Extension: Good Intentions and Hard Realities," *The World Bank Research Observer*, 2004, *19* (1), 41–60.

Besley, T. and A. Case, "Modeling Technology Adoption in Developing Countries," *American Economic Review*, 1993, *83* (2), 396–402.

____ **and** ____ , "Diffusion as a Learning Process: Evidence from HYV Cotton," *Working Papers*, 1994.

Bordalo, P., N. Gennaioli, and A. Shleifer, "Salience Theory of Choice Under Risk," *Quarterly Journal of Economics*, 2012, *Forthcoming.*

Brandner, L. and M.A. Strauss, "Congruence Versus Profitability in the Diffusion of Hybrid Sorghum," *Rural Sociology*, 1959, *24* (4), 381–383.

Conley, T.G. and C.R. Udry, "Learning About a New Technology: Pineapple in Ghana," *American Economic Review*, 2010, *100* (1), 35–69.

Croppenstedt, A., M. Demeke, and M.M. Meschi, "Technology Adoption in the Presence of Constraints: The Case of Fertilizer Demand in Ethiopia," *Review of Development Economics*, 2003, *7* (1), 58–70.

Duflo, E., M. Kremer, and J. Robinson, "Why Don't Farmers Use Fertilizer? Experimental Evidence from Kenya," *Massachusetts Institute of Technology and Harvard University, working paper*, 2008.

Evenson, R.E., "Economic Impacts of Agricultural Research and Extension," *Handbook of Agricultural Economics*, 2001, *1*, 573–628.

Fliegel, F.C. and J.E. Kivlin, "Attributes of Innovations as Factors in Diffusion," *American Journal of Sociology*, 1966, pp. 235–248.

Foster, A.D. and M.R. Rosenzweig, "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, 1995, pp. 1176–1209.

____ **and** ____ , "Microeconomics of Technology Adoption," *Annual Reviews of Economics*, 2010, *2* (1), 395–424.

Gabaix, X., "A Sparsity-Based Model of Bounded Rationality," 2011.

Gautam, M., *Agricultural Extension: The Kenya Experience: An Impact Evaluation*, World Bank Publications, 2000.

Gerhart, J.D., "The Diffusion of Hybrid Maize in Western Kenya," 1975.

Gittins, J.C., "Bandit Processes and Dynamic Allocation Indices," *Journal of the Royal Statistical Society. Series B (Methodological)*, 1979, pp. 148–177.

Griliches, Z., "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica*, 1957, pp. 501–522.

Jovanovic, B. and Y. Nyarko, "Learning by Doing and the Choice of Technology," *Econometrica*, 1996, *64* (6), 1299–1310.

Koszegi, B. and A. Szeidl, "A Model of Focusing in Economic Choice," 2011.

Nelson, R.R. and E.S. Phelps, "Investment in Humans, Technological Diffusion, and Economic Growth," *American Economic Review*, 1966, *56* (1/2), 69–75.

Niehaus, P., "Filtered Social Learning," *Journal of Political Economy*, 2011, *119* (4), 686–720.

Rogers, E.M. and F.F. Shoemaker, "Communication of Innovations; A Cross-Cultural Approach.," 1971.

Rosenzweig, M.R., "Why are There Returns to Schooling?," *American Economic Review*, 1995, *85* (2), 153–158.

Schultz, T.W., *The Economic Value of Education*, Vol. 63, Columbia University Press New York, 1963.

——, "The Value of the Ability to Deal with Disequilibria," *Journal of Economic Literature*, 1975, *13* (3), 827–846.

Schwartzstein, J., "Selective Attention and Learning," *Unpublished Manuscript, Dartmouth College*, 2012.

Sims, C.A., "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Suri, T., "Selection and Comparative Advantage in Technology Adoption," *Econometrica*, 2011, *79* (1), 159–209.

Tornatzky, L.G. and K.J. Klein, "Innovation Characteristics and Innovation Adoption-Implementation: A Meta-Analysis of Findings," *IEEE Transactions on Engineering Management*, 1982, *29* (1), 28–45.

A Proofs

Before turning to the proofs, we introduce some notation and a useful lemma.

No matter the regime or the farmer's choice in the first period, in the second he will compare

$$\mu^{+*} \equiv \max_{s' \in \mathcal{S}} E[f(s')|\hat{h}] = \max_{s' \in \mathcal{S}} \pi E[f(s')|\hat{h}, P] + (1-\pi) E[f(s')|\hat{h}, NP] \tag{12}$$

with

$$\bar{\mu}^{+} \equiv \frac{1}{S} \sum_{s'} E[f(s')|\hat{h}] \tag{13}$$

and will choose the input level s^* that achieves μ^{+*} if $\mu^{+*} - \bar{\mu}^{+} > m$ and will not measure otherwise. Label the event that $\mu^{+*} - \bar{\mu}^{+} > m$ by G and the event that $\mu^{+*} - \bar{\mu}^{+} \leq m$ by NG. Further, let

$$\mu_M^{+*} \equiv E[f(s^*)|\hat{h}, M] \tag{14}$$

and

$$\bar{\mu}_M^{+} \equiv \frac{1}{S} \sum_{s'} E[f(s')|\hat{h}, M] \tag{15}$$

for $M = P, NP$. Note that $\mu_{NP}^{+*} = \bar{\mu}_{NP}^{+}$ given any history and that $\mu_P^{+*} - \bar{\mu}_P^{+} > m/\pi$ conditional on event G.

Lemma A.1. *Independent of the regime:*

1. *If the farmer does not keep track of the input level in the first period, then $\mu^{+*} = \bar{\mu}^{+}$ along any history.*

2. *If the farmer keeps track of the input level in the first period, then the probability of event G tends towards 1 as m tends towards 0. Further, $\mu^{+*} > \bar{\mu}^{+}$ with probability 1.*

Proof. **Part 1.** We have

$$\begin{aligned}\mu^{+*} &= \max_{s'} \pi E[f(s')|\hat{h}, P] + (1-\pi)E[f(s')|\hat{h}, NP] \\ &= \max_{s'} \pi E[\theta(s')|\hat{h}, P] + (1-\pi)E[\theta|\hat{h}, NP],\end{aligned} \tag{16}$$

and

$$\begin{aligned}\bar{\mu}^{+} &= \pi \left(\frac{1}{S} \sum_{s'} E[f(s')|\hat{h}, P] \right) + (1-\pi) \left(\frac{1}{S} \sum_{s'} E[f(s')|\hat{h}, NP] \right) \\ &= \pi \left(\frac{1}{S} \sum_{s'} E[\theta(s')|\hat{h}, P] \right) + (1-\pi) \left(\frac{1}{S} \sum_{s'} E[\theta|\hat{h}, NP] \right).\end{aligned} \tag{17}$$

Comparing (16) with (17), we see that it suffices to show that

$$E[\theta(s')|\hat{h}, P] = E[\theta(s'')|\hat{h}, P] \tag{18}$$

for all $s', s'' \in \mathcal{S}$ and $\hat{h}$ consistent with the farmer not keeping track of the input in the first period. But this equality follows from the fact that the farmer's marginal prior distributions over $\theta(s')$ and $\theta(s'')$ are the same under model P, as are the likelihoods $\Pr(\hat{h}|\theta(s') = \tilde{\theta}, P)$ and $\Pr(\hat{h}|\theta(s'') = \tilde{\theta}, P)$ for all $s', s'', \tilde{\theta}$, and $\hat{h} = (y_{p1}, \varnothing)_{p=1}^{N}$.

