To Specialize or Diversify: Agricultural Diversity and Poverty Persistence in Ethiopia

Jeffrey D. Michler[1] and Anna Leigh Josephson[2]

[1] *Department of Agricultural and Consumer Economics, University of Illinois, Urbana, IL*
[2] *Department of Agricultural Economics, Purdue University, West Lafayette, IN*

Abstract

This article answers the empirical question: what is the relationship between the choice to specialize or diversify in crop production and household poverty status? We use household panel data from Ethiopia and a recently developed parametric method for estimating dynamic binary response models with endogenous contemporaneous regressors. Our results provide evidence that households which grow a diverse set of crops are less likely to be poor. Additionally, crop diversity reduces the probability that a household will fall into poverty and reduces the probability that a household will remain in poverty. We conclude that policies which encourage households to specialize in cash crops may be counter-productive while policies which encourage crop diversification may reduce poverty.

1 Introduction

Why do some farming households choose to specialize in crop production while other households seek to diversify their production? *Ex post*, households seeking to maximize their welfare would always choose to specialize in a single crop. However, there are numerous constraints and uncertainties in the agricultural production process that, *ex ante*, may result in households choosing to cultivate a diverse crop portfolio. Which strategy is welfare maximizing, specialization or diversification, is an empirical question and depends on the economic and agronomic environment in which a household cultivates its land. Understanding the relationship between cropping decisions and household welfare is an important first step in developing effective anti-poverty policies. Furthermore, the effect of crop mix on poverty has implications for policies designed to mitigate food insecurity, manage production risk, and assist in adaptation to climate change.

In this study, we estimate the effects of diversification in crop cultivation on household welfare in Ethiopia. We use poverty as our measure of household welfare because it provides insight regarding the distributional effects of crop diversification - whether crop choice can pull poor households out of poverty. Furthermore, the Millennium Development Goals make poverty reduction the central objective of development. Consistent with this, we follow Christiaensen et al. (2011) in focusing our analysis on poverty reduction and not household income growth. We seek to answer two specific research questions on this issue. First, does growing a diverse set of crops reduce the probability that a household is below the poverty line? Second, does an increase in crop diversity reduce the probability that a poor household will rise out of poverty or that a non-poor household will fall into poverty?

Assessing the impact of crop diversity on poverty is not straightforward. Estimation is complicated by state dependence in the binary outcome and by two potential sources of endogeneity. First, it is likely that there are unobserved household characteristics (e.g., skill, education, entrepreneurship) related to both cropping decisions and poverty status. Second, the decision to diversify or specialize may be driven by negative shocks that also increase the probability of a household being poor. To account for the endogeneity in cropping decisions we employ a control function approach applied to a static response model by Papke and Wooldridge (2008) and recently extended to a dynamic setting by Giles and Murtazashvili (2013). To account for the initial conditions problem and the existence of unobserved heterogeneity we employ a random effects model developed by Wooldridge (2005).

We find that growing a diverse set of crops reduces the probability of a household being in poverty. Specifically, a 10 percent increase in crop diversity reduces the probability of being poor by 5.2 percent. Furthermore, a 10 percent increase in crop diversity reduces the probability that a household will remain in poverty by 4.8 percent. Finally, a 10 percent increase in crop diversity reduces the probability that a household will fall below the poverty line by 5.5 percent. We conclude that agricultural diversification, and not specialization, is associated with poverty reduction. While many households in Ethiopia have switched to cultivation of a single cash crop like chat, coffee, or sesame, we find that households which cultivate several different types of crops are less likely to be poor.

Our results provide much needed evidence regarding the increasingly common policy prescription of agricultural diversification. With the slowing of crop production growth rates in many parts of the world, development agencies have begun to shift focus from promotion of a few staple grain crops to policies designed to encourage diversification. Food and Agriculture Organization (FAO) policy supports crop diversification in the belief that it is an effective strategy in dealing with issues as varied as food and nutrition security, poverty alleviation, employment generation, judicious use of land and water, sustainable agricultural development, and environmental and ecological management (FAO, 2012). A series of country level case studies undertaken by the FAO recommend methods to increase diversification and discuss various constraints to diversification but provide no quantitative evidence to support these positions (Hazra, 2001; Mengxiao, 2001; Kaguongo et al., 2013). Similarly, the International Food Policy Research Institute (IFPRI) has begun to move away from policies that encourage the focused production of staple crops. Recent IFPRI publications have speculated that growth in agricultural incomes will require more diversification by farming households but provide no evidence of this (Taffesse et al., 2011). Our research provides the first clear quantitative evidence in support of these policy positions.

In addition to our contribution to these recent shifts in policy, our study contributes to two separate strands of literature. The first strand is on the relationship between diversification, risk mitigation, and income. The second strand of literature to which we contribute is on household coping strategies to increase food security and adapt to climate change.

The literature on the relationship between diversification, risk, and income generally seeks for determinants of diversification.[1] Several studies find a positive relationship between household income and agricultural diversification (Ellis, 1998, 2000; Barrett et al.,

[1] This literature can be divided into two subsets: studies which focus on the relationship between diversification and risk and studies that focus on the relationship between diversification and poverty. Numerous seminal studies have focused on the relationship between risk and diversification. These include Rosenzweig and Binswanger (1993), Alderman and Paxson (1994), Dercon (1996), Little et al. (2001), and Di Falco and Perrings (2005). We focus on the less studied relationship between diversification and income.

2001; Caviglia-Harris and Sills, 2005). Contrary evidence exists, however, indicating that diversification may be more associated with poverty. Feder et al. (1985) argues that income drives diversification, generating income gains for the already wealthy and resulting in a poverty trap for those at the bottom. Similar evidence is presented in Weinberger and Lumpkin (2007). Additionally, a study of crop choice in the Western Punjab beginning with the early twentieth-century found that farmers specializing in crops in which they had a comparative advantage resulted in higher productivity and increased income (Kurosaki, 2003).

Instead of estimating the determinants of diversification, as much of the previous literature has done, we analyze the role diversification plays as a determinant of poverty. Few studies have taken this approach. Among studies that do, most treat diversification as an independent, exogenous variable (Bigsten and Tengstam, 2011; Bezu et al., 2012; Baird and Gray, 2014). By failing to control for endogeneity in the choice to diversify, or control for the initial condition of households, these studies provide only suggestive results regarding the association between diversification and poverty. Our econometric methodology, which includes instrumenting for diversification, resolves these issues and provides clear evidence that diversification reduces poverty.

A second body of literature to which we contribute is on household coping strategies and adaptation to food insecurity and climate change. Despite evidence that farms are becoming less diversified with time (Bradshaw et al., 2004), diversification has come to be viewed as an important way to increase food security. This is particularly true when faced with increasing variability in production due to climate change. Several studies conducted in Ethiopia find that combinations of different farming techniques, including greater crop diversity, may mitigate food insecurity and help farmers cope with climate change (Di Falco et al., 2010, 2011; Bezabih and Sarr, 2012; Di Falco and Veronesi, 2013). Our results provide further evidence that diversity, not specialization, is a viable way to deal with the exegencies of being poor.

2 Data

Our empirical analysis uses four years of panel survey data collected in the Ethiopian Rural Household Survey (ERHS) by the Economics Department, Addis Ababa University, the Centre for the Study of African Economics at Oxford University, and IFPRI. The data covers approximately 1,500 households in 15 villages over the ten year period from 1994 to

2004. The villages were selected to represent the diversity of the farming systems in the country, including the ox-plough cereal areas in the northern and central highlands and the enset-growing areas in the southern lowlands (Dercon and Hoddinott, 2011). However, as the data only considers rural, non-pastoral households, it is not considered to be nationally representative. We use four waves of the survey: 1994, 1995, 1997, and 2004. We use a balanced panel of 1,250 households from 14 villages.

2.1 Poverty Status and Household Characteristics

Our dependent variable is a binary indicator that measures if the household was below the poverty threshold. Our decision to use a binary indicator is motivated by three factors. First, the primary concern of many development agencies is raising households out of poverty. By using a binary poverty indicator, our results are easily interpreted and speak directly to the mandate of many development stakeholders. Second, consumption expenditure data in the ERHS is incomplete.[2] Due to heterogeneity in age and quality of durable and non-durable goods (as well as an inability to establish market prices for these goods), consumption data in the ERHS is limited to only food items and non-investment non-food items. By using a binary indicator for poverty we are able to minimize measurement error in calculating our dependent variable. Third, while use of a continuous dependent variable might provide more precision in coefficient estimates, our use of a binary dependent variable does not require any sacrifice in the accuracy of coefficient estimates. Thus, our use of a binary poverty indicator instead of a continuous consumption variable allows us to reduce measurement error in our dependent variable, makes our results easily interpretable, and has no cost in the accuracy of our estimates.

