

Social Learning Through Networks:

The Adoption of New Agricultural Technologies in Ghana*

The adoption of new technology is a central feature of the transformation of farming systems during the process of economic development. There is a rich tradition of social science research, therefore, on the adoption of new technologies (Evenson and Westphal; Feder, Just and Zilberman). Much of this research has focused on the problem faced by individual farmers as they decide whether to adopt a potentially profitable new technology. How do farmers learn about a new technology?

There are many possible sources of information about the new technology (Rogers). A farmer may learn from his or her own experimentation with the technology. Advice and technical information may be available from the extension service or the media. If there are many farmers in somewhat similar circumstances, then the process of learning about the new technology may be social. Farmers may learn about the characteristics of the new technology from their neighbors' experiments.

There is a natural way of thinking about this process of 'social learning' that has motivated much of the existing work on the topic (Besley and Case; Foster and Rosenzweig; Munshi). Consider the village (or other appropriate social group) as a unit of learning engaged in a process of collective experimentation. Each farmer in the village observes the farming activities of each of the other farmers, including of course those who are experimenting with the new technology. Each farmer then updates his or her own opinion regarding the technology using this information, makes decisions regarding cultivation for the next season,

and the learning process continues. There are two important assumptions about the nature of social learning in this story. First, each farmer receives information on the outcomes of experiments from every other farmer in the village. Second, each farmer observes other farmers' experiments with no loss of information. Few would argue (and none of the cited authors argue) that these assumptions are literally true, but they seem to be a reasonable starting point for a discussion of learning about agricultural technology in a village context. However, both assumptions are strongly contradicted by data from pineapple farmers in a set of Ghanaian villages.

Learning About a New Technology in Ghana

Over the past decade, in part of Ghana's Eastern region, an established system of maize and cassava intercropping for sale to urban consumers has begun to be replaced by intensive production of pineapple for export to European markets. An important component of this transformation is the adoption of agricultural chemicals that were not used in the previous farming system. In this paper, we consider how farmers in this area might learn about the appropriate use of fertilizer in this new farming system.

A survey was conducted of approximately 450 individuals in four clusters of villages in Ghana's Eastern Region over a period of 21 months in 1996-1998.¹ Two aspects of the data are relevant here. First, plot level data on inputs and outputs was collected at frequent intervals from the respondents. Second, a variety of data on farmer interactions was collected. For example, data was collected on respondents' knowledge of inputs and outputs on the plots of other respondents and on respondents' conversations about farming (and specifically about fertilizer) with other farmers. These data provide direct evidence regarding the

assumptions that underlie the canonical model of social learning in village settings.

It does not appear to be the case that farmers learn from all other farmers in their village. Each respondent was matched randomly with 10 other farmers in his/her village. In only 11 percent of these matches had one of the two individuals ever received advice about farming from the other. In 30 percent of the matches, the respondent indicates that he could approach the other farmer for advice about fertilizer. Respondents are able to provide some information on harvests and inputs used on approximately 7 percent of random matches between respondents and pineapple plots cultivated by other farmers in the village. Thus, it is better to characterize information regarding farming as flowing through a relatively sparse social network rather than being freely available in the village. Perhaps this is most effectively demonstrated through figure 1, which depicts the communication network among our sample from one of the survey villages. Information links between farmers are indicated by the edges of the graph.² Similar graphs can be drawn for the other villages and for other definitions of information exchange and the general pattern is unchanged. Farmers learn through social networks and the constituent links are not based solely on geographic proximity.

Farmers often have only limited information regarding the inputs used and output harvested by the limited number of other farmers with whom they communicate. A communication link is defined to exist in the approximately seven percent of random matches between respondents and pineapple plots cultivated by other farmers in which the respondent is able to provide some information regarding inputs or outputs on the other plot. In less than a third of these cases is the respondent able to provide a quantitative estimate of the amount

of any chemical input applied, or of the number or value of pineapples harvested. In most cases, the respondent is able to provide only *relative* information regarding the other farmers' actions. For example, the respondent often does not know the other farmer's actual harvest or input use; instead he knows only if the other farmer harvested more or less pineapples than the village average, or than the respondent himself.