Part 2.

We will consider the learning by doing regime. The proof of this result for the other regimes is similar but slightly less involved because, in those regimes, $n(s)$ is known rather than stochastic.

Define $\mu_P^s = E[f(s)|\hat{h}, P]$ for all $s \in \mathcal{S}$, so $(\mu_P^1, \mu_P^2, \ldots, \mu_P^S)$ is a random variable. Re-order these μ_P^j in terms of decreasing values, where $\mu_P^{(i)}$ denotes the i-th highest value, so $\mu_P^{(1)} \geq \mu_P^{(2)} \geq \ldots \geq \mu_P^{(S)}$. Clearly, $(\mu_P^{(1)}, \mu_P^{(2)}, \ldots, \mu_P^{(S)})$ is also a random variable.

By definition, event NG is the event that $\mu^{+*} - \bar{\mu}^{+} \leq m$ which is equivalent to the event that

$$\mu_P^{+*} - \bar{\mu}_P^{+} \leq m/\pi. \tag{19}$$

It turns out that (19) holds only if

$$|\mu_P^s - \mu_P^{s'}| \leq \frac{Sm}{\pi} \text{ for all } s, s'. \tag{20}$$

To see this, suppose that there exists some s, s' such that $|\mu_P^s - \mu_P^{s'}| > Sm/\pi$. Then, for some $j > 0, k > 0$, we must have $\mu_P^{(j)} - \mu_P^{(j+k)} > Sm/\pi$. Now, expanding $\mu_P^{+*} - \bar{\mu}_P^{+}$, we have

$$\begin{aligned}
\mu_P^{+*} - \bar{\mu}_P^{+} &= \mu_P^{(1)} - \frac{1}{S}\left(\mu_P^{(1)} + \mu_P^{(2)} + \ldots + \mu_P^{(j)} + \ldots + \mu_P^{(j+k)} + \ldots + \mu_P^{(S)}\right) \\
&\geq \mu_P^{(1)} - \frac{1}{S}\left(\underbrace{\mu_P^{(1)} + \ldots + \mu_P^{(1)}}_{(j-1)\text{-times}} + \underbrace{\mu_P^{(j)} + \ldots + \mu_P^{(j)}}_{k\text{-times}} + \underbrace{\mu_P^{(j+k)} + \ldots + \mu_P^{(j+k)}}_{S-(j-1)-k\text{-times}}\right) \\
&= \frac{S-(j-1)}{S}\mu_P^{(1)} - \frac{k}{S}\mu_P^{(j)} - \frac{S-(j-1)-k}{S}\mu_P^{(j+k)} \\
&\geq \frac{S-(j-1)-k}{S}(\mu_P^{(j)} - \mu_P^{(j+k)}) \\
&> \frac{S+1-(j+k)}{S}\left(\frac{Sm}{\pi}\right) \\
&\geq \frac{1}{S}\left(\frac{Sm}{\pi}\right) \\
&= \frac{m}{\pi}.
\end{aligned}$$

Having established that (19) holds only if (20), we will now show that

$$\lim_{m\to 0} \Pr\left(|\mu_P^s - \mu_P^{s'}| \leq \frac{Sm}{\pi} \text{ for all } s, s'\right) = 0. \tag{21}$$

Note that

$$\Pr\left(|\mu_P^s - \mu_P^{s'}| \leq \frac{Sm}{\pi} \text{ for all } s, s'\right) \leq \Pr\left(|\mu_P^s - \mu_P^{s'}| \leq \frac{Sm}{\pi} \text{ for some } s, s' \text{ with } n(s) > 0\right).$$

But, conditional on $n(s) > 0$, $|\mu_P^s - \mu_P^{s'}|$ is a continuous random variable (under the true distribution) on $[0, \infty)$, so

$$\lim_{\tilde{m} \to 0} \Pr\left(|\mu_P^s - \mu_P^{s'}| \leq \tilde{m} | n(s) > 0\right) = 0,$$

which establishes (21), and hence that the probability of event G tends towards 1 as m tends towards 0.

Finally, we also have that

$$\Pr\left(|\mu_P^s - \mu_P^{s'}| = 0 | n(s) > 0\right) = 0,$$

so $\mu^{+*} > \bar{\mu}^{+}$ with probability 1.

■

Proof of Proposition 1 (sketch). The farmer has three options in the first period:

1. Do not measure or keep track of the input level.
2. Do not measure, but keep track of the input level.
3. Measure some level $s' \in \mathcal{S}$.

There are effectively three different actions the farmer can take in the first period since the expected utility from measuring level $s' \in \mathcal{S}$ is the same as the expected utility from measuring any other level $s'' \in \mathcal{S}$. Label the farmer's first period choice of action by a. The farmer's first period choice of action, together with his prior and second period strategy (as detailed at the beginning of the Appendix) induces a probability distribution over outcomes. Label this distribution by $\Pr^a$.

The farmer's first period expected (per-piece of land) flow utility is

$$\mu - \delta_m m - \delta_e e, \tag{22}$$

where δ_m is an indicator for whether the input is measured and δ_e is an indicator for whether the farmer keeps track of the input.

Conditional on the farmer's first period action, his second period expected (per-piece of land) flow utility under model NP is

$$\begin{aligned} u^a_{NP} &= \Pr^a(G|NP)(E^a[\mu^{+*}_{NP}|G, NP] - m) + \Pr^a(NG|NP)E^a[\bar{\mu}^+_{NP}|NG, NP] \\ &= E^a[\bar{\mu}^+_{NP}|NP] - m\Pr^a(G|NP) \text{ (since } \mu^{+*}_{NP} = \bar{\mu}^+_{NP}) \\ &= \mu - m\Pr^a(G|NP) \text{ (by the law of iterated expectations).} \end{aligned} \tag{23}$$

The farmer's second period expected (per-piece of land) flow utility under model P is

$$u^a_P = \Pr^a(G|P)(E^a[\mu^{+*}_P|G, P] - m) + \Pr^a(NG|P)E^a[\bar{\mu}^+_P|NG, P]. \tag{24}$$

It is easy to see that

$$\mu - m \leq u^a_{NP} \leq \mu \tag{25}$$

and

$$\mu \leq u^a_P \leq E^a[\mu^{+*}_P|P]. \tag{26}$$

A less obvious, but important inequality is that, conditional on the agent keeping track of the input level in the first period, there exists an $\varepsilon > 0$ such that

$$u^a_P > \mu + \varepsilon \tag{27}$$

for sufficiently small m. We will take (27) as given for now, and will establish this inequality at the end of the proof.

