

Learning About a New Technology: Pineapple in Ghana*

Timothy G. Conley
University of Chicago
tim.conley@gsb.uchicago.edu

Christopher R. Udry
Yale University
udry@yale.edu

May 24, 2004

Abstract

This paper investigates the role of social learning in the diffusion of a new agricultural technology in Ghana. We use unique data on farmers' communication patterns to define each individual's information neighborhood, the set of others from whom he might learn. Our empirical strategy is to test whether farmers adjust their inputs to align with those of their information neighbors who were surprisingly successful in previous periods. We present evidence that farmers adopt

*The authors have benefitted from the advice of Richard Akresh, Federico Bandi, Dirk Bergemann, Larry Blume, Adeline Delavande, Steven Durlauf, Ana Fernandez, Garth Frazer, Lars Hansen, Ethan Ligon, Charles Manski, Francesca Molinari, Stephen Morris, Jeff Russell, Chris Taber, Otto Toivanen, Giorgio Topa, and seminar participants at a number of seminars. We especially thank Yaw Nyarko, Mark Rosenzweig and three anonymous referees for their advice. The data used in this paper were collected by Ernest Appiah, Robert Ernest Afedoe, Patrick Selorm Amihere, Esther Aku Sarquah, Kwabena Moses Agyapong, Esther Nana Yaa Adofo, Michael Kwame Arhin, Margaret Harriet Yeboah, Issac Yaw Omane, Peter Ansong-Manu, Ishmaelina Borde-Koufie, Owusu Frank Abora, and Rita Allotey under the direction of the second author and of Markus Goldstein, who were visiting scholars at the Institute of Statistical, Social and Economic Research at the University of Ghana. This research has received financial support from the NSF (grants SBR-9617694 and SES-9905720), International Food Policy Research Institute, World Bank Research Committee, Fulbright Program, Social Science Research Council and the Pew Charitable Trust. The authors are of course responsible for any errors in this paper.

surprisingly successful neighbors' practices, conditional on many potentially confounding factors including common growing conditions, credit arrangements, clan membership, and religion. The relationship of these input adjustments to experience further supports their interpretation as resulting from social learning. In addition, we apply our methods to input choices for another crop with known technology and they correctly indicate an absence of social learning effects.

1 Introduction

The transformation of technology is fundamental to the development process. For a new technology to be adopted by an agent, particularly in agriculture, it must be adapted to the circumstances faced by that agent. Its characteristics usually will not be transparent to the new user (Evenson and Westphal (1995)). Consequently, an investment in learning about the new technology is associated with its adoption. If there are multiple adopters of the new technology in similar circumstances, as is often the case with an innovation in agriculture, then the process of learning about the new technology may be *social*. New users of the technology may learn its characteristics from each other.

The role of social learning in promoting growth and technology diffusion has been featured in the endogenous growth literature (Romer (1986); Lucas (1988); Aghion and Howitt (1998)). It is also an integral part of current practice in agricultural research and extension systems in developing countries. New technologies are introduced either by farmers' own experimentation or through formal sector intervention and the process of social learning encourages their diffusion (Bindlish and Evenson (1997); Rogers (1995)). Only recently, however, have economists made efforts to measure the quantitative importance of learning from others.¹

In this paper we investigate learning about a new agricultural technology by farmers in the Akwapim South district of Ghana. Over the last decade, an established system of maize and cassava intercropping for sale to urban consumers has begun a transformation into intensive production of pineapple

¹In contrast, there is a long tradition of empirical studies by economists of the adoption of new technologies in agriculture. Griliches (1957) is the seminal work. For reviews see Feder *et al* (1985) and Evenson and Westphal (1995). This important literature, however, does not isolate the role of learning processes from other determinants of adoption.

for export to European markets (Obeng (1994)). This transformation of the region's farming system involves the adoption of a set of new technologies, in particular the intensive use of fertilizer and other agricultural chemicals.

Measuring the extent of social learning is difficult for two major reasons. First, the set of neighbors from whom an individual can learn is difficult to define. Second, even with a proper definition of this set, distinguishing learning from other phenomena that may give rise to similar observed outcomes is problematic. In the absence of learning, individuals may still act like their neighbors as a result of interdependent preferences, technologies, or because they are subject to related unobservable shocks.

Direct data on information interconnections is typically unavailable to economists.² Consequently, economic investigations of the process of social learning have typically made assumptions that relate observed relationships between individuals - such as geographic proximity - to unobserved flows of information. This set of assumptions is critical for the measurement of the extent of social learning, but can rarely be tested because of data limitations.³ For example, Foster and Rosenzweig (1995) provide tabulations indicating that 'friends and neighbors' are an important source of information about fertilizer use, but must use village aggregates as the relevant information set for social learning.

We have unusually rich data that allows us to address the concerns of neighbor definition more directly. Our approach draws on the classic work by Coleman *et al* (1957) which related adoption of new antibiotics to the network of social interconnections between the doctors. We collected detailed

²Exceptions include Woittiez and Kapteyn (1998) and Kapteyn (2000) who use individuals' responses to questions about their 'social environments' to describe their reference groups. Romani (2003) uses information on ethnicity and membership in cooperatives in Côte d'Ivoire to infer the probability of information flows. Another exception is Bandiera and Rasul (2002), who have information on the number (though not the identities) of people using a new technology known by particular farmers. Rauch and Casella (2001) is a very useful collection of papers that use direct information on social interactions more generally.

³In many investigations of learning in developing country agriculture, the reference group is taken to be all farmers in the village (Foster and Rosenzweig (1995), Besley and Case (1994), Yamauchi (2002)). Munshi and Myaux (1998) take exceptional care in the construction of reference groups for social learning by using external evidence on communication barriers arising from religion. See Manski (1993) for a concise discussion of the importance of reference group designations in identification of endogenous social effects.

information on who individuals know and talk to about farming. Hence we follow Coleman *et al* by defining information links between agents using responses to questions about which other agents they turn to for information.⁴

Once neighborhoods are defined, the identification of learning is still a formidable problem. The classic problem of omitted variables prevents us from inferring that learning effects must be present simply from observations on, say, the diffusion process of a new technology. The fact that a farmer is more likely to adopt a new technology soon after his neighbors have done so might be a consequence of some unobserved variable that is spatially and serially correlated, rather than learning. We believe that correlated unobservables are a general problem in the literature on agrarian technology, and it is apparent that they are important in the sample region (see sections 3.2 and 6.3). We have collected data to mitigate this problem. Our data contains detailed geographic and soil information as well as information on credit and family relationships, allowing us to control for many potentially confounding factors.

Our identification problem can be thought of as a special case of the general problem of identification in social interactions models studied by Manski (1993, 1997), Moffitt (1999), Brock and Durlauf (1999) and others.⁵ This literature is concerned with the problem of inferring whether an individual's behavior is influenced by the behavior of those in his neighborhood or reference group. Manski (1993) demonstrates that identification of social effects in a cross section is generally tenuous. With panel data, prospects for identification improve, subject to the important caveat that the assumptions regarding the timing of social interactions must be properly specified. We are able to make progress because we are examining a specific form of social interaction: learning from each others' experiments. Our strategy for identifying learning effects has two prongs: first, we use data on both the geographic and the informational relationships between farmers to help distinguish the effects of learning from those of unobserved spatially-correlated shocks; and second, we use the time dimension of our data to isolate the

⁴Rogers (1995) and Birkhaeuser *et al* (1991) provide valuable surveys of research that describes and characterizes the set of neighbors from whom agents learn about new innovations in a wide variety of settings. Van den Bulte and Lilien (2001) show that the social contagion effects found by Coleman *et al* vanish once marketing effort is taken into account.

⁵See Brock and Durlauf (1999) for a survey of the literature on social interactions models.

impact of new observations of the productivity of fertilizer on innovations in a farmer’s fertilizer use. We investigate whether farmers change their input decisions to align with those of their information neighbors with similar circumstances who were previously surprisingly successful, earning higher than expected profits. The timing of responses to neighbors’ actions arises naturally from the staggered revelation of information from preceding plantings.

We model farmers’ learning about the productivity of inputs. Each harvest opportunity gives the farmer an observation on output for a given amount of input, and thus reveals information about the productivity of that input level. Our primary method to test for social learning is to estimate how farmers’ input decisions respond to the actions and outcomes of their neighbors. We know the inputs used and output harvested by each farmer, and thus can infer aspects of the information conveyed by each ‘experiment’ with the new technology by each respondent. We use our data on information flow between farmers to trace the impact of the information revealed by each experiment on the future input decisions of other farmers who are in the information neighborhood of the cultivator who conducted the experiment.

We find strong effects of news about fertilizer productivity in the information neighborhood of a farmer on his future innovations in fertilizer use.⁶ Specifically, we find for a given farmer: (1) he is more likely to change his fertilizer use after his information neighbors who use similar amounts of fertilizer achieve lower than expected profits; (2) he increases (decreases) his use of fertilizer after his information neighbors achieve unexpectedly high profits when using more (less) fertilizer than he did; (3) his responsiveness to news about the productivity of fertilizer in his information neighborhood is much greater if he has only recently begun cultivating pineapple; and (4) he responds more to news about the productivity of fertilizer on plots cultivated by experienced farmers and farmers with wealth similar to his. These conclusions hold when conditioning on the fertilizer use of farmers who are physically nearby and who therefore experience unobserved growing conditions that are highly correlated with his. In addition, they are robust to a variety of different definitions of information flow between farmers, and conditional on the fertilizer use of farmers with whom he has financial ties. Finally, we apply analogous methods to labor input choices to provide ev-

⁶We use the male pronoun to refer to farmers because the large majority of pineapple farmers in our data are men.

idence that pineapple farmers also learn about the labor productivity from their information neighbors, but that there is no evidence of similar learning in an established maize-cassava technology in these villages.

The remainder of this paper is organized as follows. In Section 2 we present the simple learning model that motivates our empirical specifications. Section 3 describes the empirical setting and our data. The implementation of the empirical model is discussed in Section 4 and Sections 5 and 6 describe our results.

2 A Learning Model

This section describes a simple model of learning about a new technology that we use to guide our empirical work. The basic form of this model is that farmers are trying to learn about the responsiveness of output $y_{i,t+1}$ on plot i to a discrete-valued input $x_{i,t}$ which we will call fertilizer:

$$y_{i,t+1} = w_{i,t}f(x_{i,t}) + \varepsilon_{i,t+1}. \quad (1)$$

$\varepsilon_{i,t+1}$ is an expectation zero productivity shock that is IID across farmers and time. Farmers do not know the function f ; it is the object of learning. The variable $w_{i,t}$ is a positive, exogenous growing conditions variable influencing the marginal product of $x_{i,t}$ that is correlated across farmers and time. This is motivated by the fact that agricultural production is often affected by shocks to the marginal product of inputs that are spatially and serially correlated (e.g., soil moisture, weeds, or pests).⁷ We assume the $w_{i,t}$ are observable to farmers but not the econometrician. The time indices reflect the revelation of information: in particular, input decisions might be influenced by the revelation of the growing conditions. To simplify notation, we suppose that each farmer i has only one plot, also designated i . The price of the input $x_{i,t}$ is a constant p .⁸

The essential feature of the farmer’s problem is that the responsiveness of output to fertilizer use is unknown. However, the farmer has beliefs about the technology f that evolve over time. We will not be specific about how the farmer’s beliefs evolve, but simply posit that farmer i ’s subjective expectations, $E_{i,t}(\cdot)$, can be defined. We begin by deriving implications of

⁷ $w_{i,t}$ could include a forward looking component, e.g. a rain forecast.

⁸In our study area, fertilizer prices are common across farmers and essentially constant throughout the sample time span.

social learning for observable actions in a very simple model. This model abstracts from some important aspects of learning about a new technology, in particular, from farmers' experimentation with different levels of fertilizer use in order to refine their knowledge of the technology and from potential strategic motivations. It is, however, sufficiently rich to clarify the essential identification problem that plagues attempts to measure the extent of social learning about agricultural technologies.

The time t problem faced by a myopic farmer i is to choose the input, x , to maximize his subjective expectation of time $t + 1$ profits for this input choice, $\pi_{i,t+1}(x, w_{i,t}) \equiv (w_{i,t}f(x) + \varepsilon_{i,t}) - px$. Since inputs are discrete, this is nothing more than choosing input level $x_{i,t}^*$ so that:

$$E_{i,t}\pi_{i,t+1}(x_{i,t}^*, w_{i,t}) \geq E_{i,t}\pi_{i,t+1}(x, w_{i,t}), \quad (2)$$

all $x \in \text{supp}(x_{i,t})$.

Learning influences the farmer's choices by changing the farmer's subjective expectation of $f(x_{k,t})$ which can be interpreted as his expectation of 'growing-conditions-adjusted' output:

$$\begin{aligned} E_{i,t} \{f(x_{k,t})|x_{k,t}, w_{k,t}\} &= E_{i,t} \left\{ f(x_{k,t}) + \frac{\varepsilon_{k,t+1}}{w_{k,t}} |x_{k,t}, w_{k,t} \right\} \\ &= E_{i,t} \left\{ \frac{y_{k,t+1}}{w_{k,t}} |x_{k,t}, w_{k,t} \right\} \equiv g_{i,t}(x_{k,t}). \end{aligned} \quad (3)$$

Examination of equation (2) reveals two main reasons for farmers' choices of $x_{i,t}^*$ to be interrelated. First, farmers may face similar realizations of growing conditions and thus make similar decisions about the optimal amount of fertilizer use. Therefore spatial or serial correlation in growing conditions will tend to induce a corresponding correlation in farmers' fertilizer choices. Second, farmers with related subjective information sets would have similar subjective expectations of growing-conditions-adjusted output, $g_{i,t}(\cdot)$, leading to similar input choices when they face similar prices and growing conditions. Learning will have an impact on fertilizer choices through changes in $g_{i,t}(\cdot)$ across farmers. If farmers learn from others, there is social learning and this will induce a systematic dependence of $g_{i,t}(\cdot)$ across farmers to the extent that farmers have common information. This dependence in $g_{i,t}(\cdot)$ series will induce a correlation in innovations in actions for farmers who share information.

The empirical task we face is to distinguish between these two reasons for correlations across farmers' innovations in fertilizer use. Shocks to growing conditions are a major concern as they are positively spatially and serially correlated (see e.g. Carter 1997). Even if farmers have perfect knowledge of the production function (so $g_{i,t}(x_{i,t}) = g(x_{i,t})$), if the $w_{i,t}$ are spatially correlated then from (2) the changes in fertilizer use ($x_{i,t}^* - x_{i,t-1}^*$) are also spatially correlated. We are able to distinguish the effects of spatially-correlated growing conditions from social learning because we collected data on *both* the geographic location of plots and the flows of information between farmers.

Local Learning

We do not want to focus attention on a particular model of learning or updating. Therefore, we do not start with primitive assumptions specifying a particular learning model. Instead, we consider empirical implications for a set of models where farmers learn from observations of inputs and outputs about their production technology locally, in the vicinity of the inputs they observe.

We choose to model learning as *local* in order to capture the notion that the information content of an outcome depends upon the input choice: in order to learn about a particular part of the production function, farmers (or their neighbors) must experiment with inputs in that portion of the function. This corresponds both to sample farmers' own descriptions of their experimentation with fertilizer inputs and to a substantial descriptive literature (Richards 1985; Amanor 1994). It is in contrast to some models of learning in which the same information about the production function can be deduced regardless of the portion of the production function that is used (Prescott 1972; Jovanovic and Nyarko 1996; Foster and Rosenzweig 1995).

To be more precise, consider a farmer's reaction to a new observation of inputs, growing conditions, and output that the farmer will use to update his subjective expectations from $g_{i,t}(\cdot)$ to $g_{i,t+1}(\cdot)$. We consider learning rules that are local in the sense that a new observation with input level x will only affect $g_{i,t+1}(x)$ and not $g_{i,t+1}(\tilde{x})$ for $\tilde{x} \neq x$. We assume in addition that farmers are *learning* in the sense that $g_{i,t+1}(x) - g_{i,t}(x) > 0$ when they observe higher shock-adjusted output than expected ($\frac{y}{w} > g_{i,t}(x)$), and likewise reduce their subjective expectation when they observe lower than expected $\frac{y}{w}$. This local learning rule is consistent with multiple kinds of learning, including Bayesian learning with independent priors over the elements of the support of x .