The evidence from these Ghanaian villages indicates that farmers know little about inputs or outputs on the plots of most farmers in the village. Instead, information flows through relatively restricted channels, with each farmer learning about new technology from a few sources. Moreover, even with respect to this limited number of direct contacts information is not always perfect: farmers are more likely to know broad facts (e.g., the other farmer had a particularly good harvest) rather than specific details. Our purpose in the remainder of this paper is to examine some of the implications for optimal learning behavior when (a) information is incomplete and (b) flows through networks.

Examples of Social Learning

This section describes a simple model of learning about optimal inputs. The basic form of this model is that farmers know the production technology up to the distribution of an optimal or target input. Suppose that each farmer has the following technology and that each operates it on one plot, for simplicity in notation. On farmer i 's plot, a single input (fertilizer) $f_{i,t}$ is chosen at time t and the following period output $q_{i,t+1}$ is produced. The amount of output is random because the *ex post* optimal level of the input is random, given by a 'target input' variable $\theta_{i,t+1}$ that for simplicity we take to be IID across agents and time. We assume the loss associated with suboptimal inputs is quadratic: $q_{i,t+1} = \lambda - (f_{i,t} - \theta_{i,t+1})^2$.

Further suppose that the distribution of the target input shock is known by all farmers to be normally distributed with known variance V and unknown expected value m . So farmers have only the expected value of the target input shock, m , to learn.

The farmer's problem is to choose input $f_{i,t}$ given his beliefs. We model farmers as acting simultaneously in choosing $f_{i,t}$ so that others' choices are only observable in the following period. In this target input model, there are no strategic nor experimental motivations for a farmer to vary his use of inputs. Farmers are assumed to know the form of the production function, therefore the same information is revealed by any level of input use, so the optimal strategy is to choose $f_{i,t}$ to maximize profits. For simplicity of exposition, we assume that inputs are costless. The assumption of quadratic loss from having suboptimal inputs implies that the optimal choice of inputs is the expectation of θ with respect to his subjective prior distribution. So the optimal input choice in period 1 for farmer i is $f_{i,1} = E_{i,1}(\theta_{i,2})$.

In this note, we model farmers' beliefs about m in a Bayesian framework so that they have a subjective prior distribution over values of m and learning is identified with updating these priors. We consider first an example of a two-farmer village where both farmers are in direct communication with each other. We also suppose that each farmer observes without error the realization of θ for himself as well as his information neighbor. The second example illustrates the process of social learning when networks matter, that is, when not all farmers are in direct communication with each other. Finally, the last example illustrates the impact of limited observability on updating rules when all farmers directly communicate with each other.

Full Information

In example one, suppose there are two farmers A and B in communication with each other. Assume that each farmer has full information *ex post* about his own and the other's realization of θ . It should be noted that the form of the production function implies that knowledge of inputs and outputs is not sufficient to deduce θ ; it must also be the case that each farmer knows on which side of the optimal θ his fertilizer use fell. This is a general problem with production functions with an interior optimal input choice.³ Suppose each farmer starts out in period zero with priors $N(\widehat{m}_0, \sigma_0^2)$. Farmer i 's period zero actions will be $f_{i,0} = \widehat{m}_0$. In period 1, each farmer will update beliefs with the information revealed that period: $\theta_{A,1}$ and $\theta_{B,1}$. The updated $\widehat{m}_{A,1}$ will be a weighted average of the θ observations and \widehat{m}_0 where the weights depend on the known variance of θ , V , and the precision of the prior, σ_0^{-2} . Optimal actions will then be equal to this updated expectation. For farmer A:

$$1. \quad f_{A,1} = \widehat{m}_{A,1} = \frac{\sigma_0^2}{\sigma_0^2 + \frac{V}{2}} \left[\frac{1}{2}(\theta_{A,1} + \theta_{B,1}) \right] + \frac{\frac{V}{2}}{\sigma_0^2 + \frac{V}{2}} \widehat{m}_0$$

This example is simple to generalize to N farmers, and illustrates that with full information, neighbors' experiments are just as informative as a farmer's own experiments and so they have a symmetric impact upon beliefs. It illustrates the principle that with full information, a farmer's beliefs in any period depend only upon his priors and the realizations of θ by anyone in his information neighborhood. Neither the identity of the farmer nor the ordering of the realizations has any effect on posterior beliefs.