The farmer's expected utility given action a is

$$U^a = N[\mu - \delta_m m - \delta_e e + (1-\pi)u^a_{NP} + \pi u^a_P]. \tag{28}$$

When the farmer does not measure or keep track of the input level ($a = NE$), Lemma A.1 gives us that $\Pr^{NE}(G) = 0$, which implies that $u_P^{NE} = u_{NP}^{NE} = \mu$. Thus, the farmer's expected utility when he does not keep track of the input level equals

$$U^{NE} = 2N\mu. \tag{29}$$

If the farmer takes some other action (i.e., one in which he does keep track of the input level) his expected utility is bounded above by

$$N\{\mu - e + (1-\pi)\mu + \pi \max_a E^a[\mu_P^{+*}|P]\} \tag{30}$$

by (25), (26), and the fact that $e \leq m$. It is easy to see that (30) is less than U^{NE} for sufficiently small π. Thus, the farmer chooses not to measure or keep track of the input level for sufficiently small π. When the farmer does not measure or keep track of the input level, Lemma A.1 tells us that his posterior expectation of $f(s)$ will not vary in s.

We now want to show that the farmer keeps track of the input level in the first period so long as π is sufficiently large and e and m are sufficiently low. The utility the farmer achieves through a strategy in which he keeps track of the input level in the first period is bounded below by the utility from the strategy that includes not measuring, yet attending to the input level in the first period. Label the expected utility from this strategy by U^E. We have that

$$\begin{aligned} U^E &= N[\mu - e + (1-\pi)u_{NP}^E + \pi u_P^E] \\ &\geq N[\mu - e + (1-\pi)(\mu - m) + \pi u_P^E] \text{ (by inequality (25))} \\ &> N[2\mu - e - (1-\pi)m + \pi\varepsilon] \text{ for sufficiently low } m \text{ (by inequality (27)).} \end{aligned} \tag{31}$$

The right hand side of the above inequality is increasing in π and decreasing in m and e. Further, the right hand side tends towards

$$2N\mu + \varepsilon > U^{NE} \tag{32}$$

as $(\pi, e, m) \to (1, 0, 0)$, thus establishing that the farmer keeps track of the input level in the first period so long as π is sufficiently large and e and m are sufficiently low. When he keeps track of the input level, his posterior expectation almost surely varies in s (by Lemma A.1) and he measures the input level in the second period with positive probability. Further, when he keeps track of the input level, his posterior expectation for a given input level, s, almost surely tends to $\pi f_0(s) + (1-\pi)\frac{1}{S}\sum_{s'} f_0(s')$ as $N \to \infty$, which is maximized at $s = s_0$.

It only remains to show that (27) holds. Let G^δ be the event that $\mu^{+*} - \bar{\mu}^+ > m + \delta$ and $G^{\neg\delta}$ be the event that $m < \mu^{+*} - \bar{\mu}^+ \leq m + \delta$. We can write

$$\begin{aligned} u_P^a &= \Pr^a(G^\delta|P)(E^a[\mu_P^{+*}|G^\delta, P] - m) + \Pr^a(G^{\neg\delta}|P)(E^a[\mu_P^{+*}|G^{\neg\delta}, P] - m) \\ &\quad + \Pr^a(NG|P)E^a[\bar{\mu}_P^+|NG, P] \\ &\geq \mu + \Pr^a(G^\delta|P)\delta, \end{aligned} \tag{33}$$

with strict inequality whenever $\Pr^a(G|P) > 0$. The inequality follows from the fact that $\mu_P^{+*} - \bar{\mu}_P^+ > m + \delta$ conditional on event G^δ, together with the law of iterated expectations.

Lemma A.1 implies that, when the farmer keeps track of the input level in the first period, there exists some $p, \delta, \bar{m}$ such that $\Pr^a(G^\delta|P) \geq p$ for all $m \leq \bar{m}$, so (33) implies that $u_P^a > \mu + p\delta$ for all $m \leq \bar{m}$. Setting $\varepsilon = p\delta$ completes the proof.

■

Proof of Proposition 2. The farmer has two options in the first period:

1. Keep track of the input level.

2. Do not keep track of the input level.

Using arguments and notation analogous to those used in the proof of Proposition 1, the farmer's expected utility if he does not attend to the input level equals

$$U^{NA} = 2N\mu \tag{34}$$

and his expected utility if he attends to the input level equals

$$U^{A} = N[\mu - e + (1-\pi)u^{A}_{NP} + \pi u^{A}_{P}], \tag{35}$$

which is bounded above by

$$N[\mu - e + (1-\pi)\mu + \pi E^{A}[\mu_P^{+*}|P]]. \tag{36}$$

Since (36) is less than U^{NA} for sufficiently small π, the farmer does not attend to the input level for sufficiently small π in the learning by observing regime. Thus, a farmer who would not have attended to the input level absent exogenous experimentation also may not in this regime. Finally, such a farmer's posterior expectation of $f(s)$ does not vary in s by Lemma A.1.

■

Before proving Proposition 3, we state and prove an intermediate Lemma.

Lemma A.2. *Given any history h, the farmer will make the same decision in the second period whether he knows h, or just $((y_{p1}, \varnothing)_{p=1}^{N}, \bar{y}_{s^*}, s^*) \equiv \tilde{h}$.*

Proof. Recall that, in the second period, the farmer will choose a level s^* that achieves μ^{+*} if $\mu^{+*} - \bar{\mu}^{+} > m$ and will not measure otherwise. It will be helpful to write μ^{+*} and $\bar{\mu}^{+}$ as functions of information known by the farmer, so

$$\mu^{+*}(h) = \max_{s' \in \mathcal{S}} E[f(s')|h],$$

$$\mu^{+*}(\tilde{h}) = \max_{s' \in \mathcal{S}} E[f(s')|\tilde{h}],$$

etc. It suffices to show that $\mu^{+*}(h) = \mu^{+*}(\tilde{h})$ and $\bar{\mu}^{+}(h) = \bar{\mu}^{+}(\tilde{h})$ for all h.

We have

$$\begin{aligned}\mu^{+*}(h) &= \max_{s' \in \mathcal{S}} \pi \left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + n\sigma^2}\mu + \frac{n\sigma^2}{\sigma_\varepsilon^2 + n\sigma^2}\bar{y}_{s'} \right) + (1-\pi)\left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + N\sigma^2}\mu + \frac{N\sigma^2}{\sigma_\varepsilon^2 + N\sigma^2}\bar{y} \right) \\ &= \pi \left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + n\sigma^2}\mu + \frac{n\sigma^2}{\sigma_\varepsilon^2 + n\sigma^2}\bar{y}_{s^*} \right) + (1-\pi)\left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + N\sigma^2}\mu + \frac{N\sigma^2}{\sigma_\varepsilon^2 + N\sigma^2}\bar{y} \right).\end{aligned} \tag{37}$$

From (37), we see that, to calculate $\mu^{+*}(h)$, the farmer only needs to know $(s^*, \bar{y}_{s^*}, \bar{y})$, which he can calculate from $\tilde{h}$. Thus, $\mu^{+*}(h) = \mu^{+*}(\tilde{h})$.