Implications for Actions

We now outline the implications for actions of a given farmer i in response to a new piece of information: $(x_{k,t}, w_{k,t}, y_{k,t+1})$. To make the argument more transparent, suppose for the moment that this farmer faces constant growing conditions. This new piece of information allows him to locally learn about f at input level $x_{k,t}$ and has an impact upon his subjective expectation $g_{i,t}(x)$ for $x = x_{k,t}$. With constant growing conditions, only the expected profits associated with input level $x_{k,t}$ shift; good news, higher than expected profits, inducing an upward shift and bad news, lower than expected profits, inducing a downward shift. Responses to changes in expected profits will be different for farmers who used the level $x_{k,t}$ in their previous planting versus those who used an alternative input level.

First consider farmers who previously used $x_{k,t}$. Good news, increases in expected profits, at input level $x_{k,t}$ will reinforce these farmers' beliefs that $x_{k,t}$ is the optimal choice for their growing conditions and they will continue to use it. Bad news, with a sufficiently large decrease in expected profits at input $x_{k,t}$, may induce farmers who used this level to switch to an alternative. However, bad news about $x_{k,t}$ provides no information about which among multiple alternatives the farmer will choose if he changes inputs in response to it. Farmers will revert to whatever option was second-best before the bad news arrived. Thus, our model offers no prediction for either the magnitude or direction of an input change induced by bad news about profits at $x_{k,t}$.

Now consider farmers who used an alternative to $x_{k,t}$. If such farmers receive bad news about $x_{k,t}$ they will not change their action as they already considered $x_{k,t}$ an inferior option. Good news about level $x_{k,t}$ may persuade these farmers to change to $x_{k,t}$, if it induces a sufficiently large increase in the associated expected profits. The magnitude and direction of a change in response to a single observation of good news about $x_{k,t}$ is also perfectly forecasted by the difference between $x_{k,t}$ and the farmer's previously used input level.

These implications hold for farmers with any model of local learning where new information $\{x_{k,t}, w_{k,t}, y_{k,t+1}\}$ effects only $g_{i,t}(x)$ for $x = x_{k,t}$. This includes models where more prior information is combined with the farmer's sample to update, perhaps in a Bayesian fashion, the means of each conditional distribution. They will not hold, in general, for models in which learning is global. For example, if farmers know that the production function lies in a particular parametric family, then it is possible to construct examples of learning algorithms in which an observation of the event

$\left\{ \left[\frac{y_{k,t+1}}{w_{k,t}} - g_{i,t}(x_{k,t}) \right] > 0 \right\}$ induces the farmer to move away from $x_{k,t}$.

In a more general setting in which growing conditions are random, the responsiveness of actions to news holds in a probabilistic sense. To see this, consider the special case in which x can take on two values $\{H, L\}$ with $H > L$. Suppose farmer i chooses $x_{i,t}^* = L$. He will change his input level to $x_{i,t+1} = H$ if $w_{i,t+1}g_{i,t+1}(H) - pH \geq w_{i,t+1}g_{i,t+1}(L) - pL$.⁹ So i chooses $x_{i,t+1} = H$ for any realization of period t growing conditions such that

$$w_{i,t+1} \geq w_{i,t+1}^* \equiv \frac{p(H - L)}{g_{i,t+1}(H) - g_{i,t+1}(L)}.$$

The threshold $w_{i,t+1}^*$ is strictly declining in $g_{i,t+1}(H)$, so the probability that i will choose $x_{i,t+1}^* = H$ conditional on all his past choices of inputs and growing conditions shocks is increasing in $g_{i,t+1}(H)$. Thus, the probability that i will change inputs to $x_{i,t+1}^* = H$ is higher if he observes a positive value of $\left[\frac{y_{k,t+1}}{w_{k,t}} - g_{i,t}(H) \right]$. If i 's neighbor achieves surprisingly high growing-conditions-adjusted output and profits at some level of fertilizer, he will be more likely to use that level in the next period.

However, correlated growing conditions shocks can induce a very similar pattern even in the absence of learning. Consider the two input level example when the production function is known. Farmer i will choose $x_{i,t}^* = H$ when

$$w_{i,t} \geq \frac{p(H - L)}{g(H) - g(L)}$$

and this will cause $y_{i,t+1}$ to be high relative to its unconditional expectation. If $w_{i,t}$ is positively spatially and serially correlated, then farmers' choices of $x_{i,t}^*$ and hence $y_{i,t+1}$ will be positively correlated across space and time as well. In particular, if growing conditions at small lags in time are highly positively correlated for physically proximate plots, then farmer i 's choices $x_{i,t}^* = H$ and associated likely-to-be-high output $y_{i,t+1}$ will tend to be followed by choices of H and higher outputs of i 's physical neighbors, solely due to the positive dynamic correlations in growing conditions. To the econometrician who does not observe $w_{i,t}$, a higher than long-run average realization of yields and profits by farmer i using the high fertilizer level tends to be followed in near future by an increased use of that quantity of fertilizer by his physical

⁹Of course this requires $g_{i,t+1}(H) > g_{i,t+1}(L)$: the farmer must expect H to be more productive than L if he is ever to use H .

neighbors who also tend to achieve a higher than long-run average profit. Therefore, it will be important in our empirical work to adequately condition on growing conditions.

This is a central problem in the identification of social learning in agriculture. It is difficult to gather sufficiently rich data on local conditions to be confident that learning effects can be distinguished from spatially- and serially-correlated shocks. Because growing conditions variables like weather and weeds are spatially correlated, the use of geography to identify learning neighborhoods is particularly problematic. Our strategy is to use geography (and timing) to identify spatially- and serially-correlated growing conditions, and direct data on information linkages to distinguish the effects of learning from those of similar growing conditions.

The discussion thus far has focused on myopic agents who learn from their neighbors' experience. In Appendix 1 we show that the core implications of local learning for actions can persist with forward-looking agents who have an incentive to experiment to learn about the characteristics of the new production function.

Empirical Strategy and Obstacles

Our empirical approach is to use a set of regressions to examine how innovations in farmer i 's fertilizer use respond to news about fertilizer productivity from his information neighbors' experiments, which we will refer to as *plantings*. First, we estimate a logistic regression where the binary outcome is an indicator that i changed his level of fertilizer from that used on his own previous planting. For simplicity, we drop stars and let $x_{i,t}$ be the value of fertilizer used per new plant applied by farmer i for the pineapples planted in round t . Our model implies a farmer who used $x_{i,t-j}$ at time $t-j$ and whose next planting was at time t is more likely to change if he observes (between $t-j$ and t) good news about $x \neq x_{i,t-j}$, or bad news about $x = x_{i,t-j}$; and is less likely to change if he observes good news about $x = x_{i,t-j}$ or bad news about $x \neq x_{i,t-j}$. Our approach to testing these implications is to construct indices meant to reflect the number of good and bad news experiments observed by farmer i at both $x = x_{i,t-j}$ and $x \neq x_{i,t-j}$; and to use these indices as regressors in a logistic regression predicting changes. The key variable that we use to control for the alignment of actions due to common growing conditions, $w_{i,t}$ is an average of the absolute deviation of $x_{i,t-j}$ from the inputs used by set of farmers whose plantings are close enough in time and space to farmer i 's time t planting that they had very similar growing

conditions. We also include several additional control regressors including ones for wealth, clan, religion, and experience farming pineapple .

After examining the binary description of input changes, we proceed to regressions predicting changes themselves: $\Delta x_{i,t} = (x_{i,t} - x_{i,t-j})$. This is motivated by our model’s implications for the magnitude and direction of changes in response to an observation of good news about an input level. In the model, if farmer i changes inputs in response single observation of good news about $x_{k,\tau}$, $\Delta x_{i,t}$ will equal $(x_{k,\tau} - x_{i,t-j})$, he will switch to $x_{k,\tau}$. Of course, farmer i may not always choose to switch in response to a good news observation at $x_{k,\tau}$, but $(x_{k,\tau} - x_{i,t-j})$ should be positively correlated with $\Delta x_{i,t}$. When farmer i observes good news about more than one input level, he could of course move to any of the good news levels. Therefore, we construct an index that is a weighted average of $(x_{k,\tau} - x_{i,t-j})$ across all the relevant good news experiments observed by farmer i . This index is meant to reflect whether the predominant good news information is associated with inputs above, below, or equal to his previous level $x_{i,t-j}$ and should be positively correlated with $\Delta x_{i,t}$. Again, a key variable in this regression is a control for common growing conditions. To create this regressor we construct another weighted average of $(x_{k,\tau} - x_{i,t-j})$. As for the logits, this growing conditions control is a weighted average across a set of farmers whose plantings are close enough in time and space to farmer i ’s time t planting that they had very similar growing conditions. Again, additional control regressors include wealth, clan, religion, and pineapple experience.

We modify these baseline regressions to explore the role of farmers’ experience and the source of information in determining their changes in actions. In any learning model in which agents’ beliefs converge, farmers’ reactions to new information will, in some sense, be decreasing in their accumulated information.¹⁰ This motivates an investigation of whether experienced farmers react less than inexperienced ones to new information. It is also easy to think of scenarios where the source of information influences how much farmers respond. For example, if we relax the assumption that the ε shocks are identically distributed across farmers then responses to observation of an experiment, (x, w, y) , might depend on the identity of the farmer associated with it. Suppose that the conditional mean of $\frac{y}{w}$ given x does not vary across farmers but its conditional variance does. Then farmers may

¹⁰It is clearly not necessary that farmers’ $g_{i,t}(\cdot)$ converge to the true conditional expectations, only that they do converge.

weight observations based on their relative variances under many learning rules. In particular, we investigate whether actions of experienced or large-scale farmers, or those within the observers' soil or wealth category are more informative than actions of their counterparts.

There are four components of this empirical strategy that need to be operationally defined in order to proceed. (1) We must define an appropriate set of individuals from whom farmer i can learn, his information neighbors. (2) We must construct measures of good and bad news. This requires an measure of whether an information neighbors' outcome is above or below farmer i 's subjective expectation of productivity given growing conditions and is complicated by the fact that we do not observe growing conditions. (3) We need to create indices measuring good/bad news and input levels associated with good news, because farmers commonly observe multiple experiments between plantings. (4) We must we construct indices to provide a way to condition on growing conditions to disentangle changes in inputs that are due to patterns in growing conditions from those predicted by our measures of good/bad news and good news inputs. Step one is detailed in Section 3, which describes our available data and the empirical counterparts to the variables in our stylized model of production above. Steps two, three and four are detailed in Section 4, which describes our empirical measures of subjective expectations and our construction of the indices that we use to measure good/bad news, good news input levels, and to control for growing conditions. A complete specification of our regressions and presentation of our results follows in Section 5.

3 Empirical Setting and Data

This section describes the data we use in our empirical work. First, we discuss the measure of farmers' communication that we use to define neighborhoods. Then we describe the economic and agronomic context of the farmers' learning problem and the mapping from our data to the variables in the learning model.

The data are drawn from a two-year survey (1996-98) of approximately 240 households in southern Ghana. A fairly comprehensive set of individual and agronomic data was collected. As some of the information was quite sensitive, we limited the size of the sample in order to maintain close oversight of the interview process. The sample was constructed in two stages. The

process began with the purposive selection of four ‘villages’ near the towns of Nsawam and Aburi.¹¹ This region is the center of the recent growth of intensive vegetable cultivation in the Eastern Region. The second stage was a random sample of married individuals: 60 couples (or triples, when there are two wives) were chosen by a simple random sample in each village. Two enumerators lived in or near each village and interviewed each respondent in 15 rounds at intervals of approximately six weeks.

In addition to the data on pineapple production, communication, knowledge and social networks described in Sections 3.1 and 3.2, we make use of data on the characteristics of farmers.¹² Wealth is defined as the value of the non-land assets held by the farmer at the start of the survey period. The clan indicator variables denote membership in a particular *abusua*, or matrilineal clan. The church indicator denotes membership in a particular charismatic church. Table 1 reports summary statistics.

3.1 Communication and Knowledge Data

One of our main innovations is that we are able to use the survey data to define information neighborhoods. We base our measure of information availability on direct data about conversations between individuals.

Each respondent was questioned about a random sample (without replacement) of seven other individuals in the same village, and with three other predetermined individuals who appear to be focal in the village. The samples of individuals produced responses to the question: “Have you ever gone to ___ for advice about your farm?”. In this case, we say an information link exists between farmers i and j if either i responded ‘yes’ to this question about j or if j responded ‘yes’ to this question about i . We use responses to this question as our benchmark definition of information neighbors because during the field research it appeared reliably-answered and it is transparently related to the learning process under study. The median number of information neighbors among pineapple farmers is 2, the maximum is 33.

¹¹We use data from the three of these ‘villages’ where pineapple is farmed. We will refer to the units as ‘villages’ although, in fact, only two are single villages. The other two locations are a pair of adjacent villages and a village with a set of outlying hamlets.

¹²A detailed description of survey procedures, copies of the survey instruments and the data archive are available at <http://www.econ.yale.edu/~cru2/ghanadata.html>.

To illustrate how these information links vary across individuals, we estimate a logit model of the probability of a link between i and j based on their underlying joint characteristics. Specifically, we estimate the probability that the pair (i, j) is linked as a function of the physical distance between their farms, the absolute differences in their ages and wealth, and indicator variables for whether they share a common religion, belong to a common clan, have a traditional office, have similar soil, and are of the same gender. The sample consists of all pairs of pineapple farmer respondents within each village.

Logit estimates of link probabilities are presented in Table 2. There is evidence of spatial correlation in link patterns as the marginal effect of proximity is to increase link probability, but distance is not the sole determinant of links. Individuals are more likely to have information links if they are of the same gender, the same clan, and similar ages. Individuals with different levels of wealth are more likely to be linked, reflecting the strong vertical patron-client ties that exist in these villages. There is no evidence that religion influences information links.

The importance of the statistically significant predictors differs dramatically in terms of absolute changes in probabilities. For the sake of comparison, take as a base pair one with the mean values of wealth difference, age difference, and distance (2.9, 10.9, and 1.25 respectively) with the same gender and soil but different clans, religions, and where neither party holds an office. The point estimate of the link probability for this base pair is 22%. This point estimate would shift to 14% if one of the parties held some office and up to 31% if instead the only difference from base was that they were from the same clan. A reduction in estimated probabilities to around 15% would accompany an approximate doubling of the base pair’s distance or age difference, individually. Likewise, an approximate doubling of the wealth difference would result in an increase to 31%. If the pair is not of the same gender, the predicted probability of one asking the other for advice drops dramatically to 5%.

In section 6, we check the robustness of our main results to varying definitions of the information neighborhood by using three alternative measures of information flow. Two of these measures are based on lists of interactions between respondents during the course of the survey (buying or selling goods, hiring labor, exchanging gifts, etc.), and the third is based on predicted links given the characteristics of pairs of farmers from the logit model described above. All of these measures are defined in Appendix 2.

3.2 Pineapple Production in Ghana

We focus on farmers' decisions about the use of fertilizer. While many aspects of the pineapple growing technology are new to these farmers, the most salient departure from traditional techniques is the use of this new chemical input. There is agronomic evidence that pineapple yields are very responsive to fertilizer and that the impact varies with local soil conditions and moisture patterns (Abutiate (1991); Purseglove (1972)). In informal interviews, individuals in the sample villages expressed conflicting views regarding the optimal quantities of fertilizer. There are official recommendations on fertilizer use available from the extension service of the Ministry of Agriculture, but these far exceed the levels of application in these villages.¹³

The specific input decision we consider is the application of fertilizer per plant during the period from six weeks after planting through six months after planting. During this period, pineapples are extremely sensitive to nutrient availability (Bartholomew and Kadzimin (1977); Soler(1992)). Pineapple is not strongly seasonal because it can be chemically forced to flower and thus fruit at any time during the year in southern Ghana. Hence we observe pineapple being planted at each round in our survey data. For each observed 'planting', our measure of inputs $x_{i,t}$ is the per-plant value of fertilizer applied during the reference period after planting.