Communication Networks

In Ghana, learning occurs through social networks rather than in the context of the collective experiment illustrated above. What are the implications for learning when not all farmers

are in communication with all other farmers? Suppose there are three farmers A,B and C, and that A communicates only with B, C communicates only with B, and B communicates with both A and C. Assume as in the first example that each farmer has full information about the experiments of farmers with whom he communicates.

Consider the optimal period 1 actions of farmers A and B after they update their beliefs with the information available to them after the first period's θ s are revealed. A updates his beliefs to:

$$2. \quad f_{A,1} = \hat{m}_{A,1} = \frac{\sigma_0^2}{\sigma_0^2 + \frac{V}{2}} \left[\frac{1}{2} (\theta_{A,1} + \theta_{B,1}) \right] + \frac{\frac{V}{2}}{\sigma_0^2 + \frac{V}{2}} \hat{m}_0$$

while B now believes

$$3. \quad f_{B,1} = \hat{m}_{B,1} = \frac{\sigma_0^2}{\sigma_0^2 + \frac{V}{3}} \left[\frac{1}{3} (\theta_{A,1} + \theta_{B,1} + \theta_{C,1}) \right] + \frac{\frac{V}{3}}{\sigma_0^2 + \frac{V}{3}} \hat{m}_0.$$

Each farmer takes actions $f_{i,1}$ and then observes outcomes $\theta_{i,2}$. A knows all the determinants of $f_{B,1}$ with the exception of $\theta_{C,1}$; hence as soon as A observes that B applies $f_{B,1}$, A can deduce what B must have learned from C , and hence the outcome of C 's experiment in period 1. Therefore,

$$4. \quad f_{A,2} = \hat{m}_{A,2} = \frac{\sigma_0^2}{\sigma_0^2 + \frac{V}{5}} \left[\frac{1}{5} (\theta_{A,2} + \theta_{B,2} + \theta_{A,1} + \theta_{B,1} + \theta_{C,1}) \right] + \frac{\frac{V}{5}}{\sigma_0^2 + \frac{V}{5}} \hat{m}_0.$$

A and C , therefore, learn of each others' experiments (with a lag) through the impact of these experiments on the beliefs (and thus actions) of B .

If Farmer A were to use only the information contained in $\{\theta_{A,1}, \theta_{B,1}, \theta_{A,2}, \theta_{B,2}\}$, he would not have an optimal update of his beliefs because he would omit the information generated by $\theta_{C,1}$. In general, when learning occurs through networks, observation of neighbors' optimal decisions contains a source of information that is not captured by their target input shock histories. Moreover, the particular history of realizations of θ now matters because farmer

A must model farmer B's updating process in order to deduce farmer B's observation of $\theta_{C,1}$. If the network structure were to be extended, for example by adding farmer D who communicates only with farmer C, farmer A would have to extend the depth of his memory in order to deduce $\theta_{D,1}$ and $\theta_{C,2}$.

Limited Communication

We now consider how limited information will affect optimal learning behavior. For the sake of argument, consider two simple ways in which the realization of θ by one farmer may be imperfectly observable by other farmers. First, it may be observed with noise, so that the sum of θ plus an additive observation error is observed. Alternatively, only an indicator of the range of θ may be observed, for example that θ is above or below a certain value.⁴ In both cases we consider the simple network structure in which two farmers communicate with each other.

Suppose farmer A observes $\theta_{A,t}$ perfectly but his observations of farmer B allow him only to recover $\theta_{B,t} + u_{B,t}$ where $u_{B,t}$ is independent measurement error with distribution $N(0, \Sigma)$. Farmer A will update his beliefs and take action $f_{A,1}$:

$$5. \quad f_{A,1} = \hat{m}_{A,1} = \frac{\sigma_0^2(V+\Sigma)}{\sigma_0^2(V+\Sigma)+\sigma_0^2V+V(V+\Sigma)} [\theta_{A,1}] \\ + \frac{\sigma_0^2V}{\sigma_0^2(V+\Sigma)+\sigma_0^2V+V(V+\Sigma)} [\theta_{B,1} + u_{B,1}] + \frac{V(V+\Sigma)}{\sigma_0^2(V+\Sigma)+\sigma_0^2V+V(V+\Sigma)} \hat{m}_0.$$

This is similar to the updating rule in period 1 for the full information case, with the difference being that the information from neighbors is downweighted relative to the information available from one's own experiments. Farmer B's actions in period 1 are symmetric to those of farmer A.