We also have

$$\begin{aligned}\bar{\mu}^{+}(h) &= \frac{1}{S}\sum_{s'} \left\{ \pi \left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + n\sigma^2}\mu + \frac{n\sigma^2}{\sigma_\varepsilon^2 + n\sigma^2}\bar{y}_{s'} \right) + (1-\pi)\left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + N\sigma^2}\mu + \frac{N\sigma^2}{\sigma_\varepsilon^2 + N\sigma^2}\bar{y} \right) \right\} \\ &= \pi \left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + n\sigma^2}\mu + \frac{n\sigma^2}{\sigma_\varepsilon^2 + n\sigma^2}\bar{y} \right) + (1-\pi)\left(\frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + N\sigma^2}\mu + \frac{N\sigma^2}{\sigma_\varepsilon^2 + N\sigma^2}\bar{y} \right).\end{aligned} \tag{38}$$

From (38) we see that, to calculate $\bar{\mu}^{+}(h)$, the farmer only needs to know $\bar{y}$, which can be calculated from $\tilde{h}$. Thus, $\bar{\mu}^{+}(h) = \bar{\mu}^{+}(\tilde{h})$. ∎

Proof of Proposition 3. The farmer has four options in the first period:

1. Keep track of the input level.
2. Keep track of the input level and attend to recommendation.
3. Don't keep track of the input level, but attend to recommendation.
4. Don't keep track of the input level or attend to recommendation.

By Lemma A.2, options 1 and 2 are dominated by 3, so the farmer compares 3, attending to the recommendation ($a = R$), with 4, not attending to the recommendation ($a = NR$).

The farmer's expected utility if he does not attend to the recommendation is

$$U^{NR} = 2N\mu \tag{39}$$

by now familiar arguments.

If the farmer attends to the recommendation, then the probability distribution over his second period action is the same as if he kept track of pod size by Lemma A.2. However, the cost of attending to the recommendation is only e, while the cost of keeping track of the input level is Ne. Thus, the farmer's expected utility if he attends to the recommendation is given by

$$U^R = U^A + (N-1)e, \tag{40}$$

where U^A is the farmer's expected utility if he keeps track of the input (defined in the proof of Proposition 2).

While the farmer would keep track of the input level in the experimentation regime if

$$U^A > 2N\mu, \tag{41}$$

the farmer attends to the recommendation in the advice regime if

$$U^A + (N-1)e > 2N\mu. \tag{42}$$

Proposition 3 follows from comparing (41) to (42). ■

Proof of Proposition 4. (Sketch)

Part 1: Suppose the technology is not filter congruent. This means that the farmer does not attend in the first period. In the second, he then continues to farm so long as $\bar{\mu}^+ > 0$. But since the technology is not filter congruent, $\bar{\mu}^+ \to \bar{f}_0 < 0$ almost surely as $N \to \infty$. As a result, the probability that the farmer continues to farm in the second period tends towards 0 as $N \to \infty$.

Part 2: Suppose the technology is filter congruent. In the second period, the farmer continues to farm so long as

$$\max\{\mu^{+*} - m, \bar{\mu}^+\} > 0. \tag{43}$$

There are two cases to consider. In the first case, the farmer does not measure but attends in the first period. In this case, the left hand side of (43) tends almost surely to

$$\max\{\pi f_0(s_0) + (1-\pi)\bar{f}_0 - m, \bar{f}_0\} \tag{44}$$

as $N \to \infty$, where $\bar{f}_0 \equiv \frac{1}{S}\sum_{s'} f_0(s')$. Note that (44) is positive so long as π is sufficiently large by the assumption that the technology is profitable when optimized.

In the second case, the farmer does not measure or attend in the first period. In this case, the left hand side of (43) tends almost surely to $\bar{f}_0$ as $N \to \infty$, which is positive since the technology is assumed to be filter congruent.

We see then that, in both cases, the probability that the farmer continues to farm in the second period tends towards 1 as $N \to \infty$ so long as π is sufficiently large. ■

Proof of Proposition 5. (Sketch).

The farmer will decide to farm in the second period so long as

$$\max\{\mu^{+*} - m, \bar{\mu}^{+}\} > 0. \tag{45}$$

If the farmer attends in the first period, then (45) tends to $\bar{f}_0 > 0$ almost surely as $N \to \infty$ if the best practice is to not measure (i.e., if $\bar{f}_0 = \max\{f_0(s_0) - m, \bar{f}_0\}$). Otherwise, (45), tends almost surely to a value bounded below by $f_0(s_0) - m > 0$. Thus, if the farmer attends, the likelihood that he farms in the second period tends to 1 as $N \to \infty$. The probability that he is profitable tends towards 1 as $N \to \infty$ so long as $f_0(s_0) > \mu$ and m is sufficiently small because, in this case, he will choose size s_0 in the second period.[17]

[17]When $f_0(s_0) > \mu$, μ^{+*} almost surely tends to $f_0(s_0)$ and $\bar{\mu}^{+}$ almost surely tends to $\bar{\mu}^{+}_{\infty} \equiv \pi\left(\frac{f_0(s_0)}{S} + \frac{S-1}{S}\mu\right) + (1-\pi)f_0(s_0) < f_0(s_0)$. The farmer then almost surely selects s_0 in the second period so long as $f_0(s_0) - m > \bar{\mu}^{+}_{\infty}$, which is true for sufficiently small m.

If the farmer does not attend in the first period, then (45) tends to

$$\max\{f_0(s_0), \bar{f}_0\} > 0$$

almost surely as $N \to \infty$, so the likelihood that an inattentive farmer farms in the second period tends to 1 as $N \to \infty$. A farmer who did not attend in the first period chooses not to measure in the second. If the technology is filter congruent, then the farmer is almost surely profitable. If the technology is filter incongruent, then the farmer is almost surely unprofitable.

■

Table 1: Characteristics of Farmers: Demographic Characteristics

	Mean (1)	Standard Deviation (2)	N (3)
Ln(Monthly Per Capita Expenditures)	12.54	0.91	487
Age of HH Head	43.08	11.87	474
Number of Assets	8.09	3.23	487
HH Head is Literate	0.83	0.38	480
Years Farming Seaweed	18.36	7.15	475
Learned to Farm Seaweed from Parents	0.50	0.50	487
Has a Loan from Person to Whom Sells Seaweed	0.28	0.45	353

Notes: This table provide sample statistics on the demographic characteristics of farmers from the baseline survey.

