Does Unobserved Land Quality Bias Separability Tests?

Jarrad Farris
U.S. Department of Agriculture
Economic Research Service
Kansas City, MO

DRAFT Version: July 1st, 2020

Selected Paper Prepared for Presentation at the 2020 Agricultural & Applied Economics Association Annual Meeting, August 10 - August 11.

I. Introduction

Modeling agricultural household decision making is integral to the design and evaluation of development programs and policies. A key breakpoint in such models is whether farm households make their production decisions separately from their consumption decisions. The existence of separability affects household production responses to new opportunities and shocks and provides an indication of the completeness of markets (Benjamin, 1992; Singh, Squire, & Strauss, 1986).

Numerous studies implicitly or explicitly assume separability of agricultural production decisions from consumption decisions (e.g. Conley & Udry, 2010; Foster & Rosenzweig, 1995; Sheahan, Black, & Jayne, 2013; Suri, 2011). When production decisions are non-separable from consumption decisions, ignoring this non-separability may vastly misrepresent household production decisions and the implications of policies (LaFave & Thomas, 2016; Singh et al., 1986). For instance, a separable model cannot predict the preferential adoption of a new agricultural technology by larger households with greater availability of family labor. Similarly, such models cannot account for autarkic decision making based on the consumption preferences of household members. The sensitivity of models of agricultural household behavior to this hypothesis begets the importance of accurate separability tests.

This paper addresses a major identification challenge for tests of this separability hypothesis– the potential endogeneity of household demographic characteristics with unobserved land quality.[1] Using a unique plot-panel dataset, I test the separability of rural Rwandan agricultural household production decisions while controlling for the endogeneity of household demographic characteristics with land quality. I then use simulations to assess the susceptibility of standard tests based on household fixed effects to ignoring unobserved heterogeneity in land quality.

[1]This study defines land quality in the broadest sense to include all land characteristics which affect agricultural productivity (e.g. soil type, nutrients, organic matter content, slope, etc.).

The difficulty of controlling for land quality and its likely correlation with household demographic characteristics has long been a key identification challenge in the separability literature (Benjamin, 1992; Udry, 1999). Early, seminal work based on cross-sectional data relies on observable proxies for land quality and tends to fail to reject separability (Benjamin, 1992; Pitt & Rosenzweig, 1986). More recent work relies on household or farm fixed effects and tends to come to the opposite conclusion (e.g. Dillon, Brummund, & Mwabu, 2019; Kopper, 2018; LaFave & Thomas, 2016).

This paper makes two main contributions to this literature. First, using a recent dataset from Rwanda, I control for potentially confounding unobserved land characteristics by leveraging intra-plot variability in agricultural input demand. Common tests of separability using household panel data control for factors fixed at the household or farm level, such as the quality of household decision making. These tests, however, are threatened by the likely correlation of household characteristics with land quality and other unobserved land characteristics when farmland is not static across survey waves. I find that the non-separability result in Rwanda is robust to controlling for land quality and other unobserved time invariant plot characteristics. This emphasizes the need to integrate consumption characteristics into models of production decision making and support programs and policies designed to alleviate market failures in agricultural settings.

Second, I use simulations to examine a future with well-functioning markets where the separability hypothesis holds, but consumption traits are correlated with unobserved plot characteristics. Using these simulated datasets, I show that separability tests based on household fixed effects are prone to bias, and that ignoring unobserved land quality can lead to false rejections of separability. Furthermore, this bias is exacerbated as the land market becomes more active.

This relationship to land market activity is particularly important given the close link between the separability of agricultural household behavior and the existence of well-functioning markets (Benjamin, 1992; Singh et al., 1986). Separability tests are more useful in contexts with active land markets, as this attribute increases the likelihood that a separable agricultural household model (AHM) may accurately describe household production responses; the simulation results, however, suggest that standard tests based on household panel data are likely to perform worse in these contexts.

These findings highlight the need for additional research that incorporates more robust means of controlling for unobserved land quality, such as plot panel data which enable the use of plot fixed effects. In areas with functioning land markets where some households change operated land area between survey waves, inadequate control of land quality in reduced form separability tests based on household fixed effects could bias inferences on agricultural household decision making.

The remainder of the paper is structured as follows. In the first and second sections, I describe the theoretical framework underlying the AHM and the empirical strategy underpinning reduced-form separability tests, with a focus on plot-level characteristics. In the third and fourth sections, I describe the rural Rwandan plot-panel dataset used in the empirical application and present the Rwanda results. In the fifth section, I simulate data to illustrate the potential bias from unobserved plot-characteristics and how it is exacerbated by shifts in cultivated land between periods. I conclude in the final section with a summary of the key findings and implications.

II. Theoretical Model

In this section, I illustrate the intuition behind reduced form separability tests and highlight the role of unobserved land quality. I do so by incorporating unobserved land quality à la Udry (1999) into the LaFave and Thomas (2016) and Dillon et al. (2019) dynamic extensions of the static AHM in Singh et al. (1986).

A household's objective is to maximize expected discounted utility as follows:

$$\max\ E\left[\sum_{t=1}^{\infty}\beta^{t-1}U(\mathbf{x_{mt}},\mathbf{x_{at}},x_{lt};\ \mathbf{D_t},\mu_t)\right] \tag{1}$$

where household utility in time period t is captured by a time-separable, concave, strictly increasing utility function, $U(\cdot)$, over a vector of market goods, $\mathbf{x_{mt}}$, a vector of agricultural goods, $\mathbf{x_{at}}$, and leisure, x_{lt}. The utility derived from these goods differs according to household consumption preferences observed by the analyst (e.g. demographic characteristics), $\mathbf{D_t}$, and a composite of characteristics unobserved by the analyst, μ_t. Utility derived in future time periods is discounted at the rate β^{t-1}.

The household's budget constraint in period t is:

$$\mathbf{p_{mt}}\cdot\mathbf{x_{mt}}+\mathbf{p_{at}}\cdot\mathbf{x_{at}}+w_t x_{lt}+\frac{1}{1+\tau_t}W_{t+1}=w_t E_t+\pi_t+W_t \tag{2}$$

where the prices of market goods, agricultural goods, and leisure are $\mathbf{p_{mt}}$, $\mathbf{p_{at}}$, and w_t respectively, W_{t+1} is wealth in the next period, which is negative if the household is in debt and positive otherwise, τ_t is the interest rate for borrowing or lending, E_t is the household's total time endowment, and π_t is total farm profit.