To construct our poverty indicator we follow Dercon et al. (2009) in using a cost-of-basic-needs approach. Food poverty is considered to be consumption of a bundle of food items that provide less than 2,300 kcal per adult per day. To this is added a bundle of

[2] While the ERHS includes a rich set of household characteristics and agricultural production variables, income and expenditure data are problematic. Previous research using the ERHS has noted that income data is generally underreported. While underreporting of income is a common feature of surveys in developing countries, Bezu et al. (2012) note that underreporting in the ERHS is severe. Average household consumption expenditure per adult equivalent is $125 while household income per adult equivalent is $68. (Income and expenditure are given in USD at 2000 constant prices). Additionally, income data was collected at four month intervals which do not perfectly align with growing seasons, resulting in great likelihood of measurement error for households engaged in seasonal employment. This issue is especially acute for the 1997 round which, unlike the 1994, 1995, and 2004 rounds, was collected in the immediate post-harvest period. Due to the issues with household income data, many studies using the ERHS rely on consumption expenditure data to determine household well being (Dercon et al., 2009; Dercon and Christiaensen, 2011; Bezu et al., 2012).

non-food items as in Ravallion and Bidani (1994).[3] In our sample, 42 percent of households are below this threshold. Relevant literature on the topic finds about 40 percent of rural Ethiopian households live below the poverty line (Bigsten et al., 2002; Bogale et al., 2005). This suggests that our poverty term is broadly representative. The share of households living in poverty in each village is highly variable. Imdibir has the largest share of poor households, with 73 percent of households falling below the poverty line. Sirbana Godeti has the smallest share of poor households, with only 13 percent of households living below the poverty line (See Table 1).

In this study we are particularly interested in the dynamics of poverty - how poverty responds to changes in crop diversity. However, given our binary poverty indicator and with only four observations per household, informative measures of household poverty dynamics are difficult to construct. To that end, we focus on poverty dynamics at the village level. Figure 1 displays the bivariate kernel density contours of the mean poverty level in each village in each survey year compared with the previous survey year.[4] To this we have added a 45° line. Villages that, from one survey year to the next, have experienced an increase in household poverty are below the 45° line. Encouragingly, much of the mass of the poverty distribution lies above the 45° line, indicating that most villages saw a reduction in poverty over the survey period.

This reduction in village level poverty does not, at first glance, appear to be correlated with changes in crop diversity. Figure 2 is a scatter plot of changes to village poverty and changes to average village crop diversity from one survey year to the next. To this we add a linear trend line whose slope is not significantly different from zero. Taken together, Figure 1 and 2 provide suggestive evidence that, on average, households in Ethiopia are becoming less poor and that this dynamic has no obvious correlation with changes in cropping decisions.

In addition to our household poverty indicator, we also use a selection of household demographic characteristics to evaluate and control for the relationship between crop diversity and poverty status. These include household size, land per capita, and the years of education obtained by the head of household. We also include an indicator variable for whether or not the head of household is female. Descriptive statistics for these variables, as well as for poverty status, can be found in Table 2.

[3]Additional details on the specifics of each consumption bundle and the various sources of price data can be found in Dercon and Krishnan (1996, 2003); Dercon et al. (2009).

[4]We use households observations from each year to calculate mean poverty levels for each village in each year.

2.2 Crop Diversity Index

To measure crop diversity we generate a crop diversity index, using detailed cropping data from the survey. Our index measures the total number of crops a household grows in a year (n_{it}), in relation to the total number of unique crops grown within the village in that year (N_{jt}).

$$div_{it} = \left(\frac{n_{it}}{N_{jt}}\right)^2 \tag{1}$$

This approach has several advantages to alternative methods of index construction. First, by using the total number of crops presently grown in the village as the denominator in our index, we can control for village specific agronomic environments. Thus, a household's crop diversity, or lack thereof, is not measured against the agricultural practices of households in other villages, but against the common practices of it's own village. Households living in agronomic zones that only allow for a limited number of crops are not penalized for only growing a few crops.[5] Second, we update the denominator each survey year to allow for changes to the environment that might increase or decrease the number of viable crops in a village. This allows us to accommodate the insight that in each village in each year a different strategy may be welfare maximizing. Third, by measuring a household's diversity in relation to the total number of crops grown in the village, we can capture the inequality between households in a given community. In a recent paper Thiede (2014) shows that adverse environmental events have heterogeneous effects on households within a village, disproportionately harming poorer households. By constructing our index in relation to village practices, we can explore the interaction between poverty status and crop diversity within the village.

As our index is a ratio, lower values indicate a more agriculturally specialized household relative to the total number of distinct crops grown in the village and higher values indicate a more diversified household relative to the village. We include in our diversity count 33 different crops, including staple crops such as teff, maize, and barely, as well as cash crops such as linseed and sesame. Several types of tree crops are also included, such as coffee, chat, and enset.[6] Table 1 shows summary statistics of crop diversity for each village as well

[5] Ethiopia exhibits distinct agronomic zones: the highlands and the lowlands. The highlands are distinguished by steady rainfall and plateaus which are conducive to a variety of crops, while the lowlands generally have shallow soils, little rainfall, and more limited crop choices (Pankhurst, 2009).

[6] We do not include livestock in our index. However, the inclusion of livestock does not fundamentally change our results.

as the maximum and minimum number of crops grown.

Similar to our examination of poverty dynamics at the village level, Figure 3 displays the bivariate kernel density contours of the mean level of diversity in each village in each survey year compared with the previous survey year.[7] Villages that, from one survey year to the next, have experienced an decrease in crop diversity are below the 45° line. Much of the mass lies above the 45° line, indicating that most villages saw an increase in crop diversity over the survey period. Despite the increase in crop diversity and decrease in poverty over the survey period, as we saw in Figure 2, there does not appear to be a strong correlation between these events. This result could be due to several reasons. One is that while, on average, poverty fell and diversity increased, the villages (and households within villages) that reduced poverty were not the same as those that increased crop diversity. A second reason is that our analysis is bivariate and fails to control for confounding factors such as endogeneity of the treatment and the initial conditions problem. Our empirical strategy addresses both of these confounding factors.

3 Empirical Strategy

Estimation of the relationship between crop diversity and poverty faces numerous econometric issues. These include two potential sources of endogeneity. The first is the potential for unobserved heterogeneity and state dependence in our dynamic setting. The second is a reverse causality problem in that poverty may be driving crop choice. In this section we discuss these issues and briefly outline our method for dealing with them, which was developed by Giles and Murtazashvili (2013).

3.1 A Dynamic Binary Response Panel Data Model

The first potential source of endogeneity is the existence of unobserved household characteristics or unobserved shocks affecting both cropping decisions and poverty status. In a dynamic panel data model, how unobserved characteristics affect the initial condition is an important problem to address. We use a control function approach introduced by Smith and Blundell (1986) and applied to a nonlinear setting by Papke and Wooldridge (2008).

We begin by assuming that for our binary response function, there is an underlying latent variable model:

[7]We use households observations from each year to calculate mean diversity within each village in each year.

$$y^*_{1it} = \mathbf{z}_{1it}\beta_1 + \beta_2 y_{2it} + \rho y_{1i,t-1} + c_{1i} + u_{1it} \tag{2}$$

where $y_{1it} = 1[y^*_{1it} \geq 0]$ for $t = 1, ..., T$, $\mathbf{z}_{1it}$ is a $1 \times (K-1)$ vector of exogenous variables, y_{2it} is an endogenous covariate, c_{1i} is an unobserved effect, u_{1it} is an idiosyncratic error term, and β_1, β_2, and ρ are parameters to be estimated.

