Heterogeneity, risk, and learning in the adoption of agricultural technologies

Andrew Zeitlin*
University of Oxford

Paper submitted to the RES 2012 Annual Conference

Abstract

In this paper we explore whether low rates of sustained technology use can be explained by heterogeneity in returns to adoption. We derive distinct predictions for models in which this relationship is driven by precautionary savings and by learning, and test these using data on the Cocoa Abrabopa Association, which provided a package of fertilizer and other inputs on credit to cocoa farmers in Ghana. High estimated average productive impacts for treated farmers are found to be consistent with negative economic profits for a substantial proportion of the treated population. By constructing an individual-specific measure of returns, we demonstrate that low realized returns among adopters are associated with low retention rates, even after conditioning on output levels and successful repayment. Patterns of returns and adoption support the hypotheses that high average returns mask substantial and persistent heterogeneity, and that farmers experiment in order to learn about their idiosyncratic returns.

*I thank Daniel Clarke, Stefan Dercon, Marcel Fafchamps, Christopher Ksoll, Kim Lehrer, Eliseus Opoku-Boamah, Marcella Vigneri, Christopher Udry, and seminar participants at the Ghana Cocoa Board, Oxford University, and the NEUDC for helpful comments and discussions. Moses Awoonor-Williams, Stefano Caria, Mark Fiorello, Petr Jansky, Christoph Lakner, and Sean O'Leary provided excellent research assistance. I am grateful to Richman Dzene, Francis Teal, Petr Jansky, and Emmanuel Opoku for their collaboration in the project on which this paper builds. The project would not have been possible without the contribution of the Ghana Cocoa Board and the Cocoa Abrabopa Association. Funding from the Bill and Melinda Gates Foundation the Ghana Cocoa Board is gratefully acknowledged.

1 Introduction

What constrains investment in agricultural technologies? This question is important - and stubbornly persistent - in development economics. Not only does agriculture continue to represent the primary source of income for many of the world's poor, but low adoption rates of agricultural technologies, such as fertilizers and improved seed varieties, have accompanied the stagnation of agricultural productivity in Africa in particular (World Bank 2008).

There is little dispute that there exist agricultural technologies with high expected returns in many Sub-Saharan contexts. This view is supported by a growing body of evidence. Notably, Esther Duflo, Michael Kremer and Jonathan Robinson report experimental evidence of a mean seasonal return of 36 percent to fertilizer use among maize farmers in the Busia District of Kenya (Duflo, Kremer & Robinson 2008). And yet rates of fertilizer use are low: fewer than 24 percent of farmers in Duflo and coauthors' study had used fertilizer in the preceding year. Even where supposedly high-return technologies do get adopted, many farmers abandon them. In a distinct sample of Kenyan maize farmers, Tavneet Suri documents that 30 percent of farmers switch into and out of the use of hybrid seeds in a given year (Suri 2011). In Ethiopia, Stefan Dercon and Luc Christiansen find that, while only 22 percent of farmers use fertilizer in a given year, a further 14 percent of farmers in the final round of their survey were not using fertilizer in spite of having done so in previous survey rounds (Dercon & Christiaensen forthcoming). Low rates of adoption and lack of sustained use of the technology combined with high rates of return to those technologies therefore present a puzzle.

Several mechanisms have been put forward to explain observed patterns of agricultural technology adoption. Processes of social learning have been much studied (Bandiera & Rasul 2006, Conley & Udry 2001, Foster & Rosenzweig 1995, Munshi 2004). If social learning is sufficiently important, low-adoption equilibria may persist in spite of potentially high returns. Alternative theories include credit and supply-side constraints (Moser & Barrett 2006, Zerfu & Larson 2010). In Kenya, Duflo and co-authors find evidence consistent with the view that time inconsistency in farmers' preferences causes inefficient adoption decisions (Duflo, Kremer & Robinson 2009).

While they do address important elements of observed adoption patterns, these theories are generally not well equipped to explain why adoption is not sustained. In the most common form of learning model, for example, farmers only adopt technologies when they know how to use them effectively, and this knowledge, once acquired, is never lost (Foster & Rosenzweig 1995, Jovanovic & Nyarko 1996). Likewise, instability in the supply of inputs alone seems an ad hoc explanation, and one incapable of explaining the widespread failure of farmers to persistently adopt profitable technologies even in cases where farmers have accessed them in the past. One exception is Dercon and Christiaensen (Dercon & Christiaensen forthcoming), who argue that year-to-year variation in the ability of households to bear risks associated with high-return technologies may explain instability in their use. Even so, if the typical farmer experiences such high returns as have been reported in the literature - a premise that we revisit in this paper - one would expect such technologies, once established, to pay for themselves. Indeed, these stylized facts lead Duflo to assert that "prima facie, neither limited liability nor risk aversion seem capable of explaining such a low level of fertilizer use" (Duflo 2006).

In this paper, we test the hypothesis that persistent differences in returns across farmers play a key role in determining sustained adoption of new technologies. We do so using a unique panel dataset that exploits the growth of a large scale non-profit initiative, the Cocoa Abrabopa Association, which alleviated credit constraints to the adoption of the *hi-tech* package of inputs (a specific combination of fertilizer, insecticide, and fungicide) among cocoa farmers in Ghana. We use the timing of the roll-out of this program to identify the effects of hi-tech adoption, and to test competing theories of the role of impact heterogeneity in determining subsequent adoption decisions.

Following the work of Heckman and coauthors, interest in models with "essential heterogeneity" (Heckman, Urzua & Vytlacil 2006). In these authors' framework, essential heterogeneity is defined by the empirical relationship between ex-ante decisions of selection into treatment and the ex-post distribution of returns. In applications where agents learn about their (possibly heterogeneous) returns to a treatment over time, both the agents' information set and their economic constraints may be affected by the outcomes of experimentation. Long-term dynamics of adoption may then be governed by the way in which agents respond to experimentation—a possibility that requires us to distinguish the mechanisms linking ex-post heterogeneity to subsequent adoption decisions. Central in this analysis will be the degree of persistence in the process governing potential outcomes.

The role of treatment effect heterogeneity in explaining low adoption rates in African agriculture is understudied in the literature, with papers by Tavneet Suri (2011) and Duflo, Kremer and Robinson (2008) providing two notable exceptions. Using an observational, panel dataset of Kenyan maize farmers to estimate average returns for discrete subgroups, Suri finds a non-monotonic relationship between the adoption rates and expected returns of these subgroups in her data. Duflo, Kremer and Robinson (2008) use treatment interactions and quantile treatment effects to characterize heterogeneity in their experimental data. They find that net returns to fertilizer adoption are negative for 13.5 percent of farmers, in spite of the estimated 36 percent seasonal return.[1]

[1]As pointed out by Foster & Rosenzweig (2010), these "net" returns do not account for labor costs. If labor or other complementary inputs also increased in response to fertilizer use, then they will estimate the true net returns, and underestimate the fraction of farmers for whom net returns are negative.

Existing studies shed only limited light on the extent to which heterogeneity in the ex-post distribution of returns is driven by persistent differences across farmers or transient sources of risk. To address this question, we draw out distinct empirical implications of two alternative theoretical mechanisms linking ex-post returns to subsequent adoption.

In addition to estimating mean returns for observable subgroups, we use quantile treatment effects and construct a proxy for farmers' information about returns at the individual level. The latter in particular allow us to test mechanisms of prudence and learning that may relate realized returns to subsequent adoption decisions.

