

Foster, Andrew D.; Rosenzweig, Mark R.

Working Paper

Microeconomics of technology adoption

Center Discussion Paper, Economic Growth Center, No. 984

Provided in Cooperation with:

Economic Growth Center (EGC), Yale University

Suggested Citation: Foster, Andrew D.; Rosenzweig, Mark R. (2010) : Microeconomics of technology adoption, Center Discussion Paper, Economic Growth Center, No. 984

This Version is available at:

<http://hdl.handle.net/10419/59155>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

ECONOMIC GROWTH CENTER
YALE UNIVERSITY
P.O. Box 208629
New Haven, CT 06520-8269
<http://www.econ.yale.edu/~egcenter/>

CENTER DISCUSSION PAPER NO. 984

MICROECONOMICS OF TECHNOLOGY ADOPTION

Andrew D. Foster
Brown University

and

Mark R. Rosenzweig
Yale University

January 2010

Notes: Center Discussion Papers are preliminary materials circulated to stimulate discussions and critical comments.

This paper can be downloaded without charge from the Social Science Research Network electronic library at: <http://ssrn.com/abstract=1540038>

An index to papers in the Economic Growth Center Discussion Paper Series is located at: <http://www.econ.yale.edu/~egcenter/publications.html>

Microeconomics of Technology Adoption

Andrew D. Foster and Mark R. Rosenzweig

Abstract

There is an emerging consensus among macro-economists that differences in technology across countries account for the major differences in per-capita GDP and the wages of workers with similar skills across countries. Accounting for differences in technology levels across countries thus can go a long way towards understanding global inequality. One mechanism by which poorer countries can catch up with richer countries is through technological diffusion, the adoption by low-income countries of the advanced technologies produced in high-income countries. In this survey, we examine recent micro studies that focus on understanding the adoption process. If technological diffusion is a major channel by which poor countries can develop, it must be the case that technology adoption is incomplete or the inputs associated with the technologies are under-utilized in poor, or slow-growing economies. Thus, obtaining a better understanding of the constraints on adoption is useful in understanding a major component of growth.

Keywords: technology adoption review

JEL Codes: O10, O13, O33

1. Introduction

There is an emerging consensus among macro-economists that differences in technology, or TFP, across countries accounts for the major differences in per-capita GDP and the wages of workers with similar skills across countries of the world (Caselli and Coleman, 2001; Comin and Hobijn, 2004; Rosenzweig, forthcoming). Accounting for differences in technology levels across countries thus can go a long way towards understanding global inequality. One mechanism by which poorer countries can catch up with richer countries is through technological diffusion, the adoption by low-income countries of the advanced technologies produced in high-income countries (Nelson and Phelps, 1966). In this survey, we examine recent micro studies that focus on understanding the adoption process. By technology we mean the relationship between inputs and outputs, and by adoption of new technologies we mean both the use of new mappings between inputs and outputs and the corresponding allocations of inputs that exploit the new mappings.

The last major survey of technology adoption focused on agriculture in low-income countries (Feder *et al.*, 1985). As most of the world's poor work in agricultural occupations and agriculture is an important industry in most poor countries, this focus is well-placed. However, to understand fully the determinants of technological adoption, it is useful to examine adoption behavior in a variety of settings for a variety of technology types. We will thus look at studies examining the adoption of new seeds, use of fertilizer, improved bed nets, pills, boats, water purifiers, contraceptives, menstrual aids, and other innovations that are presumed to either augment profits or human welfare directly. Most studies, however, still focus on agriculture, in part because it is easier to measure inputs and outputs, although this advantage is not always well-exploited, and partly because agriculture continues to be important and there have been a flow of important innovations in agriculture, including most prominently, new high-yielding variety (HYV) seeds. And, as fertilizer is a key input for maximizing the potential of many of these new seeds, there are many studies of this farm input.

If technological diffusion is a major channel by which poor countries can develop, it must be the case that technology adoption is incomplete or the inputs associated with the technologies are under-utilized in poor, or slow-growing economies. Thus, obtaining a better understanding of the constraints on adoption and input allocations are useful in understanding a major component of growth. Documentation of such underutilization of existing technologies and inputs in, for example, agriculture in the form of unusually high rates of returns outside of

experimental plots and laboratories, however, is almost nonexistent, a topic we will discuss in more detail below.¹

What are the principal determinants of technology adoption? Table 1 reports estimates from a simple, cross-sectional regression of the probability that farmers in India in 2007 were using any HYV seeds on any of their plots of land in terms of variables that are typically looked at in studies of adoption.² And, the estimates are also typical of the major descriptive findings in the literature: First, adoption and schooling are positively correlated, net of wealth. Second, larger and wealthier farmers are more likely to adopt new technologies than are poorer households, and the effects may be non-linear. Third, the adoption by an individual farmer is positively correlated with the extent of prior adoption by his “neighbors”, in this case measured by the number of adopters in the village. What is not revealed by these estimates is the underlying causes. Does the schooling relationship reflect the fact that the more schooled have superior knowledge about the technology? Are the poorer farmers less likely to adopt the new technologies because of credit constraints, or are they more risk averse and less protected from risk than richer farmers? Or are there important economies of scale in adoption? Or are wealthy farmers wealthy because they have adopted HYV’s? Does the correlation with neighbors’ prior adoption behavior reflect learning externalities, or is it simply a reflection of common unobservables that make HYV returns higher for the farmer and his neighbors. Indeed, missing as a determinant in Table 1 is the return to adoption, which may be correlated with all of the right-hand side variable.

Studies that have taken place since the 1985 survey have gone a long way towards answering many of these questions, using new data, new empirical methods, and new theoretical approaches. We will discuss those studies that have advanced our understanding in this area, or that raise new questions about our understanding. We begin with a discussion of measurement issues that pertain to evaluating the returns to technology adoption, and then go on to discuss the role of learning, individual and group, the role of education, and the roles of operational

¹Such direct evidence for under-investment in schooling in poor countries is similarly lacking, but given the possible complementarities between schooling and technology and its change, understanding the barriers to technology adoption may provide insights into the importance of schooling as a determinant of growth in low-income countries. We discuss this link below.

²The data are from the sixth round of the Rural Economic and Development Survey (REDS), which is a probability sample of rural households in the major states of India.

scale, credit markets and insurance markets in explaining the wealth-adoption relationship. We also discuss recent studies that explore non-standard models of human behavior, and end with our conclusions about what we think we have learned and where we need to learn more.

2. Returns, Input Use and Adoption

a. Measurement issues.

An important determinant of the adoption of a new technology is the net gain to the agent from adoption, inclusive of all costs of using the new technology. Under-adoption is defined as a situation in which there are substantial unrealized gains to the use of a new technology or expansion of input use. It is thus generally reflected in a high return to adoption or input use at the relevant margin. Measures of the marginal return to input use or a marginal expansion in technology are thus informative about whether there are market or other problems that constrain adoption. Measurement of outcomes is also a prerequisite for assessing to what extent agents are responsive to variation in the returns to the use of inputs or technologies. Measurement of outcomes, however, is not straightforward.

In the case of technologies used by profit-maximizing entities, it is clear that technology profitability is the key measure. For technologies that improve an agent's utility, such as those that improve health, measurement of returns is less straightforward. Agents choose to use a technology based on the gain in welfare, which cannot be directly measured. In the case of the adoption of contraceptive technologies, for example, the return depends importantly on couples' preferences for family size (Rosenzweig and Schultz, 1989) or social norms about family size (Munshi and Molyneaux, 2006). For medical technologies such as improved bed nets, curative pills or water purifiers, adoption will depend on how agents value health and other attributes of the technology (e.g., taste, side-effects, style), which will depend on both preferences and on the returns to health in the economy. In Miguel and Kremer's (2004) study of the adoption of wormicide pills among school-age children in Kenya, for example, school attendance and scholastic test scores are used as indirect measures of the gains from pill use. However, these may understate the utility gains. First, to the extent that the pills increase vigor, pill adoption for a child will also raise the return to activities outside of school (like working or playing) and thus may increase the opportunity cost of schooling. In that case schooling may increase or decrease even when pill use improves health and welfare. Second, schooling, even if it increases, may not be efficacious in

increasing learning (or test scores measure learning poorly) and/or the returns to schooling in the labor market may be low. In fact, while Miguel and Kremer do find that schooling time is increased, test scores do not rise. More importantly, Miguel and Kremer find in their follow-up study (2007) that use of the pill declined with increased knowledge about it. Thus, although the pill clearly is effective in reducing worm infection, the net private gain in utility to the children was evidently not high.³ What we cannot know, given the difficulty of measuring outcomes, is why.

Even in the case of technology used by profit-maximizing entities, there are few studies that carefully estimate the returns to profits arising from increased input use or from the adoption of new technologies. There are two problems. The first is that profits, while conceptually straightforward, are not easy to measure. Information is needed on the costs of all inputs, but data on many inputs and the relevant cost of these inputs is not easy to collect. Typically in studies of farms, information is obtained on paid-out input costs, such as for seeds and fertilizer, but there is rarely information on labor use, particularly use of family labor that dominates labor inputs in many low-income countries. In Duflo *et al.* (2008), for example, the “returns” to fertilizer use from field experiments in which farmers in Kenya were randomly assigned fertilizer amounts were based solely on measures of crop output, not farmer or plot-specific profits. Individual farmer or plot data on labor inputs, for example, were incompletely collected. While the authors report that there was no increase in weeding labor based on informant observation, it is not possible to have increased crop output without at least some increase in harvest labor, which was not measured. If fertilizer usage did increase harvest labor, the returns to fertilizer use in terms of farm profits are biased upward by their output measure.⁴

b. Optimal input use and the returns to inputs: heterogeneity and perfect input

³This finding may be due to the fact that the returns to health in either the labor market or in school may be low in the specific context in which the field experiment was carried out.