Correct estimation of the model requires several assumptions. First, we assume that the dynamics in the model are correctly specified and $\mathbf{z}_{1it}$ is strictly exogenous conditional on the unobserved effect, c_{1i}. This assumption implies that the error term is serially uncorrelated. Second, we assume that we can model the endogenous covariate as a linear function of the following variables:

$$y_{2it} = \mathbf{z}_{1it}\delta_1 + \mathbf{z}_{2it}\delta_2 + c_{2i} + u_{2it} \tag{3}$$

where $\mathbf{z}_{2it}$ is a set of instrumental variables and u_{2it} is an idiosyncratic error term also free of serial correlation. Third, consistent with Mundlak (1978), we assume that the unobserved effect in the first stage equation, c_{2i}, can be replaced with its projection onto the time averages of all exogenous variables such that

$$c_{2i} = \bar{\mathbf{z}}_i\lambda + a_{2i} \tag{4}$$

where $\bar{\mathbf{z}}_i$ is a vector of time averages of $\mathbf{z}_i = (\mathbf{z}_{1it}, \mathbf{z}_{2it})$. Following Papke and Wooldridge (2008), we can use equation (4) to rewrite equation (3) as the linear reduced form equation:

$$y_{2it} = \mathbf{z}_{1it}\delta_1 + \mathbf{z}_{2it}\delta_2 + \bar{\mathbf{z}}_i\lambda + v_{2it} \tag{5}$$

where $v_{2it} = a_{2i} + u_{2it}$. Fourth, we assume that (u_{1it}, u_{2it}) has a bivariate normal distribution with mean zero and is independent of $\mathbf{z}_i$. This assumption allows us to write the error term in equation (2) as a function of the error term in equation (3):

$$u_{1it} = \theta u_{2it} + \epsilon_{1it} = \theta(v_{2it} - a_{2i}) + \epsilon_{1it} \tag{6}$$

where $\theta = \text{Cov}\,(u_{1it}, u_{2it})/\text{Var}\,(u_{2it})$ and ϵ_{1it} is an idiosyncratic error term free from serial correlation due to our first and second assumptions.

Given our four assumptions, we can rewrite equation (2) as:

$$\begin{aligned} y_{1it} &= 1\left[\mathbf{x}_{1it}\beta + c_{1i} + \theta(v_{2it} - a_{2i}) + \epsilon_{1it} \geq 0\right] \\ &= 1\left[\mathbf{x}_{1it}\beta + \theta v_{2it} + c_{0i} + \epsilon_{1it} \geq 0\right] \end{aligned} \tag{7}$$

where $\mathbf{x}_{1it} = (\mathbf{z}_{1it}, y_{2it}, y_{1i,t-1})$ contains our data, $\beta = (\beta_1', \beta_2, \rho)'$ is a vector of coefficients to be estimated, and $c_{0i} = c_{1i} - \theta a_{2i}$ is the composite unobserved effect. By including v_{2it} we have controlled for the endogeneity of y_{2it} in time period t. However, there may be feedback loops such that y_2 in other time periods may affect y_{1it}. Thus, while we have controlled for the endogeneity in y_{2it} caused by the unobserved effect c_{2i}, we have not yet controlled for the unobserved effect c_{0i}.

To control for c_{0i}, we adopt an approach similar to that used in equation (4). We assume the composite unobserved effect, c_{0i}, is independent of the initial condition, y_{1i0}, and the exogenous covariates, $\mathbf{z}_i$, but not v_{2i}:

$$c_{0i} = \alpha\bar{v}_{2i} + a_{1i} \tag{8}$$

where $\bar{v}_{2i}$ is a vector of time averages. This final assumption regarding the independence of the initial condition and the composite unobserved effect allows us to rewrite equation (7) as:

$$y_{1it} = 1\left[\mathbf{x}_{1it}\beta + \alpha\bar{v}_{2i} + a_{1i} + \epsilon_{1it} \geq 0\right]. \tag{9}$$

By including $\alpha\bar{v}_{2i}$ and a_{1i} equation (9) controls for the unobserved effects of c_{0i} and c_{2i} and is now free of endogeneity caused by unobserved heterogeneity.

We follow the two-step estimation procedure outlined in Giles and Murtazashvili (2013). First, we estimate equation (5) using pooled OLS. We save the residuals, $\hat{v}_{2it}$, from this reduced form equation and calculate $\bar{\hat{v}}_{2i} = \frac{1}{T}\sum_{t=1}^{T}\hat{v}_{2it}$. Next we estimate our probit model in equation (9) using the conditional MLE and including the residuals and their time averages as right hand side regressors. We bootstrap our standard errors since our second stage regression involves first stage residuals.

3.2 Identification of Crop Diversity

The second potential source of endogeneity in our regression equation is a reverse causality problem in that poverty status may affect crop choice. We control for the potential endogenous regressor by choosing instrumental variables which allow us to estimate equation (3).

To identify crop diversity we use the distance from each village to the nearest agricultural cooperative interacted with the lag of household land per capita.[8] Agricultural cooperatives in Ethiopia are vital conduits for the dissemination of seed, technology, and information. Cooperatives also operate as a home base for extension agents. Given the evidence on the role extension agents have in technology adoption in Ethiopia (Asrat et al., 2011; Krishnan and Patnam, 2013; Di Falco et al., 2011; Di Falco and Veronesi, 2013), proximity to a co-operative is likely to be associated with crop choice. The left panel of Figure 4 shows that changes in village crop diversity are negatively associated with the distance to an agricultural cooperative.[9] However, the non-parametric lowess plot makes clear that there are nonlinearities in the bivariate relationship. To account for potential nonlinearities in the relationship between distance to a cooperative and crop diversity we also include higher order terms as instruments.

While distance to the nearest agricultural cooperative is likely to be correlated with crop diversity, proximity to a cooperative is likely to be uncorrelated with household characteristics, such as household poverty. This is because of a government policy to establish complete geographic coverage of rural areas by cooperatives. While cooperative location is not random, neither is the government's choice of location determined by village size, village wealth, or local land quality, let alone a household's poverty level. Among the 14 villages used in our study, four villages host agricultural cooperatives while two villages are more than 15 km from a cooperative. Of the four villages with cooperatives within the village, two of those villages have poverty levels of over 50 percent while the other two villages have poverty levels below 40 percent. The right panel of Figure 4 provides further evidence by showing the relationship, or lack of, between changes in village poverty and the distance to an agricultural cooperative. The non-parametric lowess plot provides inconclusive evidence regarding the bivariate relationship. However, a linear trend line makes clear that there is

[8] We believe the use of distance to an agricultural cooperative is an improvement on instruments used in recent studies of agricultural households and adaptation strategies (such as crop diversification) in Ethiopia (Asrat et al., 2011; Di Falco et al., 2011; Di Falco and Veronesi, 2013; Krishnan and Patnam, 2013). In order to control for potential endogeneity in the relationship between adaptation strategies and outcome, Di Falco et al. (2011) and Di Falco and Veronesi (2013) use extension services. They argue that use of extension services is correlated with the decision to choose an adaptation strategy but is not correlated with the outcome of the strategy (output or revenue). We feel that correlation may exist between unobserved household characteristics (and therefore poverty) and the propensity to take advantage of extension services. If such a relationship exists, extension services would no longer be a valid instrument for agricultural diversity. Therefore, we prefer distance to agricultural cooperatives as an instrument over the more commonly used extension services instrument.

[9] The slope of the trend line is -0.0011 which is statistically significantly at the 99 percent level.

no significant correlation between changes in poverty and distance to a cooperative.[10]

Since distance to agricultural cooperative is a village level variable, we interact distance with the lag of household land per capita. This provides us with household level variance in our IV. We verify the validity of the our instrument by performing a simple falsification test: if the variable is a valid instrument, it will affect the crop choice decision, but it will not affect household poverty. To determine that our instrument is correlated with crop diversity, we estimate the reduced form equation (5) as:

$$div_{it} = AG_{it}\delta + \mathbf{x}_{it}\alpha + \bar{\mathbf{x}}_i\lambda + \mathbf{d}_t + \mathbf{v}_t * \mathbf{t}_t + e_{it} \tag{10}$$

where AG_{it} is a set of instruments based on the distance to the nearest government run agricultural cooperative, and $\bar{\mathbf{x}}_i$ are the time averages of the household variables in $\mathbf{x}_{it}$. We test several specifications for our instrument, including the distance to an agricultural cooperative, the distance squared and cubed, and distance interacted with the lagged amount of household land per capita. Results from these specifications are presented in Table 3. While only one of the quadratic terms is significant, an F-test fails to reject the null hypothesis that the four instruments are jointly equal to zero. Therefore, in our subsequent analysis we use the specification in column (3) of Table 3 for our first stage regression. We calculate the residuals and add them, along with their averages, as control variables in the binary response model.