We take as our starting point that the literature is far from conclusive on three questions: (1) How heterogeneous are the rates of return to agricultural technologies, such as fertilizer? (2) Does heterogeneity in returns affect the sustained adoption of such technologies, beyond farmers' initial experimentation? And (3) is this heterogeneity in realized returns caused by prudence in response to transient shocks, or learning about persistent differences in the suitability of a technology across farms and farmers? We find that heterogeneity in the returns to fertilizer use is substantial, it affects continued adoption, and it reflects persistent differences across farmers.

The remainder of this paper proceed as follows. Section 2 provides a theoretical basis for the empirical work, by drawing out the distinct implications of mechanisms of prudence and learning for the relationship between ex post returns and subsequent adoption decisions. Section 3 describes the specific intervention studied, the data and the quasi-experimental setting that will be used to estimate impacts. Section 4.1 presents our estimates of the average treatment effect, demonstrating robustness to a range of identifying assumptions. Having established large positive returns on average, we turn in Section 4.2 to demonstrating the heterogeneity of these treatment effects. In Section 5, we test hypotheses of the relationship between experienced treatment effects and decisions to continue adoption. Section 6 concludes.

2 Theoretical framework

When the effects of agricultural technologies are heterogeneous, at least two theoretical mechanisms relate the ex post distribution of returns to experimenting farmers' decisions to continue adoption. On the one hand, such farmers may update their beliefs about the distribution of outcomes that they face based on these realizations.[2] Farmers experiencing low returns in a given period may become more pessimistic and decide to abandon the

[2] This may be true regardless of whether returns are homogeneous across farmers, but unknown in the population, or when the returns are farmer-specific. We return to this distinction below.

technology. On the other hand, even if the distribution of output under available technologies is known perfectly by all farmers, we may still see a relationship between ex post realizations of returns and decisions to continue adoption. As shown below, farmers who exhibit prudence—that is, who build a buffer stock in anticipation of the adoption of risky technologies (Kimball 1990)—may be deterred by low realized returns to adoption simply because it affects their ability to insure themselves against adverse events in the future. In this section we propose a stylized model of prudence and learning mechanisms, which can be used to distinguish empirically between the two.

These two mechanisms linking output realizations and adoption decisions have very different policy implications. When prudence drives the relationship between realized returns and subsequent choices, policies that insure farmers against downside risks can induce higher levels of adoption and consequent welfare gains. By contrast, in the learning case, and particularly when returns are homogeneous but unknown, there may be a role for policy in resolving uncertainty by providing information about returns. Where social learning is possible, experimentation—a public good—is undersupplied.[3] If returns are both heterogeneous and unknown, policies designed to encourage all farmers to adopt a given technology may be inefficient. Instead, a second-order question becomes central to policy: are the returns themselves amenable to intervention?

These points are readily understood in a stylized model of a farmer's decisions over consumption and technology adoption. Anticipating our empirical application, we consider the farmer's decision between traditional, 'low-tech' technology, L, and a new, 'high-tech' technology, H, where the latter is characterized by a higher mean and higher variance.

At the beginning of the first period, the farmer is endowed with assets A, and receives a draw of income, y_{1H} from technology H. Given a concave utility function u, the farmer's decision problem is to choose both a consumption level, c, and a second-period technology, $w \in \{H, L\}$, in order to maximize expected lifetime utility

$$V(c, w) = u(c) + \mathrm{E}u(A + y_{1H} - c + y_{2w}), \tag{1}$$

where the expectation is taken with respect to the (possibly subjective) probability distribution for the chosen technology of second-period income.

Consider first the case in which the distribution of outcomes under each technology is known by the farmer, so that prudence governs the relationship between the realization of y_{1H} and the farmer's technology choice w. Initial assets and first-period income affect decisions only through their sum.

[3] This undersupply is exacerbated when farmers engage in strategic delay (Bardhan & Udry 1999, Bandiera & Rasul 2006). In the simple model that follows, however, we focus only on learning by doing, so that such strategic considerations do not apply.

Proposition 1. *Initial income and technology choice under prudence*

In the farmer's maximization problem in equation (1) *an increase in first-period income* y_{1H} *increases the expected utility of the high-tech technology* H *relative to the traditional technology* L *if and only if, for a given initial income and assets and second-period technology choice optimal consumption under* H, c^*_H, *is lower than optimal consumption under* L, c^*_L.

Proof. We are interested in the sign of $\partial(V^*_H, V^*_L)/\partial y_{1H}$, where V^*_w represents the value of the value function V, with consumption chosen optimally for a given technology choice $w = H, L$.

Note that, for a given technology choice, w, the first order condition from the choice of c is given by

$$u'(c) - \mathrm{E}_w\left[u'(A + y_{1H} - c + y_{2w}\right] = 0 \tag{2}$$

which implicitly defines the c^*_w as a function of $A + y_{1H}$.

To understand what happens when y_{1H} increases, we differentiate the value function for a given prospect, w, to obtain

$$\frac{\partial V^*_w}{\partial y_{1H}} = \mathrm{E}_w\left[u'(A + y_{1H} - c^*_w + y_{2w})\right] = u'(c^*_w), \tag{3}$$

where the first equality makes use of the envelope theorem and the fact that c^*_w is chosen optimally, and the second equality is obtained by substituting in the first order condition. Consequently we have

$$\frac{\partial(V^*_H - V^*_L)}{\partial y_{1H}} > 0 \iff u'(c^*_H) > u'(c^*_L), \tag{4}$$

which is equivalent to the condition that $c^*_H < c^*_L$, by the concavity of u. □

Whether this condition will hold depends on features of the utility function and on the properties of the distribution for the counterfactual outcomes y_{2H}, y_{2L}. A few conditions can be stated in general (Kimball 1990). First, the condition in equation (4) will only hold if the farmer's utility function exhibits *prudence*—that is, if the third derivative of the utility function is positive.[4] Second, if the distribution of y_H first order stochastically dominates y_L, then equation (4) will not hold. Third, if the distribution of y_H is a mean-preserving spread of the distribution under y_L, then condition (4) will hold. Intuitively, then, an increase in first-period income will have a greater effect on the likelihood of choosing a higher-mean, higher-variance technology, H, when farmers exhibit higher prudence, and where the increase in variance caused by moving from the distribution L to H is large relative to the increase in expected value.

[4]Specifically, Kimball (1990) proposes u'''/u'' as a measure of absolute prudence.

The prudence mechanism draws no distinction between the effect of (liquid) assets, A and first-period income as determinants of the choice between second-period technologies. Consider an extension of the example above, in which in period 0 the farmer first realizes an income y_{0L} drawn from the traditional technology, before deciding a period 0 consumption level, c_0 such that the assets at the beginning of period 1 are given by $A = y_{0L} - c_0$. Starting from an interior solution in which the constraint $c_0 \leq y_{0L}$, the marginal propensity to consume out of an increase in the realization of y_{0L} will be less than one. In this case, even past income realized under the traditional technology can be conducive to the adoption of the new technology.

This result contrasts with the implications of a Bayesian learning mechanism. Under this alternative mechanism, farmers update beliefs about returns to a new technology on the basis of output realizations. To make this contrast as stark as possible, suppose that there is no savings technology, so that consumption and income in each period are equal. Consider again the decision of a farmer who has just experimented with technology H for the first time. This farmer holds a (subjective) prior distribution for high- and low-technology options, with cumulative distribution functions $F_{H,0}, F_{L,0}$, respectively, at the end of period 0. She decides whether to use technology H on this basis. If the farmer is myopic (i.e., if she does not take into account the value of learning acquired through experimentation), then she will use technology H in period 1 if $\int u(y)\mathrm{d}F_{H,0}\mathrm{d}y > \int u(y)\mathrm{d}F_{L,0}\mathrm{d}y$.