⁴The increase in yields from the small-dosage fertilizer treatments was in fact small, and the associated increments in complementary inputs may have been undetectable without more resource-intensive survey methods. If, on the other hand, harvest labor in fact did not increase because, say, of labor market barriers to the use of hired labor, then the returns to fertilizer use measured in terms of crop output value may understate the returns to profitability if the labor market were to be more efficient.

markets

The second problem in inferring the returns to technology adoption, or its associated inputs, given correctly-measured outcomes, is that adoption and input use are the outcomes of optimizing by heterogeneous agents. In particular, it cannot be inferred from the observation that farmers using high levels of fertilizer earn substantially higher profits than farmers who use little fertilizer that more farmers should use more fertilizer. Consider first the farmer problem of optimal input use for a given technology θ , which describes a concave mapping from inputs of fertilizer (f), labor (l) into an output good y_{it} at location i for crop season t on land of quality u_i

$$(1) \quad y_{it} = g_{\theta}(f, l, u_{it}),$$

where u_{it} is a time (season)- and location-specific exogenous environmental variable that affects output. We assume that markets are well-functioning in that each farmer can purchase the amount of fertilizer or labor he wants at a given price that is invariant to quantity.⁵

Equation (1) defines the technology-specific profit function, which pertains to a setting in which agents maximize profits within “perfect” input markets and access to credit at rate $\rho - 1$ conditional on a technology with known u_{it}

$$(2) \quad \pi_{\theta}(p_f, p_l) = \max_{f, l} g_{\theta}(f, l, u_{it}) - \rho p_f f - \rho p_l l,$$

Profit maximization yields the standard result that the marginal contribution of both inputs to discounted output value is just equal to their marginal cost (price), or the marginal returns to profits from increasing the value of each input is zero. Thus, the marginal returns to profits from, say, fertilizer use will be the same across all farmers. Fertilizer use and the average returns to fertilizer use, however, will differ across farmers in a given season, varying in particular with the u_{it} . If u_{it} and fertilizer are complements (for each level of f the marginal product of fertilizer is higher on “better” land), for example, then there will be a positive correlation between average profit returns and fertilizer use. But the cross-sectional variation between fertilizer use and average profits does not identify the marginal returns to fertilizer use and therefore whether fertilizer is under-used in a setting in which prices do not vary. With little price variation and substantial farm heterogeneity it

⁵The price faced by the farmer may be subsidized. What is important for the implications of the model is that he not be quantity constrained.

is difficult to identify the returns to an input from cross-sectional variation in input use and farm profits even though in such a setting there will be variation in input use across farmers.

One natural solution to the problem of inferring returns in the presence of farm heterogeneity is to exploit observations on farmers using different levels of inputs at different points in time on the same land, say, as a result of changes in input costs. There are two problems with using panel data to infer returns. First, the environmental variable u_{it} may vary over time. In particular, it is important to recognize that a point estimate of the *ex post* profitability of input use, even if profits and inputs are extraordinarily well measured, does not necessarily reflect the information that was available to the farmer at the time that an adoption decision was made. If the econometrician does not measure u_{it} but the farmer observes it and acts on it then again the profitability of inputs use will be mis-measured.

The second problem in inferring profit returns using panel data arises if credit markets (and insurance markets) are imperfect so that lagged shocks to profits affect current input choices. We consider the issue of credit and insurance markets in more detail below, but here we examine its implications for inferring the returns to inputs from panel information. In particular, decompose u_{it} into two additive components: $u_i + e_{it}$, where e_{it} is the time varying component that is also specific to the farmer or plot. Then, the relationship between the difference in profits $\Delta\pi_{it}$ over time for the same land (and farmer) and changes in input use is given by:

$$(3) \quad \Delta\pi_{it} = \beta_f \Delta f_{it} + \beta_{pf} \Delta p_{ft} + \beta_{plf} \Delta p_{lt} + \Delta e_{it} + \Delta \varepsilon_{it},$$

where ε_{it} is an additional shock to profits in period t that occurs *ex post*. If farmers know e_{it} and that affects the returns to the input then there will be covariation between input use and the compound error term in the differenced equation, causing bias in the coefficient measuring the effect of input use on profits β_f . The bias cannot be signed: if the contemporaneous *ex ante* shock observed by the farmer in making his input allocation decision is complementary with (substitute for) the input, then the bias will be positive (negative); the bias arising from credit constraints - positive shocks to profits in the previous period increase input use in the current period - is negative.

Given all of these problems in inference associated with estimating the returns to input use, the field experiment carried out by Duflo *et al.* (2008) among Kenyan farmers in which fertilizer was allocated randomly across farmers is of interest because it creates variation in input use that is orthogonal to land and farmer quality as well as to time-varying profit shocks. However, as noted, the

measure of outcomes in this study is not profits, so it is difficult to know if the correctly-measured returns (increases in output value minus all input costs) to marginal increases in fertilizer are in fact high. Although it is possible in that study that labor inputs did not detectably increase as a consequence of fertilizer adoption because of the low yield increases, it is useful to consider the general question of whether the neglect of input costs matter for inference about the returns to inputs?⁶ To answer this question requires either a repetition of the experiment with better measures of inputs or a credible method of inferring returns using observational data that deals with the inference problems discussed.

Foster and Rosenzweig (2009) exploit plot-specific data on crop yields and inputs for each of three seasons for the farmers in the sixth round of the REDS. In particular, they estimate the returns to fertilizer use by estimating the following equation for a farmer j who farms multiple plots indexed by i across seasons indexed by t :

$$(4) \quad \Delta\pi_{ijt} = \sum \beta_{fx} \Delta f_{ijxt} + \mu_{jt} + \Delta \varepsilon_{ijt},$$

where $\Delta\pi_{ijt}$ are per-acre outcome measures, including profits and the Δf_{ijxt} are dummy variables indicating intervals x of fertilizer use per-acre and the β_{fx} are the associated interval-specific coefficients. More importantly, equation (4) also differs from (3) because the dummy variable μ_{jt} - the interaction between a farmer fixed effect and a season fixed effect - absorbs any differences over time in the farmer's lagged profits that may constrain input use, any differences in *ex ante* farm-level shocks that may influence the choice of fertilizer, any differences in farm input prices over time, including farmer-specific price differences. Of course the differencing by plot eliminates the influence of heterogeneity in plot characteristics for a given farmer on input choice. Essentially this method exploits the remaining random variation in fertilizer use due to random mis-measurement of appropriate inputs by the farmer - variation uninfluenced by profit variation, input choices or *ex ante* shocks to profits.

The set of estimated coefficients describing the relationship between fertilizer use per acre and per-acre farm outcomes at the plot level; measured by "true" profits - output value minus all input costs inclusive of all labor costs, family or hired; profits gross of all but family labor costs; profits gross of all labor costs; and crop output value (no input costs); is displayed in Figure 1. As can be seen evidently some farmers for some plots/seasons used too much fertilizer, as the

⁶It is also possible that the Kenyan farmers did not appropriately increase labor effort to fully exploit the gains from fertilizer use.

estimates identify a profit maximum, at around 400 kilograms per acre. Most farmers, however are using fertilizer below this level, at around 250 kilograms.⁷ The estimates also indicate, however, that the measures of outcomes that do not completely account for labor costs, such as those used by Duflo *et al.* (2008) indicate average and marginal returns and the optimal use of fertilizer (600 kg per acre) that are much higher than those indicated by the outcome measure that nets out all costs. Evidently in India, labor and fertilizer use are strong complements, and labor costs are major component of profits.

c. Optimal technology choice, heterogeneity and the returns to technology adoption.

Analogous and additional problems afflict inferences about the returns to technology when there is heterogeneity in land across farmers. The choice of technology for each given location is described by the problem:

$$(5) \quad \theta_i = \arg \max_{\theta} \pi_{\theta}(p_f, l_f)$$

Consider a farmer with multiple plots of heterogeneous land deciding on how much of each plot will be allocated to a new technology seeds. A convenient specification is to assume for illustrative purposes assume there are two technologies $\theta = \{0, 1\}$, technology is linear in u_{it} with coefficient a_{θ} ,

$$g_{\theta}(f, l) + a_{\theta}u_{it} ,$$

and that farm productivity is uniformly distribution over the interval $[0, K]$ and ordered such that for a farmer with total area A , $u_i = iK / A$. A profit-maximizing farmer will use the same inputs on all plots using the same technology. Furthermore if he plants technology in place i he will plant that technology at place $j > i$. Thus his profit maximization problem if he is unconstrained in terms of access to fertilizer and we normalize $a_{\theta} = 0$ can be written

⁷These levels of fertilizer are substantially higher than current recommended levels of Urea, for example, of around 100kg per acre for particular grain crops that come from experimental studies. It is important to note, however, that the measure available in this data set includes all types of fertilizer, not just Urea, and that Foster and Rosenzweig are looking at overall profitability and thus internalizing the farmers decision with regard to what crop to choose rather than determining the profit maximizing level of fertilizer for a given crop.

$$\begin{aligned}
(6) \quad & \max_{h_1, f, l} \int_0^{A-h_1} (g_0(f_0, l_0) - p_f f_0 - p_l l_0) di + \int_{A-h_1}^A (g_1(f_1, l_{j1}) + a_1 i K / A - p_f f_1 - p_l l_1) di \\
& = \max_{h_1} (A - h_1) \pi_0(p_f, p_l) + h_1 \pi_1(p_f, p_l) + a_1 K (h_1 - \frac{h_1^2}{2A})
\end{aligned}$$

subject to the constraint that $h_1 \in [0, A]$. Note that given these assumptions the maximizing value of h_1 is proportional to area so all farmers would denote the same fraction of acreage to the new technology if there was no farmer-specific heterogeneity.

If there is heterogeneity across individuals in the returns to different technologies and these returns are not be easily measured by the econometrician then, for example, the finding that profits among those farmers who use a high yielding variety of a crop are substantially higher than profits among those farmers who use a traditional variety need not imply that the farmers planting traditional varieties are acting in a manner inconsistent with profit maximization or are otherwise constrained. It may simply be the case that some farmers have land (or other attributes) that are well-suited to the new variety and other farmers have land that is not well-suited to this variety. Thus in order to test whether technological choice is importantly determined by the relative profitability of the different varieties for a particular farmer it is necessary to know how profitable each technology is for that farmer on a given plot of land. But this information is not, in general, easy to obtain because at any given time farmers will only be using one technology on a given plot of land.

The problem of inferring the returns to a new technology will thus depend on how sensitive the returns to productivity are to difficult-to-measure variables like weather and soil and how variable such conditions are in the setting studied. Munshi (2004) shows that in the early stages of the green revolution in India, HYV rice was more sensitive than HYV wheat, and that rice regions were also more heterogeneous in growing conditions. He also finds that HYV rice was more slowly adopted than HYV wheat, presumably because of the difficulty that farmers had in inferring returns. We discuss learning models in more detail below.