Total farm profit, π_t, is determined by the household's agricultural input choices and is the sum of profit across all the household's plots as follows:

$$\pi_t=\sum_{i=1}^{N_t}p_{qt}f(L_t^i,\widetilde{A}^i,\mathbf{Z_t^i};\ v_t)-w_tL_t^i-r_t\widetilde{A}_t^i-\mathbf{p_{zt}}\cdot\mathbf{Z_t^i} \tag{3}$$

where N_t is the number of plots farmed by the household in the given period. The farm-production technology, $f(\cdot)$, determines agricultural output on plot i and is a function of labor input, L_t^i, quality-adjusted plot size, $\widetilde{A}^i$, a vector of other inputs, $\mathbf{Z_t^i}$, and an exogenous, community-specific shock, v_t. The agricultural output price, wage rate, land rental rate, and other input prices are given by p_{qt}, w_t, r_t, and $\mathbf{p_{zt}}$ respectively.[2]

Quality-adjusted plot size, $\widetilde{A}^i$, reflects that plots have varying qualities which influence their productivities. Two plots of the same size may produce different outputs depending on the quality of each plot, *ceteris paribus*. In determining the productivity of land input, the size of each plot is adjusted to account for quality differences as follows:

$$\widetilde{A}^i = \theta^i A^i \quad \forall\, i \in \{1, 2, \ \ldots \ , N_t\} \tag{4}$$

where θ^i and A^i are the quality and size, respectively, of plot i.

The key characteristic of this AHM problem is that farm input decisions are independent of household consumption preferences. For instance, the first order condition for labor demand on a given plot is:

$$p_{qt} \frac{\partial f(L_t^i, \widetilde{A}^i, \mathbf{Z_t^i};\, v_t)}{\partial L_t^i} = w_t \tag{5}$$

This optimality condition is independent of household demographic characteristics (and other characteristics which only affect consumption decisions). Thus, in the separable AHM, production decisions are based solely on profit maximization. Optimal household labor demand, as well as other input demands, are invariant to changes in household preferences or demographic characteristics. Reduced form tests of the separability hypothesis rely on this result to assess whether production decisions are consistent with the separable AHM.

The optimality condition in equation 5 does, however, depend on plot characteristics. For example, the analyst may observe plot size, A_t^i, but not plot quality, θ_t^i. This could be problematic for the reduced form separability test if household demographic characteristics, $\mathbf{D_t}$, are correlated with these unobserved plot quality characteristics (e.g. if larger households tend to have better quality land). In the next section, I assess the implications of this problem for empirical reduced form separability tests.

[2] For ease of exposition, this model focuses on a single, land-based agricultural output. The separability result extends to multiple outputs (Singh et al., 1986).

III. Empirical Strategy

Popularized in LaFave and Thomas (2016), the current standard reduced form test of the separability hypothesis applies the household fixed effects estimator to total farm labor demand as follows:

$$lnL_{cht} = \kappa + \delta\mathbf{D_{cht}} + \beta\mathbf{X_{cht}} + \eta_{ct} + \eta_h + \epsilon_{cht} \tag{6}$$

where L_{cht} is the total person-days of labor used during an agricultural season in year t by household h in community c, κ is the overall intercept, $\mathbf{D_{cht}}$ is a vector of household demographic characteristics, $\mathbf{X_{cht}}$ is a vector of other observed characteristics which affect labor demand and are potentially correlated with $\mathbf{D_{cht}}$, and ϵ_{cht} is the idiosyncratic error.[3] The null hypothesis of interest is $\delta = 0$ as this implies that household demographic characteristics do not influence labor input demand and that the separability hypothesis cannot be rejected (Benjamin, 1992; LaFave & Thomas, 2016).

The community-time fixed effects, η_{ct}, exploit variation within a community in a given year to control for any time varying community-level characteristics which may be correlated with household demographic characteristics; for example, community-wide shocks and prices (LaFave & Thomas, 2016).

The household fixed effects, η_h, exploit within-household variability in labor demand to allow household demographics to be correlated arbitrarily with time invariant household characteristics (both observed and unobserved). This is an important improvement over early studies of separability, which relied on observed variables to control for farm characteristics and other correlates to household demographic characteristics (Benjamin, 1992; LaFave & Thomas, 2016).

The specification in equation 6, however, does not control for plot-level unobservables, which may be correlated with household demographics within a particular community. For

[3] $\mathbf{X_{cht}}$ controls for time varying farm size and characteristics that reflect differences in farmer experience, such as household head characteristics.

example, soil quality, an important factor in input demand, is typically unobserved and likely correlated with household size and other demographic characteristics (Kopper, 2018; Udry, 1999). Failure to control adequately for such plot-level characteristics could result in spurious correlation of $\mathbf{D_{cht}}$ and ϵ_{cht}, biasing the test of separability.

Given these plot-specific characteristics and the aggregated household-level input demand in equation 6, the idiosyncratic error can be approximated as:

$$\epsilon_{cht} = \sum_{i=1}^{N_t} \eta_i + v_{cht} \tag{7}$$

where η_i are time invariant plot-level unobservables, N_t is the number of plots at time t, and v_{cht} is the remaining composite idiosyncratic error. If a household does not change its farmed plots between survey waves, then $\sum_{i=1}^{N_t} \eta_i$ will be subsumed by η_h. Thus, if none of the sampled households change the composition of their farmed land between survey waves, then the separability test given in equation 6 will control for unobserved plot characteristics.

If some portion of the households in the sample alters its farmed landholdings between survey waves, whether through newly rented in, rented out, bought, or sold plots, then the aggregate sum of time invariant plot characteristics is no longer constant between survey waves and will not be differenced out by household fixed effects. In this case, a correlation between household demographic characteristics and time invariant plot characteristics may cause a rejection of the null hypothesis of separability even if separability holds.