For such a farmer, a high realization of y_{1H} increases the expected value of choosing $w = H$ in period 2. In the case of normally distributed beliefs, for example, the farmer's subjective expected value for y_{2H} increases linearly in this realization. A positive relationship between realizations of y_{1H} and the expected value of technology H is consistent with a range of learning models and outcome distributions.

By contrast, in the absence of a precautionary savings mechanism, a learning mechanism gives no reason for past realizations of the traditional technology to be positively associated with the adoption of H in subsequent periods. On the contrary, if the distribution of y_L is also an object of learning, then high realizations of y_{0L}, defined as a draw from technology L in period 0, reduce the expected gain from choosing $w = H$ over $w = L$ in period 2. This observation provides the basis of an empirical test. While both learning and prudence mechanisms link income from experimentation with technology H to the decision to sustain its adoption, in the most plausible cases they have opposite implications for the association between adoption of H and previous yields under traditional technologies.

3 Context and data

In 2006, the Cocoa Abrabopa Association (CAA), a not-for-profit subsidiary of Wienco Ghana Ltd, began a program of distributing inputs on seasonal credit to cocoa farmers in Ghana. With the support of the Ghana Cocoa Board, CAA provided farmers with access to two acres' worth of a package of fertilizer, pesticides, and fungicides. This specific bundle of inputs, known as the *hi-tech* package, had been promoted by the Cocoa Research Institute of Ghana since 2001, though problems of poor repayment rates had limited distribution.[5] CAA provided these inputs to groups of between 8 and 15 farmers on a joint liability basis, with dynamic incentives: groups that failed to repay in full would be suspended for a minimum of one year, while those that repaid successfully would be given four acres' worth of inputs in the following year, subject to approval of a CAA field officer. In addition to these physical inputs, farmers in the first year of membership were advised on their proper application by a CAA promoter, and some business training would be provided by Technoserve Ghana. Judged by its expanding membership rolls, the program has been wildly successful: from 1,440 farmers in 2006, CAA expanded to a membership of 18,000 farmers by 2009 (Cocoa Abrabopa Association 2009).

To identify the impact of CAA membership on farmer incomes, we took advantage of the fact that much of CAA's expansion during this time was at an extensive margin: it involved expansion into new villages. CAA's expansion operates on an annual cycle, as follows. Promoters first arrive in a new village in January of a given year, and by February farmers make their decisions to opt into the program (or not), forming groups accordingly. Inputs begin to arrive in May, but the harvest does not take place until October, with repayment of inputs due by December of that year.

We visited farmers in September of 2008 and 2009 to observe the outcomes for the 2007/08 and 2008/09 seasons, respectively. In each wave of the survey, we conducted a representative sample of two types of villages: those that had been reached by CAA for the first time in the prior year, and those that had been reached by CAA for the first time in the current year, i.e., the growing season that was ongoing at the time of our visit.[6] The former had experienced one full season since the arrival of Abrabopa and had made their membership decisions for the following season, but had not

[5]Although the use of these broad categories of inputs is not new to Ghanaian cocoa farmers, the particular configuration was. Evidence from other contexts (Duflo et al. 2008) shows that economic returns can be highly sensitive to precise quantities and combinations of inputs used.

[6]In the 2009 survey, we also revisited farmers from the previous survey who had by then been exposed to the program for two seasons. To focus on a comparable set of "early adopters", and in light of the possible cumulative effects of sustained fertilizer use, we do not make use of these observations in this paper. See Opoku, Dzene, Caria, Teal & Zeitlin (2009) for further details of the survey.

Table 1: Estimating sample, by survey round and membership classification

		Observations, by survey round	
Year of first visit	Initial adoption decision	2007/08	2008/09
2007	Adopt	82	0
	Do not adopt	41	0
2008	Adopt	88	72
	Do not adopt	42	29
2009	Adopt	0	95
	Do not adopt	0	37

Note: survey round refers to the most recent completed harvest as of the time of each survey.

yet harvested any cocoa in the second year of exposure by the time of our visit. The latter had made membership decisions for the season in progress at the time of our survey, but had not yet experienced the results of those decisions.

In each type of village, we conducted representative samples of two populations of farmers: those who joined in the first year of its availability in their village, whom we call early adopters, and those who opted not to join CAA in the same year, whom we call early non-adopters. The resulting sample used in estimation is given in Table 1.

As will be described in detail in the next section, we use a cross-sectional difference-in-differences approach to estimate the average effect of the first wave of Abrabopa membership on early adopters. To do so, we pool data from both survey rounds. We then define two key variable. First, we define a dummy variable, z_{vt}, to indicate exposure to Abrabopa in village v and year t. Given the sampling strategy in Table 1, any village in our sample with $z_{vt} = 0$ must have $z_{v,t+1} = 1$. Second, we define m_{ivt} as an indicator of the adoption decision of farmer i in village v and year t. To denote early adopters, who adopt in the first year of exposure, we drop the time subscript and define $m_{iv} = \mathbf{1}\{m_{ivt} = 1, z_{vt} = 1, z_{vs} = 0\}$ for some t and for all $s < t$. Early adopters choose to adopt in the first year of availability in their village, but not necessarily in subsequent years, as we shall see. In the sample used to estimate impacts on farmer production, we restricts attention to the year of exposure or the year immediately prior. Treatment w_{ivt} occurs in this sample when the individual is an early adopter and Abrabopa is present in the village—that is, if $w_{ivt} = m_{iv} z_{vt}$.

The basis for our difference-in-differences identification strategy is illustrated in Figure 1, which shows our four categories of farmers. The first $(z_{ivt} = 1, m_{iv} = 1)$ are those who receive treatment (i.e., they have adopted the hi-tech package). The second $(z_{ivt} = 1, m_{iv} = 0)$ were offered the treatment but did not adopt. The third $(z_{ivt} = 0, m_{iv} = 1)$ are farmers who

Figure 1: Average cocoa output, by treatment status in year of study

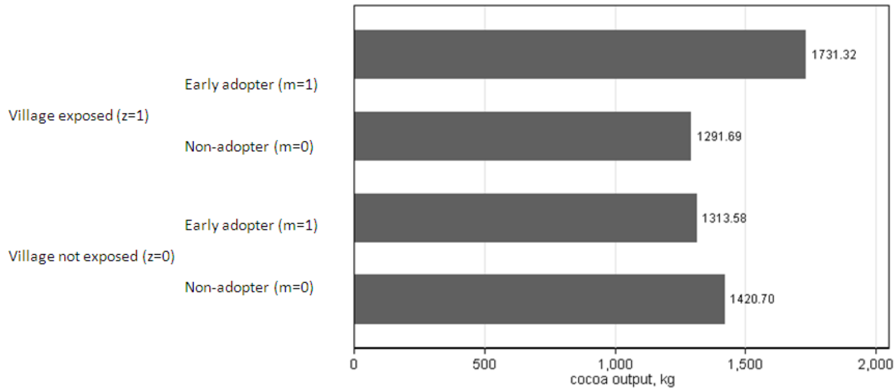

will adopt in the first year that Abrabopa reaches their village, but whose villages have not completed a season under treatment at the time of the survey round under consideration. And fourth are those farmers who will choose not to join in the first year of exposure, and whose villages have not yet been exposed to Abrabopa ($z_{ivt} = 0, m_{iv} = 0$).