One natural solution to the problem of inferring returns in the presence of farm heterogeneity is to observe farmers using different technologies at different points in time on the same land, say, as a result of changes in the cost of or access to the new technologies, as discussed for fertilizer. However, in the case of technology, this approach creates a new problem because unless the technology cost shocks are very large, only a small subset of farmers are likely to use different

technologies at different points of time and these farmers are themselves an importantly selected sample—the sample of those farmers for whom the differences in profitability of the two technologies happen to not be very large. Without imposing some additional structure it is thus impossible to assess whether those farmers, for example, who never adopt the newer technology at all are not doing so because for them the traditional technologies are more profitable.

In a recent paper, Suri (2009) tackles this problem using an approach developed to examine the operation of comparative advantage in the labor market. The basic idea is to use information on the relationship between differential productivity and the productivity of a technology among those farmers who end up using both technologies to project the differential productivity for those farmers who use only one technology. Formally, it is helpful take our basic technological model and assume that the land-specific characteristic u_i has two possibly correlated dimensions u_{1i} and u_{2i} determining productivity in the traditional and modern technologies, respectively, and that profitability is additive in the respective term:

$$(7) \quad y_{i\theta} = g_{\theta}(f, l) + u_{\theta i} .$$

It is then straightforward to decompose the productivity terms into two additive uncorrelated components

$$(8) \quad u_{1i} = v_{1i} + v_{2i}$$

$$(9) \quad u_{2i} = (1 + \phi)v_{1i} + v_{2i}$$

where the second term reflects the common component of productivity and the first term reflects the “comparative advantage” part. Thus if ϕ is positive, the difference $u_{2i} - u_{1i}$ will be positively correlated with productivity using the traditional variety and thus those areas with higher traditional crop productivity will also have a larger differential benefit of using the new technology than those with low traditional crop productivity.⁸

Because by construction v_{2i} and the difference $u_{1i} - u_{2i}$ are uncorrelated, a regression of u_{2i} on $u_{2i} - u_{1i}$ within the population of those using both technologies yields a consistent estimate of ϕ , it is possible to construct a noisy but unbiased estimate of, for example, u_{2i} , among those using only technology 1, as long as the decomposition reasonably reflects the true underlying data generating process. These differential productivity measures can then in principle be related to farmer

⁸This ignores differences in input costs across the technologies, which we can do under this structure because of the additive errors and the assumption of optimizing behavior.

choices and to measures that may influence the cost of adoption to determine whether in fact choices are consistent with a model of profit maximizing behavior. Suri in fact concludes that there are three sets of farmers: a set of farmers for which there are small differences in the profitability of traditional and modern varieties who end up using both technologies, a set of farmers who have high returns to the modern variety but who do not adopt due to the difficulty of accessing the new technologies, and a third set with moderate returns to the new technology that always adopt it.

While this approach is simple enough in principle, there are important complications in practice. In particular, productivity shocks are not likely to be additive with respect to inputs as assumed above. This means that the process of extracting the u_{θ_i} from profit data is not straightforward. Suri makes use of information on fertilizer as an input but assumes that labor does not vary across traditional and high-yielding varieties. But, as noted, at the very least harvest labor must be higher if output is higher and this will create a wedge between yield differentials and profitability differentials. Second, the approach ignores the possibility that there are *ex ante* shocks, specifically in this case technology-specific shocks to profits influencing technology choice. Third, the approach assumes that the relationship between differential productivity and the level of productivity is consistent across farmers. It is certainly restrictive to assume that the joint distribution over the u_{2i} and u_{1i} is such that the v_{1i} and v_{2i} will be independent for any given φ in both the set of farmers who adopt both technologies and those who adopt only one—though as a first-order approximation this may be reasonable. Nonetheless, the paper makes an important advance on the existing literature in terms addressing the problem of heterogeneity in returns in trying to draw inferences about the process of technological adoption.

d. Do estimated returns indicate that farmer's are under-investing?

We have seen that the experimental evidence of Duflo *et al.* (2008) is interpreted by the authors as suggesting that the returns to small quantities of fertilizer are high in Kenya, although these returns may be over-estimated due to the lack of cost data. Thus they conclude that in Kenya fertilizer use is too low. Duflo *et al.* also find that there is variation in returns across farmers residing in different regions, as they measure them, but they do not assess whether differences in returns are related to actual input use. Suri's data indicate that some farmers with high returns to adopting hybrid seeds do not adopt, and her methods suggest that some of this is due to supply constraints associated with poor infrastructure.

The Kenyan environment, from which both Suri's and Duflo *et al.*'s data are from, is one in which the technology has been stable for some years, so that any under-investment in technology or inputs is not likely to reflect lack of knowledge about returns. And both studies find no evidence consistent with learning. Udry's examination of fertilizer use among Ghanaian pineapple farmers, which shows farmers switching in and out of fertilizer use in response to new information about profits, clearly shows that expected profitability also affects input use. However, the estimated profitability of fertilizer in the Ghanaian environment net of costs is low, suggesting that on average under-investment in fertilizer is not high, with learning costs actually playing only a small role in constraining fertilizer use.

3. Learning and Technology Adoption

a. Definition of learning.

Under-investment in an input or a new technology could arise, when true returns are high, because of ignorance about returns or about how to manage the technology in order to receive high returns. This is more likely a cause of under-investment in a setting in which a technology is new. We define learning as taking place when new information affects behavior and results in outcomes for an individual that are closer to the (private) optimum.⁹ Thus, in an environment where there is no new information, learning is unlikely, while in a setting in which a new technology or input is introduced, learning should be important. The finding of Duflo *et al.* that farmers do not obtain information from their neighbors about fertilizer use in Kenya is not evidence contradicting models in which agents learn from their neighbors because it is an important assumption of learning models that there be something new to learn, and the existing technology in Kenya at the time of the study was not new. Thus, this evidence supports learning models! However, the authors find that there is learning associated with own input use - farmers who adopted fertilizer in the first round of the experiment were more likely to use fertilizer in subsequent rounds. It is important to note that the definition of learning does not imply that learning increases the use of an input or a new technology. It may be that what is learned is that the new technology is in fact not efficacious. Thus, if more experience with a new technology leads to less use, as in Miguel and

⁹It is possible that learning leads to a move away from a social optimum. Agents may learn that free riding is optimal, as is consistent with the findings in Miguel and Kremer (2007), described below.

Kremer's (2007) follow-up study of pill use, that is also evidence in favor of learning.¹⁰

Learning may not be important for all new technologies - some technologies are simple to learn, others not. The complexity of new technology matters as well. Thus, technologies like the contraceptive pill, which simplified contraception relative to traditional methods, were very rapidly adopted in the United States by all couples with small family goals (Rosenzweig and Schultz, 1987). The salience of learning, reflected in slower take up (or discarding) of a new technology is thus very context-specific. Finally, learning may depend importantly how technological returns vary with individual attributes and what is known about the structure of this relationship by those considering adopting the technology. Thus the fact that some individuals adopt a technology and others do not is not necessarily evidence that learning effects are not important.

The "green revolution" produced new, high yielding variety (HYV) seeds that were more sensitive to soil and water conditions in the initial years compared with traditional seeds. Thus, farmers in the early stages of the green revolution were faced with both a potentially more profitable but more complex technology with an uncertain return. Moreover, new seeds, with different properties, are marketed almost continuously in many areas of the world, so that learning may be an important component of seed adoption in agriculture contemporaneously. Thus, many investigators have studied the take-up of HYV seeds both in the early stages of the green revolution and subsequently using learning models. In the early stages of the revolution, a farmer's choice was essentially between HYV or traditional seeds. In the current period, many farmers choose among different vintages of HYV seeds. The choice is whether to adopt the newest seed or one that has been well-tested by actual farmers. Most studies, however, tend to look at seed adoption choice as HYV versus traditional, rather than the choice of seed vintage.

b. Individual learning

Learning about the returns to new technologies and their associated inputs can be captured with the above specification with the additional assumption that the u_{it} are generated by a distribution that is fixed up some unknown (to potential adopters) parameters. Because realizations of u at time t can be used to draw

¹⁰Some of the learning may have been about the returns to free riding due to health externalities.

inferences about the unknown parameters, say through a process of Bayesian updating, past use of the technology provides a basis on which to better forecast the u at time $t+1$ and thus make more profitable choices with respect to technology and/or input use at that time. Papers on learning differ in terms of how the u_{it} enter the production function, the parametric structure of the underlying distribution, and the extent to which information that is acquired is specific to a particular agent.

In a learning-about-productivity model it is assumed that individuals learn about the overall profitability of a new technology and compare this to the profitability of the existing technology that is well established. This is the approach used, for example, by Munshi (2004) in his study of individual versus social learning in the context of agriculture and by Besley and Case (1992) in their study of HYV cotton. Consider, for example, a profit function generated from the simple quasi-linear production function applied to a given land area A ,

$(\pi_0 + a_0 u_i)A$, with u_i fixed and known to be drawn from a

distribution $N(\mu, \sigma_u^2)$ and $a_0=0$. Under these assumptions, and if u_i were known, the new technology ($\theta=1$) would be chosen if

$$(10) \quad u_i > (\pi_0 - \pi_1) / a_1.$$

In this simple model, one obtains the result that new technology will be overutilized initially because of ignorance. Given that u_i is not known the new technology will be used, assuming expected profit maximization and that the priors over u_i reflect the true data generating process, only if $\mu > (\pi_0 - \pi_1) / a_1$. Thus for all values of u_i between $(\pi_0 - \pi_1) / a_1$ and μ there will be a loss of profits associated with the use of a technology that is in fact relatively unprofitable.

We now introduce learning. With learning, one again obtains the results that a new technology may be adopted by too many farmers. Let π_L denote the expected loss per unit of land in profit from not knowing the true value of u_i . The information technology is such that in the first period if the farmer plants the new technology to at least h_1 units of land he will know u_i for sure but that otherwise he will not gain any new information about the new technology. Ignoring discounting he will adopt the new technology if $\mu > (\pi_0 h_1 - \pi_1 h_1 - \pi_L A) / a_1$.