This threat to identification is addressed by plot fixed effects:[4]

$$ln\widetilde{L_{chit}} = \kappa + \delta\mathbf{D_{cht}} + \beta\mathbf{X_{cht}} + \eta_{ct} + \eta_i + \epsilon_{chit} \tag{8}$$

[4] This identification strategy is contingent on collecting repeated labor use data for the same plot operated by the same household. If the land market is extremely volatile in a particular context and nearly all households change all their farmed plots over the period studied, then the plot fixed effects approach is infeasible.

where $\widetilde{L_{chit}}$ is the total person-days of labor used during an agricultural season of year t on plot i of household h in community c, η_i are plot fixed effects, and ϵ_{chit} is the idiosyncratic error for plot i.[5]

By using plot panel data with plot fixed effects, any unobserved time invariant plot characteristics correlated with the household demographic characteristics are differenced away. Similarly, as time invariant household and community characteristics are fixed over time for a given plot, these characteristics are also subsumed by the plot fixed effects (Udry, 1999).

Importantly, whether a plot is observed in the plot panel or not can be arbitrarily correlated with $\mathbf{D_{cht}}$, $\mathbf{X_{cht}}$, η_{ct}, and η_i without affecting the consistency of the test in equation 8; for instance, a household's decision to farm (or not farm) a given plot in a particular period can be correlated with their demographic characteristics in $\mathbf{D_{cht}}$, a community-level shock, or time invariant plot or household characteristics without affecting the consistent estimation of δ (Wooldridge, 2010, pg. 829).[6]

IV. Data

I assess separability in Rwanda using a two-round, panel survey conducted in 2014 and 2017. The initial survey wave used a three-stage cluster sampling method within Rwanda's Northern, Southern, and Eastern provinces where the sector, village, and household were the primary, secondary, and tertiary sampling units, respectively. The survey is not representative of Rwanda as a whole; rather, the sampling frame focused on promoting food security and nutrition in rural areas and targeted households with a child less than five years of age

[5]Separable labor demand on a given plot is a function of plot size which is controlled for via η_i; therefore, $\mathbf{X_{cht}}$ omits total farm size in this specification.

[6]Let $\mathbf{s_i} = [s_{i1}, s_{i2}, \ldots, s_{iT}]$ where s_{it} is an indicator variable equal to one if a given plot i was observed in period t and T is the number of survey waves. Similarly, denote the elements in equation 8 as $\mathbf{Z_{it}} = [\mathbf{D_{cht}}, \mathbf{X_{cht}}, \eta_{ct}]$ and $\mathbf{Z_i} = [\mathbf{Z_{i1}}, \mathbf{Z_{i2}}, \cdots, \mathbf{Z_{iT}}]$. Estimation of δ is consistent if $E(\epsilon_{chit} \mid \mathbf{s_i}, \mathbf{Z_i}, \eta_i) = 0 \; \forall \; t$ and the outer product of the covariate matrix is nonsingular. This requires $\mathbf{s_i}$ to be uncorrelated with ϵ_{chit} after controlling for $[\mathbf{Z_i}, \eta_i]$, but does not restrict the relationship between $\mathbf{s_i}$ and $[\mathbf{Z_i}, \eta_i]$. Wooldridge (2010, pg. 829) provides a formal proof.

or a pregnant woman. The survey collected detailed plot-level agricultural information as well as household demographic information.

This analysis focuses on total plot-specific labor demand during Rwanda's major agricultural season of February to June. Total labor demand is measured as the sum of the labor days of family and hired labor for land preparation, planting, and field management after planting. Harvest labor is excluded from consideration, as harvest labor requirements are typically proportional to output rather than being a production choice variable. Child labor days, defined as labor provided by household members under 15 years old, are scaled by 0.5 to reflect productivity differences relative to adult labor (Dillon & Barrett, 2017).

During the 2017 survey round, plots were linked to the 2014 survey round by the main household member responsible for agricultural decisions. After describing a given 2017 plot, the respondent was read a list of the household's unique plot descriptions and plot sizes reported in the 2014 survey round. The respondent then either linked the given 2017 plot to a unique plot description provided in the 2014 round, reported the 2017 plot to be a new plot, or reported that they did not know and could not identify a match. Using this method, about half (51%) of plots observed in 2017 were successfully linked to 2014 plot observations.[7]

Table 1 provides household level characteristics for the 1,494 households with a least one plot in the plot panel subsample relative to the full analytical sample of 1,800 households. Although observed characteristics are similar between groups and a vast majority of households have at least one plot in the plot panel sample, I restrict the household level separability analysis to the subsample of 1,494 households with at least one plot in the plot panel sample. This reduces concerns that households without a plot in the plot panel may

[7]Respondent matched plots above the 95th percentile for the absolute value of plot size difference between survey waves are trimmed from the plot panel subsample. The remaining 95% of matched plots have a mean absolute plot size difference between survey waves of 0.098 hectares with a 0.998 correlation in plot size between survey waves. The trimmed plots have an average absolute plot size difference between survey waves of 16.407 hectares and a -0.057 correlation in plot size between survey waves, suggesting that these plots were erroneously matched.

follow a markedly different decision framework.

Not all plots farmed by these 1,494 households were observed in both survey waves. Table 2 provides summary statistics for the 4,580 plot-wave observations in the plot panel subsample relative to the full sample of 7,303 plot-wave observations. As the plot panel plots are not a random subset of each household's farmed plots, they are unlikely to be representative of households' total landholdings. For example, plots in the plot panel sample are smaller on average. As discussed previously, consistent estimation of the separability test based on the plot fixed effects estimator is not reliant on balanced plot characteristics (Wooldridge, 2010, pg. 829).

V. Results

Separability tests based on the household fixed effects specification in equation 6 reject the null hypothesis of separability in both the parsimonious regression of the natural log of household size and the expanded regression with shares of household members by age group (Table 3 columns 1 and 2). The validity of these findings is reliant on the exogeneity of household demographic characteristics given controls for the natural log of farmed land area, community-wide time-varying shocks, and unobserved time-invariant household or farm specific heterogeneity. These household fixed effects results are sufficiently robust to account for potential productivity differences in child household members (Appendix Table 7 columns 1 and 2), control for a land quality proxy (Appendix Table 8 columns 3 and 4), and including households without a plot in the plot panel sample (Appendix Table 9).