Plot inputs and outputs were recorded at approximately six-week intervals over the two-year survey period. In terms of these six-week periods, pineapple cultivation in southern Ghana has the following approximate sequence. A plot planted at time t receives its crucial fertilizer inputs during periods $t + 1$ through $t + 4$. Chemically forced flowering occurs approximately at $t + 5$, at which point the eventual size of the crop begins to become apparent. The pineapple harvest is generally complete by $t + 9$. At this time, an exporter applies a treatment of chemicals that prompt ripening, harvests the fruits

¹³The recommendation is 400 Kg. of fertilizer/hectare, which is more than 10 times the mean fertilizer application observed in our sample. Only 4 of the 208 plantings we observed exceeded the recommended level of fertilizer application. Farmers argue that the recommended level is too high because fertilizer is 'too expensive', which we believe reflects an opportunity cost of capital to these farmers that is much higher than the interest rate used to calculate the recommendation. In our theoretical discussion, we incorporate the opportunity cost of capital in the price of fertilizer and assume it is the same across farmers. In the empirical application, we include farmer wealth as a conditioning variable, because this is likely to strongly associated with any variation in this opportunity cost across farmers.

and ships them by air to European markets. Therefore, a plot planted in period t begins to reveal its eventual profitability by $t + 5$, and the outcome is fully known by about $t + 9$. A plot planted in period $t + 1$ continues to receive fertilizer inputs through period $t + 5$. Hence, fertilizer inputs by a farmer on the plot planted in $t + 1$ would be influenced by the experience of his information neighbor on plots planted in period t , though the full impact of time t plantings would not be realized until $t + 9$. For the remainder of the paper, our notation will be consistent with these six week periods: $x_{i,t}$ will refer to the fertilizer applied during periods $t + 1$ through $t + 4$ on a plot planted at time t , $w_{i,t}$ refers to the growing conditions for that planting, and output is designated $y_{i,t+9}$.

In our empirical work, we focus on profits rather than yields because quality variation in pineapple output implies there is no natural unit in which yield can be measured. In addition, the cost of other inputs (most importantly labor) should be deducted from the value of output. We calculate the profits earned on these plantings, again on a per-plant basis.¹⁴ We calculate profits by deducting the value of all inputs, including family labor valued at the relevant gender-specific wage from the value of output.

We expect $w_{i,t}$ to be positively correlated across both space and time. Periods are sufficiently short that there is substantial correlation in soil moisture, weeds, and pest conditions on a given plot over time. Concern about spatial correlation is motivated by the observation in these villages that growing conditions vary spatially on the scale of hundreds of meters. Soil types and topographical features are highly correlated across neighboring plots, but vary within villages. Therefore, common village-level weather shocks can have varying impact across the village. Moreover, rainfall realizations can be different on opposite sides of a single village. Finally, weeds spread in a broadly continuous manner across space, and soil moisture and pest and disease environments are often much more similar on nearby plots than on more distant plots within villages. The nine period growing cycle also re-

¹⁴Actual harvests are observed only for those plantings which occurred early in our fieldwork. Plantings after round 5 were not fully harvested before the fieldwork ended. In the last round of the survey, respondents were asked for the price at which they could sell the crops that were currently standing on their plots. This is a normal transaction for crops such as cassava, in which plots full of partially grown plants are commonly sold to traders, who then hire labor to complete cultivation and harvesting of the plots. This kind of transaction is rare for pineapple, but respondents' familiarity with such sales of other crops made it easy for them to place values on their standing crops.

sults in substantial correlation for physically close plots at different but near points in time, due to the overlap in much of the environmental conditions they experience.

We also use information about the plots themselves. We have information on soil type and (for approximately 80% of plots) measures of soil pH and organic matter. In addition, all plots were mapped using global positioning system equipment. This procedure yields much more accurate measures of plot size and location than are available in most surveys in LDCs and makes it possible for us to distinguish explicitly between the effects of information connections and those of geographic proximity. To do so, we define *geographic* neighbors of a given plot to be those within 1 kilometer of the geographic center of all their pineapple plots.¹⁵ The median number of geographic neighbors is 12, the maximum is 25.

The scope for experimentation within plots is severely limited because the plots in our sample are close to the minimal viable scale. The median plot size in our data is approximately .5 hectares. The inexperienced pineapple farmers who exhibit the most evidence of learning have a median plot size of .25 hectares; exporters are reluctant to harvest and export crops from plots this small (only 5 plantings in our data were as small as .125 ha.). Plots have to be harvested on a single day for efficient export of the fresh fruit by air to Europe. It is essential, therefore, that the fruits mature simultaneously, which requires common treatment across plants within the plot.

Our main estimation sample is constructed as follows. We begin with information on pineapple being grown on 406 plots by 132 farmers. Of these plots, 288 were planted during our survey. Plot input data is missing on 3 of these plots, leaving 285. 77 of these were planted in round 12 or later, so they lack complete information on fertilizer use, leaving 208 plantings. We are missing data for some rounds on 8 of these, leaving 200 plantings. 87 of these are initial plantings, leaving 113 observed changes in fertilizer use. GIS information is missing on 6 of these plots, leaving a final estimation sample of 107 plantings by 47 farmers.

Our main sample is comprised of pineapple farmers. Figure 1 shows the pattern of adoption of pineapple in our sample villages: from less than 10% in 1990, pineapple spread very rapidly until more than 46% of farmers were cultivating pineapple in 1997. For some of the subsequent analysis, we

¹⁵Our results are not very sensitive to the radius chosen here, qualitatively identical results obtain with a range of bandwidths from 500 to 1500 meters.

divide pineapple farmers into two groups: experienced farmers who adopted pineapple before 1994, and inexperienced farmers who adopted more recently. This divides the sample of pineapple farmers approximately in half.¹⁶

Those who cultivate pineapple are different from those who do not in three important respects. First, pineapple farmers are much wealthier, with an average non-land wealth of 1.1 million cedis, compared to .4 million cedis for those who do not cultivate pineapple.¹⁷ Second, pineapple farmers are much more likely to be male: 70% of pineapple farmers are male, while only 27% of those who do not cultivate pineapple are male. Third, pineapple farmers – especially experienced pineapple farmers – are more likely to be in each others’ information neighborhood. Table 3 provides a summary of our baseline information link distribution by experience. Over 20% of experienced pineapple farmers (within each village) have approached each other for advice about farming, while only 6% of non-pineapple farmers are in each others’ information neighborhood. A similar pattern is observed using our other information metrics. It may be the case that these information connections were important determinants of the adoption process; however, we have too few instances of new adoption during our sample period to address this question formally. In section 5, we discuss the possibility that farmers vary in their ability to learn from others, and in particular that the pineapple farmers who comprise our sample are selected along that dimension.

¹⁶We chose not to use a continuous measure of experience because a few farmers have been cultivating pineapple for more than a decade. However, this early experience is unlikely to be relevant, because fertilizer was rarely used. Hence, there is little relevant difference in experience between those who have cultivated pineapple for, say, a decade and those who have cultivated for four years.

¹⁷Cedis are small units. The exchange rate during the sample period ranged from 1700-2300 cedis/US dollar.

4 Construction of Indices

This section details our measures of good and bad news, index of input levels associated with good news, and controls for growing conditions.

Measuring Subjective Expectations, and Good/Bad News

Farmer i increases (decreases) his expectation of the productivity of fertilizer at $x_{k,t}$ according to whether $\frac{y_{k,t+9}}{w_{k,t}} - g_{i,t}(x_{k,t})$ is positive (negative). Without knowledge of $w_{k,t}$ we obviously can not calculate the magnitude of this difference. However, it must have the same sign as $\{\pi_{k,t+9}(x_{k,t}, w_{k,t}) - E_{i,t}[\pi_{k,t+9}(x_{k,t}, w_{k,t})]\}$. High, positive spatial and serial correlation in growing conditions shocks permits estimation of the ‘rational expectations’ analog of $E_{i,t}\pi_{k,t+9}(x_{k,t}, w_{k,t})$: the objective expectation of $\pi_{k,t+9}(x_{k,t}, w_{k,t})$ for k within i ’s information neighborhood for a particular but unknown (to the econometrician) value of $w_{k,t}$. Suppose that within farmer i ’s information neighborhood there is a set of H farmers using input level $x_{k,t}$ sufficiently recently (say, at $t - 1$) and sufficiently geographically close to k that they have common objective growing conditions. For plots h in this set, $w_{h,t-1} = w_{k,t}$ and $x_{h,t-1} = x_{k,t}$ hence the expectation of $\pi_{k,t+9}(x_{k,t}, w_{k,t})$ given farmer i ’s information can be estimated with the sample median of realized profits for these H farmers, defined as $\hat{E}_i[\pi_{k,t+9}(x_{k,t}, w_{k,t})]$. Comparison of an individual’s realization, $\pi_{k,t+9}(x_{k,t}, w_{k,t})$ to this median is feasible and allows us to estimate whether $\pi_{k,t+9}(x_{k,t}, w_{k,t})$ is higher or lower than expected, even though we do not know $w_{k,t}$.

We use a local regression version of this approach. Dividing inputs into coarse categories of $x = 0$ and $x > 0$, we approximate farmer i ’s expectation of $\pi_{k,t+9}(x_{k,t}, w_{k,t})$ with the median of others’ profits for his information neighbors’ plantings using inputs in the same category as $x_{k,t}$ that are ‘close enough’ to the plot k at time t . When there is no planting in i ’s information neighborhood that is close enough, we use the median profits from all plots that are close enough to approximate i ’s expectation.¹⁸ Our operational definition of ‘close enough’ is plantings on plots within a distance of 1 kilometer and from rounds $t - 1$ and $t - 2$. We recycle notation from the above

¹⁸Qualitatively identical result obtain using the median of close enough plots in all instances.

Another check on the robustness of our results to other methods for imputing subjective expectations is provided by our results for predicted information links in Section 6.1. \hat{E}_i for this metric is a weighted average of all ‘close enough’ plantings using the predicted probabilities of links based on the estimates in Table 2 as weights.

paragraph and use $\widehat{E}_i [\pi_{k,t+9}(x_{k,t}, w_{k,t})]$ to denote this local regression estimate. Our estimate of whether $\pi_{k,t+9}(x_{k,t}, w_{k,t})$ is above farmer i 's subjective expectation is given by:

$$d_{i,k,t} \equiv 1 \left\{ \pi_{k,t+9}(x_{k,t}, w_{k,t}) > \widehat{E}_i [\pi_{k,t+9}(x_{k,t}, w_{k,t})] \right\}. \quad (4)$$

Of course observation of the indicator $d_{i,k,t} = 1$ will result in farmer i adjusting g at input level $x_{k,t}$ upwards. We classify each observation within farmer i 's information neighborhood as ‘good news’ if $d_{i,k,t} = 1$ and ‘bad news’ if $d_{i,k,t} = 0$. We discuss the robustness of our results to the input category definitions and our definition of ‘close enough’ in Section 6.

Construction of Indices Reflecting Innovations in Information

A given farmer plants pineapple only occasionally, rather than every period. Hence we do not always see a farmer’s reactions to individual observations of (x, w, y) . Instead, we often observe the cumulative effect of information from several neighbors revealed between input application opportunities for farmer i . In addition, information about the success of a planting at round t is revealed gradually: as we noted in Section 3.2, fertilizer inputs are applied in a window of $t + 1$ through $t + 1$. The plants are forced to flower at approximately $t + 5$ and although a good signal about harvest size is available at $t + 5$, the harvest is not complete until $t + 9$.

Our solution to observing only cumulative responses to a set of shocks that are revealed gradually is to create indices reflecting the predominate nature of change in the farmer’s stock of information associated with observing the set of outcomes since his last input decision. These indices can be roughly interpreted as the number of pineapple plants with higher/lower than expected profits observed by the farmer since his previous planting, normalized by a measure of his overall observations. We construct measures of good and bad news at the farmers previous input choice and at alternative levels. We also construct an index meant to reflect the disparity between a farmers previous input choice and the input levels associated with good news about profits. For the sake of exposition, we describe these indices for a farmer with two plantings, the first at time t_0 and the second at time t_1 .

Our index of good news at i 's previous input level is

$$G_{i,t_1}(x = x_{i,t_0}) \equiv \frac{1}{TotalPlants_{i,t_1}} \sum_{k \in N_i} \sum_{\tau \in [t_0-4, t_1)} 1\{x_{i,t_0} = x_{k,\tau}\} d_{i,k,\tau} Plants_{k,\tau} \psi(\tau; t_0, t_1). \quad (5)$$

This is an index of the amount of good news that i has received between his previous planting (at t_0) and his current planting (at t_1) about the level of fertilizer he previously used. It is a weighted average of the number of instances of good news ($d_{i,k,\tau}$ as defined in (4)) about plantings by farmers in i 's information neighborhood (N_i) who used the input level last used by i (hence the indicator function). The information conveyed by an information neighbor's planting depends on the number of plants involved, so we weight each observation by $Plants_{k,\tau}$. The change in beliefs associated with a piece of information depends on the stock of previous experiments that i has observed, so we normalize by the total number of plants i has observed from the beginning of our data until t_1 . Finally, $\psi(\tau, t_0, t_1)$ is a piecewise linear weighting function that reflects the gradual revelation of information about the harvest.¹⁹ Some - but not all - information about a neighbor's planting that began less than five periods before t_1 would be revealed in time for i to adjust fertilizer input on his t_1 planting. $\psi(\cdot)$ reflects the idea that 1/5 of the relevant information about the eventual harvest of a period t planting is revealed in each of the periods $t + 5$ through $t + 9$. Obviously this weighting scheme is *ad hoc*; we examine the robustness of our results to variations in the weighting scheme in Section 6. Full details of the construction of the indices are provided in Appendix 3.

We define a good news index at alternatives to the farmer's previous input level, $G_{i,t_1}(x \neq x_{i,t_0})$, by simply adjusting the indicator function $1\{x_{i,t_0} = x_{k,\tau}\}$ to $1\{x_{i,t_0} \neq x_{k,\tau}\}$. Indices of bad news at the farmers past input choice and at alternatives, $B_{i,t_1}(x = x_{i,t_0})$ and $B_{i,t_1}(x \neq x_{i,t_0})$ are constructed analogously substituting the corresponding indicators for input alignment and $(1 - d_{i,k,\tau})$ in expression (5).

When estimating the determinants of the magnitude of fertilizer innovations we have the prediction from our model that, conditional on changing inputs, a farmer who received one piece of good news information at input $x_{k,\tau}$ would change his input use to $x_{k,\tau}$. The difference between the input associated with the good news and his previous input level, $[x_{k,\tau} - x_{i,t_0}]$, should perfectly predict his change in inputs. This implication motivates the following index to describe the fertilizer amounts corresponding to good news experiments in a farmer's information neighborhood. Again taking farmer

¹⁹ $\psi(\tau; t_0, t_1) \equiv \min(\frac{t_1 - \tau}{5}, 1) - \max(\min(\frac{t_0 - \tau}{5}, 1), 0)$. See Appendix 3 for a more complete discussion.

i 's plantings at t_0 and t_1 , this index is defined as:

$$M_{i,t_1} = \frac{1}{TotalPlants_{i,t_1}} \sum_{k \in N_i} \sum_{\tau \in [t_0-4, t_1)} \psi(\tau; t_0, t_1) d_{i,k,\tau} Plants_{k,\tau} [x_{k,\tau} - x_{i,t_0}]. \quad (6)$$

Each neighbor with higher than expected profits, ($d_{i,k,\tau} = 1$) adds a term $\psi(\tau; t_0, t_1) Plants_{k,\tau} [x_{k,\tau} - x_{i,t_0}]$ which will be positive if farmer k used more fertilizer than farmer i did at time t_0 . This term's contribution to the index is negative if k used less fertilizer than i . Normalizing by the inverse of $TotalPlants_{i,t_1}$ dampens the magnitude of this index for those farmers with relatively higher numbers of previously observed plants. If all of farmer i 's good news observations are higher than his previously used level x_{i,t_0} , M_{i,t_1} is positive; if all of i 's good news observations used less than x_{i,t_0} , it is negative; and it will be near zero if all good news input levels are near x_{i,t_0} . Thus we expect that M_{i,t_1} should be a good predictor of both the direction and magnitude of changes in inputs.