To show that our distance, distance squared, and the lagged land interaction terms satisfy the exclusion restriction, we test for their significance in determining household poverty status. None of our instruments are significant factors in determining the probability that a household is below the poverty line (See Table 4). Thus, our instruments satisfy the simple falsification test: they are correlated with crop diversity while also being uncorrelated with household poverty status other than through their effect on crop diversity.

3.3 Estimating the Impact of Crop Diversity on Poverty

We estimate the dynamic binary response model for the likelihood that household i in village j falls below the poverty line at time t as:

$$pov_{it} = 1[\beta_1 pov_{it-1} + \beta_2(pov_{it-1} * div_{it}) + \beta_3 div_{it} + \mathbf{x}_{it}\alpha + \mathbf{d}_t + \mathbf{v}_t * \mathbf{t}_t + u_t + \epsilon_{it} \geq 0] \tag{11}$$

[10] The slope of the trend line is 0.0001 which is not statistically significantly different from zero.

where pov_{it} is a binary indicator for whether the household is poor. A household's poverty status is affected by poverty status in the previous period, pov_{it-1}, our measure of crop diversity, div_{it}, a vector of household characteristics, $\mathbf{x}_{it}$, year dummies, $\mathbf{d}_t$, and village time trends, $\mathbf{v}_t * \mathbf{t}_t$. Our approach allows us to address potential correlation between unobserved household heterogeneity, u_t, and the other covariates. We also control for endogeneity of crop diversity with our village level instrument interacted with the lag of household level land per capita.

In applications similar to ours, Wooldridge (2005) and Giles and Murtazashvili (2013) present a correlated random effects model. The correlated random effects approach relaxes the assumption that the unobserved heterogeneity is uncorrelated with the exogenous variables. However, the need to adopt a correlated random effects model, instead of a pure random effects model, is not mandatory. We test for the existence of correlation in our errors by using a standard ANOVA test. We find that our error terms are free of serial correlation ($p\text{-}value = 0.557$) and so proceed with our empirical analysis using a pure random effects approach.

To test the robustness of our results, we estimate our model across several specifications (See Table 5). First, we treat crop diversity as exogenous and show results from both the linear probability (LPM) and probit estimations. We next control for the potential endogeneity of crop diversity by introducing our instrumental variable and show results for both the LPM and control function (CF) implementations.

4 Discussion

We estimate equation (11) but first treat diversity as exogenous. We give estimation results from both the LPM and the probit model in columns (1) and (2) of Table 5. In these specifications the relationship between household crop diversity and household poverty is not statistically significant. This implies that crop choice has no effect on whether a household is in poverty. Given the descriptive evidence shown in Figure 2, it is unsurprising that we find no correlation between crop diversity and poverty status in the exogenous specifications.

We next treat diversity as endogenous and introduce the distance to cooperative instrument, which allows us to identify the relationship between changes in crop diversity and changes in household poverty status. Columns (3) and (4) in Table 5 give results for the LPM and the control function approach. In these specifications, the relationship between household crop diversity and household poverty is negative and significant at the 1 and 5 percent levels, respectively. The upward bias in the coefficients when diversity is treated as exogenous indicates the need for an estimation strategy which permits the identification in a dynamic binary response model where there are endogenous regressors. This suggests an improvement on previous studies, which have often treated crop choice as exogenous (Bezu et al., 2012; Bigsten and Tengstam, 2011; Baird and Gray, 2014).

In Table 6 we report average partial effects (APEs) for both the exogenous and endogenous specifications of the LPM, probit, and control function approach. In most cases, APEs are averaged across both the cross-section of the covariates and time. However, due to the presence of the interaction term between the diversity index and lagged poverty status, calculation of the appropriate APEs requires some attention. The APE for lagged poverty status is calculated at the average across observed values of the diversity index in the data. To explore the effects of crop diversification on poverty persistence, we calculate the APE of crop diversity when lagged poverty status equals zero for all households, when lagged poverty status equals one for all households, and the average across the observed values of poverty status in the data.

As one might expect, the coefficients on lagged poverty are significant and positive in all specifications. Households that are in poverty in one period are more likely to remain in poverty in the next period. This indicates a strong persistence in poverty among the sample households that is robust to various specifications and estimation techniques. Here again, the probit model and the LPMs exhibit upward bias in their estimation of coefficients. Further, the endogenous LPM overestimates the effects of diversity, relative to the control function specification. This suggests that models which fail to explicitly control for the initial conditions problem overstate the importance of poverty persistence.

A somewhat surprising result from our models is that the interaction term between the diversity index and poverty status is not significant. This suggests that changes in diversity do not disproportionately impact wealthy households compared to poor households. There are numerous examples in the literature regarding treatments that have heterogeneous effects on households and that such heterogeneity is driven by differences in household wealth levels. Giles and Murtazashvili (2013) find that households in poverty are more likely to be impacted by growth in village migrant networks compared to households that are not in poverty. Thiede (2014) finds that rainfall shocks have a larger detrimental effect on poor households compared to wealthy households. Conversely, in our sample, we find no heterogeneous effects of crop diversity that can be attributed to differences in poverty status.

To answer our first research question, does crop diversity affect poverty status, we focus our analysis on results from the control function approach (see column (4) in Table 6) since the model controls for both sources of endogeneity and therefore does not overestimate the values of the coefficients. We find strong evidence that increased diversity decreases the probability that a household will be below the poverty line. On average, a 10 percentage point increase in the crop diversity index reduces the probability of being in poverty by 5.18 percentage points. Regarding our second research question, we find strong evidence that an increase in crop diversity increases the probability that a poor household will rise out of poverty and reduces the probability that a non-poor household will fall into poverty. Specifically, for a household already above the poverty line, we find that a 10 percentage point increase in crop diversity reduces the probability of falling into poverty by 4.79 percentage points. For a household already below the poverty line, we find that increasing diversity by 10 percentage points reduces the probability of remaining in poverty by 5.51 percentage points.

We also find that household size and lag of land per capita have a statistically significant relationship with poverty status. Household size has a positive relationship while lag of land per capita has a negative relationship. These results are unsurprising; households with more members are more likely to be in poverty than those with fewer members while households with more land are less likely to be in poverty than those with less land.

Synthesizing these results, a clear trend emerges: increasing crop diversity for rural households can help to mitigate poverty, by both raising and keeping households above the poverty line. The key result is that households which grow a more diverse set of crops are less likely to be in poverty. Thus, agricultural diversification, and not specialization, is associated with poverty reduction among households in our study.

5 Conclusion

In order to answer the empirical question regarding the effects of crop diversity on poverty status we adopted a recently developed dynamic binary response model that controls for state dependence, unobserved heterogeneity, and endogeneity of the treatment. This approach represents an improvement over previous studies which have failed to control for potential reverse causality in the relationship between crop choice and household welfare.

Results from our empirical strategy provide evidence that growing a diverse set of crops decreases the probability of being in poverty. Furthermore, crop diversity reduces the probability that a household will remain in poverty or will fall into poverty in the future. We conclude that agricultural diversity, and not specialization, is associated with poverty reduction among surveyed households. These results do not disproportionately impact wealthy households compared to poor households but are consistent across wealth levels.

Our conclusions help to elucidate a potential path out of poverty for the rural poor. Although the motivating factor to diversify may not be clear, and may range from a general desire to mitigate risk to a method of adaptation to climate change, it is clear that the specific economic and agronomic environment in Ethiopia means that diversification can reduce household poverty. Policies should be directed to encourage and increase household level crop diversity, rather than to promote specialization. In the case of Ethiopia, this means a greater focus on biodiversity and a lesser focus on encouraging households to specialize in cash crops such as coffee, sesame, or chat. Because the interaction term between the diversity index and lagged poverty status is not significant, we conclude that such policies will not have a disproportionate impact on one group over another. Therefore, actions taken to encourage crop diversity will generally be beneficial to all households; those which are presently in poverty will improve their likelihood of moving out of poverty, and those who are presently not in poverty will improve their likelihood of staying out of poverty.

Understanding the effects of household cropping decisions on poverty is an important first step in developing effective policies for household risk management. In generating strategies to address rural poverty, promoting and extending services which encourage crop diversification should be an important component. Ultimately, our research provides clear evidence in support of policies that attempt to help households mitigate food insecurity and adapt to climate change through diversification of crop production because such policies may create additional benefits by also reducing poverty.

References

Alderman, H. and C. H. Paxson (1994). In E. L. Bacha (Ed.), *Economics in a Changing World: Development, Trade and the Environment, Vol. 4*. London: St. Martin Press.