Separability tests based on the plot fixed effects specification in equation 8 suggest that the non-separability result in this sample of Rwandan households is also robust to controlling for unobserved land quality and other unobserved time-invariant plot characteristics (Table 3 columns 3 and 4). While the parsimonious plot fixed effects specification fails to reject the null of separability, separability is rejected once the specification is expanded to include shares of household members by age group. These plot fixed effects results are robust to account for potential productivity differences in child household members (Appendix Table 7 columns 3 and 4). The findings suggest that this sample of small-scale, agricultural households integrates demographic characteristics into their production decisions.

Restricting the household fixed effects specification to only plots in the plot panel sample provides another useful check on these results. Unlike the summation of labor demand over all plots farmed by a household in a given year, summing household labor demand over only plots in the plot panel subsample forces land quality and other unobserved time-invariant land characteristics to remain fixed at the household level; this enables the household fixed effects specification to control for unobserved land quality in a similar manner to the plot fixed effect specification. The results, presented in Table 4, are consistent with those based on plot fixed effects; the parsimonious regression of log household size fails to reject separability, but separability is rejected in the expanded regression with shares of household members by age.

Restricting the analysis to the subsample of households where no person left or joined the household apart from children born between survey waves provides a further check on the main results (Appendix Table 10). This subsample analysis reduces the likelihood that endogeneity of the demographic variables with the remaining idiosyncratic error drives the non-separability result as changes in the household demographic variables exclude migrants in or out of the household.[8] This restriction has the disadvantage, however, of an inability to assess separability for households with a non-birth related change to their household rosters and a large reduction in sample size. In the parsimonious plot fixed effects specification on this subsample, separability is rejected (at the 1% level). Furthermore, the estimated coefficient on the log of household size is negative, indicating that the birth of a child between

[8]LaFave and Thomas (2016) further restrict the sample to households with static rosters where changes in the demographic variables are driven solely by aging. I do not have adequate power to further restrict the sample in this way as the survey sampling frame targeted Rwandan households with a child under five years of age or a pregnant woman.

survey waves reduces labor demand.[9] This is consistent with a reduction in a new mother's own-farm labor supply following childbirth which is not offset by hired labor.

While the more robust plot fixed effects tests corroborate those based on household fixed effects for this sample of Rwandan households, the latter should not be relied on as a primary indicator of whether household decision making is consistent with separability when land quality is unobserved and farmed land is not fixed. Next, I demonstrate this in a controlled environment via simulation, assessing the performance of tests based on the household fixed effects and plot fixed effects estimators as land markets become more active.

VI. Simulation

VI.1 Simulation Methods

In this section, I simulate the models outlined in Section III to analyze the performance of tests in a future with well-functioning markets where separability holds, but consumption traits are correlated with unobserved plot characteristics. The findings from this simulation show that separability tests based on the household fixed effects estimator are prone to bias, and that ignoring unobserved land quality can lead to false rejections of separability. The simulation results further show how tests based on the plot fixed effects estimator address these issues.

What follows is a description of the simulation procedure for a single replication. I repeat this process 1,000 times and compare the performance of the household fixed effects and plot fixed effects specifications under different scenarios.

First, I generate three-level panel data (community, household, and plot) over two time

[9]The other robustness specifications also have negative household size coefficients, but the estimated coefficients are not significant.

periods by the following process:

$$Y_{chit} = exp(\kappa + \eta_c + \eta_h + \eta_i + \epsilon_{chit}) \tag{9}$$

where Y_{chit} is analogous to labor demand in time t on plot i of household h in community c. This process was chosen to mirror, in a simplistic way, the linear-in-logs specification of this study's empirical strategy.[10]

Using this data generating process, data are simulated for 250 communities and 25,000 households, 100 households per community. Community and household-level unobservables are drawn independent and identically distributed (i.i.d.) as $\eta_c \sim N(0, 4)$ and $\eta_h \sim N(0, 4)$ respectively. Each household in the same community is assigned a common η_c and each plot-observation within the same household is assigned a common η_h. Plot-level unobservables, η_i, are drawn as:

$$\eta_i = aX_{c,h,t=1} + bZ_{chi} \tag{10}$$

where $X_{c,h,t=1} = \widetilde{X}_{c,h,t=1} + 1,\ \widetilde{X}_{c,h,t=1} \sim Poisson(4),$ and $Z_{chi} \sim N(0, 4)$.

$X_{c,h,t=1}$ is analogous to a household demographic variable in the first period. This structure is chosen to simulate correlation between observable household demographics and unobserved time-invariant plot characteristics. Given this structure and the independence of $X_{c,h,t=1}$ and Z_{chi}, a can be derived as follows:

$$a = \frac{Cov(X_{c,h,t=1}, \eta_i)}{Var(X_{c,h,t=1})} = \frac{Cov(X_{c,h,t=1}, \eta_i)}{4} \tag{11}$$

[10] This simulation focuses on the correlation between plot unobservables and a household demographic variable. Thus, other characteristics are not included in the data generating process and community level unobservables are simulated as time constant and thus absorbed by the household or plot fixed effects.

Similarly, b can be derived as:

$$b = \frac{Cov(Z_{chi}, \eta_i)}{Var(Z_{chi})} = \frac{Cov(Z_{chi}, \eta_i)}{4} \tag{12}$$

The population correlation between $X_{c,h,t=1}$ and η_i is then given by:

$$Corr(X_{c,h,t=1}, \eta_i) = \frac{Cov(X_{c,h,t=1}, \eta_i)}{\sigma_X \sigma_\eta} = \frac{Cov(X_{c,h,t=1}, \eta_i)}{4\sqrt{a^2 + b^2}} \tag{13}$$

Defining $Cov(X_{c,h,t=1}, \eta_i)$ and $Cov(Z_{chi}, \eta_i)$ as one and four, respectively, sets a population correlation between the first period household demographic variable and time invariant plot unobservables of approximately one quarter (0.2425). This population correlation is larger than that of plot unobservables and X_{cht} once temporal variation in X_{cht} is introduced.