Controls for Similar Growing Conditions

Our control for input adjustment in response to similar growing conditions is formed in a manner analogous to that used for M . We construct a plant-weighted index of input differentials from x_{i,t_0} across a neighborhood, N_i^{Geo} , defined strictly based on geographic and temporal proximity rather than information flows between farmers. This index measures how the farmer's previous input level differs from those recently used by his physical neighbors whose growing conditions are highly related to w_{i,t_1} :

$$\Gamma_{i,t_1} = \frac{1}{\sum_{k \in N_i^{Geo}} \sum_{\tau \in [t_1-4, t_1-1]} Plants_{k,\tau}} \sum_{k \in N_i^{Geo}} \sum_{\tau \in [t_1-4, t_1-1]} Plants_{k,\tau} [x_{k,\tau} - x_{i,t_0}], \quad (7)$$

Plantings within the four rounds preceding t_1 are included in this average across N_i^{Geo} which is defined as the set of farmers whose plot centers are within one kilometer from that of farmer i . Analogous indices are constructed to measure the absolute difference between geographic neighbors' inputs and x_{i,t_0} as well as to summarize input innovations within financial and gift networks - the 'financial neighborhood.'

5 Estimation Results

Base Regression Specifications

This section presents our base regression specifications. Alternative specifications and a discussion of robustness are considered in Section 6. We present these specifications as though each farmer has a time t and $t - 1$ planting for simplicity. We let the characteristics of i and his plot that we use for conditioning be contained in a vector $z_{i,t}$. These characteristics include the farmer's wealth, soil characteristics, and indicators for religion, clan, village, round of the planting and an indicator that is one if the farmer has been farming pineapple for export for less than three years at the start of the survey. Using the notation $\Delta x_{i,t}$ for the first difference of inputs, $(x_{i,t} - x_{i,t-1})$, and the notation $\Pr\{\Delta x_{i,t} \neq 0\}$ to refer to the probability of changing fertilizer use conditional on observable (to the econometrician) information at time $t - 1$, we estimate a logistic specification for this conditional probability:

$$\Pr\{\Delta x_{i,t} \neq 0\} = \Lambda \left[\begin{array}{l} \alpha_1 G_{i,t}(x = x_{i,t-1}) + \alpha_2 G_{i,t}(x \neq x_{i,t-1}) \\ + \alpha_3 B_{i,t}(x = x_{i,t-1}) + \alpha_4 B_{i,t}(x \neq x_{i,t-1}) \\ + \alpha_5 \frac{1}{N_i^{Geo}} \sum_{j \in N_i^{Geo}} |x_{i,t-1} - x_{j,t-1}| + z'_{i,t} \alpha_6 \end{array} \right] \quad (8)$$

The first four terms reflect the nature of new information to the farmer. The first term is our index of good news about farmer i 's time $t - 1$ input choice $x_{i,t-1}$. The second term is the index of good news about alternatives to $x_{i,t-1}$. The third and fourth terms are the indices of bad news at $x_{i,t-1}$ and at alternatives to $x_{i,t-1}$. Local learning implies that $\alpha_1, \alpha_4 < 0$ and $\alpha_2, \alpha_3 > 0$. The fifth term is the average of absolute deviations of the farmer's $t - 1$ input from those of his geographic neighbors, our control for growing-conditions-induced changes in input levels. We expect unobserved shocks to growing conditions to be positively spatially and serially correlated and thus $\alpha_5 > 0$.

Our baseline regression predicting changes in fertilizer use is:

$$\Delta x_{i,t} = \beta_1 M_{i,t} + \beta_2 \Gamma_{i,t} + z'_{i,t} \beta_3 + v_{i,t}. \quad (9)$$

$M_{i,t}$, defined in (6), is our index of input levels associated with good news that should predict the magnitude and direction of changes, and should therefore have a positive coefficient β_1 . The sign and magnitude of changes in response to bad news are ambiguous since they are determined by the previously second-best input level. Therefore, we confine our attention to the effect of good news on innovations. $\Gamma_{i,t}$ defined in (7) is our crucial control for

movements in $x_{i,t}$ induced by correlated growing conditions. As in (8), $z_{i,t}$ includes wealth, soil characteristics, and indicators for religion, clan, village, round of the planting and experience indicator. In addition, it includes a regressor analogous to $\Gamma_{i,t}$ with the neighborhood definition based on financial rather than geographic neighborhoods. This is motivated by alternate explanations that would suggest that significant β_1 estimates might be caused by omitted variable bias because information neighbors share common access to credit arrangements. Finally, the error term $v_{i,t}$ is permitted to be conditionally heteroskedastic and spatially correlated across plots as a general function of their physical distance using the spatial GMM approach of Conley (1999).

Farmers certainly learn from their own experience as well as from the experience of other farmers in their information neighborhood. But our panel's short time span prohibits us from including outcomes of lagged experiments by farmer i in (8),(9) (only 14 farmers have more than 2 fertilizer changes in the time span of our data). However, we are able to control for some aspects of learning from own experience directly by including an experience regressor and indirectly through the denominator of our indices G , B , and M .

The parameters of equations (8) and (9) are not parameters of a structural model and are best interpreted as those for a prediction of innovations in behavior. As such, measurement error in definitions of N_i is not a major concern, because this would not lead to an incorrect inference that news about fertilizer productivity is a useful predictor of innovations in fertilizer use. Endogeneity of neighborhoods that arises from individuals choosing neighbors who would be good to learn from is also not a concern. However, endogeneity of neighborhoods arising from unobservables (like wealth were it an omitted variable) that influence neighbor choices and profitability as well as Δx are, of course, an important concern as these unobservable effects could result in the same correlations as learning from neighbors with higher than expected profits/growing-conditions-adjusted yields. In section 6 we replicate our main results using the predicted information neighborhood, which is less subject to this concern because the predictions are based on variables that reflect deeper social connections that are plausibly unrelated to unobserved determinants of innovations in fertilizer use.

Results

Table 4 presents the coefficient and spatial standard error estimates from equation (8).²⁰ The dependent variable in the logistic regression is an indi-

²⁰The standard errors in all our specifications use limiting results for cross section esti-

cator equal to one if the farmer changed input use at time t . The metric for information neighbors in this Table is that derived from the question “Have you ever gone to ___ for advice about your farm?”

In column A, we see that the direction of the influence of our information and experience variables upon the predicted probability of changing is as anticipated. Observations of bad news at the farmer’s lagged input level strongly increase the predicted probability that he will change input levels. The estimated coefficient is positive, significantly different from zero (p-value of 0.001) and large. The standard deviation of bad news at the lagged fertilizer use is .12. Thus the estimated impact of a one standard deviation increase in this index of lower than expected profits at the old level of fertilizer used by i on the logit index is 1.5. The estimated range from 75th to the 25th percentiles of the distribution of estimated logit indices is 2.8. Observations of bad news at alternative levels of fertilizer use decrease the predicted probability of changing. The coefficient is negative and significantly different from zero at the 12% level. A one standard deviation increase (about .13) in the index of bad news at alternative fertilizer levels is associated with a decline in the logit index of approximately .15. The point estimate of the effect of good news at alternative levels of fertilizer use on the probability of changing is positive, and that of the effect of good news at the lagged level of use is negative, as expected, but these coefficients are very small and statistically indistinguishable from zero. As would be expected in virtually any model of learning, inexperienced farmers seem to be more likely to change input levels, but while the size of the coefficient is large, it is not estimated with precision (it is different from zero at only approximately the 20% level). The imprecision in this estimate may reflect the difficulty of distinguishing

mation with spatial dependence characterized by physical distance between the centroids of each farmer’s set of plots. Serial dependence is allowed for only by use of time (round) dummies. Specifically, spatial standard errors are calculated using the estimator in Conley (1999) with a weighting function that is the product of one kernel in each dimension (North-South, East-West). In each dimension, the kernel starts at one and decreases linearly until it is zero at a distance of 1.5 km and remains at zero for larger distances. This estimator is analogous to a Bartlett (1950) or Newey-West (1987) time series covariance estimator and allows general correlation patterns up to the cutoff distances. Note that plantings by the same farmer are allowed to be arbitrarily correlated as they are all distance zero from each other.

The inferences reported below are robust to cutoff distances between 1 km and 2 km. This is largely due to the fact that there is little spatial correlation in our regression errors because we are conditioning upon indices involving geographic neighbors’ actions.

it from our indices of news, which are scaled by $TotalPlants_{i,t}$ to directly capture an experience effect.

The estimated probability of changing fertilizer levels is strongly increasing in the average absolute deviation of farmer i 's lagged inputs from his geographic neighbors, providing evidence of the importance of positively serially- and spatially-correlated unobserved shocks to the productivity of fertilizer. None of the unreported round dummies are statistically different from zero nor are they jointly different from zero at conventional significance levels.

Column B of Table 4 adds an indicator for whether the respondent has ever received advice from the local extension agent. We do not know when any such conversation occurred. Those who have received advice from the extension agent are less likely to adjust their fertilizer use – this is consistent with the idea that farmers who have had extension contact are more confident in their beliefs.

Table 5, column A presents the results of estimating equation (9). The coefficient on the index of good news in the farmer's information neighborhood is positive, clearly statistically significant and large. A one standard-deviation increase in M (about 4) is associated with an increase in fertilizer use of approximately 4 cedis per plant, which is greater than the median level of fertilizer use per plant of those farmers who use fertilizer.

There is no aggregate trend in fertilizer use over our sample period. None of the round indicators is individually significant, nor are they jointly significantly different from zero. There is no evidence that changes in input use are significantly related to inputs used by financial neighbors. It is apparent that the growth of fertilizer use is much slower in village 1 than in the other villages. We believe that this is a reflection of the fact that pineapple cultivation is much less well-established in village 1 than in the omitted village 3, where (on a per-capita basis) six times more pineapple was cultivated. In this as in the following columns, changes in fertilizer use are strongly in the direction of the lagged use of one's geographic neighbors.

In column B, we examine relationship between experience and a farmer's responsiveness to information on the profitability of fertilizer. The coefficients on M for experienced and inexperienced farmers are statistically different from each other at the 3% level. The definition of experience is the same as that in Table 4. There is no evidence that experienced pineapple farmers respond at all to good news about alternative levels of fertilizer use. For inexperienced farmers, in contrast, a one standard-deviation increase in M (about 4) is associated with an increase in fertilizer use of approximately

4 cedis per plant. We raised the possibility in section 3.2 that farmers might be heterogeneous in their ability to learn from others, and in particular that lower ability farmers adopt pineapple more slowly (or not at all) and do not react to information from their neighbors. If this is the case, then our use of a sample of current pineapple farmers overstates the responsiveness of farmers in general to information from neighbors. The results in Column B provide some evidence on the importance of this kind of selection. If there is selection such that later adopters of pineapple are less responsive to news, then this selection is sufficiently weak that recent adopters are still very responsive to new information.

Columns C-F present the results of an investigation of the influence of the source of information on farmer i 's reactions. In alternate specifications, we use variants of M defined on partitions of farmer i 's information neighborhoods based on i 's information neighbors' experience, farm size, relative wealth and relative soil type. Our experience indicator is as described above. We define large farms as those with plantings of at least 60,000 total pineapple plants over our sample period (27 % of farmers have large farms).²¹ Finally, we define a classification of wealth as rich/poor with rich as those whose non-land wealth at the start of the survey is greater than the mean non-land wealth (30% of farmers are rich by this definition).

Column C defines M separately for inexperienced and experienced farmers in i 's information neighborhood. The coefficient of the index for experienced farmers' results is large and significant and for inexperienced farmers it is not. Column D presents a partition depending on whether i 's information neighbor is in i 's wealth category (both rich or both poor). The index is an important and significant predictor for same category neighbors but not for different category neighbors. For each of the pairs of M partition coefficients C-D, their difference is statistically significant with a p -value under 2%. Column E presents analogous estimates with a partition of M depending upon the size of the farms in i 's information neighborhood. Both coefficient estimates are large, positive, and statistically significant. Point estimates suggest that the responsiveness of input use to news from large farmers may be stronger than it is to similar news from small farmers, over and above the per-plant weighting of the M index itself. However, these estimates are not statistically different from each other. Finally, Column F presents estimates

²¹Median and mean numbers of plants planted by farmers in our sample are 22,000 and 41,000, respectively.

with M defined using a partition based on whether i 's neighbor has the same soil type as i (sandy or clay). These estimates provide no significant evidence in support of that hypothesis. In summary, inexperienced farmers appear to be the ones reacting to good news and tend to react to information revealed by neighbors who are experienced and who have similar wealth.²²

6 Robustness Checks and Extensions

6.1 Alternative Information Neighborhoods

In Table 6 we examine whether our finding that M predicts innovations in fertilizer is robust to changes in the definition of an information link. Full definitions of each of these alternatives are provided in Appendix 2. In column A, j is considered to be in i 's information neighborhood if j is named by i when asked a series of open-ended questions about who taught them to farm and from whom they have received farming advice (or *vice versa*). In column B, we use the broader definition of an information link if either i or j is listed anywhere in the other's roster of interactions with other sample members. In column C we define information neighborhoods based on the predicted probabilities for going to another farmer for advice (corresponding to the estimates in Table 2).²³

Regardless of the precise definition of the information neighborhood, the coefficient on M is statistically significant and large for inexperienced farmers (the standard deviation of M is approximately 3.5 for the first two information neighborhoods, and about 1 for the "Predicted Advice" neighborhood). In each case we find that when i is inexperienced, good news experiments in i 's information neighborhood tend to be followed by i changing his fertilizer use in the direction of those experiments, conditional on our growing conditions control Γ , village and round effects, and i 's wealth, clan and church membership. In contrast, there is evidence of responsiveness to information by experienced pineapple farmers only for one metric: predicted ask-for-advice.

²²Of course an important caveat to this summary is that our small sample size leads us to examine these partitions one variable at a time rather than jointly.

²³When using the predicted information neighborhoods as described in Section 3.1, M_{i,t_1} is a plant-weighted average of $\psi(\tau; t_0, t_1)d_{i,k,\tau} [x_{k,\tau} - x_{i,t_0}]$ across the entire village with weights for pair (i, k) equal to the predicted probability of being linked.

Our logit result that changes in inputs are predicted by bad news at previous input choice is also robust to the three alternate metrics in Table 6.²⁴ Indices of bad news at farmers’ previous input levels calculated with these alternate metrics remain statistically significant predictors of whether $\Delta x_{i,t} \neq 0$. Increases in the indices of bad news about previous levels increase the predicted probability of changing. Our result that bad news about alternate levels is associated with an statistically significant decrease in the probability of changing obtains for both roster of contacts measures, but not that based on predicted neighborhoods.

6.2 Robustness to Assumptions on Specification

In Table 7, we examine some of the assumptions we have made about the timing of learning, the categories of fertilizer used in constructing our proxy for subjective expectations, the size of the geographic neighborhood, and conditioning on soil characteristics. In each case we look at the impact of the specification change for our regression of Δx on M interacted with our experience indicator.

In columns A and B, we examine the assumptions embedded in ψ regarding the timing of information flows from neighbors’ pineapple plots. Recall that in our base specifications we assumed that information about the eventual size of the harvest of a period t planting began to be revealed after flowering in period $t + 5$, with the outcome fully known at harvest in $t + 9$. Under this base assumption, information about a time t planting begins to be available in time to influence fertilizer choice on a period $t + 1$ planting (because on the latter planting, fertilizer continues to be applied through period $t + 5$). In column A, we assume that no information is available from a period t planting until period $t + 6$, but maintain the idea that information is linearly revealed until harvest at $t + 9$. Thus, information from a time t planting can influence plantings starting at period $t + 2$. In column B, we assume that the flow of information is further restricted, so that no information about the eventual harvest of a period t planting is available until period $t + 7$; hence can influence plantings starting at period $t + 3$.²⁵ In both

²⁴Full results for our logistic regressions for these alternate metrics are available upon request, we omit a full discussion here to conserve space.