This expression has several implications. First, the presence of learning means that some farmers may adopt the technology in the first period who would have found it unprofitable to do so in the absence of the second period. In expected value they lose money in the first period through the process of expectation more than makes up for it in the second. Second, it is possible that some farmers for

whom the new technology is profitable given u_i will not plant the crop at all because their priors are such that they do not expect to receive sufficiently “good news” from their experience with the new technology. Third, there are likely to be important scale effects associated with learning. A large farmer faces the same cost of learning but, in the event he receives good news about the new technology he can adopt that new technology on a larger scale and thus receive higher expected profits. Thus, large farmers are more likely to adopt a new technology initially independent of any relationship between landholding and costs of inputs.

But uncertainty about the profitability of a new technology is not the only challenge overcome by learning. Learning may also involve acquiring information about how to optimally manage the new technology. Foster and Rosenzweig (1996) argue that in fact this idea has particular salience in the context of agriculture. Agricultural research organization and extension agents carry out controlled experiments on new seeds and can thus determine the maximal possible yields and even, for given set of prices, maximal profitability. What they cannot necessarily do is provide information on how best to achieve these yields given the specific characteristics of the soil and climate of a particular farmer. They argue, for example, that the optimal level of fertilizer use may depend on the nitrogen content of the soil as well as permeability and rainfall that may be specific to a particular plot. As such a farmer may have to experiment with a crop on his own land in order to sort out how much fertilizer to use.

A simple and analytically convenient implementation of this idea used by Foster and Rosenzweig is the target-input model in which the u_i enters the production function in the following fashion:

$$(11) \quad g_\theta(l) - a_\theta(u_i - f_i)^2$$

The basic point here is that maximal yields are achieved for given l if $f_i = u_i$. If the u_i is known then that presents no difficulty. But if u_i is not known then the farmer will on average tend to miss the target and thus receive less than optimal profits. By learning about u_i the farmer is able to better target fertilizer use to the conditions on his particular land and thus receive better outcomes. Assuming that the traditional technology does not use fertilizer or is sufficiently well established so that it can be properly managed, this specification generates many of the same predictions as the learning-about-productivity model. In particular, a farmer may have incentive to try out the new technology using his best guess of u_i even if he loses money at first in doing so. He also may not experiment with the technology at all even though he knows for sure that if he were to learn to properly manage the new technology it would be more profitable than would the older technology: the

short-term cost may outweigh the long term benefit.

The model is easily adapted to cases in which the *ex post* optimal u_i varies randomly over space and across time around some mean, so that the farmer will not be able to fully determine the optimal use from a single realization of u_i and thus will need to aggregate information across space and/or time in order to determine the optimal level of input use. It also creates a relatively straightforward way of thinking about the process of learning from the experience of neighbors, which has been a prominent focus of the empirical literature on learning as discussed below.

An attractive feature of the target-input model, at least in the special case where the cost of the input is zero, is that future profits do not depend on the revealed values of the farm-specific parameters.¹¹ This means that a farmer will know for sure the relationship between experimentation and his future profits. Thus it is possible in the learning by doing model to write down formal testable implications for the relationship between past experimentation and current profitability. These propositions are formally tested in Foster and Rosenzweig. In particular, they find evidence that the profitability of the new technology is increasing but concave in past experience. Note that this need not be the case in a learning about profitability model. Experimentation in that case may influence adoption but it should not affect profitability given adoption. Similarly, the learning by doing model provides predictions about how experimentation should be related to adoption, but the learning about productivity model does not necessarily do so. In the latter case experimentation may lead one to conclude that a new technology is inferior and thus lead to lower adoption.

c. Learning from others

As noted, a primary focus of the learning literature has been on the issue of whether and by how much agents learn from others. Although learning from others clearly facilitates the acquisition of knowledge compared to a world in which one has to learn only from own experience, such learning externalities can give rise to sub-optimal adoption of new technologies. In Foster and Rosenzweig (1995) it is assumed that the input target is the same on every farmer's land but the actual input decision of neighboring farmers is observed with error so that a given neighbor's experience contributes less information than does one's own

¹¹If inputs are costly then the realized value of u_i does affect future profitability because it affects the cost of achieving profit-maximizing input use.

experience. On the other hand it is recognized that very farmer has multiple neighbors so that the overall effect of average neighbors' experience could be greater or less than own experience in terms of area planted to the new technology. Using a data set collected in rural India starting from the time that high-yielding variety seeds were first introduced into the country, they establish that, as with own experience, the profitability of the new technology is increasing in neighbors' experience at a decreasing rate. Consistent with the idea that own and neighbors' experience are substitutable they further establish that the rate of decrease in returns to experimentation is the same for oneself and one's neighbors. They also show that experimentation both by oneself and one's neighbors increase adoption.

The fact that own and neighbors' experience are substitutes creates the potential for free riding behavior. In particular, a farmer who knows his neighbor is likely to experiment with the new variety may have an incentive to reduce his own experimentation and then benefit from the increased information. Obviously this problem might be overcome if farmers can license or otherwise market their information to others or if other institutions are in place to reward those farmers who bear a disproportionate share of the cost of experimentation. But whether or not this internalization of the potential externality takes place may be critical for policy. If there is free-riding behavior there may be inefficient under-provision of information and even, in the extreme case, non-adoption of a new technology that would be profitable, inclusive of the cost of experimentation, from a social perspective. To test for the presence of free riding behavior, Foster and Rosenzweig make use of the idea suggested above that the returns to experimentation are increasing in scale when this experimentation is not specific. Given within-farmer heterogeneity in the suitability of the crop for the new technology, these scale effects imply not only that farmers with large land area and other fixed assets will experiment more but also that, from a social planner's perspective, those farmers with neighbors with larger operational scale should experiment more. In the absence of coordination, however, a farmer will know that his larger neighbor has a private incentive to experiment more. Thus as long as costs of experimentation are sufficiently high, a farmer "close to" a larger farmer will under-adopt a new technology, conditional on own and neighbors' experience.

The empirical evidence presented by Foster and Rosenzweig confirms the presence of free-riding behavior and thus may provide a case for socially optimal subsidy of experimentation on new technologies at the level of the village. However, it is important to point out that in fact the presence of learning spillovers, particularly in the case of heterogeneity in land ownership, also means that the diffusion of a new technology once adopted by the largest farmers, can be quite

rapid. In short, larger farmers with the highest incentive to experiment do so and then this information is transferred to the smaller farmers who then adopt without bearing the full cost of experimentation. In the particular case of the early stages of the green revolution in India, the estimates from Foster and Rosenzweig suggest that diffusion of information about the new technologies within a given village is more or less complete in 3-4 years following the first adoption of that new technology.

The Foster and Rosenzweig learning study focuses on a model in which farmers benefit from the cumulative experience of all farmers in the village. The key assumptions that deliver this result are that the optimal management of the new technology does not vary much within the village; that information flow within the village is not importantly constrained by networks based on kin or social status, and that individual farmers have a good sense of the structure of the technology. The first seems plausible to the extent that soil and moisture conditions typically do not vary substantially within a village. The second seems more restrictive particularly in the light of recent work showing the importance of caste networks in determining, for example, access to credit (Munshi and Rosenzweig, 2009). However, issues of trust and reciprocity that may be critical in the context of credit provision would seem to be less important in the context of transferring information about agricultural input use. Nonetheless it would be useful to know whether learning about agricultural technology flows more effectively within caste or other groups.

The third assumption about knowledge of the true structure would seem to be of particular interest. A key assumption of the target-input model is that farmers have a full understanding of the relationship between input use and profits given the unobservable u_{it} and they can make use of this understanding to back out estimates of u_{it} from their experience. A surprising implication of this model, taken literally, is that the information that can be gleaned from the planting of the new technology is the same regardless of the level of input actually used. Thus, one can learn as much from acreage planted to the new technology with good information as one can from farmers with relatively weak information. While this is analytically convenient as noted, one might imagine a more complicated signal extraction problem in which there is an unknown non-parametric relationship between input use and profitability. This idea that farmers “know the structure” may not seem particularly limiting in the context of inputs like pesticides or fertilizer where one would expect to see a smooth relationship between inputs and outputs. It may be quite unreasonable however in a case in which individuals are choosing among a variety of high-yielding seeds with, say, differential resistance

to different types of pests. In this case experience with one seed may have little relevance for experience with other seeds. The key prediction of such a model is that one is likely to try and replicate the input use of those who were relatively successful and to not make the same choices as those who are not.

Udry and Conley (forthcoming) examines a model of this form in the context of pineapple farmers in Ghana. The basic proposition that is tested is that when learning about a new technology farmers will adopt the behavior of those farmers who were unexpectedly successful, in the sense that they had high profits given other observables that influence profitability. Despite the similarity of the question being asked, it is notable that the data used for this study are quite different in scope than the data used by Foster and Rosenzweig's study of Indian farmers. Instead of being a large nationally representative survey in which the village provides an important source of variation, Udry and Conley look at a small number of villages and farmers in which there is more detailed information on the network structure of the village that might influence information flows. Nonetheless, Udry and Conley's results are also supportive of the notion that individuals are learning from the experience of their neighbors. Not only does one see movement towards the fertilizer behavior of those who are successful and away from those who are unsuccessful, but also, as expected, people move toward the input use of more experienced neighbors and are more likely to move if they have little experience of their own. Such patterns are not observed in the case of traditional crops with which most farmers have a great deal of experience. It is unclear, however, whether these results are inconsistent with a more parametric approach to learning. In particular, one would expect in a parametric model for those farmers who have more precise information about the proper management of a technology given local characteristics (i.e., they know u_{it}) to in fact be using inputs that generate higher profits. Thus unless one can fully condition on the information set of the farmers in question one would expect movement towards the behavior of successful farmers under either model.