Asrat, S., G. Berhane, G. Getachew, J. Hoddinott, F. Nisrane, and A. S. Taffesse (2011). Sources of inefficiency and growth in agricultural output in subsistence agriculture: A stochastic frontier analysis. Ethiopian Strategies Support Program II Working Paper No, 19, International Food Policy Research Institute, Washington, D.C.

Baird, T. D. and C. L. Gray (2014). Livelihood diversification and shifting social networks of exchange: A social network transition? *World Development 60*(1), 14–30.

Barrett, C., T. Reardon, and P. Webb (2001). Nonfarm income diversification and household livelihood strategies in rural Africa: Concepts, dynamics, and policy implications. *Food Policy 26*(4), 315–331.

Bezabih, M. and M. Sarr (2012). Risk preferences and environmental uncertainity: Implications for crop diversification decisions in Ethiopia. *Environmental and Resource Economics 53*(4), 483–505.

Bezu, S., C. B. Barrett, and S. T. Holden (2012). Does the nonfarm economy offer pathways for upward mobility? Evidence from a panel data study in Ethiopia. *World Development 40*(8), 1634–46.

Bigsten, A., B. Kebede, A. Shimeles, and M. Taddesse (2002). Growth and poverty reduction in Ethiopia: Evidence from household panel surveys. Working Papers in Economics, Goteborg University.

Bigsten, A. and S. Tengstam (2011). Smallholder diversification and income growth in Zambia. *Journal of African Economies 20*(5), 781–822.

Bogale, A., K. Hagedorn, and B. Korf (2005). Determinants of poverty in rural Ethiopia. *Quarterly Journal of International Agriculture 44*(2), 101–120.

Bradshaw, B., H. Dolan, and B. Smith (2004). Farm-level adaptation to climate variability and change: Crop diversification in the Canadian Prairies. *Climatic Change 67*(1), 119–141.

Caviglia-Harris, J. L. and E. O. Sills (2005). Land use and income diversification: Comparing traditional and colonist populations in the Brazilian Amazon. *Agricultural Economics 32*(3), 221–37.

Christiaensen, L., L. Demery, and J. Kuhl (2011). The (evolving) role of agriculture in poverty reduction - an empirical perspective. *Journal of Development Economics 96*(2), 239–54.

Dercon, S. (1996). Risk, crop choice, and savings: Evidence from Tanzania. *Economic Development and Cultural Change 44*(3), 485–513.

Dercon, S. and L. Christiaensen (2011). Consumption risk, technology adoption and poverty traps: Evidence from Ethiopia. *Journal of Development Economics 96*(1), 159–173.

Dercon, S., D. O. Gilligan, J. Hoddinott, and T. Woldehanna (2009). The impact of agricultural extension and roads on poverty and consumption growth in fifteen Ethiopian villages. *American Journal of Agricultural Economics 91*(4), 1007–1021.

Dercon, S. and J. Hoddinott (2011). The Ethiopian Rural Household Surveys 1989-2009: Introduction. International Food Policy Research Institute, Washington, D.C.

Dercon, S. and P. Krishnan (1996). A consumption based measure of poverty in Ethiopia: 1989-1995. In M. Taddesse and B. Kebede (Eds.), *Poverty and Economic Reform in Ethiopia*. Addis Ababa: Proceedings Annual Conference of the Ethiopian Economics Association.

Dercon, S. and P. Krishnan (2003). Changes in poverty in rural Ethiopia. In A. Booth and P. Mosley (Eds.), *The New Poverty Strategies*. London: Palgrave McMillan.

Di Falco, S., M. Bezabhi, and M. Yesuf (2010). Seeds for livelihood: Crop biodiversity and food production in Ethiopia. *Ecological Economics 69*(8), 1695–1702.

Di Falco, S. and C. Perrings (2005). Crop biodiveristy, risk management, and the implications of agricultural assistance. *Ecological Economics 55*(4), 459–466.

Di Falco, S. and M. Veronesi (2013). How can African agriculture adapt to climate change? A counterfactual analysis from Ethiopia. *Land Economics 89*(4), 743–766.

Di Falco, S., M. Veronesi, and M. Yesuf (2011). Does adaptation to climate change provide food security? A micro-perspective from Ethiopia. *American Journal of Agricultural Economics 93*(3), 829–846.

Ellis, F. (1998). Household strategies and rural livelihood diversification. *Journal of Development Studies 35*(1), 1–38.

Ellis, F. (2000). *Rural Livelihoods Diversity in Developing Countries.* Oxford: Oxford University Press.

FAO (2012). Crop diversification for sustainable diets and nutrition: The role of FAO's Plant Production and Protection Division. Technical report, Plant Production and Protection Division, Food and Agriculture Organization of the United Nations, Rome.

Feder, G., R. E. Just, and D. Zilberman (1985). Adoption of agricultural innovations in developing countries: A survey. *Economic Development and Cultural Change 33*(2), 255–298.

Giles, J. and I. Murtazashvili (2013). A control function approach to estimating dynamic probit models with endogenous regressors. *Journal of Econometric Methods 2*(1), 69–87.

Hazra, C. (2001). Crop diversification in India. In M. K. Papapdemetriou and F. J. Dent (Eds.), *Crop Diversification in the Asia-Pacific Region*. Bangkok: Food and Agriculture Organization for the United Nations.

Kaguongo, W., A. Nyangweso, J. Mutunga, J. Nderitu, C. Lunga'ho, N. Nganga, D. Kipkoech, J. Kabira, M. Gathumbi, P. Njane, J. Irungu, A. Onyango, D. Boru, and E. Schutlte-Geldermann (2013). A policymakers' guide to crop diversification: the case of the potato in Kenya. Technical report, Plant Production and Protection Division, Food and Agriculture Organization of the United Nations, Rome.

Krishnan, P. and M. Patnam (2013). Neighbors and extension agents in Ethiopia: Who matters more for technology adoption? *American Journal of Agricultural Economics 96*(1), 308–327.

Kurosaki, T. (2003). Specialization and diversification in agricultural transformation: The case of the West Punjab, 1903-92. *American Journal of Agricultural Economics 85*(2), 372–386.

Little, P., K. Smith, A. Cellarius, D. Coppock, and C. Barrett (2001). Avoiding disaster: Diversification and risk management among East African herders. *Development and Change 32*(3), 401–433.

Mengxiao, Z. (2001). Crop diversification in China. In M. K. Papapdemetriou and F. J. Dent (Eds.), *Crop Diversification in the Asia-Pacific Region*. Bangkok: Food and Agriculture Organization for the United Nations.

Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica 46*(1), 69–85.

Pankhurst, A. (2009). *Moving People in Ethiopia: Development, Displacement, and the State*. Rochester, New York: Bodell and Brewer.

Papke, L. E. and J. M. Wooldridge (2008). Panel data methods for fractional response variables with an application to test pass rates. *Journal of Econometrics 145*(1-2), 121–33.

Ravallion, M. and B. Bidani (1994). How robust is a poverty profile? *World Bank Economic Review 8*(1), 75–102.

Rosenzweig, M. and H. Binswanger (1993). Wealth, weather risk, and the consumption and profitability of agricultural investments. *The Economic Journal 103*(416), 56–78.

Smith, R. J. and R. W. Blundell (1986). An exogeneity test for a simultaneous equation Tobit model with an application to labor supply. *Econometrica 54*(3), 679–85.

Taffesse, A. S., P. Dorosh, and S. Asrat (2011). Crop production in Ethiopia: Regional patterns and trends. Ethiopian Strategies Support Program II Working Paper No, 16, International Food Policy Research Institute, Washington, D.C.

Thiede, B. C. (2014). Rainfall shocks and within-community wealth inequality: Evidence from rural Ethiopia. *World Development 64*(1), 181–193.

Weinberger, K. and T. Lumpkin (2007). Diversification into horticulture and poverty reduction: A research agenda. *World Development 35*(8), 1464–1480.

Wooldridge, J. M. (2005). Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of Applied Econometrics 20*(1), 39–54.