²⁵The definitions of previous planting respect these cutoffs as well, resulting in a reduction in sample size when $t - 2$ plantings are not allowed to be the initial period in

cases, the coefficient on M is positive and significantly different from zero for inexperienced farmers, and not for experienced farmers.

In column C, we modify our categorization of fertilizer use. Expectations about profitability of fertilizer use had been defined over the two coarse categories of $x = 0$ and $x > 0$. For 40% of the plantings in our sample, there is at least one planting in the farmer’s information neighborhood that provides information about the profitability of $x = 0$; for 32% of the plantings there is at least one planting in the farmer’s information neighborhood that provides information about $x > 0$. We now define expectations over three categories of input intensity: $x = 0, 0 < x \leq x_h, x_h < x$, where $x_h = 2.5$ (the 80th percentile of fertilizer use is about 2.5). As can be seen in column C, this change in specification has no qualitative effect on the results. This conclusion is robust for any x_h less than the 85th percentile of fertilizer use. For larger x_h , the precision of the estimates falls enough that the coefficient on M is not significant at conventional levels. It does not appear to be feasible to define expectations over more than three meaningful categories given the size of our dataset.

In column D, we alter the definition of the geographic neighborhood so that only plots within 500 meters fall within a plot’s geographic neighborhood. Again, we find that inexperienced farmers change their fertilizer use in the direction of inputs associated with good news experiments by their information neighbors, but that experienced farmers do not. Very similar results are obtained when geographic neighborhoods are defined as within 1500 meters.

In column E we include information on soil characteristics in the conditioning set. We lose some observations by doing so, because soil testing was not completed on all plots, but once again the core result is unchanged: the coefficient on M is positive, large and statistically significant for inexperienced farmers but not for experienced farmers.

Finally, in column F we examine the possibility that our results are an artifact of mean reversion in fertilizer use. Lagged own fertilizer use appears both in the dependent variable and in $M_{i,t}$ raising the possibility that mean reversion in fertilizer use, perhaps due to large measurement error, might lead us to find a spuriously significant coefficient on $M_{i,t}$. In the absence of any learning effects, the average $x_{k,t-1}$ across good news observations would be an estimate of the conditional mean of $x_{k,t}$ given $\pi_{k,t}$ was above its expectation.

determining changes in x

So $M_{i,t}$ could be interpreted as a noisy, biased estimate of whether $x_{i,t-1}$ is above or below its unconditional expectation which might be positively correlated with $\Delta x_{i,t}$ solely due to mean reversion. *A priori*, we think this is an unlikely source of spurious results as we include $\Gamma_{i,t}$ in (9) in addition to $M_{i,t}$. The sample size within geographic neighborhoods is considerably larger than that in information neighborhoods. Despite a higher spatial correlation within geographic neighborhoods, averages within this larger neighborhood will have a smaller variance than averages within information neighborhoods. Therefore, if mean reversion were driving correlations, $\Gamma_{i,t}$ should be a much less noisy measure of whether $x_{i,t-1}$ is above its long run mean. Once $\Gamma_{i,t}$ is conditioned upon, $M_{i,t}$ should offer little or no predictive power for $\Delta x_{i,t}$ resulting from mean reversion.²⁶ However, to be confident our results are not an artifact of mean reversion, in column F we add the lagged fertilizer level $x_{i,t-1}$ to the regression. The coefficients on $M_{i,t}$ change by a magnitude comparable to some of our other alternative specifications and for inexperienced farmers it remains a significant predictor. The coefficient on $\Gamma_{i,t}$ changes the most dramatically; this is unsurprising since the lagged input levels are spatially correlated and so provide an alternate control for spatially correlated growing conditions.

Our logit results for responses to bad news at previous input choice are robust to analogs of the robustness checks in Table 7.²⁷ Indices of bad news at farmers' previous input levels remain statistically significant predictors of

²⁶A special case of mean-reverting $x_{i,t}$ would result if our data on inputs were dominated by large amounts of classical measurement error. However we think this case is unlikely to have occurred as the field research was specifically designed to collect accurate data on farming inputs (including the number of plants planted) and output by sacrificing sample size in exchange for frequent and thorough visits to respondents.

We also examined the special case of measurement error by performing Monte Carlo experiments (available upon request) adding artificial measurement error to our fertilizer data. The mean of per-plant fertilizer use is 4, and its standard deviation is 7; to our data we added a mean zero normal draws with standard deviation 1 to 7 (truncated so that measured fertilizer use is never negative). The estimated coefficient on M becomes insignificant at conventional levels when the standard deviation of the added noise is 4 while Γ remains a significant predictor.

²⁷Columns A through F have analogs for our logistic regressions and results (omitted here to save space) are available upon request. The robustness check in Column G has no analog as the concern of potentially spurious results due the construction of M does not apply to our logits.

whether $\Delta x_{i,t} \neq 0$; more bad news about previously used levels increases predicted changes.

6.3 Learning about Optimal Labor Use in Pineapple and in Established Crops

We have focused on learning about the productivity of fertilizer because this is the most dramatically new aspect of pineapple production in the study area. However, it is apparent that a single-input production function is a considerable simplification. Labor is also an important component of the cost of pineapple production and there is at least some uncertainty about its optimal level. In focus group interviews, pineapple farmers reported uncertainty regarding the optimal amount of labor-intensive weeding; they were not sure about how important it is to keep weeds clear of the plot. In this section, we use the methods developed above to examine the hypothesis that pineapple farmers learn from their neighbors about the productivity of labor in pineapple production.

For comparison and as an informal specification test we also estimate labor input regressions with data on an established crop for which there should be no learning.²⁸ A standard maize-cassava intercrop pattern has been the foundation of the rural economy in these villages since the local decline of cocoa cultivation in the 1930s. The characteristics of the maize-cassava production function are well-known to farmers in these villages and there should be no uncertainty about the role of labor in the production of maize-cassava. We estimate a regressions of changes in labor inputs for pineapple plots and for maize-cassava plots with a specification analogous to (9):

$$\Delta \tilde{x}_{i,t} = \delta_1 \tilde{M}_{i,t} + \delta_2 \tilde{\Gamma}_{i,t} + z'_{i,t} \delta_3 + u_{i,t}. \quad (10)$$

Where \tilde{x} is the labor input per plant for pineapples and per hectare for maize-cassava. $\tilde{M}_{i,t}$ and $\tilde{\Gamma}_{i,t}$ are constructed exactly as above for this labor input.²⁹ We expect a positive δ_1 for pineapple plots if pineapple farmers are learning from their neighbors about the productivity of labor. A nonzero δ_1

²⁸We cannot undertake a parallel analysis for chemical inputs, because these are not used for the cultivation of maize-cassava.

²⁹Labor inputs include both the value of hired labor and that provided by the farmer's household. The labor input range was divided into two categories (above and below median) for determining whether profits were unusually high given inputs.

for maize-cassava cannot be attributed to learning, because this technology is well-established.

We estimate (10) for the same sample of pineapple plots examined above. Labor inputs are measured over the crucial period early in the life cycle of the pineapple during which fertilizer inputs also occur (for a plot planted at time t , this is periods $t + 1$ through $t + 4$). All farmers change their labor inputs across plantings, so there is no need to estimate an analog of the first stage logit (8). \tilde{M} is defined using the benchmark (asked for advice) information neighborhood. Column A of Table 8 presents the results of estimating (10). We condition on the average deviation of i 's lagged labor use from the lagged labor used by his geographic neighbors, $\tilde{\Gamma}$, and its analog for his financial neighbors.

Good news experiments in i 's information neighborhood are followed by i changing his labor use in the direction of those experiments' labor, conditional on geographic and financial neighbors' lagged labor use, plot characteristics, village and round effects, and i 's wealth, clan and church membership. The coefficient is also large: a one standard deviation increase in \tilde{M} (which is 396) is associated with an increase in labor use of approximately 376 cedis per sucker, which is one-fifth of the median labor use per sucker on pineapple plots (which is 1,845). Pineapple farmers appear to be learning about the productivity of labor in the cultivation of pineapple from the experiences of their information neighbors. This conclusion should be treated with caution, because labor and fertilizer inputs are not likely to be separable in pineapple production (most obviously, labor is used to apply fertilizer). The adjustment of labor inputs to news about profitability, therefore, may not be a result independent of learning about the productivity of fertilizer. In addition, the results shown in column A are not as robust to changes in specification as are the results on fertilizer innovations in Tables 4 – 7.

In contrast, in column B, we see that there is no evidence that maize/cassava farmers adjust labor inputs to information from the cultivation of their information neighbors.³⁰ The coefficient of the learning index \tilde{M} is virtually zero (at the point estimate, a one standard deviation increase in \tilde{M} is associated

³⁰There are two differences in specification between the pineapple and maize-cassava regressions: first, in contrast to pineapple, the maize-cassava intercrop system is seasonal. Hence we compare inputs across successive seasons and replace the round indicators in A with season indicators in B. Second, the maize-cassava mixture is grown in all four of our survey villages, while pineapple is grown in only three villages, hence there is an additional village indicator for maize-cassava.

with an increase in labor use of 11 thousand cedis per hectare, while the mean change in labor use is 153 thousand cedis, and its standard deviation is 700 thousand). \tilde{M} has no significant predictive power for innovations in labor use in maize-cassava production, just as we expect given the familiarity of this farming system in the study region.

For both pineapple and maize-cassava, we find that there is a strong geographic correlation in innovations in labor use. There seem to be important spatially- and serially-correlated shocks to the productivity of inputs. This underscores the value of direct data on communication for defining information neighborhoods. In the more typical case in which we had data only on geographic proximity, it would be tempting to rely on this to proxy for information links. The consequences of this are presented in column C. We construct a new variable, analogous to \tilde{M} but based instead only on geographic proximity. We see in column C that maize-cassava farmers adjust labor inputs in the direction of successful ‘experiments’ in their geographic neighborhood. The coefficient is large (a one standard deviation increase in the index of experiments in the geographic neighborhood is associated with an increased labor input of 590,000 cedis/ha., compared to mean labor input of 650,000 cedis/ha.) and on the conventional boundary of statistical significance. This result has nothing to do with learning; it is induced entirely by the strong correlations in growing conditions. However, without our direct data on communication we might incorrectly infer the existence of social learning about labor productivity in this well-established farming system.

7 Conclusion

This paper presents evidence that social learning is important in the diffusion of knowledge regarding pineapple cultivation in Ghana. We take advantage of data that combines agronomic and conventional economic information with details regarding relationships between farmers to address the challenge of identifying learning effects in an economy undergoing rapid technological change. We trace the effect of a farmer’s successful experiment with fertilizer on the innovations in fertilizer use by other cultivators with whom the experimenter shares information. Our findings suggest that farmers are more likely to change input levels upon the receipt of bad news about the profitability of their previous level of fertilizer use, and less likely to change when they observe bad news about the profitability of alternative levels of fertil-

izer. Looking at the magnitude of innovations in fertilizer use, we find that a farmer increases (decreases) his use after someone with whom he shares information achieves higher than expected profits when using more (less) fertilizer than he did. These findings hold when controlling for common input usage due to similar unobservable growing conditions via conditioning on actions of geographic neighbors, for credit arrangements using a notion of financial neighborhoods, and across several information metrics. Support for the interpretation of our results as indicating learning effects is provided by the fact that it is inexperienced farmers who are most responsive to news in their information neighborhoods. Further support is provided by our finding no evidence of learning when our methodology is applied to a known maize-cassava technology.

There is evidence that social learning plays a role in the cultivation decisions of these farmers. Information, therefore, has value in these villages, as do the network connections through which that information flows. This raises the possibility that farmers consider the consequences for the availability of information when forming the connections that underlie their information neighborhoods. If so, measurement of the extent of social learning is not sufficient for adequate evaluation of policy regarding the diffusion of technology. It is necessary, in addition, to understand the endogenous process of information network formation. For example, consider the impact of a subsidy offered to one farmer in a village that induces him to use an optimal large amount of fertilizer and (with high probability) get high profits. The speed with which this information spreads, and hence the value of the subsidy, depends upon the choices of the subsidized farmer and others in the village to make and maintain information linkages. These choices may depend upon the value of the information to each farmer and upon the costs of information links, which may depend upon a rich array of characteristics of the farmers and the social structure of the village. In some contexts, differing religions may be an effective barrier to communication. In others, gender, wealth or family ties may be the most salient determinants of the shape of the information network.

One of the drawbacks of having data early in the learning process is that the value of learning cannot be estimated well. Inexperienced farmers often make mistakes by switching to what is truly a suboptimal input level after seeing it perform surprisingly well in a small number of experiments. Farmers' responses to both signal and noise provide a superb context for detecting evidence of learning, but is a poor one for determining the long-

run profits associated with optimal input use. In the short run those that change inputs include large segments moving both towards and away from the optimal input levels.³¹ Therefore we should not expect to, and in fact do not find increases in realized profits for those who change their input levels.

The next step in this research program is to model the choices of farmers regarding the formation of information links in these villages. A large literature examines network efficiency, with the goal of characterizing the network configuration that maximizes a value function (Bolton and Dewatripont (1994); Hendricks, Piccione *et al.* (1995); Economides (1996)). This is appropriate for a planner (such as a telecommunications monopoly) but not for the decentralized process that governs the formation of a social network like those in the sample villages. Accordingly, future work should focus on the incentives of the individuals who build the links that define the network ((Coleman (1966); Granovetter (1973); Granovetter (1992); Fafchamps (1999); Fafchamps and Minten (1999)). Bala and Goyal (1999), Jackson and Wolinsky (1996) and Montgomery (1996) provide alternative theoretical schema that can underpin an empirical analysis of link formation, but we know of no empirical work in economics that examines the formation of decentralized networks.

³¹Recall that almost all farmers in our sample are using far less fertilizer than recommended by the Ministry of Agriculture. However, approximately the same number of farmers are reducing their level of fertilizer input in response to good news about low levels of fertilizer use as are increasing fertilizer use in response to good news about high levels of use.

References

- [1] Abutiare, W. S. (1991) *Guide to the Commercial Production of Pineapple, Papaya and Mango in Ghana*.
- [2] Amanor, K. (1994) *The New Frontier*. London, Zed Books.
- [3] Aghion, P. and P. Howitt (1998) *Endogenous Growth Theory*. Cambridge, MIT Press.
- [4] Bala, V. and S. Goyal (1999) "A Non-Cooperative Theory of Network Formation." *Econometrica* forthcoming.
- [5] Bandiera, O. and Rasul I. "Social Networks and Technology Adoption in Northern Mozambique" Working paper LSE.
- [6] Bartlett, M.S. (1950) "Periodogram Analysis and Continuous Spectra." *Biometrika*, 37, 1-16.
- [7] Bartholomew, D. and S. Kadzimin (1977). "Pineapple." in *Ecophysiology of Tropical Crops* by P. d. T. Alvim and T. T. Kozlowski. New York, Academic Press: 113-156.
- [8] Berry, D.A. and B. Fristedt (1985) *Bandit Problems*. New York, Chapman & Hall.
- [9] Besley, T. and A. Case (1994). "Diffusion as a Learning Process: Evidence from HYV Cotton." Working Paper.
- [10] Bolton, P. and M. Dewatripont (1994). "The Firm as a Communication Network." *Quarterly Journal of Economics* 109(4): 809-840.
- [11] Bindlish, V. and R. E. Evenson (1997) "The Impact of T&V Extension in Africa: The Experience of Kenya and Burkina Faso." *World Bank Research Observer* 12(2): 183-201.
- [12] Birkhaeuser, D., R. E. Evenson, et al. (1991) "The Economic Impact of Agricultural Extension: A Review." *Economic Development and Cultural Change* 39(3): 607-650.
- [13] Brock, W. and S. Durlauf (1999) "Interactions-Based Models," Forthcoming in *Handbook of Econometrics*, J. Heckman and E. Leamer eds.