Table 2: Farming Practices and Beliefs

	Mean (1)	Standard Deviation (2)	N (3)
Panel A: Farming Practices			
Number of Days in Previous Cycle	36.74	7.75	487
Mean Distance Between Lines at Baseline (Enumerator Measured)	15.47	1.96	486
Mean Distance Between Pods at Baseline (Enumerator Measured)	15.20	1.47	486
Mean Pod Size at Baseline (Enumerator Measured)	105.74	28.72	487
Farms Cottoni Strand	0.34	0.47	487
Panel B: Self-Reported Beliefs About Farming			
Length of Typical Line	5.05	1.04	481
Distance Between Lines	16.49	3.14	482
Current Pod Size	118.11	57.01	70
Optimal Distance Between Knots	15.97	2.84	481
Optimal Distance Between Lines	16.39	3.01	481
Optimal Pod Size	148.26	248.45	63
Optimal Cycle Length	37.43	7.14	486
Panel C: Percentage Who Could Not Provide an Answer for….			
Length of Typical Line	0.02	0.13	489
Distance Between Lines	0.01	0.12	489
Current Pod Size	0.86	0.35	489
Optimal Distance Between Knots	0.02	0.13	489
Optimal Distance Between Lines	0.02	0.13	489
Optimal Pod Size	0.87	0.34	489
Optimal Cycle Length	0.01	0.08	489
Panel D: Would Change Pod Size or Try Another Size Because….			
Of own accord	0.24	0.43	37
Because it worked in another plot	0.41	0.50	37
On NGO or government recommendation	0.08	0.28	37
On friend/family recommendation	0.11	0.31	37
Panel E: What Reasons May Make You Try a New Method			
Would not want to make any changes	0.04	0.18	482
Own initiative	0.02	0.13	482
Pest or Price	0.02	0.13	482
Advice from friend	0.10	0.30	482
NGO or Government recommendation	0.11	0.31	482
Seeing results on the plots of other farmers	0.39	0.49	482
Panel F: Notice when…			
Distance when distance between knots differs	0.21	0.41	486
Pod size is significantly larger than normal	0.21	0.41	487
Panel G: Self-Reported Changes in:			
Distance Between Lines	0.01	0.08	484
Cycle Length	0.09	0.28	475
Distance Between Knots	0.04	0.21	478
Pod size	0.12	0.32	446
Strain	0.51	0.50	487
Any aspect of farming, other than change in strain	0.19	0.39	487
Any aspect of farming	0.57	0.50	487

Notes: This table provide sample statistics on the farming techniques and beliefs of farmers from the baseline survey.

Table 3: Predicted Gains From Changing Farming Methodology

Panel A: Sort Treatment	
Median Percent Gain from Moving from Worst Return to Best Return	0.30
Income from Seaweed Farming (in Dollars)	928.75
Income from All Activities (in Dollars)	1400.36
Predicted Gain in Income from Seaweed Farming (In Dollars)	279.57
Predicted Percentage Gain in Total Income	0.20
Panel B: Weight Treatment	
Median Percent Gain from Moving from Current Return to Best Return	0.48
Income from Seaweed Farming (in Dollars)	881.33
Income from All Activities (in Dollars)	1417.72
Predicted Gain in Income from Seaweed Farming (In Dollars)	421.14
Predicted Percentage Gain in Total Income	0.30

Notes: This table provides the expected returns from changing seaweed farming technology. Panel A provides the expected returns from moving from the lowest return category to the highest return category in the sort treatment. Panel B provides the returns from moving from the category in the trial that is closest to existing farming methods to the highest return. Income data is self-reported. Predicted income assumes that farmers do not change methods (have fewer cycles, harvest earlier, etc.) if pod size were changed.

Table 4: Effect of Treatment on Propensity to Experiment and Observed Pod Size

	Made any change in technique (not including strain) since last survey		Enumerator measured pod size	
	(1)	(2)	(3)	(4)
Post 1	-0.146	-0.148	-11.333	-11.661
	(0.048)***	(0.057)**	(3.003)***	(3.578)***
Treatment * Post 1	0.072	0.079	-2.051	-1.550
	(0.060)	(0.071)	(4.411)	(5.306)
Post 2	-0.145	-0.150	-13.587	-13.859
	(0.050)***	(0.061)**	(2.896)***	(3.496)***
Treatment * Post 2	0.162	0.171	6.951	7.316
	(0.069)**	(0.084)**	(4.095)*	(4.982)
Sub-village fixed effects	X		X	
Farmer fixed effects		X		X
Observations	684	684	684	684

Notes: This table provides the coefficient estimates of the effect of treatment on farming methods in follow-up 1 and follow-up 2, conditional on baseline farming methods. The treatment dummy indicates that the farmer belongs in either the sort or weight treatment group. The mean of the dependent variable in Column 1 and 2 is 0.14, and mean pod size in the baseline (Columns 3 and 4) is 100.7. All regressions are estimated using OLS and standard errors are clustered at the farmer level. Statistical significance is denoted by: *** p<0.01, ** p<0.05, * p<0.10.

Table 5: Effect of Sort versus Weight Treatments

	Made any change in technique (not including strain) since last survey		Enumerator measured pod size		Enumerator Measured Distance Between pods	
	(1)	(2)	(3)	(4)	(5)	(6)
Post 1	-0.146	-0.148	-11.333	-11.661	0.142	0.127
	(0.048)***	(0.058)**	(3.010)***	(3.586)***	(0.180)	(0.213)
Sort * Post 1	0.089	0.100	3.944	4.657		
	(0.065)	(0.077)	(4.461)	(5.310)		
Weight * Post 1	0.052	0.053	-9.257	-8.929	0.289	0.304
	(0.075)	(0.089)	(6.610)	(7.882)	(0.328)	(0.387)
Post 2	-0.145	-0.150	-13.588	-13.859	0.014	-0.000
	(0.051)***	(0.061)**	(2.903)***	(3.504)***	(0.167)	(0.199)
Sort * Post 2	0.141	0.153	10.908	11.768		
	(0.075)*	(0.091)*	(4.418)**	(5.286)**		
Weight * Post 2	0.187	0.192	2.185	2.093	0.226	0.172
	(0.095)*	(0.114)*	(5.819)	(7.002)	(0.303)	(0.362)
Sub-village fixed effects	X		X		X	
Farmer fixed effects		X		X		X
Observations	684	684	684	684	499	499

Notes: This table provides the coefficient estimates of the effect of the different treatments on farming methods in follow-up 1 and follow-up 2, conditional on baseline farming methods. The mean of the dependent variable in Column 1 and 2 is 0.14, the mean pod size in the baseline (Columns 3 and 4) is 100.7, and the mean distance between pods in the baseline (Columns 5 and 6) is 15.37. All regressions are estimated using OLS and standard errors are clustered at the farmer level. Statistical significance is denoted by: *** p<0.01, ** p<0.05, * p<0.10.