Given this setup, we can read off a difference-in-differences estimate of program impacts directly from Figure 1. We estimate the average treatment effect on the treated as

$$\begin{aligned} \tau_{ATT} \quad = \quad & \mathrm{E}\left[y_{ivt}|m_{iv} = 1, z_{vt} = 1\right] - \mathrm{E}\left[y_{ivt}|m_{iv} = 0, z_{vt} = 1\right] \\ & - \left(\mathrm{E}[y_{ivt}|m_{iv} = 1, z_{vt} = 0] - \mathrm{E}[y_{ivt}|m_{iv} = 0, z_{vt} = 0]\right). \end{aligned} \quad (5)$$

The first two terms of equation (5) give the differences in output between those who were early adopters and those who were not, in villages exposed to Abrabopa. This within-village difference nets out one potential form of selection bias, arising from any common, village-level productivity variables that may be correlated with the timing of Abrabopa exposure. Still, this may be a biased estimate of the treatment effect if relatively productive farmers were more likely to become members. By subtracting the second two terms of equation (5), the pre-treatment difference between early adopters and non-adopters in villages not yet exposed to Abrabopa, we can account for this second form of selection bias, arising from within-village selection. In the pooled cross-section without any controls, the average treatment effect of Abrabopa adoption in equation (5)can be read off from Figure 1 as 546.75 kg.[7] From the figure, neither form of selection bias appears very strong: in villages not yet experiencing output under treatment, those who go on to

[7]This differs from the estimate in column (2) of Table 2 only because the estimates in that table include a control for the year of survey.

become early adopters are, if anything, worse off than no-adopters. Moreover, non-adopters in villages under treatment are, if anything, worse off than non-adopters in villages not yet treated.

Sampled farmers also provided information on a range of socio-economic characteristics and agricultural practices. These data are summarized in Appendix Table A1, where we present summary statistics by survey round and treatment status. Table A1 includes the prima facie evidence of Abrabopa's impacts that underlie Figure 1: early adopters' output exceeds that of non-adopters in exposed villages in both survey waves, while early adopters' output levels in the season before their village is exposed to Abrabopa are higher than non-adopters in 2008 (1615 vs 1460 kg) but lower in 2009 (1034 vs 1376 kg). Rates of fertilizer use are low (less than 50 percent) among farmers not yet reached by Abrabopa. Farms in this sector are typically small: farmers have an average of approximately four hectares of land devoted to cocoa trees. Education levels are low, just under 40 percent of farmers having completed education beyond primary level, and just over 20 percent of farmers in the sample are female. Given the quasi-experimental setting, we defer a discussion of covariate balance across treatment and control groups to the following section, where we present a more detailed explanation of our identification strategy.

4 Average and heterogeneous treatment effects

4.1 Average treatment effect on the treated

To make clear the identifying assumptions underlying our estimates of average returns, consider the following model for potential outcomes of gross output under two counterfactual scenarios - with and without the *hi-tech* inputs ($w = 1, 0$ respectively):

$$y_{0ivt} = \mu_0 + \eta_i + \lambda_{vt} + \varepsilon_{0ivt} \tag{6}$$

$$y_{1ivt} = \mu_1 + \eta_i + \lambda_{vt} + \varepsilon_{1ivt} \tag{7}$$

for farmer i in village v and year t. For the time being, we ignore the role of observed, farmer-specific covariates. The η_i give farmer-specific, time-invariant unobserved characteristics, while the λ_{vt} capture village-year unobserved shocks to productivity; we will be concerned about the potential correlation of both with treatment status.[8] Without further loss of gener-

[8]For counterfactual states, $w = 0, 1$, we can decompose the state-specific error term ε_{wivt} into two parts, $\varepsilon_{wivt} = \alpha_{wiv} + u_{wivt}$, where the first component represents a time-invariant, individual-specific return. We will refer to this as "essential heterogeneity"; it is the object of learning by the farmer. The second component, u_{wivt} is time-varying and captures a source of objective risk. Both can impact future technology decisions: the former affects subjective perceived returns, while the latter affects farmers' liquidity, buffer stocks, etc.

ality, we assume that $\mathrm{E}[\eta_i] = \mathrm{E}[\lambda_{vt}] = \mathrm{E}[\varepsilon wivt] = 0$, so that the difference $\tau_{ATE} = \mu_1 - \mu_0$ gives the average treatment effect in this population. The quantity $\tau_{ATT} = \mu_1 - \mu_0 + \mathrm{E}[\varepsilon_{1ivt} - \varepsilon_{0ivt} | w_{ivt} = 1]$ gives the average treatment effect on the treated (ATT). We focus on identification of the ATT in this section, since this is identifiable under a more plausible set of assumptions.

Observed outcomes are given by the switching regression, $y_{ivt} = y_{0ivt} + (y_{1ivt} - y_{0ivt})w_{ivt}$. Substituting in equations (6) and (7) yields

$$y_{ivt} = \mu_0 + (\mu_1 - \mu_0 + \varepsilon_{1ivt} - \varepsilon_{0ivt})w_{ivt} + \mu_0 + \eta_i + \lambda_{vt} + \varepsilon_{0ivt} \tag{8}$$

Examination of equation 8 clarifies the nature of the selection problem that must be addressed in estimating the ATT. A regression of y_{ivt} on w_{ivt} returns a consistent estimate of τ_{ATT} only if $\mathrm{E}[\eta_i + \lambda_{vt} + \varepsilon_{0ivt} | w_{ivt}] = \mathrm{E}[\eta_i + \lambda_{vt} + \varepsilon_{0ivt}] = 0$. This reflects the weaker identifying assumptions required for the ATT than the average treatment effect in the population as a whole; the former requires only that, conditional on covariates, average outcomes among the untreated are a consistent estimate for the average outcomes that treated would have obtained, in the absence of the program. Thus consistent estimates of τ_{ATT} are possible even if individual-specific returns to adoption $(\varepsilon_{1ivt} - \varepsilon_{0ivt})$ are correlated with adoption choices, w_{0ivt}. But the assumption required for identification of the ATT from a naive regression of output on treatment alone will fail if adoption is correlated with either village-level differences in productivity or with the idiosyncratic productivity of farmers. We take up these problems below.

We first formalize the process by which membership is determined. Recall that we define treatment, w_{ivt}, as the effect of the first year of use of the *hi-tech* package, in light of the potential for accumulation of impacts over years. In our data, we restrict attention to village-years, vt, in which either (a) Abrabopa has never had any members in that village before, and Abrabopa will have its first members in village v in year $t + 1$; or (b) Abrabopa has its first ever members in village v in year t. Consequently we examine first-year impacts only on the subpopulation of individuals who join Abrabopa in the first year that it is available in their village.

In this set of village-years, use of Abrabopa's hi-tech inputs is the product of two factors: firstly, that Abrabopa visits the individual's village, v, in year t, and secondly, that the individual joins in that year. Let z_{vt} be a dummy variable indicating the presence of Abrabopa in village v in year t, and let m_{iv} be a dummy variable indicating that individual i in village v is the 'type' who joins Abrabopa in the first year in which it is available in their village. Thus in these villages, $w_{ivt} = z_{vt}m_{iv}$.

Our identification strategy rests on two key features of our data. The first of these is the ability to observe the *future* membership decisions of individuals in villages that have not yet been visited by Abrabopa at the time of the output realization y_{ivt}.[9] The second of these is the ability to observe the productive outcomes for a representative sample of those who do *not* join Abrabopa in any given village-year.