Bandiera and Rasul (2006) also make use of a data set involving the adoption of a new technology, in this case sunflowers, in which there is more detailed information available on social and other networks that may influence the flow of information than was available for Foster and Rosenzweig. They explicitly focus on the question of how adoption varies by network structure. Their work highlights the point made above that, in the presence of free-riding effects, social network effects on adoption may be positive as well as negative. In particular, they conclude that when relatively few people in one's network have adopted the marginal effect of increased adoption by one's neighbor is positive. However,

when a large number of people in one's network have adopted the marginal effects may be negative as there is less incentive for an individual to undertake costly experimentation on his own. The results from their analysis during the first-year that the new technology was introduced do show this u-shaped effect. However, as the authors recognize, there are significant problems of inference that arise given that they cannot take advantage of dynamic effects. In particular, there may be important unobservables that are common within network groups, and network groups may themselves be importantly constructed with respect to willingness to adopt new technologies. Arguably these effects contribute to the estimates of the positive "effects" of adoption, but it is difficult to see how these effects could also create the negative marginal effects at high levels.

While the models of Udry and Conley, Foster and Rosenzweig, and Bandiera and Rasul assume that information about a new technology is largely non-specific, at least within the village, Munshi considers the possibility that different technologies have different degrees of specificity. In particular, if there is variation across farmers in the u_{it} then the acquisition of information by one farmer need not be useful to his neighbor. Of course to the extent that the differences can be predicted by observable characteristics, this need not be too much of an obstacle. A farmer may know, for example, that the yield of a new technology is dependent on the porosity of the soil. As such in determining whether a neighbor's high yields of a new technology are likely to be obtainable, he will try in some way to make adjustment for differences in porosity. But in cases where there are importance unobservable differences in farm or farmer characteristics he may not use that information at all and thus will have to rely disproportionately on his own experimentation. As noted, Munshi (2004) shows, in particular, that rice and wheat are quite different in this regard – yields of new wheat technologies are similar across farmers while yields on new rice technologies tended to be quite dependent on local conditions. As result one would expect stronger evidence of learning from neighboring regions and/or farmers in the case of wheat relative to rice. This contrast is in fact clearly evident in the data.

In an early draft of this paper Munshi addressed the issue raised above about learning with and without a knowledge of the structural relationship between farm attributes and profitability of the new technologies. In particular, he argued that when farmers have a good sense of the underlying structure, the experience of farmers who were quite different from them in terms of observable attributes could be useful in terms of predicting yields of a new technology for that farmer. On the other hand, if a farmer has little sense of the structure relating yields to these attributes then he can only learn from those who share similar characteristics. For

example, it may be possible to make some adjustment for the permeability of one's own soil by comparing the yields of farmers on permeable soil with those on soil that is not permeable. But if one has red soil than it is not clear what can be learned by comparing yields of those with black soil and those with sandy soil.

A common theme of this literature on learning from others is the difficulty of inference that arises when predicting adoption behavior based on the adoption decisions of one's neighbors because of the presence of common unobservables that jointly affect everyone's decisions (Manski, 1993). Dynamic data can be quite useful in this regard by allowing one to trace out a sequence of decisions over time. Another way to meet this challenge, at least on a small scale, is to exogenously and randomly induce some group of individuals to adopt a new technology and then determine if others who in some way are connected to these individuals are subsequently more likely to adopt this new technology. An early application in the economics of a randomized intervention aimed at identifying social learning is Duflo and Saez (2003), which examines the spillovers associated with attendance at a meeting on a tax deferred pension account in the US.

The first application of a randomized design to look directly at technology adoption in a development context is Kremer and Miguel (2007) based on a follow-up to their original experimental intervention (Miguel and Kremer, 2003). In the later study the authors took advantage of the randomized treatment used to evaluate the effects of a school-based de-worming intervention to assess subsequent adoption among students with direct or indirect contact (though attending a school where others have social ties). As discussed above they found those students who had contacts exposed to the de-worming intervention were less likely to use or to continue to use the pill, a result that appears consistent with the presence of learning effects given that the private returns to de-worming, inclusive of costs due to side effects, appear to have been small, particularly given technological externalities, as discussed below. The reduction in pill use over time also suggest that the population was initially overestimating the private returns to pill use, perhaps because of NGO ideology.

A recent paper Dupas (2009) uses data from an experiment with a similar design, describing the outcome of a field experiment in which households were randomly assigned vouchers that allowed them to purchase a high quality bed net at various subsidized rates (Cohen and Dupas, forthcoming). Dupas examined whether households that did not receive the initial subsidy were more likely to purchase the bed nets if they lived near other households that received relatively favorable subsidies. The results show this effect quite clearly and, given the likely sign of the technological externality, as discussed below, may provide the some of

the most compelling evidence to date that learning plays an important role in the adoption of health technologies in a low-income setting. Another example of a similar randomized trial that provides evidence of social learning, by Oster and Thornton (2009), studies the adoption of menstrual cups among school girls in Nepal. Their results suggest that friends of those randomly given access to the cup are more likely to adopt the cup subsequently. However, in this study the authors were able to further show that the take-up effect was driven by communication illuminating how to use the cup, consistent with the learning-by-doing hypothesis, rather than by changes in the value attached to the cup.

d. Technological externalities and learning

As noted, learning externalities can in some cases inhibit individual adoption and experimentation with new technologies, but the presence of learning spillovers that create these externalities can also in the long-term help to ensure the adoption of socially profitable technologies. A potentially more complex case is one in which there are externalities that are technological. There can be either positive externalities - the benefits to an individual of adopting a particular technology are increasing in the fraction of the population using these technologies - or negative externalities - the individual adoption of a technology may be less profitable if a large fraction of people adopt this technology. Formally, given the above structure, the production function might be

$$(12) \quad y_{\theta i} = g_{\theta}(f_{it}, l_{it}, u_{it}, h_{\theta\{-i\}})$$

where $h_{\theta\{-i\}}$ denotes the average adoption of technology θ by farmers other than i .

Technological externalities of this type are uncommon in agriculture, though one example might be in cases where widespread adoption of a particular type of seed in a particular area increases exposure to pests or leads to depletion of local common resources such as ground water. A more likely source of externalities would be those that arise through the prices of inputs or outputs. An influential example of the former comes from Griliches (1957), who argued that differences in market density in different parts of the US lead to differences in the supply of hybrid corn seed supply and thus to different rates of adoption in different regions.

But technological externalities may play an important role in the context of health where certain types of intervention may create herd immunity. Miguel and Kremer (2007) argue that such effects have reduced uptake of a de-worming medicine in Kenya. In this case, while de-worming treatments help protect individual students, it also tends to reduce the exposure of non-treated students in

the same schools or classroom. The net benefits to the individual of the treatment may be small—or even negative given possible unpleasant side effects—as long as the fraction of other students accepting the treatment is reasonably high. Thus it is likely that equilibrium adoption levels will be below a socially efficient levels.¹²

When both technological externalities and learning spillovers are in place it can be difficult to distinguish these two processes in practice, even in the presence of experimental variation. In the Miguel and Kremer study of the consequences of the distribution of de-worming pills (2004, 2007), while individuals learn from the experience of their friends, they are also exposed to these friends physically and this, in turn, should influence their chance of being infected and thus also the incentive to use the medication. Thus the experience of one's friends influences one's behavior directly through the technological externality as well as indirectly through its effect on perceptions about private returns. To break apart these effects would require a setting in which one learns based on the experience of people with whom one does not have direct physical contact. It is thus important to distinguish between studies of learning, particularly in the context of health, in which technological externalities are likely to be small and those in which such externalities may be large. In the case of the Dupas' study of the adoption of higher quality mosquito nets, for example, medical evidence suggests that the overall likelihood that an infected mosquito will bite someone without a high quality net is significantly affected by the fraction of households using a high quality net, so that it is difficult to quantify the amount of learning in that context. However this medical externality is likely to reduce adoption and thus is likely to offset a positive learning effect if the technology is perceived to be advantageous by those who use it. This result contrasts with MK in which the externality is negative but the learning effect may also be negative due to there being unexpectedly low private returns.

4. Education, Learning and Technology Adoption

As noted, a common finding in the adoption literature is that more educated

¹²Externalities that influence adoption in both positive and negative directions may also arise in the case of social preferences. Munshi and Molynaux (2006) show that values regarding family limitation evolve differently within different networks as might be expected if the returns to a given behavior (contraceptive use) are influenced by the fraction of people within one's social network that adopt this behavior.

agents are more likely to adopt new technologies. For example, Skinner and Staiger (2005), who examine the adoption of new, effective technologies across US states over the course of the Twentieth Century, including hybrid corn, *beta*-blockers, tractors, and computers, find that education (measured by high-school enrollments) and measures of social networks were the only variables positively associated with the adoption rates for all four innovations. We also saw that education is positively correlated with HYV seed choice among Indian farmers in 2007 in Table 1. There are three mechanisms that have been hypothesized in the literature to explain the education-adoption link: 1. More educated agents are wealthier, and thus the education-adoption relationship represents an income effect. Most of the descriptive studies linking schooling to adoption, however, include controls for income or wealth, as in Table 1 and in the Skinner and Staiger study. 2. More educated agents have better access to information. 3. More educated agents are better able to learn - to decode new information faster and more efficiently. The third mechanism has been the principal focus of economists. As noted by Nelson and Phelps (1966), the income gap between rich and poor countries could be attenuated if poor countries can catch up to rich countries by adopting new technologies developed in rich countries faster and more efficaciously. Thus, if schooling augments learning, increasing educational levels can be an effective development policy in a world in which there is technological diffusion.

There are a number of testable implications that arise from the hypothesis that the more educated are superior learners. The first is that more educated agents will have higher incomes where there are profitable and complex new technologies to understand. We can modify the profit function (6) for a farmer deciding on how much to plant of a new technology to allow profits under the different technologies to be a function of the farmer's schooling E :

$$(13) \quad \pi_\theta = \max_{h_1} (A - h_1)\pi_0(p_f, p_l, E) + h_1\pi_1(p_f, p_l, E) + a_1K(h_1 - \frac{h_1^2}{2A})$$

If schooling augments profits more under the new technology than under the old technology, then (a) more educated farmers will tend to adopt more of the new technology,

$$(14) \quad dh_1 / dE = (A / aK)[\partial\pi_1 / \partial E - \partial\pi_0 / \partial E],$$

and (b) more educated farmers will have higher earnings where there are

advantageous newer technologies available.