The simulated dependent variable is finalized by defining the overall intercept, κ, at 20 and drawing i.i.d. idiosyncratic errors from $\epsilon_{chit} \sim N(0, 1)$. Y_{chit} is then computed by combining all its component parts, as shown in equation 9.

In order to simulate exogenous temporal variation in X_{cht}, which is analogous to changes in a household demographic variable between survey waves, each household is assigned an i.i.d. draw from $\delta_h = \widetilde{\delta_h} + 1$ where $\widetilde{\delta_h} \sim Poisson(2)$. For each household, a draw from $Uniform[0, 1]$ determines how δ_h is allocated to $X_{c,h,t=2}$. One third of simulated households are randomly assigned X_{cht} increases in period two as given by $X_{c,h,t=2} = X_{c,h,t=1} + \delta_h$. Similarly, one third of simulated households are randomly assigned X_{cht} decreases in period two as given by $X_{c,h,t=2} = X_{c,h,t=1} - \delta_h$.[11] The remaining one third of households are assigned $X_{c,h,t=2} = X_{c,h,t=1}$.

The number of plot observations, n_{ht}, varies by household and time period. In the first period, each household's number of farmed plots is determined by i.i.d. draws of $N_{ht} = \widetilde{N}_{ht} + 1$ where $\widetilde{N}_{ht} \sim Poisson(2)$.

[11] $X_{c,h,t=2}$ is set to one if the simulated decrease would cause $X_{c,h,t=2}$ to fall below one.

Several different scenarios, summarized in Table 5, simulate varying degrees of farmed land changes between survey waves. The first scenario mimics the unlikely case of no household changing farmed plots between survey waves by maintaining the first period plot allocations. This scenario is chosen to demonstrate a case where separability tests based on the household fixed effects estimator are unbiased.

The other three scenarios demonstrate an active (and increasingly active) land market. They do so by simulating more realistic cases where some of the households' farmed plots vary between survey waves (whether due to newly rented in, rented out, bought, or sold plots). For these scenarios, each of a household's plots is assigned an i.i.d. draw from $Uniform[0, 1]$. Similarly, for each household, five potential new plots are simulated and assigned i.i.d. draws from $Uniform[0, 1]$. These draws are used to determine which plots are farmed by each household in the second period.

Under Scenario A, a given household plot farmed in the first period is maintained in the second period with probability 0.99. In addition, each household has a small chance of farming one or more new plots. Each of the five potential new plots are incorporated into a household's farmed plots in the second period with probability 0.03. On average across the 1,000 replications, this results in 16% of households experiencing a change in farmed plots in the second period.

Under Scenario B, a given household plot farmed in the first period is maintained in the second period with probability 0.95. Similarly, the chance of farming each of five potential new plots increases to 0.05. On average across the 1,000 replications, this results in 30% of households experiencing a change in farmed plots in the second period.

Under Scenario C, the probability of a household maintaining a given first-period plot is kept unchanged at 0.95. The chance of farming each of the five new plots, however, is increased slightly to 0.07. On average across the 1,000 replications, this results in 37% of households experiencing a change in farmed plots in the second period.[12]

For each of these scenarios, I run the household fixed effects and plot fixed effects specifications and store the results of the reduced form separability tests. Having stored the results from the first replication of the simulation, I then repeat this entire process 1,000 times. In each replication, I take new draws from the respective distributions of η_c, η_h, $X_{c,h,t=1}$, Z_{chi}, ϵ_{chit}, δ_h, N_{ht}, and the Uniform[0,1] variables. I then use these new draws to compute η_i, Y_{chit}, and $X_{c,h,t=2}$. Using the given replication's simulated dataset, I then estimate the household fixed effects and plot fixed effects specifications under each of the four scenarios and record the results of the reduced form separability tests.

VI.2 Simulation Results

The simulation results demonstrate the susceptibility of the separability test based on the household fixed effects estimator to unobserved heterogeneity in land quality. The bias in the household fixed effects estimator increases as the land market becomes more active. In contrast, the plot fixed effects estimator is robust to a correlation of the household demographic variable with unobserved, time invariant land quality, regardless of the level of land market activity.

Table 6 reports the number of Type I errors under different levels of land market activity for each estimator over the 1,000 replications. As the data are simulated, the data generating process is known, and the separability hypothesis holds. That is, as X_{cht} is not a causal determinant of Y_{chit}, systematic rejections of the null of hypothesis of separability are indicative of bias in the estimator. Across 1,000 replications, an unbiased estimator would incorrectly reject the null of separability at the 5% and 10% levels approximately 50 and 100 times, respectively.

When none of the simulated households change their farmed land between survey waves,

[12] These scenarios are conservative relative to the Rwanda sample where slightly more than half of all households reported a change in farmed plots between the 2014 and 2017 survey waves.

the Type I error rates of both the household fixed effects and plot fixed effects estimators are indicative of their unbiasedness. The plot fixed effects estimator incorrectly rejects the null hypothesis of separability at the 5% and 10% levels 50 and 94 times respectively. Similarly, the household fixed effects estimator incorrectly rejects the null hypothesis of separability at the 5% and 10% levels 44 and 92 times respectively.

The performance of the household fixed effects estimator worsens as the percentage of households with a land change increases. Under Scenario A, where 84% of simulated households have static farms and 16% gain or lose a second period plot, tests based on the household fixed effects estimator reject the null of separability 117 and 190 times at the 5% and 10% levels, respectively. When 30% and 37% of simulated households gain or lose a second period plot, the Type I errors at the 5% level increase to 133 and 245, respectively. These Type I error rates represent more than 100% increases relative to an unbiased estimator.

The empirical distributions of the coefficient estimate on X_{cht} presented in Figure 1 further illustrate the Type I error rate of tests based on the household fixed effects estimator. When all households have static farmland, the distribution of the coefficient estimates of the demographic variable using household fixed effects is correctly centered on zero. In all other scenarios, however, the distribution is shifted rightward. The magnitude of this shift increases as the land market becomes more active and a larger percentage of households change farmed land between survey waves.

In contrast, separability tests based on the plot fixed effects estimator are unaffected by the percentage of households that gained or lost a plot in the second survey wave. When 16%, 30%, or 37% of households change one or more plots between survey waves, the plot fixed effects estimator incorrectly rejects separability 47, 49, and 49 times, respectively. The empirical distributions of the coefficient estimate on X_{cht} based on plot fixed effects illustrate this consistency (Figure 2); the empirical distributions are tightly centered around zero under each scenario.