Table 1: Village Level Descriptive Statistics

	Poverty Share		Crop Diversity			
	Mean	St. Dev	Mean	St. Dev	Max	Min
Haresaw	0.40	0.49	0.16	0.20	14	0
Geblen	0.61	0.49	0.11	0.20	13	0
Dinki	0.53	0.50	0.17	0.21	11	0
Yetmen	0.34	0.47	0.22	0.21	16	0
Shumsha	0.15	0.38	0.14	0.18	16	0
Sirbana Godeti	0.13	0.34	0.17	0.19	22	0
Adele Keke	0.24	0.43	0.16	0.17	18	0
Korodegaga	0.55	0.50	0.27	0.21	11	0
Trirufe Ketchema	0.35	0.48	0.20	0.19	22	0
Imdibir	0.73	0.45	0.21	0.19	21	1
Aze Deboa	0.53	0.50	0.24	0.19	20	0
Adado	0.53	0.49	0.16	0.14	18	2
Gara Godo	0.57	0.50	0.28	0.19	16	1
Doma	0.40	0.49	0.14	0.16	18	1
Total	0.42	0.49	0.19	0.19	22	0

Note: Poverty share is percentage of households in village that are poor averaged across years. Mean crop diversity is the mean of the diversity index in each village averaged across years. Max and Min are observed maximum and observed minimum number of crops grown by a household in each village.

Table 2: Household and Village Characteristics

	Mean	St. Dev.
Poverty Status	0.43	0.49
Poverty Status in 1993	0.43	0.49
Crop Diversity Index	0.19	0.19
Household Size	4.89	2.45
Land per Capita	0.31	0.33
Years of Education	1.90	2.80
Female Headed Household	0.19	0.40
Distance to Ag Coop	6.91	6.25
Observations	5000	
Number of Households	1250	

Table 3: Distance to Ag Coop and Crop Diversity (First Stage)

	Dependent Variable: Diversity Index			
Model	(1)	(2)	(3)	(4)
Household Size	−0.001	−0.001	−0.001	−0.001
	(0.003)	(0.003)	(0.003)	(0.003)
Lag of Land per Capita	−0.011	−0.042*	−0.049*	−0.031
	(0.020)	(0.024)	(0.027)	(0.026)
Years of Education	−0.004**	−0.004***	−0.004***	−0.004***
	(0.002)	(0.002)	(0.002)	(0.002)
Female Headed Household	0.006	0.005	0.006	0.007
	(0.010)	(0.010)	(0.010)	(0.010)
Distance to Ag Coop	0.004***	0.002**	0.010**	0.008
	(0.001)	(0.001)	(0.005)	(0.009)
Distance to Ag Coop*Lag Land per Capita		0.007**	0.015	−0.018
		(0.003)	(0.013)	(0.023)
Distance to Ag Coop^2			−0.005*	−0.001
			(0.0003)	(0.001)
Distance to Ag Coop^2*Lag Land per Capita			−0.001	0.006
			(0.001)	(0.004)
Distance to Ag Coop^3				−0.00001
				(0.0001)
Distance to Ag Coop^3*Lag Land per Capita				−0.0003*
				(0.0002)
Observations	5000	5000	5000	5000
R^2	0.216	0.218	0.219	0.224
F-Stat on IV, Interaction, & Average	20.47	8.32	4.82	4.21

Note: Fully robust standard errors clustered at the household are in parentheses (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$). Regressions include time averages of explanatory variables, year dummies, and interactions between village dummies and a time trend.

Table 4: Distance to Ag Coop and Poverty Status

Model	Dependent Variable: Poverty Status
Lag Poverty Status	0.147***
	(0.016)
Household Size	0.051***
	(0.009)
Lag of Land per Capita	0.039
	(0.045)
Years of Education	0.002
	(0.005)
Female Headed Household	−0.022
	(0.036)
Distance to Ag Coop	0.012
	(0.010)
Distance to Ag Coop*Lag Land per Capita	−0.033
	(0.038)
Distance to Ag Coop^2	−0.001
	(0.001)
Distance to Ag Coop^2*Lag Land per Capita	0.002
	(0.003)
Observations	5000
R^2	0.162

Note: Fully robust standard errors clustered at the household are in parentheses (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$). Regressions include explanatory variables in each year, year dummies, and interactions between village dummies and a time trend.

Table 5: Poverty Status and Crop Diversity (Second Stage)

	Dependent Variable: Poverty Status			
	Exogenous Diversity Index		Endogenous Diversity Index	
	LPM	Probit	LPM	CF
Model	(1)	(2)	(3)	(4)
Lag Poverty Status	0.134***	0.355***	0.135***	0.226***
	(0.021)	(0.057)	(0.021)	(0.078)
Diversity Index*Lag Poverty Status	0.076	0.236	0.072	0.211
	(0.078)	(0.223)	(0.075)	(0.254)
Diversity Index	−0.072	−0.207	−0.562***	−1.611**
	(0.046)	(0.152)	(0.188)	(0.781)
Household Size	0.029***	0.082***	0.041***	0.127***
	(0.003)	(0.010)	(0.006)	(0.023)
Lag of Land per Capita	−0.132***	−0.514***	−0.053*	−0.311*
	(0.021)	(0.089)	(0.032)	(0.177)
Years of Education	−0.002	−0.003	0.001	0.005
	(0.003)	(0.009)	(0.003)	(0.014)
Female Headed Household	0.035*	0.105**	0.011	0.036
	(0.018)	(0.052)	(0.021)	(0.087)
Observations	5000	5000	5000	5000
Number of Households	1250	1250	1250	1250
R^2	0.16		0.16	
Replications for Bootstrapped Errors	1000	1000	1000	1000

Note: Bootstrapped standard errors in parentheses (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$). Regressions include explanatory variables in each year, year dummies, and interactions between village dummies and a time trend. We also include household random effects in each specification. Regressions (1) and (2) treat crop diversity as exogenous. Regressions (3) and (4) include first stage residuals free of serial correlation and their time averages.

Table 6: Average Partial Effects of Determinants of Poverty Status

	Exogenous Diversity Index		Endogenous Diversity Index	
	LPM	Probit	LPM	CF
Model	(1)	(2)	(3)	(4)
Lag Poverty Status	0.148***	0.399***	0.148***	0.088***
	(0.016)	(0.042)	(0.016)	(0.015)
Diversity Index when Lag Poverty=0	−0.072	−0.207	−0.562***	−0.551**
	(0.046)	(0.152)	(0.188)	(0.276)
Diversity Index when Lag Poverty=1	0.005	0.029	−0.490***	−0.479*
	(0.064)	(0.170)	(0.196)	(0.280)
Diversity Index Averaged	−0.065	−0.188	−0.556***	−0.518*
	(0.043)	(0.141)	(0.188)	(0.274)
Household Size	0.029***	0.082***	0.041***	0.043***
	(0.003)	(0.010)	(0.005)	(0.008)
Lag of Land per Capita	−0.132***	−0.514***	−0.053*	−0.106*
	(0.021)	(0.087)	(0.032)	(0.062)
Years of Education	0.002	−0.003	0.001	0.002
	(0.003)	(0.009)	(0.003)	(0.005)
Female Headed Household	0.035*	0.105**	0.011	0.012
	(0.018)	(0.052)	(0.021)	(0.031)
Observations	5000	5000	5000	5000
Number of Households	1250	1250	1250	1250
Replications for Bootstrapped Errors	1000	1000	1000	1000

Note: Bootstrapped standard errors in parentheses (* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$). We calculate the APEs of the diversity index in three ways: 1) when lagged poverty status equals zero for all households, 2) when lagged poverty status equals one for all households, and 3) averaged across the observed values of poverty status in the data. The APE for lagged poverty status is averaged across the observed values of the diversity index in the data. All other APEs are averaged across both the cross-section of the covariates and time.