We use the first of these features to address potential correlation between the individual-specific unobservables and treatment status, arising through individual selection into Abrabopa. The second allows us to address the potential correlation between village-level characteristics and treatment status, arising through the non-random roll-out of Abrabopa coverage.

To do so, we assume that the process by which farmers are selected into membership is constant over time, with respect to unobserved characteristics that differentiate them from village-mean productivity:

$$\mathrm{E}[\eta_i + \varepsilon_{0ivt} | m_{iv} = 1, z_{vt} = 1] = \mathrm{E}[\eta_i + \varepsilon_{0ivt} | m_{iv} = 1]. \tag{9}$$

This assumption implies that those who join Abrabopa when it first reaches their village in year t are the best farmers in those villages, then farmers who join Abrabopa upon its arrival in their villages in year $t+1$ are also the best farmers in those villages.[10]

If we were also willing to assume that the roll-out of Abrabopa availability was as good as random with regard to village productivity levels, so that $\mathrm{E}[\lambda_{vt}|z_{vt}] = 0$, then a comparison of current and future program members would suffice to identify the ATT. However, effective randomness of roll-out is a strong assumption. We are able to relax this by making use of data on non-members in program villages. In essence, we can use mean outcomes of those who do not join Abrabopa in current and future program villages to estimate the village-specific effect. Our estimates under this identification strategy are then a form of difference-in-difference estimates: we compare within-village differences between those who join Abrabopa and those who do not, in villages that have just been reached by Abrabopa in year t and those that will only be reached by Abrabopa in year $t+1$.

This strategy requires an auxiliary assumption that there are no externalities from the presence of the program on non-members—akin to the standard *stable unit treatment value* (SUTVA) assumption. Formally, we require that

$$\mathrm{E}[\eta_i + \varepsilon_{0ivt} | m_{iv} = 0, z_{vt} = 1] = \mathrm{E}[\eta_i + \varepsilon_{0ivt} | m_{iv} = 0, z_{vt} = 0] \tag{10}$$

This implies that the idiosyncratic component of the outcome observed for the untreated is uncorrelated with village-level exposure, *after* conditioning on the village-year effect, λ_{vt}. Since we are interested here only in impacts in the first year of the program's presence in a village, we believe this to be a

[9] For a similar use of future adoption decisions to address selection problems in a pipeline evaluation, see, e.g., Field (2005).

[10] We are able to provide a partial test of this assumption by exploiting the fact that a subset of farmers are seen both before and after adoption to estimate a fixed effects model.

plausible assumption. Neighboring farmers will not have had an opportunity to observe program impacts and make any corresponding adjustments in their own production.[11]

Given assumptions (9) and (10), we can operationalize our difference-in-differences estimator in two ways. Most directly, we regress output on dummy variables for adoption in the first year of exposure, m_{iv}, current exposure to Abrabopa, z_{vt}, an their interaction, $w_{ivt} = m_{iv}z_{vt}$. Under assumptions (9) and (10), a regression of the form

$$y_{ivt} = \beta_0 + \beta_z z_{vt} + \beta_m m_{iv} + \beta_w m_{iv} z_{vt} + \varepsilon_{ivt} \tag{11}$$

consistently estimates the ATT as the coefficient β_w. An alternative is to use a fixed estimator for equation (11), in which case the village-level exposure variable, z_{vt}, is not identified, but the treatment effect for early adopters remains identified by the interaction of indicators for the early adoption decision of the individual and village-level exposure.

Results of these estimates are presented in Table 2. Column (1) provides estimates of the basic difference-in-difference specification of equation (11). The estimated average treatment effect experienced by Abrabopa members can be read off from the first row as 437kg—an estimate approximately 1.9 times the output equivalent of the cost of the input package (230 kg). In column (2), we further include a vector of potentially confounding farm and farmer characteristics.[12] Columns (3) and (4) include village and village-year indicator variables, respectively. The estimated treatment effect is substantively unchanged by these additional robustness measures. Finally, column (4) includes farmer fixed effects, identifying the treatment effect off of changes in membership observed for a subset of 195 farmers who were interviewed both before and after treatment.

We cannot reject the hypothesis that the estimates of column (1) are consistently estimated in any of these specifications. Estimated coefficients on exposure at the village and individual levels (variables z_{vt} and m_{iv}, respectively) confirm the lack of evidence for selection bias by village or individual. The estimated treatment effect of 391 kg is substantial: it represents a 30% increase in the pre-treatment output of those who go on to adopt.

[11] Conley & Udry (2010) demonstrate that farmers do adapt their technology choices in response to 'news' about their neighbors' levels of production. As in their estimation strategy, we rely on farmers' inability to react until after output realizations have occurred

[12] We do not control for productive inputs, on the grounds that changes in labor and non-labor inputs mediate the causal effect of program membership on production (Foster & Rosenzweig 2010). More generally, controlling for post-treatment variables may introduce biases into estimates of causal effects (Heckman & Navarro-Lozano 2004, Pearl 2009). We include controls for farmer gender, an indicator variable for whether the farmer has attained junior secondary or higher education, and household size, as well as (quadratic functions of) farmer age and cultivated farm size. Given that it typically takes three years for cocoa trees to reach bearing age, we consider it reasonable to take our measure of cultivated farm size, which explicitly excludes trees too young to bear cocoa, as exogenous in this context.

Table 2: Estimates of average impact of CAA membership on members in the first year

	(1) OLS	(2) Controls	(3) Village FE	(4) V×t FE	(5) Ind. FE
current CAA member, w_{ivt}	437.1*	450.0**	482.8**	461.5***	391.6**
	(214.63)	(168.62)	(175.49)	(164.18)	(143.27)
villaged exposed to CAA, z_{vt}	-54.15	-201.6	-274.4		-138.6
	(278.08)	(190.15)	(163.92)		(122.52)
early adopter, m_{iv}	-130.9	-35.75	-116.7	-90.87	
	(200.30)	(185.66)	(158.42)	(145.02)	
Individual controls	No	Yes	Yes	Yes	Yes
Wave indicators	Yes	Yes	Yes	No	Yes
N	606	577	577	577	577
Controls: p-value		0.00	0.00	0.00	0.02

Dependent variable is cocoa output, in kg. Robust standard errors in parentheses, clustered at village-level. Estimates in columns (2)–(4) contain controls for farmer gender, log age, and post-primary education, and for log farm size and log tree age. Resulting p value from F test for joint significance of control variables presented.

A similar identification strategy can be used to test whether participation in Abrabopa increases use of complementary inputs. If such impacts were to be found, then these should be taken into account in calculating returns to participation and technology use. For any input, x_{ivt}, used by farmer i in village v and year t, the effect of Abrabopa membership on input use can be implemented, as in column (3) of Table 2 and analogous assumptions to (A1) and (A2), by regressing x_{ivt} on early adoption, m_{iv}, and its interaction with current exposure, $w_{ivt} = m_{iv} z_{vt}$, in a specification with village-year fixed effects. We do so in particular for three types of labor input: household labor, hired labor, and *nnoboa* labor, a traditional labor-sharing arrangement. These results are reported in columns (1) to (3) of Table **??**. In each case, point estimates are economically small, and we are not able to reject the hypothesis that these complementary input levels remain constant under participation in the program . In the absence of any such indirect costs, an impact of 391 kg corresponds to an economic return of 70 percent on the value of the inputs received.