Figure 1A: Baseline Pod Sizes

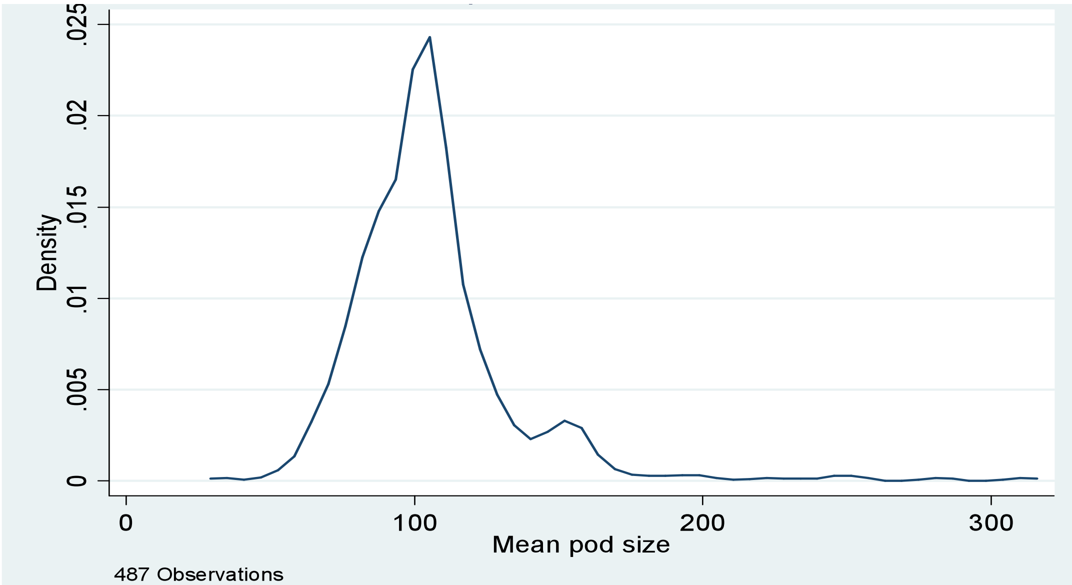

Figure 1B: Baseline Pod Sizes for Cottoni Growers

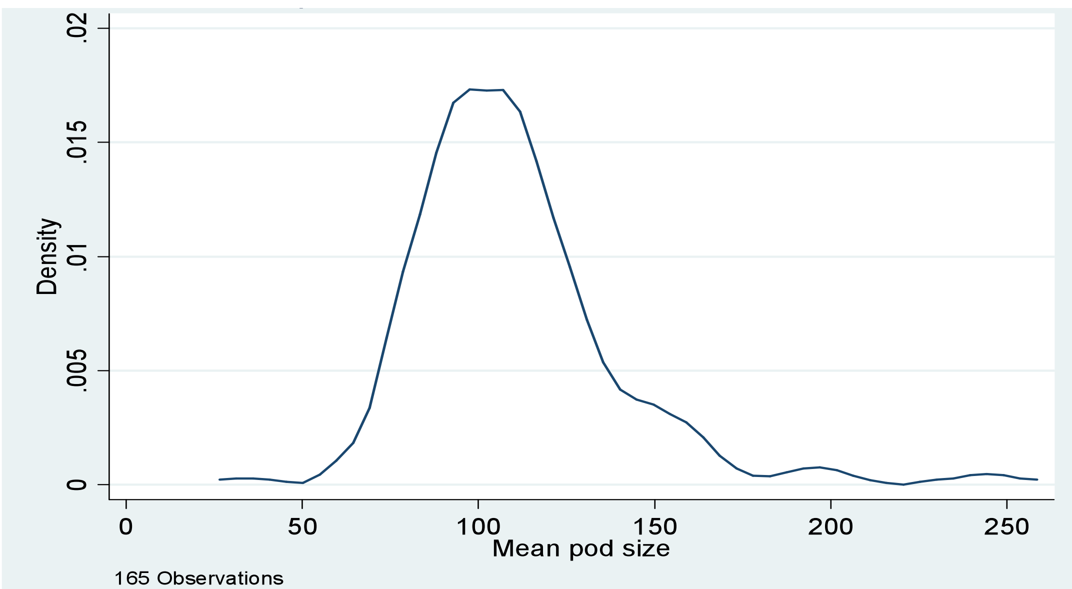

Figure 1C: Baseline Pod Sizes for Spinosim Growers

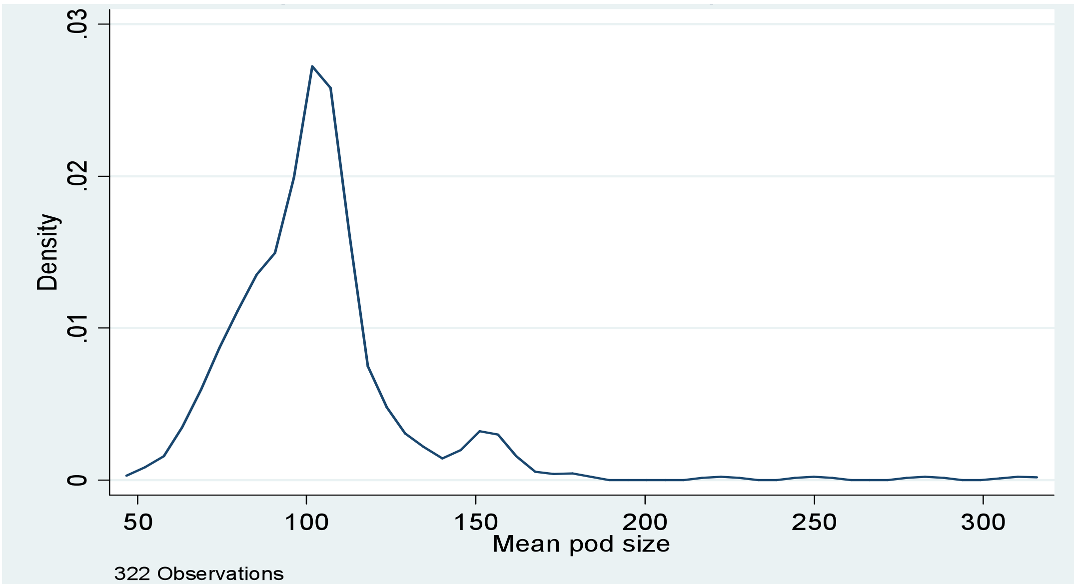

Figure 2: Recommendations from Trials, by Treatment

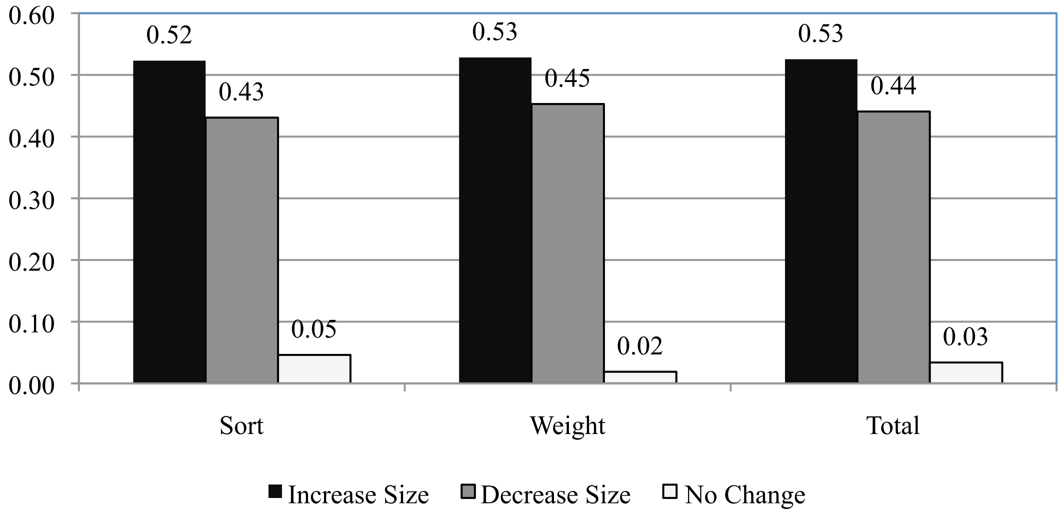

Notes: This figure gives information on the percentage of farmers told to increase or decrease their size as a result of the trials, by trial type. 65 farmers are included in the sort treatment, while 52 farmers are included in the weight treatment.