VII. Conclusion

Whether agricultural households make their production decisions separately from their consumption decisions is key to the design and evaluation of development programs and policies. Non-separability affects household production responses to new opportunities and shocks; its existence also indicates market failures.

This paper confronts an important identification challenge in common tests of this separability assumption–the endogeneity of household demographic characteristics with unobserved land quality. Leveraging intra-plot variability in labor demand, I find that the non-separability result in Rwanda is robust to controlling for land quality and other unobserved time-invariant land characteristics. Furthermore, using simulated data where the separability hypothesis is known to hold, I find that tests based on intra-household variability in labor demand are prone to false rejections of separability and that the likelihood of a false rejection increases as the land market becomes more active.

The Rwandan results are limited by the short, two-wave plot panel data which reduces observed intra-plot variability. This short timing, however, reduces the likelihood that selection into the plot panel leads to inconsistent estimation of the separability test (Wooldridge, 2010, pg. 830). I also cannot control for unobserved land characteristics when assessing the separability of agricultural households without a plot in the panel. The inclusion or exclusion of households not in the plot panel subsample, however, does not affect results based on household fixed effects.

The simulation results are also limited by the applicability of the underlying assumptions. In particular, the consequence of ignoring unobserved land quality in separability tests is dependent on the correlation with household demographic characteristics. The applicability of the plot-panel-based test is also dependent on some stability in the plots operated by the households. If the land market is volatile and few plots are maintained by the same household over time, then capturing plot panel survey data may not be feasible.

Despite these limitations, the findings in this paper provide important implications for the evaluation of households' responses to agricultural policies and programs. The robustness of the non-separability result in Rwanda begets the importance of agricultural development programs and policies which reduce market inefficiencies and consider the role of households' consumption preferences in their production decisions. Furthermore, the simulation results suggest that if the land market is not active and the vast majority of households have static farmland across all survey waves, then the appropriateness of studying households' production responses to policies in isolation of consumption decisions may be accurately assessed with household panel data. However, household production responses to policies are more likely to be separable when markets are working well, increasing the value of separability tests in contexts with relatively active land markets. When the land market is active over the sampling period, the findings suggest that a reliance on household panel data may misinform inference on the separability of agricultural households' production decisions. This could lead to a mischaracterization of household production responses to agricultural policies. Future work should incorporate plot panel data in a variety of different contexts to provide an updated view of agricultural household decision making which is robust to unobserved heterogeneity in land quality.

Tables and Figures

Table 1: Household Summary Statistics

	Plot Panel Households	All Households
Household size	5.05	5.05
	(1.78)	(1.78)
Share of female members 0 to 14 years old	0.24	0.24
	(0.18)	(0.18)
Share of male members 0 to 14 years old	0.24	0.24
	(0.18)	(0.18)
Share of female members 15 to 60 years old	0.26	0.26
	(0.12)	(0.12)
Share of male members 15 to 60 years old	0.22	0.22
	(0.13)	(0.14)
Share of female members 61 years or older	0.02	0.02
	(0.09)	(0.09)
Share of male members 61 years or older	0.02	0.01
	(0.06)	(0.06)
Farmed land area across all plots (ha)	1.08	1.27
	(12.92)	(13.08)
Number of farmed plots	2.44	2.34
	(1.40)	(1.40)
Number of household-wave observations	2,988	3,600

Note: Means with standard deviations in parentheses. Plot panel households refers to the subsample of households with at least one plot successfully linked across survey waves. Farmed land area and number of farmed plots correspond to the February-to-June agricultural season.

Table 2: Plot Summary Statistics

	Panel Subsample	All Plots
Plot size (ha)	0.22	0.44
	(2.96)	(8.09)
Total labor demand (labor-days)	24.81	24.21
	(36.22)	(35.29)
Family labor demand (labor-days)	21.30	20.76
	(25.23)	(25.38)
Hired labor demand (labor-days)	3.50	3.45
	(24.16)	(22.17)
Number of plot-wave observations	4,580	7,303

Note: Means with standard deviations in parentheses. These statistics are restricted to plots farmed by the 1,494 households with at least one plot in the plot panel. Panel subsample refers to the plots successfully linked across survey waves. "All plots" refers to any plot farmed by the 1,494 households in at least one period. Labor demand consists of labor days used for land preparation, planting, and field management after planting for the February-to-June agricultural season. Labor from children is scaled by 0.5.

Table 3: Farm and Plot Labor Demand (Log of Person Days per Season) on Household Characteristics

	Household FE		Plot FE	
	(1)	(2)	(3)	(4)
ln(Household size)	0.242**	0.356***	-0.012	-0.021
	(0.111)	(0.134)	(0.112)	(0.140)
Share of female members 0 to 14 years old		0.222		1.842***
		(0.687)		(0.609)
Share of male members 0 to 14 years old		0.221		2.057***
		(0.675)		(0.621)
Share of female members 15 to 60 years old		0.909		2.298***
		(0.650)		(0.612)
Share of male members 15 to 60 years old		0.683		2.221***
		(0.659)		(0.596)
Share of female members 61 years or older		0.940		2.096***
		(0.672)		(0.599)
F-test p-value for demogr. vars. joint signif.	0.030	0.059	0.916	0.001
Number of observations	2,988	2,988	4,580	4,580
Number of FE groups	1,494	1,494	2,290	2,290

Note: Coefficient estimates with village-level cluster-robust standard errors in parentheses. Dependent variable in (1) and (2) is the natural log of the sum of pre-harvest person days of labor used by a household across all operated plots during the February-to-June agricultural season of a given year. Dependent variable in (3) and (4) is the natural log of pre-harvest person days of labor used on a plot during the February-to-June agricultural season of a given year. The omitted demographic share in (2) and (4) is male members 61 years or older. All models control for the age, education, and gender of the household head and community-time fixed effects. (1) and (2) include household fixed effects and the natural log of farmed land area in a given year. (3) and (4) include plot fixed effects. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 4: Sum of Labor Demand Across Plot Panel Plots (Log of Person Days per Season) on Household Characteristics