The hypothesis that farmers with more schooling earn more under a new technology regime was first tested by Finis Welch (1970), who found that the relative earnings of more-educated US farmers were higher in areas where there was more farm technology R&D. Foster and Rosenzweig (1996) directly estimated the new- and old-technology profit functions, embedded in (13), to assess whether the returns to profits were higher in areas of India in which advances in agricultural technology were highest after the onset of the Indian green revolution. Using panel data on the profits of farmers from a national probability sample of rural households interviewed in 1971 and in 1982, they found that the differential in profits between illiterate and primary school graduate farmers rose from about 10% prior to the green revolution to as high as 40% in those areas of India, such as the Punjab, where the gains from agricultural technological progress were highest. Both the Welch and Foster and Rosenzweig studies assumed, and found, that labor was spatially immobile, so that it was possible to estimate a relationship between area-specific technology change and wages. Bartel and Lichtenberg (1987), using a three year decadal panel of US manufacturing industries, assume that industrial wages are spatially equalized and instead look at the demand for educated workers by industry. Specifically, they estimated a restricted translog cost function to assess if the demand for more educated workers was higher in those industries using newer technologies, as proxied by the average age of the capital stock. They find that this is the case, although their estimation procedure does not take into account that unobserved shock to the education of the labor force can affect the age of the capital stock.

The finding that more educated workers (farmers) earn more or are greater demand when there is new technology merely indicates that new technology and schooling are complements. The evidence does not necessarily imply that the reason for the higher return is due to the more educated having superior learning skills. Rosenzweig and Foster (1996) use their profit-function estimates of the returns to schooling across districts of India to test if schooling investment responded to the increase in schooling returns. A key result is that school enrollments increased in response to the higher returns to schooling in agriculture only in households with land, thus suggesting that only those making allocative decisions in agriculture benefit from schooling in the high technological change areas.

A more direct way to assess if schooling enhances learning is to estimate the relationship between profitability, education and experience with a technology. As shown in Rosenzweig (1995) using an augmented version of the target-input

Bayesian learning model used in Foster and Rosenzweig (1995) the returns to experience with new technologies should be higher for the more schooled if education enhances learning. However, experience and schooling will be substitutes if schooling merely increases initial knowledge about a technology, due, say, to superior access to external information sources. Using the same three-year panel data on Indian farmers, Rosenzweig estimated how the cumulative planting of HYV seeds affected farm profits over time differentially for primary-schooled and illiterate farmers. The results showed that an additional hectare of prior HYV planting increased per-hectare farm profits 18% more for the educated farmers than for the illiterate farmers in the second round of the survey.

Lleras-Muney and Lichtenberg (1996) use data from a 1997 US sample of individuals that contains information on specific drug purchases to assess the role of education in medication choice. They find that, controlling for income and a large variety of other characteristics, more educated patients are significantly more likely to purchase newer drugs, as indicated by the date of FDA approval. More importantly, they find that the education-newness relationship is significantly greater among those with chronic conditions, as indicated by longer histories of repeated drug purchases. That is, they find that experience with medications for an illness and education are complements, consistent with the learning hypothesis.

A corollary of the assumption that schooling augments learning is that the benefits or effects of schooling will be small in settings where either there are no new technologies or the new technologies are not difficult to decipher, where the returns to learning are low. The absence of an adoption or input use relationship with schooling can thus be evidence in favor of the learning hypothesis. Duflo *et al.* (2008), for example, do not find any effect of education on the use of fertilizer, but this is in a setting in which the agricultural technology is relatively old; the same setting in which they also find that farmers do not discuss fertilizer use with neighbors. Rosenzweig and Schultz's (1989) study of contraception adoption at the onset of the contraceptive revolution shows that, controlling for desired family size, college-educated women were no more likely to adopt the pill or IUD as contraceptive methods than were high school students. They argue that the pill and IUD are simple technologies that do not present a challenge to use effectively. They indeed show that the measured efficaciousness of the new contraceptives also did not differ by schooling. In contrast, the efficacy of the traditional rhythm method was much higher for those women with a college education than who had only completed high school. In this case, the traditional method was even more complex than the new technology, so that the ability to decode information was more advantageous.

5. The Effects of Income on the Adoption of Profitable Innovations: Risk, Credit Constraints and Scale

Although income may affect the demand for technologies that augment health or well-being, the wealth of profit-maximizing enterprises should have no effect on technology adoption if markets are perfect unless, as we have shown, there are fixed costs to technology adoption (scale effects). Given that the costs of many inputs associated with a technology must be paid up front and that the returns to new technologies are uncertain (and may be riskier), imperfections in credit and insurance markets may, however, lead to a result that wealthier agents or agents with steadier alternative income streams are more likely to adopt new technologies, at least initially. Disentangling the effects of scale, credit constraints and absent insurance (combined with risk aversion) is not an easy task, yet it is relevant in formulating policies that facilitate the diffusion of technologies. We discuss studies that address the role of risk and credit constraints as they affect adoption and input use and thus create a link between income and the adoption decision.

a. Risk and insurance

We have already discussed the role of variation in profitability in creating difficulties for assessing the true returns to a new technology and thus its profitability. But variability in the returns to new technologies has also been thought to play an important role in terms of technological adoption because of aversion to risk in contexts in which insurance markets are absent. In the context of agriculture there are three reasons that one might expect new technologies to be riskier. First, the yields of new seed varieties may be more sensitive to weather or other forms of variation than are those of traditional crops. Certainly the first generation of HYV wheat crops were very dependent on having a reliable supply of water over the crop cycle. Second, imperfect knowledge about the input management, as discussed above, may not only lower yields but it may also increase variability. Third, HYV seeds often require more investment, such as in fertilizer, prior to the full realization of uncertainty thus increasing overall risk. Thus, if there is a crop-failure just before the realization of the harvest one can reduce expenses by reducing the labor force used to harvest the crop. One cannot, *ex post*, reduce one's use of fertilizer.

To incorporate this idea into our overall notation, assume that u_{it} is a random variable across space that has a known distribution that is realized after input choices are made and that farmers have concave preferences over risk and no insurance. Consider the linear in shocks technology above and utility $V(\cdot)$, so that

the maximand is

$$(15) \quad \int V((A - h_1)g_0(p_f, p_l) + h_1g_1(p_f, p_l) + h_1a_1u) f(u | h_1) du ,$$

where integration is over the average shock on the land planted to technology 1. Because technology 1 in this case has higher risk by assumption, the farmer faces a tradeoff between risk and return. With quadratic utility and a mean zero shock this may be written as

$$(16) \quad V_0((A - h_1)g_0(p_f, p_l) + h_1g_1(p_f, p_l)) - V_1a_1^2 \text{var}(h_1u).$$

One simple additional assumption is that the u_{it} are the same for all the land of a given farmer, in which case $\text{var}(h_1u) = h_1^2 \text{var}(u_{it})$, so an interior solution for the adoption of the technology is a likely outcome. Alternatively if the u_{it} are sufficiently independent across space so that $\text{var}(h_1u) = h_1 \text{var}(u_{it})$, farmers will specialize in either the old or new technology.

Despite the large empirical literature testing for and rejecting full insurance in the context and low-income countries and the theoretical literature showing how risk can in principle affect agricultural decision making, the literature evaluating the role of risk as a constraint on adoption of new technologies is thin. The likely reason for this is that the key thought experiment involves the question of whether *ceteris paribus* an increase in the *ex ante* risk of adopting a new technology affects adoption. Unless there is reason to believe that the distribution of risk is changing over time or varies across people in the same area in some well-defined way, this precludes the use of estimates of technological adoption rules that include village fixed effects. But, given the inability to use fixed effects, any test of the effects of risk on adoption are not robust to the presence of unobserved endowments (such as land quality in the case of agriculture) that may be related to both risk and the returns to the new technology. The best one can do, in general, is to determine whether households with different abilities to accommodate risk (i.e., through higher wealth) but otherwise similar endowment (i.e., quality of land) are differentially likely to be influenced by risk.

An early example of this approach is Rosenzweig and Binswanger (1993). While these authors do not specifically deal the question of technological adoption per se, they do establish that poor farmers facing increased rainfall variability tend to hold a portfolio that is less influenced by rainfall and as a consequence tend to have lower profits. Wealthy farmers facing varying exposure to risk do not, however, exhibit changing portfolios of investments. Thus to the extent that new technologies are high mean and high variance these results would confirm the presence of a risk-based barrier to technological adoption, with wealthier farmers

more likely to adopt new, riskier (at least initially) technologies. Morduch (1990) also showed, using the same data, that poorer farmers exposed to risk planted less risky crops than wealthier farmers.

In more recent work, Dercon and Christiansen (2008) address directly the question of fertilizer adoption using panel data from Tanzania. They are able to use household fixed effects by constructing a measure of the level of consumption that would obtain if rainfall would be at the 20% level. The idea is in essence that as a household's wealth changes from year to year the consequences of an adverse weather shock changes and this, in turn, may affect willingness to absorb risk. One can think of this as a model identified by an interaction between household wealth and the underlying weather risk, with household wealth changing over time, although of course changes in household wealth may themselves be responsive to technological choices. Moser and Barrett (2008) examine a new rice production technology (SRI) in Madagascar. They take advantage of individual-level variation in exposure to risk by using a measure of whether the household has a stable source of income. This measure significantly predicts adoption and continued use of the new technology.

Given the challenges associated with using natural variation in exposure to risk to look at this question it is natural to ask whether it is possible to look at experimentally induced variation. For example, recent attempts to experimentally induce better crop insurance mechanisms (Cole *et al.* 2008) may provide a useful mechanism for evaluating the role of risk in reducing technological adoption—in principle one could establish whether farmers who are given access to a successful measure for reducing exposure to weather risk are more likely to adopt new technologies. Unfortunately, the major conclusion of that paper was that very few farmers chose to make use of this mechanism even though it was designed to have very attractive returns.