- [14] Carter, M (1997) “Environment, Technology and the Social Articulation of Risk in West Africa.” *Economic Development and Cultural Change*. 45(3): 557-590.
- [15] Chamley, C (2003) *Rational Herds: Economic Models of Social Learning*. Manuscript: DELTA, Ecole normale supérieure.
- [16] Conley, T. (1999) “GMM Estimation with Cross Sectional Dependence” *Journal of Econometrics* Vol 92(1) p.1-45.
- [17] Coleman, J. E. Katz and H. Menzel (1957) “The Diffusion of an Innovation Among Physicians,” *Sociometry* 20: 253-270.
- [18] Coleman, J. (1966). *Medical Innovation: A Diffusion Study*. New York, Bobbs-Merrill.
- [19] Economides, N. (1996). “The Economics of Networks.” *International Journal of Industrial Organization* 14(6): 673-99.
- [20] Evenson, R. and L. Westphal (1995). Technological Change and Technology Strategy. Handbook of Development Economics. J. Behrman and T. N. Srinivasan. Amsterdam, North-Holland. 3A: 2209-2300.
- [21] Fafchamps, M. (1999). Ethnicity and Credit in African Manufacturing. Manuscript: Oxford.
- [22] Fafchamps, M. and B. Minten (1999). “Relationships and traders in Madagascar.” *Journal of Development Studies* 35(6): 1-35.
- [23] Feder, G., R. Just, et al. (1985). “Adoption of Agricultural Innovations in Developing Countries: A Survey.” *Economic Development and Cultural Change* 33(2): 255-298.
- [24] Foster, A. and M. Rosenzweig (1995) “Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture” *Journal of Political Economy* Vol. 103(6) p 1176-1209.
- [25] Gittins, J.C. (1989) *Multi-armed Bandit Allocation Indices*. New York, John Wiley & Sons.
- [26] Goldstein, M. and C. Udry (1999). Gender and Land Resource Management in Southern Ghana. Manuscript: Yale University.

- [27] Granovetter, M. (1973). "Strength of Weak Ties." *American Journal of Sociology* 78(6): 1360-1380.
- [28] Granovetter, M. (1992). "Economic Institutions As Social Constructions - a Framework For Analysis." *Acta Sociologica* 35(1): 3-11.
- [29] Griliches, Z. (1957). "Hybrid Corn: An Exploration in the Economics of Technological Change." *Econometrica* 25(4): 501-522.
- [30] Hendricks, K., M. Piccione, et al. (1995) "The Economics of Hubs: The Monopoly Case." *Review of Economic Studies* 62(1): 83-101.
- [31] Jackson, M. O. and A. Wolinsky (1996) "A Strategic Model of Social and Economic Networks." *Journal of Economic Theory* 71(1): 44-74.
- [32] Jovanovic, B. and Y. Nyarko (1996). "Learning by Doing and the Choice of Technology." *Econometrica* 64(6): 1299-1310.
- [33] Kapteyn, A. (2000) "Saving and Reference Groups" Tilburg University CentER Working Paper.
- [34] Lucas, R. E. Jr. (1988) "On the Mechanics of Economic Development." *Journal of Monetary Economics*. 22:3-42.
- [35] Manski, C. F. (1993) "Identification of Social Endogenous Social Effects: The Reflection Problem" *Review of Economic Studies* 60: 531-542.
- [36] Manski, C. F. (1997) *Identification Problems in the Social Sciences* Harvard University Press, Cambridge.
- [37] Moffitt, R. (1999) "Policy Interventions, Low-Level Equilibria, and Social Interactions" Johns Hopkins Working Paper.
- [38] Montgomery, James (1996) "The Structure of Social Exchange Networks: A Game-Theoretic Reformulation of Blau's Model." *Sociological Methodology* 26:193-225.
- [39] Morris, S. (1995) "The Common Prior Assumption in Economic Theory." *Economics and Philosophy* 11:227-253.
- [40] Munshi, K. and J. Myaux (1998) "Social Effects in the Demographic Transition: Evidence from Matlab, Bangladesh" Working Paper University of Pennsylvania.

- [41] Newey, W.K. and West, K.D. (1987). "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55:703-708.
- [42] Obeng, I. S. (1994). Effects of Domestic Policies on Production and Export of Non-Traditional Agricultural Commodities: A Case Study of Fresh Pineapples in Ghana, M.Phil Thesis, Department of Agricultural Economy and Farm Management, University of Ghana, Legon.
- [43] Prescott, E. (1972) "The Multi-Period Control Problem Under Uncertainty." *Econometrica* 40(6): 1043-1058.
- [44] Purseglove, J.W. (1972) *Tropical Crops: Monocotyledons I*. London: Longman.
- [45] Rauch, James E. and Alessandra Casella, eds. (2001) *Networks and Markets*. New York: Russell Sage Foundation.
- [46] Richards, P. (1985) *Indigenous Agricultural Revolution: Ecology and Food Production in West Africa*. London: Hutchinson.
- [47] Rogers, E. (1995) *Diffusion of Innovations*, 4th Edition. New York: Simon & Schuster.
- [48] Romani, Mattia. (2003) "Love Thy Neighbor? Evidence from Ethnic Discrimination in Information Sharing within Villages in Côte d'Ivoire" *Journal of African Economies* 12(4): 533-563.
- [49] Romer, Paul (1986) "Increasing Returns and Long-Run Growth" *Journal of Political Economy*.94(5): 1002-1037.
- [50] Soler, Alain (1992) *Pineapple: Quality Criteria*. CIRAD-COLEACP-OCAB, Paris.
- [51] Van den Bulte, Christophe and Gary Lilien. (2001) "Medical Innovation Revisited: Social Contagion versus Marketing Effort." *American Journal of Sociology*, 106(5), 1409-1435.
- [52] Wilson, R. (1975) "Informational Economies of Scale." *Bell Journal of Economics*: 184-195.

- [53] Woittiez, I. and Kapteyn A. (1998) “Social interactions and habit formation in a model of female labour supply” *J. Public Economics* (70)2 p. 185-205.
- [54] Yamauchi, Futoshi (2002) “Social Learning and Investment in Human Capital: Evidence from Schooling Investments in India” Working Paper Kyoto University.

8 Appendix 1: Social Learning in a Dynamic Context

In order to make progress in the dynamic setting, it is necessary for us to be more specific about a model of learning behavior, so we assume that farmers are Bayesian learners. We retain the focus on local learning by adopting a generalization of the multi-armed bandit setting. We consider a three period model where farmers choose input levels in period one, generating profits in period two. They then repeat this choice of input levels in period two, producing profits in period three. In period two, farmers choose the expected profit maximizing input level; in this period the model is a special case of the local learning discussed thus far. In period one, however, in addition to considering profitability in period two, the farmer considers the option value of the experiment with respect to final period profits. This option value may lead the farmer to choose a period one action with less than the highest expected payoff in period two.

In order to assess farmer’s reactions to observing better than expected growing-conditions-adjusted output for a particular input level (say \tilde{x}) in this dynamic model we consider the following thought experiment. Suppose that our farmer’s initial beliefs about the productivity of input level \tilde{x} were in part a result of having observed a neighbor using that input level and the corresponding output realization. How would our farmer’s beliefs be different if he had observed a slightly higher output realization from his neighbor and what consequence will this have for his period one choice?

We show below that the probability that the farmer chooses \tilde{x} in period 1 is strictly increasing in the output realization of his neighbor who used input level \tilde{x} in the previous period. This prediction replicates that of the myopic model, demonstrating that the myopic model’s predictions can survive the introduction of forward-looking behavior. If his neighbor achieves particularly

high growing-conditions-adjusted output at a particular level of fertilizer use in period t , then a farmer's beliefs about the productivity of that level of fertilizer shift upward. This increases the relative value to the farmer of choosing that level of fertilizer use, and thus increases the probability that that level of fertilizer use will be chosen by the farmer in the following period. Obviously, it remains the case that in the absence of any learning, an observer who cannot fully observe serially and spatially correlated growing conditions could observe a similar pattern: a neighbor's realization of high profits at fertilizer level \tilde{x} will be followed in the succeeding period by an increased probability of a farmer choosing fertilizer level \tilde{x} .

Details

To simplify notation, we suppress the farmer i subscript and restrict the support of x_t to consist of only two points, $x \in \{H, L\}$.³² We assume that farmers know that ε_{t+1} has a standard normal distribution and that input prices are constant at p . The farmer does not know $f(H)$ and $f(L)$, but at time t he has independent priors that are distributed $N(g_t(H), \sigma_{tH}^2)$ and $N(g_t(L), \sigma_{tL}^2)$ respectively. Consider the farmer's final choice of x_2 . Expected profits take one of two values, so his value function in the second period is

$$V_2 = \max \{E_2(w_2 f(H) - pH + \varepsilon_3), E_2(w_2 f(L) - pL + \varepsilon_3)\}.$$

Where the expectation operator E_t refers to the subjective expectation with respect to the farmer's belief at time t . So the farmer chooses $x_2 = H$ if $w_2 g_2(H) - pH > w_2 g_2(L) - pL$.³³

Now consider the choice of fertilizer in the previous period. In this period, the farmer is concerned not only with immediate expected profits, but also with the value of any information generated by experimentation in the period. Defining $\delta \in (0, 1)$ as the discount rate, the value of choosing $x_1 = H$ is

$$V_{1H} = w_1 g_1(H) - pH + \delta E_1 \max \left[w_2 \left(\frac{y_{2H}}{w_1} \frac{\sigma_{1H}^2}{\sigma_{1H}^2 + w_1^{-2}} + g_1(H) \frac{w_1^{-2}}{\sigma_{1H}^2 + w_1^{-2}} \right) - pH, w_2 g_1(L) - pL \right]$$

³²Chamley (2003) provides a good textbook presentation of Bayesian tools for modeling social learning.

³³Recall that the realization of productivity w_2 occurs before the choice of x_2 .

where we have substituted the farmer's updating rule using the observation on $\frac{y_{2H}}{w_1}$ he will obtain if H is chosen in period one. The input choice in period 2 will be determined by the realization of y_{2H} : there is critical value \hat{y}_{2H} (which is a function of w_2) such that $x_2 = H$ if $y_{2H} \geq \hat{y}_{2H}$ and $x_2 = L$ otherwise. V_{1L} and \hat{y}_{2L} are defined analogously.

This setup is closely related to conventional multi-armed bandit models (Berry and Fristedt 1985), but differs in that the random shock makes input levels (arms) more or less attractive. As a consequence, the standard index theorems (e.g., Gittins 1989, Theorem 2.3) fail to hold. The common growing conditions shock can induce the farmer to learn about input levels that might not be experimented with if growing conditions were constant. For example, a sufficiently bad growing conditions shock will induce the agent to use the lowest available fertilizer level.

Expected lifetime profits if H is chosen in period 1 are

$$\begin{aligned} V_{1H} &= w_1 g_1(H) - pH + \\ &\quad \delta E_1 \max E_2 (w_2 f(H) - pH + \varepsilon_3), E_2 (w_2 f(L) - pL + \varepsilon_3) \\ &= w_1 g_1(H) - pH + \\ &\quad \delta E_1 \max \left[w_2 \left(\frac{y_{2H}}{w_1} \frac{\sigma_{1H}^2}{\sigma_{1H}^2 + w_1^{-2}} + g_1(H) \frac{w_1^{-2}}{\sigma_{1H}^2 + w_1^{-2}} \right) - pH, \right. \\ &\quad \left. w_2 g_1(L) - pL \right]. \end{aligned}$$

We define the critical realization of y_{2H}

$$\hat{y}_{2H} = w_1 \frac{\sigma_{1H}^2 + w_1^{-2}}{\sigma_{1H}^2} \left(g_1(L) + \frac{p}{w_2} (H - L) - g_1(H) \frac{w_1^{-2}}{\sigma_{1H}^2 + w_1^{-2}} \right)$$

such that $x_2 = H$ if $y_{2H} \geq \hat{y}_{2H}$ and $x_2 = L$ otherwise. We define $\phi_{1H}(\cdot)$ as the normal density with mean $w_1 g_1(H)$ and variance $w_1^2 \sigma_{1H}^2 + 1$ and $\mu(\cdot)$ as the distribution of period 2 growing conditions, conditional on period 1 growing conditions (we drop the notation making the conditioning on w_1 explicit to save space). Thus the value of choosing $x_1 = H$ can be more

explicitly defined as

$$\begin{aligned}
V_{1H} &= w_1 g_1(H) - pH + \\
&\quad \delta \int \int \max \left[w_2 \left(\frac{y_{2H}}{w_1} \frac{\sigma_{1H}^2}{\sigma_{1H}^2 + w_1^{-2}} + g_1(H) \frac{w_1^{-2}}{\sigma_{1H}^2 + w_1^{-2}} \right) - pH, \right. \\
&\quad \left. w_2 g_1(L) - pL \right] \times \phi_{1H}(y_{2H}) dy_{2H} d\mu(w_2), \\
&= w_1 g_1(H) - pH + \\
&\quad \delta \int \int_{\hat{y}_{2H}}^{\infty} \left(w_2 \left(\frac{y_{2H}}{w_1} \frac{\sigma_{1H}^2}{\sigma_{1H}^2 + w_1^{-2}} + g_1(H) \frac{w_1^{-2}}{\sigma_{1H}^2 + w_1^{-2}} \right) - pH \right) \phi_{1H}(y_{2H}) dy_{2H} d\mu(w_2) \\
&\quad + \delta \int \int_{-\infty}^{\hat{y}_{2H}} (w_2 g_1(L) - pL) \phi_{1H}(y_{2H}) dy_{2H} d\mu(w_2).
\end{aligned}$$

Equivalently:

$$\begin{aligned}
V_{1H} &= w_1 g_1(H) - pH \\
&+ \delta \frac{1}{w_1} \frac{\sigma_{1H}^2}{\sigma_{1H}^2 + w_1^{-2}} \int w_2 \Pr(y_{2H} > \hat{y}_{2H}) E(y_{2H} | y_{2H} > \hat{y}_{2H}) d\mu(w_2) \\
&+ \delta g_1(H) \frac{w_1^{-2}}{\sigma_{1H}^2 + w_1^{-2}} \int w_2 \Pr(y_{2H} > \hat{y}_{2H}) d\mu(w_2) \\
&+ \delta (-pH) \int \Pr(y_{2H} > \hat{y}_{2H}) d\mu(w_2) \\
&+ \delta g_1(L) \int w_2 \Pr(y_{2H} < \hat{y}_{2H}) d\mu(w_2) \\
&+ \delta (-pL) \int \Pr(y_{2H} < \hat{y}_{2H}) d\mu(w_2)
\end{aligned}$$

where probabilities and expectations refer to the agent's subjective distribution. The interesting quantity in the first term is $\Pr(y_{2H} > \hat{y}_{2H}) E(y_{2H} | y_{2H} > \hat{y}_{2H})$. Standardizing, we have

$$\begin{aligned}
\Pr(y_{2H} > \hat{y}_{2H}) E(y_{2H} | y_{2H} > \hat{y}_{2H}) &= \\
&w_1 g_1(H) \Phi\left(-\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right) + v_H \phi\left(-\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right)
\end{aligned}$$

where $v_H = (w_1^2 \sigma_{1H}^2 + 1)^{1/2}$. Therefore,

$$\begin{aligned}
\frac{\partial V_{1H}}{\partial g_1(H)} &= w_1 \\
&+ \delta \frac{1}{w_1} \frac{\sigma_{1H}^2}{\sigma_{1H}^2 + w_1^{-2}} \int w_2 \left\{ w_1 \Phi\left(-\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right) \right. \\
&+ w_1 g_1(H) \phi\left(-\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right) \left[-\frac{1}{v_H} \frac{\partial \hat{y}_{2H}}{\partial g_1(H)} + \frac{w_1}{v_H} \right] \\
&+ \phi\left(-\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right) (\hat{y}_{2H} - w_1 g_1(H)) \left[-\frac{1}{v_H} \frac{\partial \hat{y}_{2H}}{\partial g_1(H)} + \frac{w_1}{v_H} \right] \left. \right\} d\mu(w_2) \\
&+ \delta \int w_2 \frac{w_1^{-2}}{\sigma_{1H}^2 + w_1^{-2}} \Phi\left(-\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right) d\mu(w_2) \\
&+ \delta \int w_2 \frac{w_1^{-2}}{\sigma_{1H}^2 + w_1^{-2}} g_1(H) \phi\left(-\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right) \left[-\frac{1}{v_H} \frac{\partial \hat{y}_{2H}}{\partial g_1(H)} + \frac{w_1}{v_H} \right] d\mu(w_2) \\
&+ \delta (-pH) \int \phi\left(-\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right) \left[-\frac{1}{v_H} \frac{\partial \hat{y}_{2H}}{\partial g_1(H)} + \frac{w_1}{v_H} \right] d\mu(w_2) \\
&+ \delta g_1(L) \int w_2 \phi\left(\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right) \left[\frac{1}{v_H} \frac{\partial \hat{y}_{2H}}{\partial g_1(H)} - \frac{w_1}{v_H} \right] d\mu(w_2) \\
&+ \delta (-pL) \int \phi\left(\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right) \left[\frac{1}{v_H} \frac{\partial \hat{y}_{2H}}{\partial g_1(H)} - \frac{w_1}{v_H} \right] d\mu(w_2)
\end{aligned}$$