Different forms of imperfect information can imply different updating rules. For example,

suppose farmer A observes only whether $\theta_{B,t}$ is larger or smaller than a known value θ^* . Suppose that in period 1 farmer A observes $\theta_{A,1}$ and that $\theta_{B,1} > \theta^*$, the likelihood of this observation will have a density term for the $\theta_{A,1}$ information and a cumulative distribution term due to observation of the event $\theta_{B,1} > \theta^*$. The resulting posterior mean does not to our knowledge have a closed form solution so we do not present it here. However, it is clear that it will result in a distinct likelihood and hence different posterior means and actions from the preceding case where $\theta_{B,1}$ is observed with an additive error.

In either case, farmer A can use farmer B's period 1 actions to deduce additional information about the distribution of θ . Returning to the additive error case, notice that $f_{B,1}$ depends on $\theta_{B,1}$ measured without any error, and on $\theta_{A,1} + u_{A,1}$, where $u_{A,1}$ is independent of $u_{B,1}$. Therefore, in period 2, farmer B's previous action in period 1 will provide farmer A with more information than was contained in $\{\theta_{A,2}, \theta_{A,1}, \theta_{B,2} + u_{B,2}, \theta_{B,1} + u_{B,1}\}$. Similarly, in the case in which A observed only that $\theta_{B,1} > \theta^*$, farmer B's action in period 1 depends upon a precisely observed $\theta_{B,1}$ and hence can provide farmer A with additional information about $\theta_{B,1}$. In each instance, farmer A (for example) can infer from B's actions some of the information that is directly available to B but not to A.

In the case of incomplete information, just as in the case of learning through networks discussed above, each farmer therefore finds it in his interest to infer additional information about the distribution of θ from the actions of the other farmer. In order to make this inference, each farmer must model the learning process of the other farmers with whom he communicates. When communication occurs through networks, farmer A must reason how farmer B updates his priors upon receipt of information from farmer C, with whom

A does not directly communicate. When communication is limited, farmer A must reason how farmer B learns from his own farming experience, because A directly observes only an imperfect signal of B's experience.

Thus we find that social learning involves higher order reasoning by farmers: they care not only about their own direct observations of realizations of θ , but also about how their neighbors learn about θ . This adds a layer of additional complexity to the empirical analysis of social learning. When learning occurs through communication networks, or when it involves less than perfect communication, optimal learning behavior by a farmer involves keeping track of long histories of actions and experimental outcomes of everyone with whom that farmer communicates. In general, there are no simple statistics that can summarize the state of knowledge of a farmer engaged in this process. In the interest of making empirical work more tractable, researchers might assume that there exist such simple summary statistics. However, when doing so, researchers should recognize that they are in effect assuming that farmers do not engage in optimal learning behavior, but that they instead use simple rules of thumb to guide their behavior when innovating.

Footnotes

*Timothy Conley, Assistant Professor, Department of Economics, Northwestern University and Christopher Udry, Professor, Department of Economics, Yale University. The authors thank Chris Barrett, Stephen Morris, and Robert Vigfusson for helpful comments. The data used in this paper were collected by Markus Goldstein, 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 while the second author was a visiting scholar 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), Clair Brown Fellowship, International Food Policy Research Institute, Institute for the Study of World Politics, World Bank Research Committee, Fulbright Program, and Social Science Research Council. The authors are of course responsible for any errors in this paper.

¹The dataset is described at <http://www.econ.yale.edu/~cru2//ghanadata.html>.

²In this graph, we have defined information links based on farmers listed in each respondent's roster of individuals with whom he/she has had a significant conversations about farming.

³In general, any target input model implies that the loss is observed, perhaps with error. Knowledge of the loss generally provides the observer only with knowledge that the realized θ fell into a set. For example, if output was observed to be x , then θ must have been $f \pm \sqrt{\lambda - x}$.

⁴It should be noted that imperfect information regarding inputs and outputs of a particular form (say, additive error in observed output) does not imply a similar form for imperfect information about θ . This is a consequence of the fact that knowledge of inputs and outputs only restricts θ to a set, rather than to a point. See footnote 3.

References

Besley, T., and A. Case. "Diffusion as a Learning Process: Evidence from HYV Cotton."