Figure 1: Bivariate Density of Mean Village Poverty

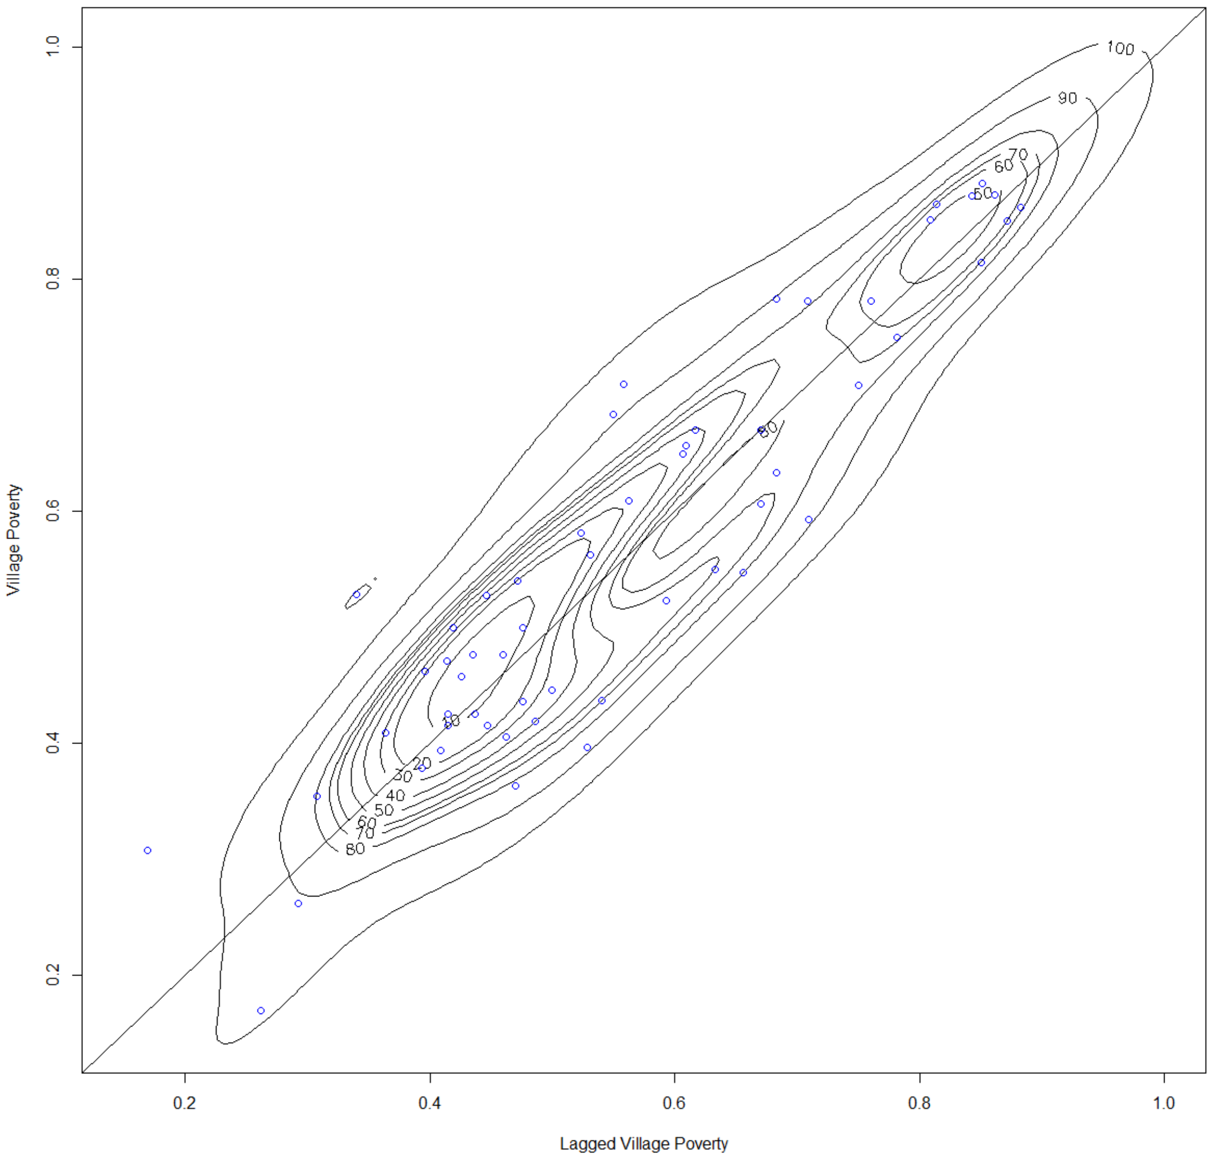

Note: Figure shows the bivariate kernel density contours of the mean poverty level in each village in each year. Poverty level is one hundred minus the percentage of households in the village that are poor. Thus, observations close to zero come from villages with high poverty levels while observations close to one come from villages with low poverty levels. Circles indicate observed data. Villages above the 45° line have fewer poor households compared to the previous year. Villages below the 45° line have more poor households compared to the previous year.

Figure 2: Change in Poverty in Village Versus Change in Crop Diversity Index

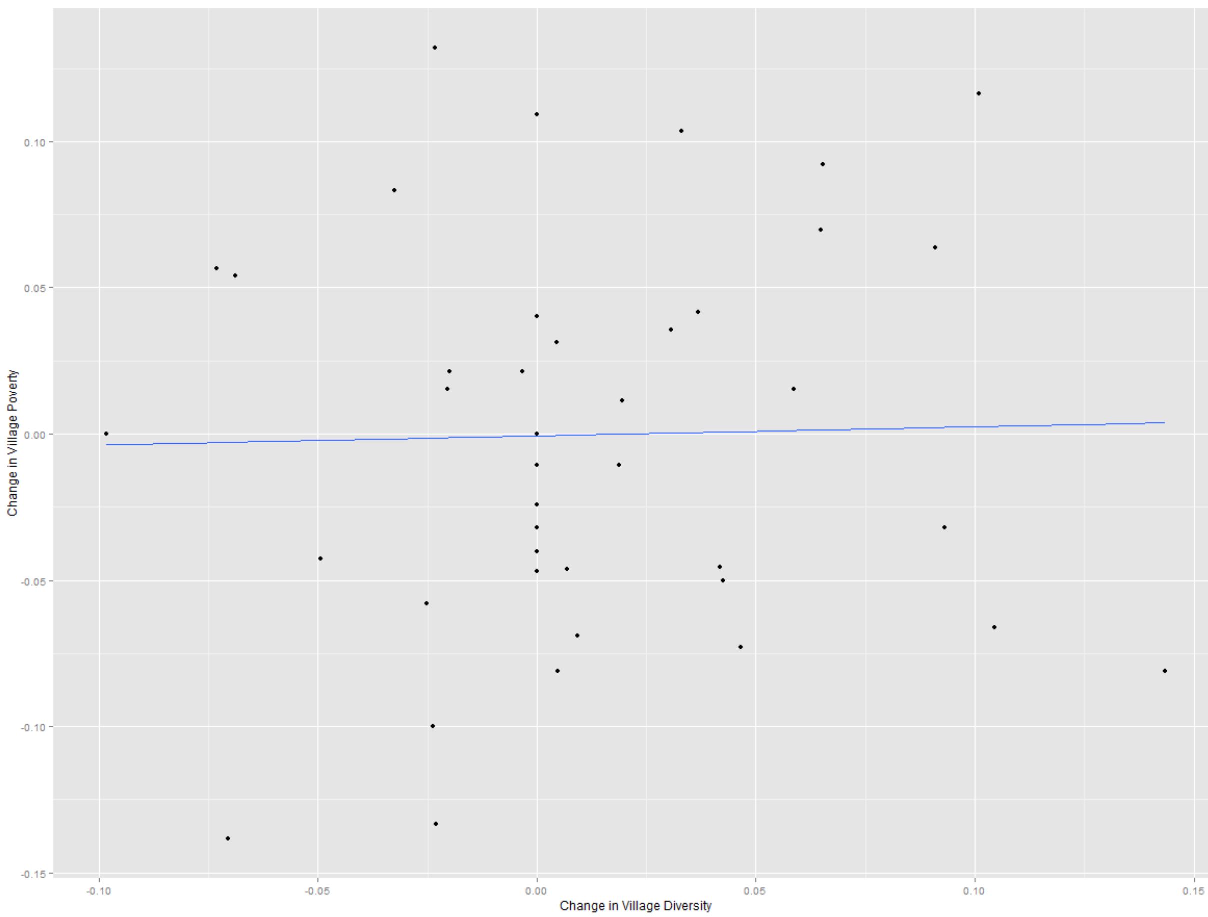

Note: Figure shows scatter plot of changes to village poverty and changes to village crop diversity from year $t-1$ to t. The figure also includes a linear trend line with slope of 0.0313, which is not statistically different from zero.