Recalling that \hat{y}_{2H} is defined so that

$$w_2 \frac{\hat{y}_{2H}}{w_1} \frac{\sigma_{1H}^2}{\sigma_{1H}^2 + w_1^{-2}} + w_2 g_1(H) \frac{w_1^{-2}}{\sigma_{1H}^2 + w_1^{-2}} - pH = w_2 g_1(L) - pL$$

and combining the remaining terms we have

$$\begin{aligned}
\frac{\partial V_{1H}}{\partial g_1(H)} &= w_1 + \delta \int w_2 \Phi\left(-\frac{\hat{y}_{2H} - w_1 g_1(H)}{v_H}\right) d\mu(w_2) \\
&= w_1 + \delta \int w_2 \Pr(y_{2H} > \hat{y}_{2H}) d\mu(w_2).
\end{aligned}$$

Now consider the value of choosing L in period 1:

$$\begin{aligned}
V_{1L} &= w_1 g_1(L) - pL + \delta E_1 V_2 \\
&= w_1 g_1(L) - pL + \\
&\quad \delta E_1 \max E_2 (w_2 f(H) - pH + \varepsilon_3), E_2 (w_2 f(L) - pL + \varepsilon_3) \\
&= w_1 g_1(L) - pL + \\
&\quad \delta E_1 \max [w_2 g_1(H) - pH, \\
&\quad w_2 \left(\frac{\sigma_{1L}^2}{\sigma_{1L}^2 + w_1^{-2}} \frac{y_{2L}}{w_1} + \frac{w_1^{-2}}{\sigma_{1L}^2 + w_1^{-2}} g_1(L) \right) - pL].
\end{aligned}$$

Similarly to above, we define a critical realization of $y_{2L} : \hat{y}_{2L}$ such that $x_2 = L$ if $y_{2L} \geq \hat{y}_{2L}$ and $x_2 = H$ otherwise so

$$\begin{aligned}
V_{1L} &= w_1 g_1(L) - pL + \\
&\quad \delta E \left[w_2 \left(\frac{\sigma_{1L}^2}{\sigma_{1L}^2 + w_1^{-2}} \frac{y_{2L}}{w_1} + \frac{w_1^{-2}}{\sigma_{1L}^2 + w_1^{-2}} g_1(L) \right) - pL \mid y_{2L} \geq \hat{y}_{2L} \right] \Pr(y_{2L} \geq \hat{y}_{2L}) + \\
&\quad + \delta E [(w_2 g_1(H) - pH) \mid y_{2L} < \hat{y}_{2L}] \Pr(y_{2L} < \hat{y}_{2L}) \\
&= w_1 g_1(L) - pL + \\
&\quad \delta \frac{1}{w_1} \frac{\sigma_{1L}^2}{\sigma_{1L}^2 + w_1^{-2}} \int w_2 \left[w_1 g_1(L) \Phi\left(-\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) + v_L \phi\left(\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) \right] d\mu(w_2) \\
&\quad + \delta g_1(L) \frac{w_1^{-2}}{\sigma_{1L}^2 + w_1^{-2}} \int w_2 \Phi\left(-\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) d\mu(w_2) \\
&\quad + \delta (-pL) \int \Phi\left(-\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) d\mu(w_2) \\
&\quad + \delta g_1(H) \int w_2 \Phi\left(\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) d\mu(w_2) \\
&\quad + \delta (-pH) \int \Phi\left(\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) d\mu(w_2).
\end{aligned}$$

Differentiating,

$$\begin{aligned}
\frac{\partial V_{1L}}{\partial g_1(H)} &= \delta \frac{1}{w_1} \frac{\sigma_{1L}^2}{\sigma_{1L}^2 + w_1^{-2}} \int w_2 \{ \\
&\quad + w_1 g_1(L) \phi\left(-\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) \left[-\frac{1}{v_L} \frac{\partial \hat{y}_{2L}}{\partial g_1(H)} \right] \\
&\quad + \phi\left(-\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) (\hat{y}_{2L} - w_1 g_1(L)) \left[-\frac{1}{v_L} \frac{\partial \hat{y}_{2L}}{\partial g_1(H)} \right] \} d\mu(w_2) \\
&\quad + \delta \frac{w_1^{-2}}{\sigma_{1L}^2 + w_1^{-2}} g_1(L) \int w_2 \phi\left(-\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) \left[-\frac{1}{v_L} \frac{\partial \hat{y}_{2L}}{\partial g_1(H)} \right] d\mu(w_2) \\
&\quad + \delta (-pL) \int \phi\left(-\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) \left[-\frac{1}{v_L} \frac{\partial \hat{y}_{2L}}{\partial g_1(H)} \right] d\mu(w_2) \\
&\quad + \delta g_1(H) \int w_2 \phi\left(\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) \left[\frac{1}{v_L} \frac{\partial \hat{y}_{2L}}{\partial g_1(H)} \right] d\mu(w_2) \\
&\quad + \delta \int w_2 \Phi\left(\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) d\mu(w_2) \\
&\quad + \delta (-pH) \int \phi\left(\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) \left[\frac{1}{v_L} \frac{\partial \hat{y}_{2L}}{\partial g_1(H)} \right] d\mu(w_2) \\
&= \delta \int w_2 \Phi\left(\frac{\hat{y}_{2L} - w_1 g_1(L)}{v_L}\right) d\mu(w_2) \\
&= \delta \int w_2 \Pr(y_{2L} < \hat{y}_{2L}) d\mu(w_2)
\end{aligned}$$

The first term of $\frac{\partial V_{1H}}{\partial g_1(H)}$ simply reflects expected profits in period one being linear in expected shock-adjusted output. The second term corresponds to the marginal change in profits, w_2 , times the subjective probability that H will be chosen in period two, averaged over prospective shock values w_2 . A period one term is of course absent from $\frac{\partial V_{1L}}{\partial g_1(H)}$ as L rather than H has been chosen and the derivative is again the marginal change in expected output times the probability that H will be chosen in period two, averaged over the values of that the growing conditions shock may realize in period 2.

Now consider the relative sizes of $\Pr(y_{2H} > \hat{y}_{2H})$ and $\Pr(y_{2L} < \hat{y}_{2L})$ for

any given value of w_2 :

$$\begin{aligned}\Pr(y_{2H} > \hat{y}_{2H}) &= \Phi \left(\frac{(w_1^2 \sigma_{1H}^2 + 1)^{1/2}}{\sigma_{1H}^2} \frac{1}{w_1} \left[\frac{p}{w_2} (L - H) + g_1(H) - g_1(L) \right] \right) \\ \Pr(y_{2L} < \hat{y}_{2L}) &= \Phi \left(\frac{(w_1^2 \sigma_{1L}^2 + 1)^{1/2}}{\sigma_{1L}^2} \frac{1}{w_1} \left[\frac{p}{w_2} (L - H) + g_1(H) - g_1(L) \right] \right)\end{aligned}$$

There are two main cases, because $\frac{(w_1^2 \sigma_{1H}^2 + 1)^{1/2}}{\sigma_{1H}^2} \frac{1}{w_1} > 0$.

Case 1 : $\frac{p}{w_2} (L - H) + g_1(H) - g_1(L) < 0$

- If $\sigma_{1H}^2 > \sigma_{1L}^2$, then

$$\frac{1}{2} > \Pr(y_{2H} > \hat{y}_{2H}) > \Pr(y_{2L} < \hat{y}_{2L}).$$

- If $\sigma_{1H}^2 \leq \sigma_{1L}^2$, then

$$\frac{1}{2} > \Pr(y_{2L} < \hat{y}_{2L}) \geq \Pr(y_{2H} > \hat{y}_{2H}).$$

Case 2 : $\frac{p}{w_2} (L - H) + g_1(H) - g_1(L) \geq 0$

- If $\sigma_{1H}^2 > \sigma_{1L}^2$, then

$$\frac{1}{2} \leq \Pr(y_{2H} > \hat{y}_{2H}) < \Pr(y_{2L} < \hat{y}_{2L}).$$

- If $\sigma_{1H}^2 \leq \sigma_{1L}^2$, then

$$\frac{1}{2} \leq \Pr(y_{2L} < \hat{y}_{2L}) < \Pr(y_{2H} > \hat{y}_{2H}).$$

Therefore, for any realization of w_2 , $\Pr(y_{2H} > \hat{y}_{2H}) > \Pr(y_{2L} < \hat{y}_{2L}) - \frac{1}{2}$.

Now consider the relative sizes of $\frac{\partial V_{1H}}{\partial g_1(H)}$ and $\frac{\partial V_{1L}}{\partial g_1(H)}$:

$$\begin{aligned}\frac{\partial V_{1H}}{\partial g_1(H)} - \frac{\partial V_{1L}}{\partial g_1(H)} &= w_1 + \delta \int w_2 [\Pr(y_{2H} > \hat{y}_{2H}) - \Pr(y_{2L} < \hat{y}_{2L})] d\mu(w_2) \\ &> w_1 - \frac{\delta}{2} \int w_2 d\mu(w_2)\end{aligned}$$

Therefore, $\frac{\partial V_{1H}}{\partial g_1(H)} > \frac{\partial V_{1L}}{\partial g_1(H)}$ if

$$E_1(w_2) < 2w_1, \forall w_1, \quad (11)$$

for any set of priors $(g_t(H), \sigma_{tH}^2)$, $(g_t(L), \sigma_{tL}^2)$, growing conditions realization w_1 and price p . If we do not expect the change in growing conditions from period 1 to period 2 to be so strong that yields are doubled, we can be sure that good news regarding the growing conditions of H has a stronger positive effect on the value of choosing H in period 1 than on the value of choosing L . If growing conditions over adjacent periods is sufficiently strongly positively correlated, this condition will hold; similarly, if the range of growing conditions outcomes is sufficiently restricted, this condition will hold even if growing conditions were not serially correlated.

If (11) holds and hence $\frac{\partial V_{1H}}{\partial g_1(H)} > \frac{\partial V_{1L}}{\partial g_1(H)}$, then the probability that H is chosen in period 1 is increasing in the farmer's value of $g_1(H)$. It can be shown as in the myopic model, there is a unique critical value of w_1 (\hat{w}_1) defined implicitly when $V_{1H}(\hat{w}_1)$ equals $V_{1L}(\hat{w}_1)$, such that if $w_1 \leq \hat{w}_1$ the farmer chooses $x_1 = L$, and otherwise $x_1 = H$.³⁴

9 Appendix 2: Alternative Measures of Information Connections

To concisely describe the definitions of our metrics, let i and j be farmers within a village. Let $l_{ij} \in \{0, 1\}$ describe the relationship between any two farmers in the village. We say that i and j are linked if $l_{ij} = 1$, and the neighborhood of i is defined as $N_i = \{j | l_{ij} = 1\}$.

We generated one set of alternate metrics for each respondent using a listing of all the individuals named by that respondent in a number of different contexts. This data includes people named in response to questions designed to record all 'significant' conversations about farming between individuals, and people who were hired by, borrowed from, lent or sold output to, or

³⁴As in the myopic model, this requires $g_1(H) \geq g_1(L)$. In addition, we need to assume that good productivity in period 1 is not too bad news about the likely productivity in period 2: a sufficient condition is that $\mu(\cdot|a)$ first-order stochastically dominates $\mu(\cdot|b)$ for $a > b$. With these conditions, $\frac{\partial V_{1H}}{\partial w_1} > \frac{\partial V_{1L}}{\partial w_1}$ and the statement follows.

exchanged gifts, transacted land or jointly held assets with the respondent.³⁵ We construct two metrics from this information, first defining $l_{ij} = 1$ if either i reports learning about farming from j , or vice versa. Because important information might be transmitted during quite casual conversation, we also define a broader information neighborhood with $l_{ij} = 1$ if i appears anywhere in j 's roster of contacts, or if j appears on i 's roster.

Both our baseline ask for advice metric and these roster of contacts based measures have potential drawbacks. The ask for advice measure is based on a random sample of other farmers, and so yields estimates of the information neighborhood of a farmer that are smaller than his actual information neighborhood. The roster of contacts measures include some pairs who probably do not discuss farming activities, and depends upon the respondents' subjective understanding of 'significant conversations about farming.' Therefore, we also construct a predicted information neighborhood based on estimates reported in Table 2 of the probability of a link (based on the question "Have you ever gone to ___ for advice about your farm?") given pair characteristics.

10 Appendix 3: Constructing Indices of News

This Appendix details our construction of good/bad news indices. The first component in our index of good news at previous inputs is meant to reflect the number of his neighbors' pineapple plants (with the given fertilizer level) that turned out to have surprisingly high profits. We start with the hypothetical case where all plantings are of a single plant each. Each neighbor's planting from t_0 up to $t_1 - 1$ using a given input level provides the farmer with some information about that level's productivity. Suppose all plantings by the farmer's neighbors were indexed at times τ that were in the interval $[t_0, t_1)$ and at least 5 periods before t_1 . In this case, all information about these plantings would have been unknown at time t_0 but would be revealed within the window of opportunity to apply fertilizer to the time t_1 planting. Fertilizer for a time t planting is the sum of that applied during periods $t + 1$ to $t + 4$, so a time $t - 5$ planting will be harvested and its outcome revealed by then end of this span (see Section 3.2). In this case, we could obtain a count

³⁵Significant conversations include, for example, discussions of techniques for using agricultural chemicals, seeds, dealing with agricultural problems, or crop choice.

of all the good news plantings whose information is relevant for the time t_1 planting by summing the appropriate $d_{i,k,\tau}$ over farmer i 's N_i neighbors. However, in our data many of the plantings are less than a full 5 periods before t_1 so some information about their outcome is known to farmer i at t_1 , but not necessarily all of it. To capture the partial revelation of good news events, we could define an information aggregate that is a weighted sum of good news event indicators for farmer i as:

$$G_{i,t_1}^*(x = x_{i,t_0}) \equiv \sum_{k \in N_i} \sum_{\tau \in [t_0, t_1)} \tilde{\psi}(t_1 - \tau) d_{i,k,\tau} 1\{x_{i,t_0} = x_{k,\tau}\}. \quad (12)$$

The first sum is over all N_i information neighbors of farmer i , the second sum over all times τ for experiments within his information neighborhood since his last planting at t_0 . The weight is specified to be zero for all non-positive lags, rise to one-fifth at lag 1 ($\tilde{\psi}(1) = 1/5$) and increase linearly until it is one at lag 5, and then remain constant at one for all further lags. Thus, G^* gives all observations 5 or more periods in the past a weight of one and indicators of good news that are not a full 5 periods in the past are linearly discounted.

We do not use G^* itself for two reasons: some farmers have plantings τ occurring soon enough before t_0 that some of their information should still be considered as news at time t_1 and our farmers do not all have the same number of pineapples on each plot nor the same amount of observed outcomes.