Figure 3: Recommended Size of Change, By Treatment

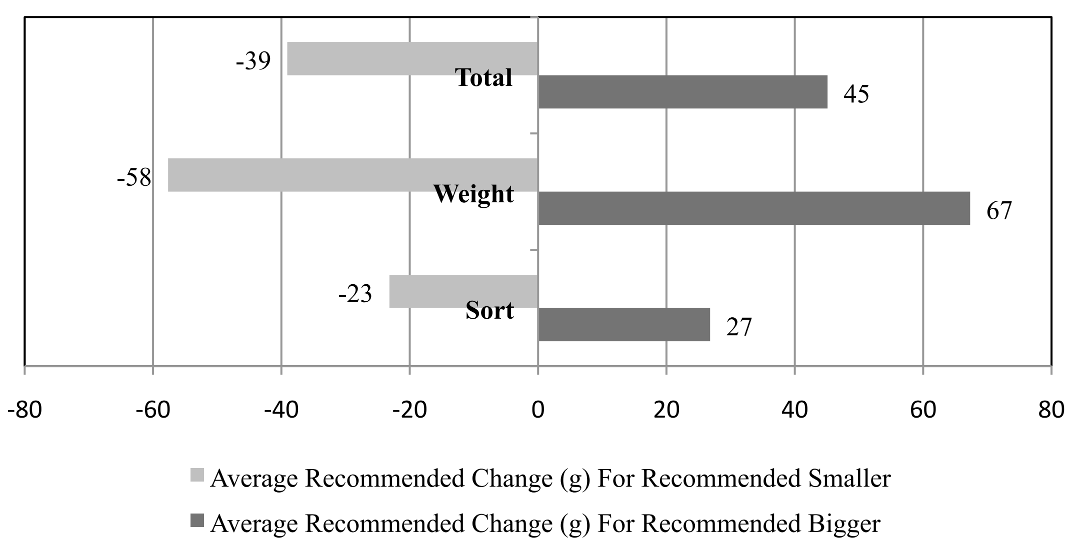

Notes: This figure provides the average recommended change in grams for farmers, by whether the farmer was informed to increase or decrease pod size.

Appendix Table 1: Regression Check

	Mean, by Treatment Group			Differences		
	Control	Sort	Weight	Col 1 - Col 2	Col 1 - Col 3	Col 2 - Col 3
	(1)	(2)	(3)	(4)	(5)	(6)
Monthly Per Capita Income	12.43	12.42	12.67	0.233	-0.016	-0.249
	(1.37)	(0.72)	(0.76)	(0.164)	(0.156)	(0.138)*
HH Head is Litarate	0.79	0.84	0.88	0.093	0.050	-0.043
	(0.41)	(0.37)	(0.32)	(0.059)	(0.060)	(0.064)
Number of Assets	8.36	7.95	8.28	-0.079	-0.409	-0.330
	(3.15)	(3.19)	(3.06)	(0.511)	(0.494)	(0.580)
Age of HH Head	43.23	43.95	43.54	0.304	0.718	0.414
	(12.22)	(12.14)	(11.33)	(1.937)	(1.905)	(2.192)
Years Farming	18.00	18.94	18.27	0.275	0.937	0.662
	(6.96)	(7.06)	(7.06)	(1.183)	(1.103)	(1.330)
Parents Farmed Seaweed	0.47	0.55	0.53	0.063	0.081	0.019
	(0.50)	(0.50)	(0.50)	(0.083)	(0.078)	(0.093)
Loans from someone sells to	0.31	0.33	0.22	-0.085	0.026	0.111
	(0.46)	(0.48)	(0.42)	(0.088)	(0.085)	(0.097)
Farms Cottoni	0.34	0.33	0.47	0.135	-0.008	-0.144
	(0.47)	(0.47)	(0.50)	(0.082)*	(0.074)	(0.091)
Podsize at Baseline	109.01	102.83	112.78	3.777	-6.174	-9.951
	(28.79)	(22.57)	(38.00)	(5.850)	(3.884)	(5.930)*
Number of Days of Previous Cycle	36.61	36.00	37.08	0.463	-0.612	-1.075
	(8.24)	(7.31)	(6.54)	(1.178)	(1.190)	(1.281)
P-Value from Joint Test				0.8895	0.1198	0.1516

Notes: This table provides a check on the randomization. Columns 1 - 3 provide the mean and standard deviation each baseline characteristic for the control group, sort group, and weight group, respectively. Columns 4 - 6 give the difference in means (and standard errors) between the noted experimental groups. Columns 7 - 9 give the difference in means conditional on sub-village fixed effects. The p-value for the joint significance of the treatment on all dependent variables is provided for Columns 4 – 6 (unfortunately, the joint test could not be estimated when including all sub-village fixed effects due to the limited sample size). Statistical significance is denoted by: *** p<0.01, ** p<0.05, * p<0.10.

Appendix Table 1: Regression Check (cont.)

	Differences, Controlling for Subvillage FE		
	Col 1 - Col 2	Col 1 - Col 3	Col 2 - Col 3
	(7)	(8)	(9)
Monthly Per Capita Income	0.378	0.049	-0.175
	(0.226)*	(0.182)	(0.162)
HH Head is Litarate	0.135	0.058	-0.148
	(0.061)**	(0.069)	(0.080)*
Number of Assets	-0.119	0.088	0.474
	(0.546)	(0.491)	(0.573)
Age of HH Head	0.193	0.919	4.396
	(2.289)	(2.063)	(2.751)
Years Farming	1.068	1.839	2.220
	(1.414)	(1.126)	(1.614)
Parents Farmed Seaweed	0.025	0.035	-0.005
	(0.098)	(0.086)	(0.126)
Loans from someone sells to	-0.025	0.081	0.029
	(0.119)	(0.092)	(0.132)
Farms Cottoni	0.107	0.105	0.001
	(0.075)	(0.067)	(0.089)
Podsize at Baseline	0.177	-5.802	-4.305
	(5.870)	(3.881)	(5.838)
Number of Days of Previous Cycle	0.459	-1.604	-1.592
	(1.306)	(1.178)	(1.566)

Appendix Table 2: Association Between Characteristics and Tendancy to Make Changes in Baseline