Foster and Rosenzweig (2009) attempt to test whether imperfect insurance leads to sub-optimal use of fertilizer, exploiting their plot-specific data on Indian farmers who have multiple plots and cultivate at least some of them across seasons. Specifically, they test whether profit shocks for a given farmer prior to planting affect his subsequent per-acre input use on a given plot of his land. If farmers are fully insured, then variation in profits will be sterilized, and have no effect on input decisions. In the absence of insurance, there will be a positive relationship between lagged income or profits π_{jt-1} for a farmer j on the use of inputs f_{ijt} on plot i in the subsequent crop season t , as farmers with a positive income shock experience an increase in wealth and thus can absorb more risk, or because credit markets are also imperfect so self-financing of inputs is necessary. Expressed as a linear

relationship:

$$(17) \quad f_{ijt} = \delta \pi_{jt-1} + \zeta_{ij} + v_j + e_{ijt} + \varepsilon_{jt},$$

where ζ_{ij} is a plot-specific fixed effect (e.g., soil quality), v_j is a farmer/farm fixed effect, e_{ijt} is an iid time-varying, plot-specific shock and ε_{jt} an iid time-varying shock that is common across all plots (e.g., farmer illness). If there is imperfect insurance, and perhaps also if there are credit constraints $\delta > 0$. Estimation of this equation by OLS would lead to a biased estimate of δ because farm profits may be correlated with the farmer and plot fixed effects, which reflect the *ex ante* return to input use (profitability). Thus, one could find that higher lagged profits and current fertilizer use are positively related simply because time-invariant land quality is complementary with fertilizer.

Differencing (17) over time for the same plot eliminates all unmeasured plot characteristics and time-invariant farmer characteristics:

$$(18) \quad \Delta f_{ijt} = \delta \Delta \pi_{jt-1} + \Delta e_{ijt} + \Delta \varepsilon_{jt}.$$

However, there two problems: First, there will be a negative covariance between the change in lagged farm profits $\Delta \pi_{jt-1}$ and the difference in the plot-specific error. This can be eliminated by taking out of the lagged difference farm profits that component associated with plot i . Thus, farmers with more than one plot are needed to identify δ . Second, the change in farm profits associated with the other farmer plots will be correlated with any shock to fertilizer common to all plots. If this is a shock common to all farmers in a village, then a village dummy can absorb this effect. If the shock is common to the farmer, such as farmer illness that prevents the use of optimal fertilizer and other inputs that affect profits, there will be a negative covariance between the change in lagged profits and this farmer-specific shock. Thus, the estimate of δ , the effect of lagged farm profits on current input use, will be biased negatively. The finding that δ is positive, however, would certainly reject full-insurance; a coefficient of zero or negative would lead to an uncertain conclusion. Estimated over 4,045 farmers cultivating more than one plot in at least two of three seasons, Foster and Rosenzweig find that δ is indeed positive and statistically significant, but only for farmers whose land size puts them in the bottom quartile of the land distribution. Consistent with other studies, the problem of lack of insurance appears to afflict the poorest farmers.

b. Credit constraints

Any input or technology that entails paying up-front costs requires that the agent have funds available prior to the realization of the gains from using the input or adopting the technology. If all agents can borrow, then whether a new technology is adopted will only depend on net returns and not on the timing of costs and benefits and therefore not on the characteristics of the agent, net of returns. If the ability to borrow, however, depends on the agent having assets that can be used as collateral or if borrowing is not an option so that the agent must supply his own funds, then such agent characteristics as wealth or the history of prior income realizations will affect current input and technology choices that have an investment element (assuming imperfect insurance). But, as we have seen, income and wealth may be correlated with the scale of operation, which affects returns, and with the ability to cope with *ex post* risk when formal insurance is unavailable. Identifying the role of credit market imperfections is thus difficult.

There are two methods used in the literature to quantify the role of credit constraints. The first is to ask agents (farmers) the primary reason(s) why they did not adopt a technology. This method was used by Miyata and Sawada (2007) in their study of the adoption of floating net aquaculture (FNA). The “reasons” can then be correlated with wealth or income to draw inferences about the importance of the credit “constraint” by income. This was done by Bhalla (1979) for farmers at the onset of the Indian green revolution; he found that 48% of large farmers and only 6% of small farmers reported that lack of access to credit was a reason for not purchasing fertilizer. The problem, of course, is that if the returns to adoption of HYV (which is fertilizer intensive) vary by farm scale, then even if all farmers faced the same credit price, small farmers would find it unprofitable to adopt while large farmers would find it profitable at the going interest rate. Small farmers would optimally adopt at lower interest rates, and thus may report that they are credit constrained in that context. Or, lenders may be unwilling to make loans to small farmers because the returns, given fixed costs, are low. Subjectively-reported credit constraints and returns may thus be highly correlated.

The second method for identifying the role or existence of credit constraints is to look for income effects that are net of returns, scale effects and insurance. Gine and Klöner (2005) attempt to isolate the role of credit with an intensive examination of the adoption of a new technology whose returns do not depend on scale. In particular, they look at the timing of the adoption of plastic reinforced fiber (FRP) boats in a fishing village in Tamil Nadu. Purchase of the boats requires up-front payments, and given that labor markets were well-functioning the new, larger-scale boats do not depend on the size of the fisherman’s household. The authors also argue that the fisherman have a well-functioning informal risk-sharing

scheme and that, because of moral hazard, boat rental was not an option. Gine and Klöpper also estimate the gains from adopting the new boat. They find that households with a higher-value house, for given returns, were more likely to purchase the boat earlier. The problem is that the variation in house value across the fisherman may reflect unobservables that affect investment returns, which are only imperfectly measured - fishermen with big houses may be more capable fishermen, and we must accept that the fishermen are fully insured, which seems unlikely. If not, then the wealth effect may again reflect risk aversion. The authors show, however, that whatever the reasons for the differential timing of adoption, the new technology was fully diffused within five years.

6. Behavioral Economics and Adoption Behavior

Given increased evidence from experimental laboratories, in the United States and increasingly in low-income countries, that individual behavior appears at times to be at variance with standard economic models it is natural to ask the question of whether behavioral models can be useful in understanding rates of technological adoption and input use in low-income countries. Of course, given the complexity of the adoption process and the difficulties associated with measurement that have been highlighted, it may be that the resolution to some of the puzzles of the technology adoption literature lies in more careful measurement and theorizing rather than taking significant steps away from the standard economic paradigm. However, to the extent that governments and NGOs are resistant to adopting insights from more conventional economic models, a finding that a behavioral mechanism is absent or of limited importance may have a constructive effect on policy design. On the other hand, if departures from standard optimizing models are salient, standard policy prescriptions based on such models should be modified or reversed.

Two recent papers that set out to test an explicit behavioral model do in the end seem to support a more conventional approach. Ashraf, Berry, and Shapiro (2007) used a randomized field experiment to study the adoption of packaged chlorine to purify drinking water. The intervention was designed to explore the idea that raising the cost of a technological device may increase the actual use of the device because agents are loss averse, so that sunk costs affect behavior, which should not be the case if agents are purely rational. Of course, in this experiment it is important to distinguish between a sunk cost effect and selection, arising from the fact that individuals who are more likely to use a device place a higher value on the device and thus are less sensitive to price. This problem is

addressed by randomly assigning a discount for an item after individuals had already agreed to purchase it at a given price. Because the discount does not introduce any additional selection into the process, response in terms of use to the discount variation captures any sunk cost effect. While the results show clear evidence of a selection effect, there is not clear evidence of the sunk cost effect.

Dupas's (2009) analysis of the field experiment randomizing the selling prices of bed nets among individuals in groups (Cohen and Dupas, forthcoming) was also initially conceived as a shedding light on behavioral hypotheses. In particular, the question asked was whether people who knew that a good was sold to neighbors at a subsidized price would, *ceteris paribus*, be less likely to purchase the good at a given price. The difficulty with this experiment is that, as noted above, those people who have neighbors who faced a favorable price in the early period may also be influenced by a learning effect – they will know more people who used the technology and thus will be more likely to adopt the technology if it is considered (by the previous adopters) to be valuable. It is clear from the results that the learning effects are far larger than any reference price effect as the reduced-form effect of neighbor prices on own adoption is negative, while the behavioral effect predicts a positive relationship. Note that if the price paid by a neighbor does in fact directly affect one's own behavior, that would invalidate the use of neighbors' prices as instruments to predict bed net adoption and thus prevent identification of pure learning effects from the experiment.

Duflo, Kremer and Robinson (2009) examine the efficacy of a commitment device designed to increase fertilizer adoption among Kenyan farmers by offering them small discounts on fertilizer when farmers are relatively liquid due to a recent harvest season that is then delivered at the time the fertilizer is used. The predictions of a standard exponential model of discounting with effective credit markets is, of course, that a farmer who is given the opportunity to purchase an investment item that will not be used until at time $t+1$ will prefer to purchase that item at time $t+1$ at a given price than to purchase that item at time t at the same price. To test whether the standard model is relevant in their context, the authors employ a randomized design, that offers three main treatments: 1. At the time of harvest the farmer is offered a contract for the free delivery of fertilizer when it is needed in the subsequent season. 2. At the top-dressing period of the subsequent season the farmer is offered free delivery of fertilizer at that time. 3. At the time of harvest the farmer is offered a choice between contracts 1 and 2.

In all of these cases, the price of fertilizer inclusive of delivery costs is the same. Thus, one would have expected that no farmer would choose to take up contract 1 if offered contract 2 and that more farmers would take up contract 2 than

would take up contract 1. However, the authors find that more farmers take up the delayed contract 1 than contract 2 (although the difference is marginally statistically significant), and half of the farmers offered contracts 1 and 2 actually chose contract 1. Therefore, the authors conclude that farmers value a commitment device, which could be due to hyperbolic discounting. One alternative explanation for these results is that cash payments up front have a negative return - loss due to sharing obligations, theft or inflation. Therefore, the authors overlaid randomly on the other treatments the subsidized sale of maize. Those receiving the subsidy had more cash and thus should have under this alternative explanation differentially preferred contract 1. They did not. These results suggest that the deadweight loss from subsidizing fertilizer can be reduced by offering subsidies along with commitment contracts if at least some farmers exhibit behavior consistent with hyperbolic discounting. However, in the absence of well-documented information on the profitability of fertilizer use (as discussed above) or of the different treatments in this setting it is difficult at this point to evaluate the full consequences for welfare of a behaviorally-enlightened subsidy program.

7. Conclusion

The adoption and efficient use of new technologies is an important feature of the development process. It is thus not surprising that there is a lively and growing literature attempting to understand whether or not there is substantial under-adoption or sub-optimal application of profitable or otherwise socially beneficial technologies, and if so why or why not. There is considerable variation, however, in what is known about different aspects of the process.