	(1)	(2)
ln(Household size)	-0.052	-0.069
	(0.110)	(0.138)
Share of female members 0 to 14 years old		1.609***
		(0.607)
Share of male members 0 to 14 years old		1.720***
		(0.604)
Share of female members 15 to 60 years old		1.938***
		(0.599)
Share of male members 15 to 60 years old		2.003***
		(0.581)
Share of female members 61 years or older		1.670***
		(0.562)
F-test p-value for demogr. vars. joint signif.	0.635	0.025
Number of observations	2,988	2,988
Number of FE groups	1,494	1,494

Note: Coefficient estimates with village-level cluster-robust standard errors in parentheses. Dependent variable is the natural log of the sum of pre-harvest person days of labor used by a household across all *plot-panel* plots during the February-to-June agricultural season of a given year. Plots observed in only one survey wave (and thus excluded from the plot panel) are omitted. The omitted demographic share in (2) is male members 61 years or older. All models control for the age, education, and gender of the household head, household fixed effects, and community-time fixed effects.
* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 5: Simulated Second Period Land Changes

Scenario	Prob. Keeping Each old plot	Prob. Gaining Each of 5 new plots	Avg. % of HHs that Lost a Plot	Avg. % of HHs that Gained a Plot	Avg. % of HHs with Land Change
Static Farmland	1	0	0%	0%	0%
Scenario A	0.99	0.03	2%	14%	16%
Scenario B	0.95	0.05	10%	22%	30%
Scenario C	0.95	0.07	10%	30%	37%

Note: The probability of keeping each old plot is the probability that a given plot farmed in the first period is maintained in the second period. The probability of gaining each of five new plots is the probability that a given new plot is incorporated into a household's farmed plots in the second period. The average percent of households that lost a plot corresponds to the percentage of households that stopped farming a first period plot in the second period averaged across the 1,000 replications. The average percent of households that gained a plot corresponds to the percentage of households that farmed a new plot in the second period that they had not farmed in the first period averaged across the 1,000 replications. The average percent of households with a land change is the percentage of simulated households that lost or gained at least one plot between survey waves averaged across the 1,000 replications.

Table 6: Number of Observed Type I Errors in 1,000 Replications

	Percentage of Simulated Households with Land Change							
	0%		16%		30%		37%	
	Significance Level		Significance Level		Significance Level		Significance Level	
Estimator	0.05	0.10	0.05	0.10	0.05	0.10	0.05	0.10
Plot Fixed Effects	50	94	47	93	49	90	49	90
Household Fixed Effects	44	92	117	190	133	205	245	353

Note: Type I error is for incorrect rejection of the null hypothesis of separability. Percentage of simulated households with land change is the percentage of simulated households that lost or gained at least one plot between survey waves averaged across the 1,000 replications. An unbiased estimator would reject the null at the 5% and 10% levels approximately 50 and 100 times respectively across 1,000 replications.