Working paper, Department of Economics, Princeton University, 1994.

Evenson, R., and L. Westphal. "Technological Change and Technology Strategy."

Handbook of Development Economics. J. Behrman and T. N. Srinivasan, eds.,

pp. 2209-2300. Amsterdam: North-Holland.

Feder, G., R. Just, and D. Zilberman. "Adoption of Agricultural Innovations in Developing

Countries: A Survey." *Econ. Develop. and Cult. Change*. 33(January 1985): 255-298.

Foster, A., and M. Rosenzweig. "Learning by Doing and Learning from Others: Human

Capital and Technical Change in Agriculture." *J. Polit. Econ.* 103(December 1995):

1176-1209.

Munshi, K. "Learning From Your Neighbors: Why Do Some Innovations Spread Faster than

Others?" Working paper, Department of Economics, University of Pennsylvania, 1999.

Rogers, E. *Diffusion of Innovations*. New York: Simon & Schuster, 1995.

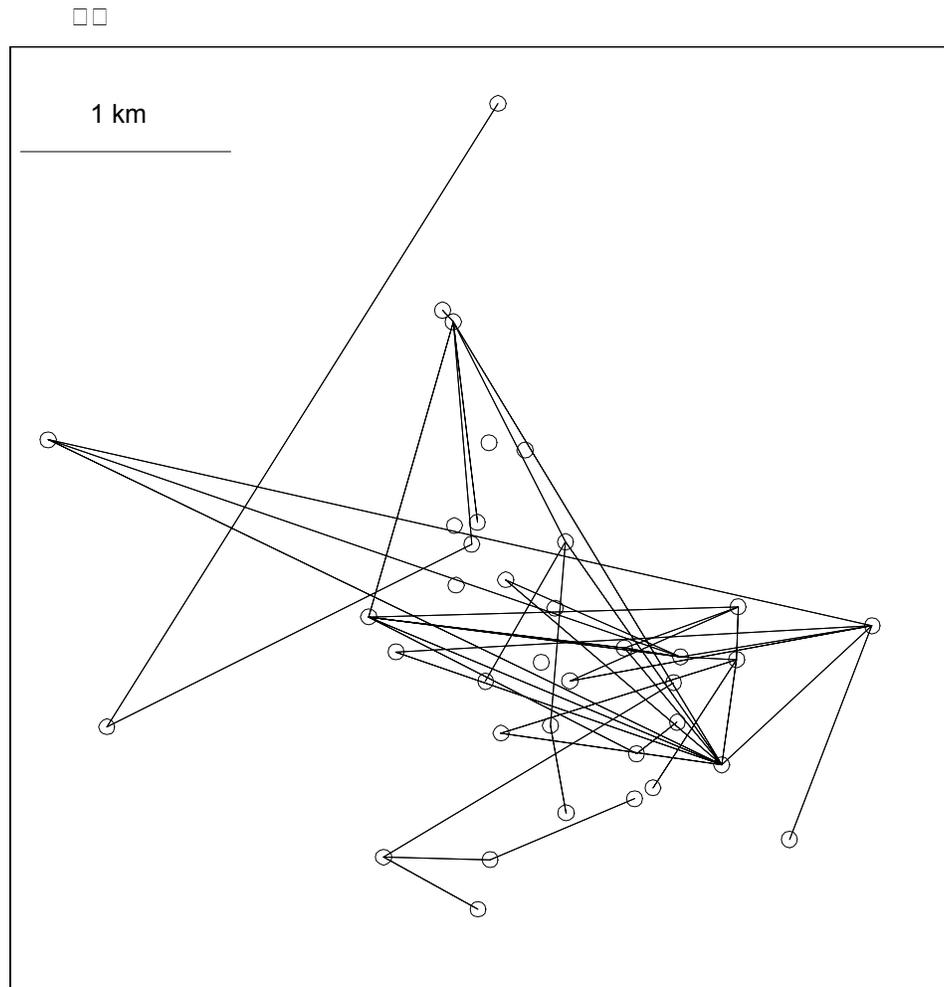


Figure 1. Roster connections and ave. plot coordinates (village 3)

Note: The circles represent the geographic center of plots for each farmer in Village 3. The lines connect farmers listed on each other's roster of contacts as having had significant conversations about farming.