	Any Change in Methods		Any Change in Methods, not including strain	
	(1)	(2)	(3)	(4)
Age of Household Head	-0.000	-0.000	0.002	0.002
	(0.002)	(0.003)	(0.001)	(0.002)
Ln(Monthly Per Capita Expendutures)	0.016	0.022	0.010	0.008
	(0.022)	(0.029)	(0.015)	(0.018)
Head of Household Literate	0.004	-0.031	0.018	0.024
	(0.060)	(0.068)	(0.046)	(0.055)
Years Farming	-0.002	-0.005	0.001	-0.001
	(0.003)	(0.004)	(0.002)	(0.002)
Number of Assets	0.004	0.003	0.016	0.014
	(0.007)	(0.008)	(0.005)***	(0.006)**
Number of Farmers Dicuss with (out of random 3 farmers)	0.120	0.127	0.057	0.052
	(0.016)***	(0.017)***	(0.014)***	(0.014)***
Taught to Farm: Family	-0.034	-0.544	-0.033	-0.503
	(0.046)	(0.083)***	(0.036)	(0.287)*
Taught to Farm: Friend	-0.033	-0.507	-0.051	-0.529
	(0.045)	(0.079)***	(0.035)	(0.286)*
Taught to Farm: NGO or Government	0.299	-0.180	0.394	-0.176
	(0.134)**	(0.167)	(0.188)**	(0.353)
Taught to Farm: Buyer	0.161	-0.341	0.109	-0.443
	(0.089)*	(0.116)***	(0.087)	(0.300)
Taught to Farm: Myself	0.009	-0.494	0.109	-0.408
	(0.097)	(0.132)***	(0.087)	(0.301)
Observations		448		448

Notes: This table provides correlations between farmer characteristics and tendency to experiment in the baseline survey. In Columns 1 and 3, each coefficient estimate (and associated standard error) comes from a separate regression. In Columns 2 and 4, the coefficient estimates come from a single regression of all characteristics on tendency to experiment. All equations are estimated using OLS and standard errors are robust. Statistical significance is denoted by: *** p<0.01, ** p<0.05, * p<0.10.

Appendix Table 3: Effect of Treatment, by Recommendation

	Enumerator measured pod size	
	(1)	(3)
Increase	-21.091	-50.976
	(4.807)***	(1.851)***
Increase * Treatment * Post	27.264	27.060
	(6.223)***	(7.627)***
Increase * Treatment * Post2	19.502	19.242
	(5.808)***	(7.149)***
Sub-village fixed effects	X	
Farmer fixed effects		X
Observations	675	675

Notes: This table provides the coefficient estimates of the effect of the recommendations from the trial on farming methods in follow-up 1 and follow-up 2, conditional on baseline farming methods. The treatment dummy indicates that the farmer belongs in either the sort or weight treatment group. Bigger is an indicator variable for being told to increase size. Mean pod size in the baseline is 100.7. All regressions are estimated using OLS and standard errors are clustered at the farmer level Statistical significance is denoted by: *** p<0.01, ** p<0.05, * p<0.10.

Appendix Figure 1: Experimental Design

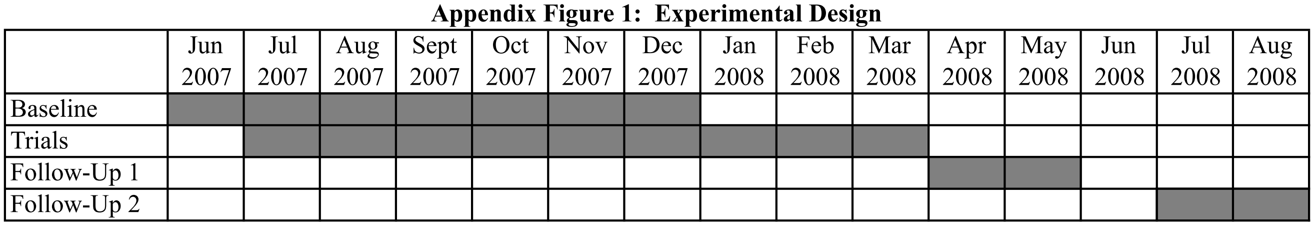

Appendix Figure 2: Sample Design

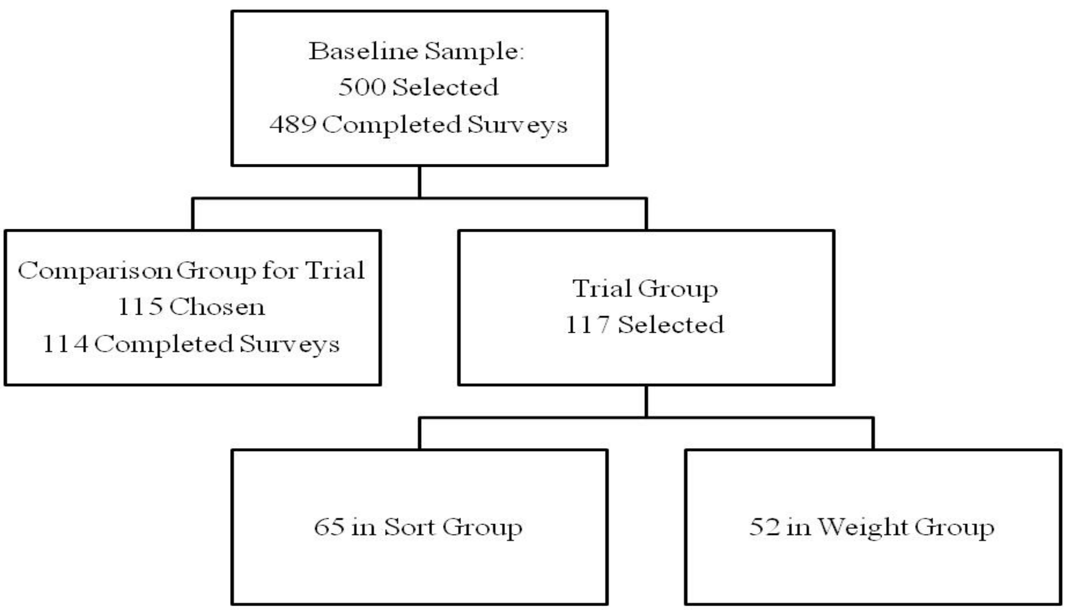

Appendix Figure 3: Example of Presented Trial Results

A. Weight Example

Pod Size	Distance	#Pods per line	Initial investment	Return per line
40	15	33	1650	4510
40	20	26	1300	3553
60	15	33	2310	1517
60	20	26	1820	1195
80	15	33	2970	1871
80	20	26	2340	1474
100	15	33	3630	1904
100	20	26	2860	1500
120	15	33	4290	597
120	20	26	3380	470
140	15	33	4950	1574
140	20	26	3900	1240

Currently

Pod Weight: 152.5	**Distance: 15**

Recommendation:

Pod Weight: 40	**Distance: 15**

B. Sort Example

Line Type	Distance	Average Pod Weight (g)	Return per pod (g)	Average pods per line	Return per line (g)
Large	14	129.92	167.65	34.1	5716.76
Medium		98.34	155.51		5302.92
Small		86.18	158.82		5415.88

Recommendation Switch to **large** pod size, with average weight **129.92**g.