First, we face the question of how to apportion the information for experiments at times τ that are 1 to 4 periods before t_0 . Some of the information they contain should be reflected in the t_0 planting but it would not all be revealed at t_0 . Some information in the τ planting would be revealed after t_0 and be news from the perspective of time t_1 . To apportion the time τ good news information between the t_1 and t_0 plantings we modify the weighting in G^* and assign this term the weight

$$\psi(\tau; t_0, t_1) \equiv (\tilde{\psi}(t_1 - \tau) - \tilde{\psi}(t_0 - \tau)) \quad (13)$$

at time t_1 . Thus we assign $\tilde{\psi}(t_0 - \tau)$ of the good news event to time t_0 and $\tilde{\psi}(t_1 - \tau) - \tilde{\psi}(t_0 - \tau)$ to time t_1 . Using $\psi(\tau; t_0, t_1)$ as our weight function modifies our measure of good news to:

$$G_{i,t_1}^{**}(x = x_{i,t_0}) \equiv \sum_{k \in N_i} \sum_{\tau \in [t_0-4, t_1)} \psi(\tau; t_0, t_1) d_{i,k,\tau} 1\{x_{i,t_0} = x_{k,\tau}\}.$$

One main issue remains: our farmers do not all plant the same number of pineapple plants per plot and they have observed differing amounts of plants' outcomes in the past.³⁶ Hence we construct a plant-weighted version of G^{**} by multiplying the indicators $d_{i,k,\tau}$ in G^{**} by the number of pineapple plants on the corresponding plot: $Plants_{k,\tau}$. Dividing this sum by the total number of plants in experiments the farmer has observed in his information neighborhood from the beginning of our data up to time t_1 , $TotalPlants_{i,t_1}$, gives us a measure of the number of 'good news plants' relative to the overall total the farmer has observed. This is our index, $G_{i,t_1}(x = x_{i,t_0})$ as defined in (5).

³⁶There is also an occasional complication caused when the application window for a farmer's previous planting is still open when a new planting starts. For example, take a farmer with plantings at periods 3, 7, and 8. There is an issue about how to assign the "previously used" input amount for the period 8 planting as the amount used on the period 7 planting has been started but not completed when the round 8 application window opens. When this type of issue occurs, we construct a weighted average of input use for the previous plantings with weights proportional to ψ^* , i.e. they increase linearly to maximum of 1 at lag 5.




Figure 1. Sample proportion of farmers cultivating pineapple by year. We define farmers who adopted before 1994 as experienced.

Table 1: Descriptive Statistics

	Mean	Std. Deviation
Fertilizer Use (cedis per sucker)	1.938	5.620
Change in Fertilizer Use	-0.315	10.335
Indicator of Change in Fertilizer \neq 0	0.496	0.502
Indicies of Successful Experiments:		
M Ask advice	0.114	4.312
M Talk frequently	-0.041	4.259
M Know each other's plots	-0.033	3.308
M Roster of contacts, farm info only	-0.174	3.467
M Roster of contacts, full list	-0.126	3.469
M Predicted ask for advice	-0.020	0.940
Avg. Dev. of Lagged Use From Geographic Neighbors' Use	0.498	7.790
Avg. Dev. of Lagged Use From Financial Neighbors' Use	-0.053	0.380
Wealth (million cedis)	2.331	2.835
Clan 1 Indicator	0.327	.
Clan 2 Indicator	0.451	.
Church 1 Indicator	0.487	.
pH	5.952	0.738
Soil Organic Matter (%)	2.927	1.110
Loamy Soil Indicator	0.434	.
Sandy Soil Indicator	0.053	.
Contact with Extension Agent Indicator	0.327	.
Inexperienced Farmer Indicator	0.230	.

Table 2: Logit Predicting Ask For Advice

	Coefficient	Standard Error
Either Party Holds Traditional Office	-0.55	0.26
Same Religion	0.04	0.33
Same Clan	0.43	0.24
Same Gender	1.73	0.78
Same Soil Type	-0.23	0.27
Absolute Age Difference (years)	-0.04	0.02
Absolute Wealth Difference (million cedis)	0.15	0.03
Distance Between Plot Centers (kilometers)	-0.46	0.16
Constant	-2.14	0.84

Logit MLE Estimates, Sample Size = 490, Pseudo R-squared = .12.

Dependent variable is one if either party answered yes to the question:

Have you ever gone to _____ for advice about your farm?

Table 3: Information Connections by Cohort of Pineapple Adoption

Proportion of pairs of individuals in each other's information neighborhood

	Not Farming Pineapple	Inexperienced Pineapple Farmer	Experienced Pineapple Farmer	Neighborhood Metric
Not Farming Pineapple	0.06	0.05	0.07	Response to "Have you ever gone to ____ for advice about your farm?"
Inexperienced Pineapple Farmer	0.05	0.09	0.13	
Experienced Pineapple Farmer	0.07	0.13	0.21	

Table 4: Predicting the Change in Input Use

Dependent Variable: Indicator of a Change in Per Plant Fertilizer Use

	A	B
Good News at Lagged Fertilizer Use	-0.13 [1.19]	0.05 [1.16]
Good News at Alternative Fertilizer Use	0.18 [0.97]	0.37 [1.02]
Bad News at Lagged Fertilizer Use	12.32 [3.72]	14.41 [4.63]
Bad News at Alternative Fertilizer Use	-2.98 [1.91]	-4.22 [2.07]
Average Absolute Deviation from Geographic Neighbors' Fertilizer Use	0.49 [0.13]	0.49 [0.14]
Inexperienced Farmer	1.14 [0.92]	1.28 [1.04]
Talks with Extension Agent		-1.39 [0.75]
Village 1	-2.34 [0.88]	-1.70 [1.06]
Village 2	-1.96 [0.99]	-2.55 [0.95]
Wealth (Million Cedis)	-0.04 [0.13]	0.04 [0.11]
Clan 1	2.05 [1.07]	2.39 [1.02]
Clan 2	2.72 [0.94]	2.88 [0.84]
Church 1	-0.18 [0.91]	-0.68 [0.91]

Logit MLE point estimates, spatial GMM (Conley 1999) standard errors in brackets allow for heteroskedasticity and correlation as a function of physical distance, see footnote 20 for details. Sample Size = 107. Pseudo R-squareds .34 and .36, columns A and B respectively. A full set of round dummies were included but not reported. Information neighborhoods defined using responses to: Have you ever gone to farmer ____ for advice about your farm?

Table 5: Predicting Innovations in Input Use, Differential Effects by Source of Information

Dependent Variable: Innovation in Per Plant Fertilizer Use

	A	B	C	D	E	F
Index of Inputs on Successful Experiments (M)	0.99 [.16]					
M * Inexperienced Farmer		1.09 [0.22]				
M * Experienced Farmer		0.10 [0.32]				
Inexperienced Farmer		4.01 [2.62]	4.20 [2.66]	4.22 [2.65]	4.19 [2.65]	4.12 [2.77]
Index of Experiments by Inexperienced Farmers			-0.13 [0.37]			
Index of Experiments by Experienced Farmers			1.02 [0.17]			
Index of Exper. by Farmers with Same Wealth				1.03 [0.18]		
Index of Exper. by Farmers with Different Wealth				-0.41 [0.32]		
Index of Experiments on Big Farms					1.10 [0.14]	
Index of Experiments on Small Farms					0.89 [0.18]	
Index of Exper. by Farmers with Same Soil						1.04 [0.16]
Index of Exper. by Farmers with Different Soil						0.91 [0.19]
Avg. Dev. of Lagged Use From Geographic Nbrs	0.54 [0.06]	0.55 [0.08]	0.58 [0.06]	0.58 [0.06]	0.58 [0.06]	0.59 [0.06]
Avg. Dev. of Lagged Use From Financial Nbrs	0.53 [0.58]	0.45 [0.58]	0.40 [0.59]	0.43 [0.55]	0.22 [0.61]	0.24 [0.60]
Village 1	-7.62 [1.16]	-7.92 [1.43]	-8.09 [1.36]	-8.24 [1.43]	-7.81 [1.31]	-7.88 [1.31]
Village 2	-0.61 [1.56]	-1.82 [2.02]	-2.15 [2.03]	-2.17 [2.11]	-1.83 [2.02]	-1.78 [2.07]
Wealth (Million Cedis)	0.13 [0.25]	0.36 [0.17]	0.41 [0.17]	0.45 [0.17]	0.29 [0.20]	0.29 [0.20]
Clan 1	-2.62 [1.29]	-2.42 [1.21]	-2.68 [1.12]	-2.62 [1.09]	-2.53 [1.11]	-2.55 [1.15]
Clan 2	-0.40 [1.44]	-0.11 [1.32]	-0.11 [1.32]	-0.15 [1.32]	-0.31 [1.30]	-0.29 [1.30]
Church 1	0.26 [1.29]	0.67 [1.12]	0.76 [1.06]	-0.60 [1.11]	0.87 [1.12]	0.88 [1.12]
R-squared	0.70	0.73	0.71	0.71	0.71	0.71

OLS point estimates, spatial GMM (Conley 1999) standard errors in brackets allow for heteroskedasticity and correlation as a function of physical distance, see footnote 20 for details. Sample Size = 107. A full set of round dummies included but not reported. Information neighborhoods defined using responses to: Have you ever gone to farmer _____ for advice about your farm?

Table 6: Alternate Definitions of the Information Network

Dependent Variable: Innovation in Per Plant Fertilizer Use

	A	B	C
Information Neighborhood Metric	Roster of Contacts: Farm Info Only	Roster of Contacts: Full Set of Contacts	Predicted Advice
M * Inexperienced Farmer	1.50 [0.28]	1.49 [0.28]	6.34 [1.14]
M * Experienced Farmer	0.19 [0.21]	0.15 [0.22]	4.52 [1.80]
Inexperienced Farmer	4.66 [2.84]	4.65 [2.84]	4.01 [2.77]
Average Deviation of Lagged Use From Geographic Neighbors' Use	0.49 [0.09]	0.49 [0.09]	0.33 [0.12]
Average Deviation of Lagged Use From Financial Neighbors' Use	0.50 [0.69]	0.51 [0.70]	0.59 [0.82]
Village 1	-7.59 [1.64]	-7.52 [1.63]	-9.25 [1.75]
Village 2	-2.09 [2.11]	-2.08 [2.10]	-1.86 [2.07]
Wealth (Million Cedis)	0.35 [0.17]	0.35 [0.17]	0.16 [0.22]
Clan 1	-2.25 [1.37]	-2.23 [1.37]	-1.66 [1.28]
Clan 2	-0.02 [1.41]	0.01 [1.40]	0.42 [1.33]
Church 1	0.62 [1.22]	0.59 [1.23]	0.75 [1.23]
R-squared	0.72	0.72	0.73

OLS point estimates, spatial GMM (Conley 1999) standard errors in brackets allow for heteroskedasticity and correlation as a function of physical distance, see footnote 20 for details. Sample Size = 107. A full set of round dummies were included but not reported. Alternative information neighborhoods are as defined in Section 3.2 and Appendix 2.

Table 7: Robustness to Changes in Specification
 Dependent Variable: Innovation in Per Plant Fertilizer Use

	A	B	C	D	E	F
	No Information from Plantings at t-1	No Information from Plantings at t-1 or t-2	Fertilizer Categories: Zero, Med, High (High > 80th percentile)	Geographic Neighborhood within 500m	Soil Characteristics	Lagged Fertilizer Use
M * Inexperienced Farmer	0.93 [0.19]	0.39 [0.15]	1.03 [0.19]	1.85 [0.20]	1.11 [0.27]	0.34 [0.13]
M * Experienced Farmer	0.16 [0.23]	-0.14 [0.24]	-0.41 [0.35]	0.04 [0.23]	-0.24 [0.45]	0.08 [0.31]
Inexperienced Farmer	4.38 [2.73]	5.83 [3.46]	4.01 [2.71]	2.87 [2.71]	5.94 [2.72]	4.05 [2.62]
Lagged Own Fertilizer Use						-0.84 [0.22]
Avg. Dev. of Lagged Use From Geographic Neighbors' Use	0.51 [0.08]	0.95 [0.09]	0.58 [0.08]	0.10 [0.06]	0.50 [0.12]	0.09 [0.17]
Avg. Dev. of Lagged Use From Financial Neighbors' Use	0.36 [0.54]	-0.17 [0.40]	0.51 [0.55]	1.06 [1.16]	0.64 [0.70]	0.59 [0.61]
Village 1	-7.51 [1.35]	-7.96 [1.58]	-8.13 [1.51]	-7.48 [2.10]	-13.46 [3.03]	-3.19 [1.55]
Village 2	-2.02 [2.09]	-2.83 [2.71]	-1.93 [2.09]	-1.40 [2.08]	-1.68 [2.37]	-2.70 [2.11]
Wealth (Million Cedis)	0.36 [0.18]	0.50 [0.18]	0.41 [0.17]	0.24 [0.20]	0.73 [0.21]	0.21 [0.19]
Clan 1	-2.30 [1.16]	-2.58 [1.35]	-2.45 [1.35]	-2.71 [1.21]	-4.04 [1.87]	-1.17 [1.32]
Clan 2	-0.07 [1.28]	0.16 [1.31]	-0.004 [1.35]	-0.62 [1.47]	0.25 [1.39]	0.71 [1.29]
Church 1	0.56 [1.14]	1.33 [1.17]	0.62 [1.13]	0.43 [1.36]	1.58 [1.58]	0.36 [1.16]
Soil Organic Matter					0.14 [0.67]	
Soil pH					4.09 [2.31]	
Soil Type = Loam					1.40 [1.14]	
Soil Type = Sandy					-5.78 [2.72]	
Sample size	107	93	107	107	89	107
R-squared	0.74	0.72	0.73	0.68	0.80	0.75

OLS point estimates, spatial GMM (Conley 1999) standard. errors in brackets allow for heteroskedasticity and correlation as a function of physical distance, see footnote 20. Round dummies included but not reported. Alternative specifications are as defined in Section 6.2.

Table 8: Predicting Innovations in Labor for Pineapple and Maize-Cassava Plots

Dependent Variable: First Difference in Labor Inputs for Pineapple and Maize-Cassava

Crop	A	B	C
	Pineapple (labor cost in cedis per plant)	Maize-Cassava (labor cost in 1000 cedis per hectare)	Maize-Cassava (labor cost in 1000 cedis per hectare)
Index of Experiments: M-tilde	0.95 [.37]	0.04 [0.14]	
Index of Experiments in the Geographic Neighborhood			0.25 [0.14]
Average Deviation of Lagged Use From Geographic Neighbors' Use	0.48 [0.22]	0.83 [0.14]	
Average Deviation of Lagged Use From Financial Neighbors' Use	-0.27 [0.23]	0.02 [0.11]	
Village 1	770.96 [238.43]	-195.98 [100.54]	-203.94 [98.52]
Village 2	567.83 [316.37]	-417.61 [119.34]	-389.13 [134.16]
Village 3		-126.29 [89.21]	-228.58 [132.30]
Wealth (Million Cedis)	47.59 [31.46]	1.09 [33.42]	-14.05 [37.69]
Clan 1	-802.58 [358.62]	213.44 [566.83]	-794.38 [450.66]
Clan 2	369.36 [232.41]	-45.33 [82.65]	-3.19 [78.85]
Church 1	93.85 [179.96]	-37.86 [77.80]	-4.69 [86.46]
Soil Organic Matter	-259.11 [202.53]	-99.26 [90.41]	-72.00 [40.58]
Soil pH	358.52 [141.46]	15.19 [131.66]	88.98 [58.55]
Soil Type = Loam	-641.90 [197.89]	-97.95 [45.75]	-105.23 [78.08]
Soil Type = Sandy	-909.36 [382.93]	128.06 [66.14]	-76.96 [126.10]
Sample size	89	346	346
R-squared	0.55	0.42	0.24

OLS point estimates, spatial GMM (Conley 1999) standard errors in brackets allow for heteroskedasticity and correlation as a function of physical distance, see footnote 20 for details. Round/season dummies included but not reported.

Information neighborhood from: Have you ever gone to farmer _____ for advice about your farm?