4.2 Treatment effect heterogeneity

Estimated returns to the CAA technological package are high among adopters, and yet 32 percent of first-year members in the sample were no longer members one year later. This appears to be a special case of the broader puzzle in agricultural technology adoption: if returns to adoption of this technology are indeed so high, and if credit constraints are not binding in this case, then why do so many members drop out?

Learning about persistent heterogeneity in returns across individuals provides one explanation. If there is persistent heterogeneity across farms or farmers—such that individuals can learn from current yields about their idiosyncratic, future returns from returns vary substantially across individuals—then incomplete adoption could be an outcome of rational choices. The average treatment effect on the treated, estimated in the previous section, is not informative about the fraction of farmers expecting positive returns - what Heckman (2010) calls the "voting criterion". In this section, we examine the extent to which high average treatment effects mask substantial heterogeneity in the ex-post distribution of impacts.

To test for such heterogeneity, we augment the village-year fixed effects empirical specification of Table 2 by interacting the indicator for current membership, w_{ivt} with a vector of observed individual characteristics, x_{ivt}.[13] These include farm size, tree ages, and the share of hybrid trees on the farm, and lagged (pre-adoption) output, as well as farmer characteristics: indicators for female farmers and farmers with post-primary (JSS) or higher education, and the log of household size.

These results are presented in column (1) of Table 3. Here we observe a negative association between lagged output and experienced returns. This is consistent with what one might expect for an intervention that provides insecticide and fungicide to farmers: it is not surprising for this to have its strongest effects among those farmers whose prior practices led to low output levels.

A potential concern with these results is that they might be biased by nonlinearities in the relationship between characteristics x and yields. Since treatment effects are identified by the interaction of an early adopter indicator, m_{iv}, with village-level exposure, such nonlinearities would be a source of bias if members have different levels of x than non-members. Under our identification strategy, this can be addressed by allowing the time-invariant membership indicator, m_{iv}, to interact with characteristics x as well. This is done in column (2). Addressing this potential source leads to a number of interesting dimensions of heterogeneity: women experience substantially lower returns, while larger farms and those with older trees experience higher returns. The p-value on a Wald test of the hypothesis that coefficients on the interaction of characteristics and membership ($m_{iv} \times x_{ivt}$) is statistically significant at the one percent level, confirming the importance of these controls.

[13] To facilitate interpretation, continuous variables are standardized prior to interaction, using means and variances among early adopters.

Table 3: Heterogeneity along observable dimensions

	(1)		(2)	
current CAA member, w_{ivt}	176.5	(181.20)	235.2	(148.47)
early adopter, m_{iv}	60.91	(102.13)	128.4	(109.10)
$w_{ivt} \times$ female	-208.8	(128.28)	-248.5*	(125.81)
$w_{ivt} \times$ JSS	-135.9	(238.67)	6.882	(199.39)
$w_{ivt} \times$ ln(adults hh members)	76.83	(92.20)	79.55	(90.31)
$w_{ivt} \times$ ln(land devoted to cocoa, ha)	-4.093	(106.64)	173.3*	(95.18)
$w_{ivt} \times$ ln mean tree age	56.20	(39.28)	99.87**	(43.78)
$w_{ivt} \times$ fraction hybrid trees	161.4	(157.47)	193.3	(149.45)
$w_{ivt} \times$ L.cocoa output, kg	-413.5*	(213.11)	-680.8***	(179.97)
Direct effects of characteristics x	Yes		Yes	
N	570		570	
Treatment interactions: p-value	0.07		0.00	
Adopter interactions: p-value			0.00	

Notes: dependent variable is cocoa output, in kgs. Individual controls for interacted characteristics included in all specifications. All columns include village-year fixed effects. Heteroskedasticity-robust standard errors reported, clustered at village level.

To provide a full characterization of the economic implications of the heterogeneity in returns, we use a quantile treatment effects approach to show that the heterogeneity in returns we do observe has economic implications. Differences in quantiles of the distribution of outcomes under treatment and control can be interpreted as quantiles of the treatment effect only under the assumption of perfect positive dependence. In this case, the treatment does not change the ranking of outcomes, so that the first quantile of the distribution without treatment, Y_0, represents the counterfactual for individuals in the first quantile of the distribution with treatment, Y_1. With treatment and control groups of different sizes, we compare impacts across quantiles of the outcome distribution, rather than directly matching individuals.

To do so, we continue with the identification strategy of Section 4.1. In particular, we estimate quantile treatment effects using a model that includes a dummy variable for presence of Abrabopa in the village in the year under study, z_{vt}, as well as a dummy identifying those individuals who join Abrabopa in the first year of exposure in their village, m_{iv}. Treatment is denoted by the interaction of early adopters with village-level exposure: $w_{ivt} = m_{iv} z_{vt}$. We now explicitly adopt a random coefficients framework to estimate the regression model,

$$y_{ivt} = \beta_{i0} + \beta_{iz} z_{vt} + \beta_{im} m_{iv} + \beta_i w z_{vt} m_{iv} + e_{ivt}. \tag{12}$$

In the potential outcomes framework of equations (6) and (7), $\beta_{iw} = \mu_1 - \mu_0 + \varepsilon_{1ivt} - \varepsilon_{0ivt}$. Under assumptions A1, A2, and PPD, quantile treatment effects on the treated are identified by the coefficient β_{iw} on this interaction term at the corresponding quantiles of the outcome distribution.[14]

[14]For quantile q, this is the treatment effect on treated individual i such that $i =$

Results are presented in Figure 2. Given that we only observe 170 individuals after one year of treatment, the figure presents estimates for vigintiles (20-quantiles) of the outcome distribution among the treated. For ease of interpretation, we include a horizontal line at an impact of 230.8 kg, the sales volume required to repay the Abrabopa loan in the years studied. In the absence of any other costs of complementary inputs, individuals make a profit if and only if they experience a treatment effect in excess of this point.

Figure 2: Quantile treatment effects

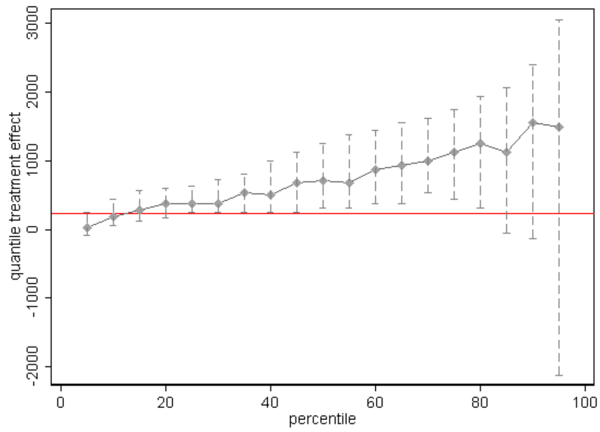

Notes: Figure illustrates estimated quantile treatment effects and associated bootstrap (90%) confidence intervals. Non-parametric, block bootstrap confidence interval based on 400 repetitions, re-sampling at village level. Horizontal line at 230.8 kg indicates cocoa output required to repay direct cost of inputs.

Confidence intervals are estimated by a block bootstrap, with resampling conducted at the village-year level to account for potential non-independence in outcomes within these sampling units. Although we can reject the null hypothesis of a treatment effect of zero for all but tails of the distribution, we fail to reject the hypothesis of a treatment effect above the break-even point for individuals at or below the 25^{th} percentile and above the 80^{th} percentile. (It should be noted that estimates at the upper end of the distribution are relatively imprecise, due to a small number of high-output farmers in our sample.)

Figure 2 suggests that, in this context, an exclusive emphasis on mean impacts misses important features of the distribution. In economic terms, there appears to be substantial heterogeneity in returns to program participation; in general, returns appear to be higher for more productive farmers.