Figure 3: Bivariate Density of Mean Village Diversity Index

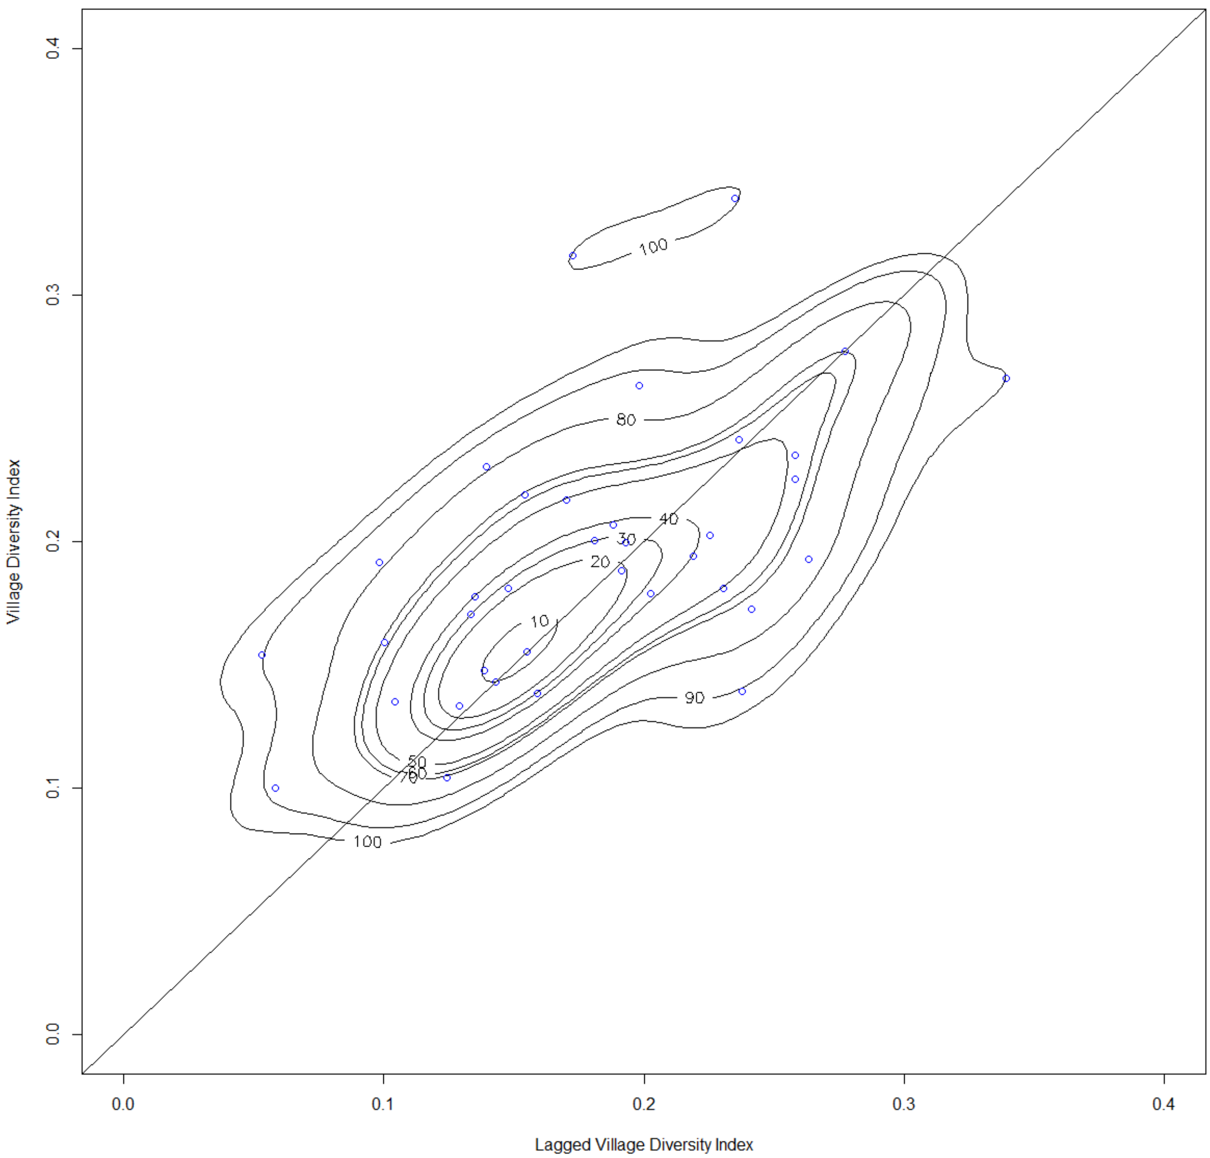

Note: Figure shows the bivariate kernel density contours of the mean diversity index in each village in each year. Observations close to zero come from villages with low levels of crop diversity while observations close to one come from villages with high levels of crop diversity. Circles indicate observed data. Villages above the 45° line have more crop diversity compared to the previous year. Villages below the 45° line have less crop diversity compared to the previous year.

Figure 4: Change in Crop Diversity Index and Poverty in Village Versus Distance to Ag. Cooperative

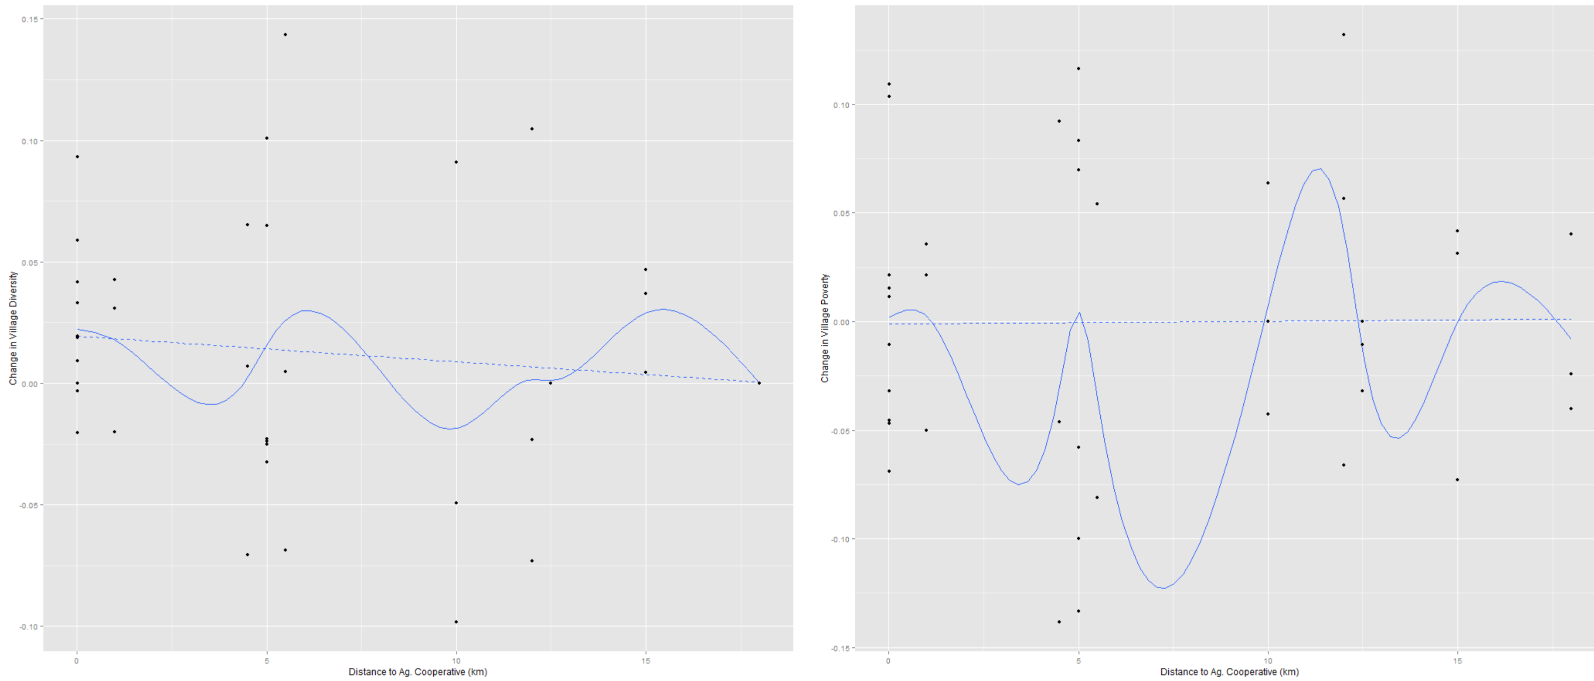

Note: Figure at left shows a lowess plot (solid line) of the relationship between distance to agricultural cooperative in kilometers and the degree of crop diversity in a village from year $t-1$ to t. The figure also includes a linear trend line (dashed line) with slope of -0.0011, which is statistically different from zero at the 99 percent level. Figure at right shows a lowess plot of the relationship between distance to agricultural cooperative in kilometers and changes to village poverty from year $t-1$ to t. The figure also includes a linear trend line with slope of 0.0001, which is not statistically different from zero.