A particular strength of the recent literature has been its focus on the role of learning in the adoption of new technologies. It is evident that as a whole learning is quite sophisticated and a key element in at least the early stages of adoption when information acquisition has large payoffs. Information about technologies that are generally beneficial tends to diffuse quite rapidly and this process appears to be well-captured by standard models of Bayesian learning. There is also evidence of active and strategic experimentation, however, which provide insights into how interventions could facilitate the adoption process. The literature also suggests that education plays an important role in facilitating the acquisition and processing of new information, which appears to account for the pervasive finding that more educated agents adopt new technologies first, and helps explain the variation in returns to schooling over time and across areas. However, although the learning literature as a whole is quite well-developed, we find the relationship

between learning and technological externalities to be complex and in need of further study. Technological externalities may be particularly important in the arena of health and therefore studies of learning behavior in the context of health may be especially difficult to conduct and interpret.¹³

There is also suggestive evidence that risk, due to the incompleteness of insurance, and credit availability play an important role in delaying the adoption of profitable new technologies and constraining the levels of inputs necessary to exploit the new technologies, particularly among the relatively poor. That wealth and income are advantageous in adoption and input use because of these institutional failures is of particular concern both because it suggests that poor countries may have difficulty developing via technology catch-up but also because it suggests that possibilities for upward mobility for poorest households are limited. However, although there is an ample literature documenting that poor households are not well protected from risk and that they may have limited access to credit, very few studies examine directly how these factors affect the process of technological adoption itself.

Perhaps a surprising gap in the literature is the paucity of studies carefully documenting the returns to inputs and technologies that are alleged to be underutilized. In some cases this is due to absent data characterizing input costs for enterprises, the remedy for which is obvious, but in others, new thinking about how to measure gains for individuals, such as from medical interventions, are needed. The salience of behavioral oddities and of particular market imperfections in observed adoption behavior may be quite different in settings where returns to input mis-allocations and distance from the technological frontier are small compared to settings where there are large gains from alternative choices. Because the same technology may have different returns for different people in different places one cannot assume that because a technology is profitable in one time or place it is also profitable in another or that the important constraints to adoption in one area generalize.

A new and promising area of research involves testing models of choice in the field that go beyond simple rationality and that are consistent with laboratory evidence on these behavioral departures. However, it is not likely that differences in technological adoption or input use across different settings are primarily the result of differences in the fundamental nature of human behavior across countries.

¹³The medical literature contains studies that are designed to illuminate externalities, such as in the case of bed nets (e.g., Hawley *et al.*, 2003).

Ultimately the interplay between behavior, market settings, traditional institutions and technology payoffs need to be addressed to more fully understand the variety of experiences over time and across countries in utilizing productive resources and adopting new technologies. A strength of micro studies of adoption is that some of these details can be incorporated into the analysis and rigorous methods of evidence adducement can be applied. However, a better understanding of differences in findings across studies requires particular attention to differences in specific conditions inclusive of those related to climate and soil coupled with difference in specific market imperfections and traditional institutions.

References

Bandiera, Oriana and Imran Rasul, 2006, "Social Networks and Technology Adoption in Northern Mozambique," *The Economic Journal*, October, 116, 869-902.

Bartel, Ann P. and Frank R. Lichtenberg, February 1987, "The Comparative Advantage of Educated Workers in Implementing New Technology," *The Review of Economics and Statistics*, pp.1 -11.

Besley, Timothy, and Anne Case, 1993, "Modelling Technology Adoption in Developing Countries," *AER Papers and Proceedings*, 83(2), pp. 396-402.

Bhalla, Surjit S., 1979, "Farm and Technical Change in Indian Agriculture." In *Agrarian Structure and Productivity in Developing Countries*, edited by R. Berry and W. Cline. Baltimore: Johns Hopkins University Press.

Cole, Shawn, Xavier Gine, Jeremy Tobacman, Petia Topalova, Rob Townsend, James Vickery, 2009, "Barriers to Household Risk Management: Evidence from India" Federal Reserve Bank of New York, May.

Caselli, Francesco and Wilbur John Coleman II, 2001, "Cross-Country Technology Diffusion: The Case of Computers," *American Economic Review*, May, 91 (2), 328-335.

Cohen, Jessica and Pascaline Dupas, forthcoming, "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment," *Quarterly Journal of Economics*.

Comin, D. and B. Hobijn, 2004. "Cross Country Technology Adoption: Making the Theories Face the Facts." *Journal of Monetary Economics* 51: 39-83.

Conley, Timothy and Christopher Udry, "Learning about a New Technology: Pineapple in Ghana," *American Economic Review*, forthcoming.

Dercon, Stefan and Luc Christiaensen, 2008. "Consumption Risk, Technology Adoption and Poverty Traps: Evidence from Ethiopia," WEF Working Papers 0035, ESRC World Economy and Finance Research Programme, Birkbeck, University of London.

Duflo, Esther and Emmanuel Saez, 2003, "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment," *Quarterly Journal of Economics* 118(3), 2003, 815-842.

Duflo Esther, Michael Kremer, and Jonathan Robinson, 2008, "How High Are Rates of Return to Fertilizer? Evidence from Field Experiments in Kenya," *American Economic Review Papers and Proceedings*, 98 (2), 482-488.

Duflo Esther, Michael Kremer, and Jonathan Robinson, 2009, "Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya," NBER Working Paper 15131, July.

Dupas, Pascaline, 2009, "Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence from a Field Experiment," August, UCLA, mimeo.

Feder, Gershon, Richard Just, and David Zilberman, 1985, "Adoption of Agricultural Innovations in Developing Countries: A Survey," *Economic Development and Cultural Change*.

Foster, Andrew D. and Mark R. Rosenzweig, 1995, "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal of Political Economy*, 103(6), 1176-1209.

Foster, Andrew and Mark R. Rosenzweig, 1996, "Technical Change and Human Capital Returns and Investments: Evidence from the Green Revolution," *American Economic Review*, September.

Foster, Andrew and Mark R. Rosenzweig, 2009, "Should Poor Farmers Use More Inputs?," mimeo.

Gine, Xavier and Stefan Klonner, 2005, "Credit Constraints as a Barrier to Technology Adoption by the Poor: Lessons from South Indian Small-scale Fishery," World Bank Policy Research Working Paper 3665, July.

Griliches, Zvi, 1957, "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica*, October, 25 (4), 501-522.

Hawley, William A. *et al.*, 2003. "Community-Wide Effects of Permethrin-Treated Bed Nets on Child Mortality and Malaria Morbidity in Western Kenya." *American Journal of Tropical Medicine and Hygiene* 68 (Suppl. 4): 121-127.

Kremer, Michael and Edward Miguel, 2007, "The Illusion of Sustainability," *Quarterly Journal of Economics* 122, 1007-1065.

Lleras-Muney, Adriana and Frank Lichtenberg, 2006, "The Effect Of Education On Medical Technology Adoption: Are The More Educated More Likely To Use New Drugs?," special issue of the *Annales d'Economie et Statistique* in memory of Zvi Griliches, No 79/80.

Manski, Charles, 1993, "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies* 60, 531-542.

Miguel, Edward and Michael Kremer, 2004, "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities," *Econometrica* 72, 159-217.

Miyata, Sachiko and Yasuyuki Sawada, 2007, "Learning, Risk, and Credit in Households' New Technology Investments: The Case of Aquaculture in Rural Indonesia," mimeo, July.

Morduch, Jonathan, 1990, "Risk, Production, and Saving: Theory and Evidence from Indian Households," mimeo, Harvard University.

Munshi, Kaivan, 2004, "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution." *Journal of Development Economics* . 73(1): 185-215.

Munshi, Kaivan and Jacques Myaux, 2006, "Social Norms and the Fertility Transition." 2006. *Journal of Development Economics* 80(1):1-38.

Nelson, Richard R. and Edmund S. Phelps, 1966, "Investment in Humans, Technological Diffusion and Economics Growth," *American Economic Review*, pp. 69-75.

Oster, Emily and Rebecca Thornton, 2009, "Determinants of Technology Adoption: Private Value and Peer Effects in Menstrual Cup Take-Up," March, mimeo.

Parente, Stephen L. and Edward C. Prescott, 1994. "Barriers to Technology Adoption and Development" *The Journal of Political Economy* 102 (April):298-321.

Rosenzweig, Mark R., 1995, "Why are there Returns to Schooling?," *American Economic Review Papers and Proceedings*, Volume 85, Issue 2, May, pp. 153-158.

Rosenzweig, Mark R. and Hans P. Binswanger, 1993, "Wealth, Weather Risk and the Profitability of Agricultural Investment," *Economic Journal*, Vol. 103 (January): 56-78

Rosenzweig, Mark R. and T. Paul Schultz, May 1989, "Schooling, Information and Nonmarket Productivity: Contraceptive Use and its Effectiveness," *International Economic Review*, pp. 457-477.

Rosenzweig, Mark R., forthcoming, "Global Wage Inequality and the International Flow of Migrants," in Kanbur, Ravi and A. Michael Spence, eds., *Equity in a Globalizing World*, New York: Routledge.

Skinner, Jonathan and Douglas Staiger, 2005, "Technology Adoption from Hybrid Corn to *Beta* Blockers," National Bureau of Economic Research Working Paper #11251, March.

Suri, Tavneet, 2009, "Selection and Comparative Advantage in Technology Adoption," mimeo, MIT Sloan School.

Welch, Finis, 1970, "Education in Production," *Journal of Political Economy*, 78(1), January/February, 35-59

Figure 1: Estimated Relationship between Fertilizer Use (kg. per Acre) and Four Outcome Measures, Indian Farmers in Crop Year 2006-2007 (N=6,273 plots)




Table 1
Determinants of the Adoption of HYV Seeds: Indian Farmers, 2007

Determinants	(1)	(2)
Maximum schooling of household (years)	.0074 (2.03)	.00589 (1.62)
Value of landholdings (Rs. x10 ⁻⁴)	.000627 (2.10)	.000447 (1.44)
Low wealth (< Rs.250,000)	-	-.126 (3.26)
Low wealth*value of landholdings	-	.00429 (2.15)
Total number of farmers using HYV in prior season	.000425 (2.50)	.000408 (2.40)
Number of farmers	4,045	4,045

Absolute value of t-statistics in parentheses.