$\arg\inf_{y_{i1}} F_1(y_1|w_{ivt} = 1) > q$.

Crucially, we fail to reject the hypothesis of zero returns, net of repayment costs, for 7 of the 20 vigintiles of the treated population.

Both time-varying and time-invariant heterogeneity in returns may matter for the sustained use of a new technology. Even shocks that farmers know to be transient can affect future technology use, if they impact on farmers' willingness to take on risk (Dercon & Christiaensen forthcoming). On the other hand, time-invariant heterogeneity may affect farmers' willingness to continue through a process of experimentation—so long as farmers are uncertain about these returns prior to membership. We address the questions of whether observed heterogeneity in returns can help to explain program retention, and through what mechanism, in the following section.

5 Heterogeneity and continued adoption

To test learning and prudence mechanisms underlying the relationship between the magnitude of the treatment effect experienced and the farmer's decision to continue adoption of the technology, we construct a proxy for treatment effects at the individual level, and show that this individual measure correlates with retention, even after conditioning on possible confounding factors. Of course, since each farmer can make one and only one membership decision in a given period, the treatment effect that she experiences is never directly observable. This is the core of the "fundamental problem of causal inference" (Holland 1986).

Under certain (necessarily untestable) assumptions, individual measures of the actual treatment effect can be constructed. Our aim here is more limited: we seek to construct a proxy for the realized treatment effect, and to argue that variation in this proxy across individuals is correlated with variation in the true treatment effect. The purpose of this is to test the hypothesis that the decision to renew membership is affected by the realization of current output. We will test this using a binary choice model of the form

$$w^*_{iv,t+1} = \phi_0 + \phi_\tau \tau_{ivt} + \phi_x x_{ivt} + u_{iv,t+1} \tag{13}$$

where membership in period $t+1$ is chosen if $w^*_{ivt} \geq 0$. The variable τ_{ivt} represents individual i's realized return in period t, and x_{ivt} represents a vector of controls for potential confounding factors, to be descried below. Under the assumption of normality of the error term, $u_{iv,t+1}$, the parameters of equation (13) can be estimated as a probit.

Because the individual's idiosyncratic return, τ_{ivt}, is not directly observable, we must estimate equation (13) with a proxy measure of the idiosyncratic return. For each individual in the treated group, we use two-period changes in cocoa output, $\hat{y}_{ivt} = y_i vt - y_{iv,t-2}$, as the basis for such a proxy for the experienced treatment effect. Following the notation of potential outcome equations (6) and (7), this will correspond to the true treatment effect,

τ_{ivt}, only if $\lambda_{vt} = \lambda_{v,t-2}$ and $\varepsilon_{0ivt} = \varepsilon_{0iv,t-2}$. This would require both that any village-level components of production are constant over time, and that individuals' idiosyncratic output in the absence of Abrabopa is the same in the pre-exposure period, $t-2$, as it is in the first period of exposure, t.

Neither of these assumptions will hold in practice: there will be village-level characteristics (such as rainfall) that vary across years, as well as shocks to potential output in the absence of fertilizer. Consequently the proxy $\hat{y}_{ivt}$ will be composed of the true treatment effect for individual i in period t, $\tau_{ivt} = y_{1ivt} - y_{0ivt}$, plus two terms that reflect the failure of these assumptions:

$$\hat{y}_{ivt} = \tau_{ivt} + \underbrace{(\lambda_{vt} - \lambda_{v,t-2})}_{\hat{\lambda}_{vt}} + \underbrace{(\varepsilon_{0ivt} - \varepsilon_{0iv,t-2})}_{\hat{\varepsilon}_{0ivt}}. \tag{14}$$

The correlation between this proxy and future membership decisions, $w_{iv,t+1}$, will reflect the effect of realized returns on individual decisions, as well as any correlation between the terms $\hat{\lambda}_{vt}$ and $\hat{\varepsilon}_{0ivt}$ and other determinants of membership. Thus measurement error in the proxy $haty_{ivt}$ creates a potential source of bias beyond the threat of omitted factors correlated with both the true, idiosyncratic treatment effect, τ_{ivt}, and the sustained adoption decision. We take up these considerations below.

Before turning to estimation of equation (13), we validate this measure by comparing the quantile treatment effects estimated in Figure (2) to the mean value of $\hat{y}_{ivt}$ for individuals in a neighborhood of the same vigintile, as was done with retention rates above. For example, we compare the quantile treatment effect at the 5^{th} percentile of the outcome distribution among the treated, $F_1(y_{ivt})$, to the mean value of the individual measure $\hat{y}_{ivt}$ between the 2.5^{th} and the 7.5^{th} percentiles of the treated population. If the perfect positive dependence assumption holds—so that quantile treatment effects can be interpreted as the average response at that quantile—and if our proxy is a relatively precise estimate of the individual treatment effect, then the two measures should correlate closely. Figure (3) shows the relationship between these outcome measures for vigintiles of the outcome distribution among the treated. The measures of these two measures of impact are indeed closely related; they have a correlation coefficient of 0.76.

Given the evidence to support this measure, $\hat{y}_{ivt}$, of individual treatment effects, we turn to estimation of equation (13). Probit coefficients are presented in Table 4, which restricts attention to the continuation decision of farmers who have joined Abrabopa in the first year of exposure in their village. The dependent variable in all specifications is the farmer's membership decision in the second year of exposure, which we observe through administrative records. In column (1), we demonstrate a positive relationship between a farmer's gain in output over the two-year period ending in their first year of membership and their membership decision in the second year. This effect is substantial, implying that a one standard deviation in-

Figure 3: Validating alternative measures of impact: quantile treatment effects versus mean two-period changes

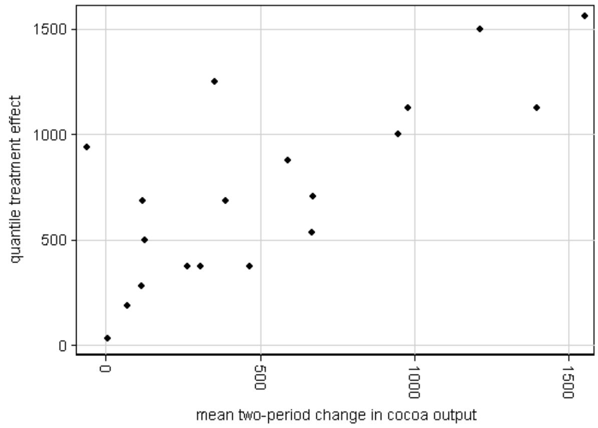

Notes: Horizontal axis is defined change in cocoa output among treated farmers. The local average of this measure within a neighborhood of a given vigintile (20-quantile) is plotted against the corresponding quantile treatment effect.

crease in this measure of the treatment effect is associated with a 13 percent increase in the likelihood of remaining in the program.

Columns (2) and (3) of Table 4 test between the precautionary savings and learning mechanisms as explanations for the relationship between realized heterogeneity and subsequent adoption. In column (2), current and past output levels are allowed to enter independently. Under the precautionary savings mechanism, both should enter the adoption decisions with the same sign, but this is not the case. The negative sign on lagged (pre-adoption) output is consistent with the view that farmers update their beliefs about the traditional technology based on recent experiences as well. In column (3), we provide an additional test by expanding the sample to include second-year as well as first-year Abrabopa members. For second-year members, their lagged output levels occur under the hi-tech package as well, so the learning hypothesis would suggest that lagged output should enter positively for second-year members, and negatively for first-year members. This is precisely what is found: the effect of past output is negative only for new members, while it is (weakly) positive for second-year members.

From these results, we conclude that the data are consistent with a model in which individuals experience heterogeneous returns to participation in the Abrabopa program, and they condition their decisions to remain within the program on these experienced returns.[15] The association between measures

[15]Sustained membership in Abrabopa is not equivalent to sustained use of the technology

Table 4: Individual output gains and program retention

	(1) New	(2) New	(3) All
Δcocoa output	0.510**		
	(0.21)		
cocoa ouput		0.433*	0.264**
		(0.23)	(0.12)
L.cocoa output		-0.591***	-0.301***
		(0.22)	(0.11)
L.member $\times$ L.cocoa			0.322***
			(0.12)
Farmer characteristics	Yes	Yes	Yes
Observations	225	225	332

Notes: Columns (1)–(3) report probit coefficients, with dependent variable **1**[continued adoption]. Column (4) reports OLS regression, with dependent variable . Standard errors clustered at village level in all specifications. Output variables rescaled by dividing by 1,000 kg prior to estimation. Farmer controls include gender, education, HH size, farm size, tree age and type.

of returns and future membership decisions does not appear to be driven by the 'supply' of the program—although repayment failures do cause expulsion in accordance with the program's bylaws. Nor is this association driven by current income alone: those who experience greater gains upon joining the program are more likely to remain, even at a given level of current production.

6 Conclusions

In many developing countries, the persistent adoption of agricultural technologies with high average returns is believed to be one of the principal policy challenges. Experimental and observational studies documenting these high average returns have led to a puzzle: why, if returns are so high, do farmers not adopt - and sustain - the use of agricultural technologies such as fertilizer or hybrid seeds?

We have presented evidence consistent with the view that high average returns mask substantial, persistent heterogeneity in realized returns. Farmers at the low end of the distribution of cocoa production exhibit particularly low returns - so much so that we are unable to reject a zero economic re-

it provides. Examination of whether treatment heterogeneity correlates with continued use of the hi-tech package is not possible in this context, since we observe only whether broad categories of inputs such as fertilizer were used, but not their exact make or proportions on a given plot; moreover, we only have these production data for 29 of the first-year members in our sample who subsequently drop out of Abrabopa. Among these 29 farmers, 14 report using fertilizer in some form and quantity.

turn for the bottom quarter of the distribution. We have shown evidence to support the view that this heterogeneity matters, economically speaking: farmers exhibiting low returns, by various measures, are less likely to continue use of the technology.

This ex post heterogeneity likely reflects a combination of persistent heterogeneity and time-varying riskiness of returns. Both may have consequences for decisions to sustain or to disadopt a given technology. Farmers may experiment in order to learn about their specific, time-varying component. Time-varying shocks to these returns may lead Bayesian farmers astray. But they may also affect adoption decisions directly, either because they lead to non-repayment of loans, or because they force farmers to deplete buffer stocks of savings, or social or other forms of collateral, in order to repay. In our data and context, there is some evidence that the latter effect is not driving the relationship between realized returns and disadoption: our measure of realized returns retains its statistical and economic significance even after conditioning on the level of past output.

The distinction between the effect of transient and persistent heterogeneity in returns seems a valuable area for future work. If persistent heterogeneity is quantitatively important, policymakers will need to be cautious in advocating widespread adoption of such technologies. Even when average returns are high, many farmers may stand to lose.

References

Bandiera, O. & Rasul, I. (2006), 'Social networks and technology adoption in Northern Mozambique', *Economic Journal* **116**(514), 869–902.

Bardhan, P. & Udry, C. (1999), *Development Microeconomics*, Oxford University Press, Oxford.

Cocoa Abrabopa Association (2009), 'Developments and challenges of the cocoa abrabopa association', Presentation to CSAE/COCOBOD workshop on "Improving productivity among Ghaniana cocoa farmers through group lending", Accra, Ghana.

Conley, T. G. & Udry, C. R. (2010), 'Learning about a new technology: Pineapple in Ghana', *American Economic Review* **100**(1), 35–69.

Conley, T. & Udry, C. (2001), 'Social learning through networks: The adoption of new agricultural technologies in ghana', *American Journal of Agricultural Economics* **83**(3), 668–673.

Dercon, S. & Christiaensen, L. (forthcoming), 'Consumption risk, technology adoption and poverty traps: Evidence from Ethiopia', *Journal of Development Economics* .

Duflo, E. (2006), Poor but rational?, *in* A. V. Banerjee, R. Bénabou & D. Mookherjee, eds, 'Understanding Poverty', Oxford University Press, Oxford, chapter 24, pp. 367–378.

Duflo, E., Kremer, M. & Robinson, J. (2008), 'How high are rates of return to fertilizer? Evidence from field experiments in Kenya', *American Economic Review* **98**(2), 482–488.

Duflo, E., Kremer, M. & Robinson, J. (2009), 'Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya', NBER Working Paper 15131.

Field, E. (2005), 'Property rights and investment in urban slums', *Journal of the European Economic Association* **3**(2/3), 279–290.

Foster, A. D. & Rosenzweig, M. R. (1995), 'Learning by doing and learning from others: Human capital and technical change in agriculture', *The Journal of Political Economy* **103**(6), 1176–1209.

Foster, A. D. & Rosenzweig, M. R. (2010), 'Microeconomics of technology adoption', Yale University, Economic Growth Center, Center Discussion Paper 984.

Heckman, J. J. (2010), 'Building bridges between structural and program evaluation approaches to estimating policy', *Journal of Economic Literature* **48**(2), 356–398.

Heckman, J. J. & Navarro-Lozano, S. (2004), 'Using matching, instrumental variables, and control functions to estimate economic choice models', *Review of Economics and Statistics* **86**, 30–57.

Heckman, J. J., Urzua, S. & Vytlacil, E. (2006), 'Understanding instrumental variables in models with essential heterogeneity', *Review of Economics and Statistics* **88**(3), 389–432.

Holland, P. W. (1986), 'Statistics and causal inference', *Journal of the American Statistical Association* **81**(396), 945–960.

Jovanovic, B. & Nyarko, Y. (1996), 'Learning by doing and the choice of technology', *Econometrica* **64**(6), 1299–1310.

Kimball, M. S. (1990), 'Precautionary savings in the small and in the large', *Econometrica* **58**, 53–73.

Moser, C. M. & Barrett, C. B. (2006), 'The complex dynamics of smallholder technology adoption: the case of SRI in Madagascar', *Agricultural Economics* **35**, 373–388.

Munshi, K. (2004), 'Social learning in a heterogeneous population: technology diffusion in the Indian Green Revolution', *Journal of Development Economics* **73**(1), 185–213.

Opoku, E., Dzene, R., Caria, S., Teal, F. & Zeitlin, A. (2009), 'Improving productivity through group lending: Report on the impact evaluation of the Cocoa Abrabopa Initiative', Centre for the Study of African Economies, technical report no. REP2008-01.

Pearl, J. (2009), 'Causal inference in statistics: an overview', *Statistics Surveys* **3**, 96–146.

Suri, T. (2011), 'Selection and comparative advantage in technology adoption', *Econometrica* **79**(1), 159–209.

World Bank (2008), *World Development Report 2008: Agriculture for Development*, The International Bank for Reconstruction and Development, Washington, D.C.

Zerfu, D. & Larson, D. F. (2010), 'Incomplete markets and fertilizer use', World Bank, Policy Research Working Paper 5235.