Figure 1: Empirical Distribution via Household Fixed Effects

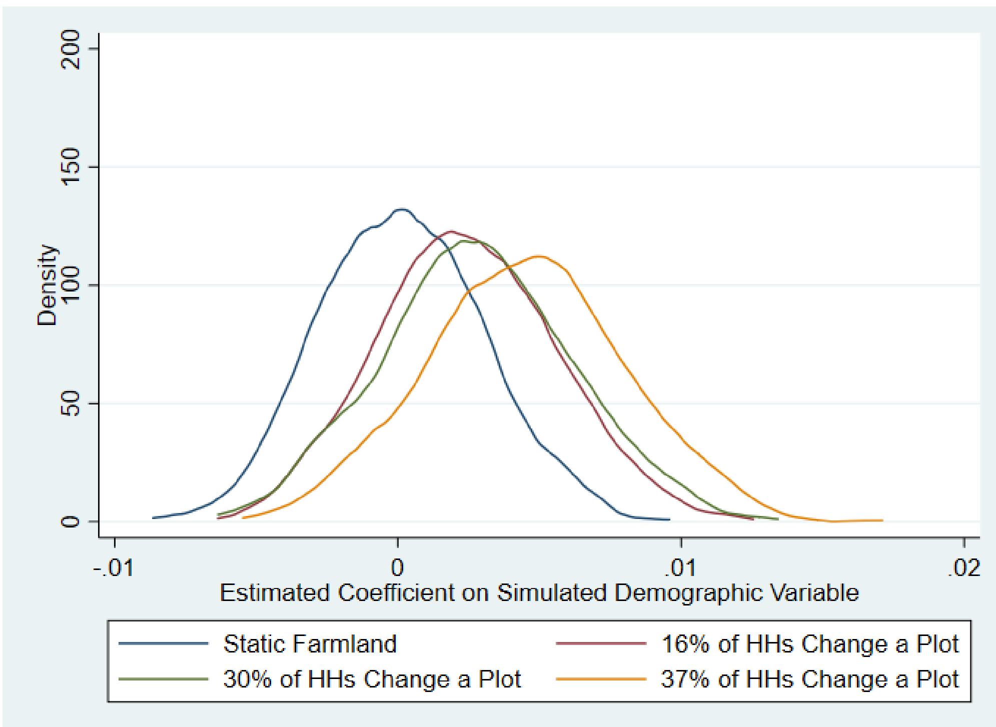

Figure 2: Empirical Distribution via Plot Fixed Effects

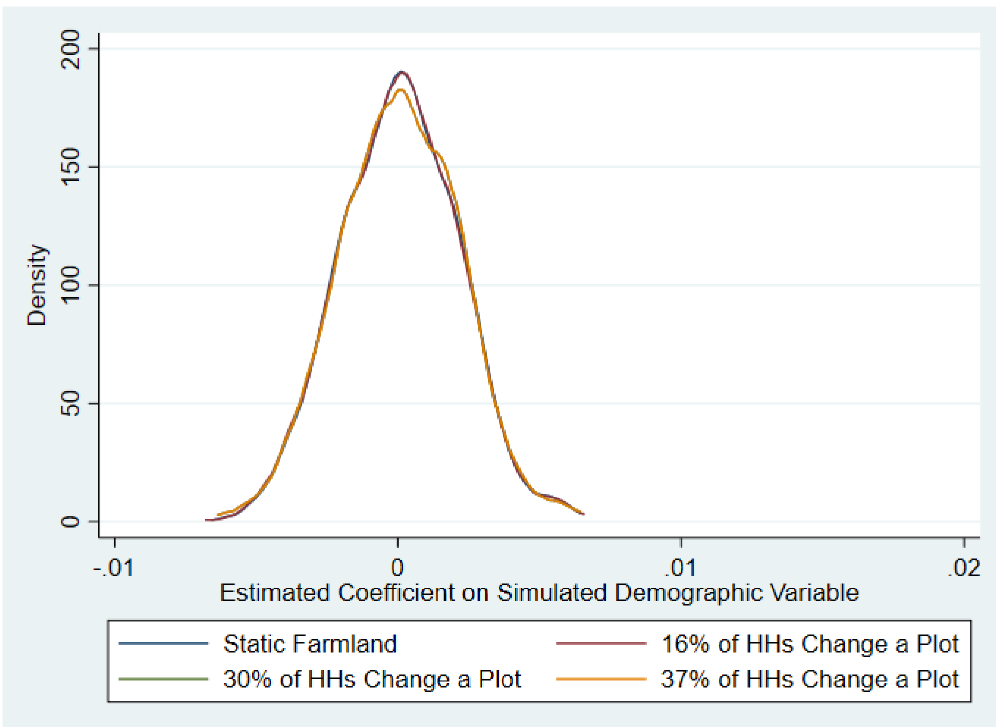

APPENDIX:

Supplementary Tables

Table 7: Labor Demand (Log of Person Days per Season) on Household Characteristics: Accounting for Potential Child Productivity Differences

	Household FE		Plot FE	
	(1)	(2)	(3)	(4)
ln(Num. of adult equivalent household members)	0.362***	0.355***	0.055	−0.019
	(0.123)	(0.133)	(0.128)	(0.139)
Share of female members 0 to 14 years old		0.442		1.825***
		(0.665)		(0.582)
Share of male members 0 to 14 years old		0.442		2.041***
		(0.656)		(0.592)
Share of female members 15 to 60 years old		0.900		2.295***
		(0.652)		(0.613)
Share of male members 15 to 60 years old		0.679		2.219***
		(0.660)		(0.596)
Share of female members 61 years or older		0.937		2.096***
		(0.671)		(0.599)
F-test p-value for demogr. vars. joint signif.	0.003	0.058	0.668	0.001
Number of observations	2,988	2,988	4,580	4,580
Number of FE groups	1,494	1,494	2,290	2,290

Note: Coefficient estimates with village-level cluster-robust standard errors in parentheses. Dependent variable in (1) and (2) is the natural log of the sum of pre-harvest person days of labor used by a household across all operated plots during the February-to-June agricultural season of a given year. Dependent variable in (3) and (4) is the natural log of pre-harvest person days of labor used on a plot during the February-to-June agricultural season of a given year. Number of adult equivalent household members counts members under 15 as half an adult to account for potential labor productivity differences. The omitted demographic share in (2) and (4) is male members 61 years or older. All models control for the age, education, and gender of the household head and community-time fixed effects. (1) and (2) include household fixed effects and the natural log of farmed land area in a given year. (3) and (4) include plot fixed effects.
* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 8: Farm Labor Demand (Log of Person Days per Season) on Household Characteristics: Controlling for Land Quality Proxy

	No Land Quality Control		Controlling for Land Quality Proxy	
	(1)	(2)	(3)	(4)
ln(Household size)	0.242**	0.356***	0.240**	0.358***
	(0.111)	(0.134)	(0.111)	(0.134)
Share of female members 0 to 14 years old		0.222		0.211
		(0.687)		(0.686)
Share of male members 0 to 14 years old		0.221		0.213
		(0.675)		(0.675)
Share of female members 15 to 60 years old		0.909		0.906
		(0.650)		(0.650)
Share of male members 15 to 60 years old		0.683		0.664
		(0.659)		(0.660)
Share of female members 61 years or older		0.940		0.991
		(0.672)		(0.667)
F-test p-value for demogr. vars. joint signif.	0.030	0.059	0.031	0.059
Number of observations	2,988	2,988	2,988	2,988
Number of FE groups	1,494	1,494	1,494	1,494

Note: Coefficient estimates with village-level cluster-robust standard errors in parentheses. Dependent variable is the natural log of the sum of pre-harvest person days of labor used by a household across all operated plots during the February-to-June agricultural season of a given year. (1) and (2) reproduce the results from columns (1) and (2) of Table 3. (3) and (4) control for household's reported farmed land area of high quality. The omitted demographic share in (2) and (4) is male members 61 years or older. All models control for age, education, and gender of the household head, the natural log of farmed land area in a given year, household fixed effects, and community-time fixed effects. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 9: Farm Labor Demand (Log of Person Days per Season) on Household Characteristics: Incorporating All Households

	Plot Panel Households		All Households	
	(1)	(2)	(3)	(4)
ln(Household size)	0.242**	0.356***	0.190**	0.332***
	(0.111)	(0.134)	(0.092)	(0.112)
Share of female members 0 to 14 years old		0.222		0.063
		(0.687)		(0.632)
Share of male members 0 to 14 years old		0.221		0.102
		(0.675)		(0.600)
Share of female members 15 to 60 years old		0.909		0.875
		(0.650)		(0.595)
Share of male members 15 to 60 years old		0.683		0.687
		(0.659)		(0.576)
Share of female members 61 years or older		0.940		0.820
		(0.672)		(0.614)
F-test p-value for demogr. vars. joint signif.	0.030	0.059	0.040	0.013
Number of observations	2,988	2,988	3,600	3,600
Number of FE groups	1,494	1,494	1,800	1,800

Note: Coefficient estimates with village-level cluster-robust standard errors in parentheses. Columns (1) and (2) are restricted to the subsample of households with at least one plot in the plot panel, reproducing the results from columns (1) and (2) of Table 3. Columns (3) and (4) include all households. Dependent variable is the natural log of the sum of pre-harvest person days of labor used by a household across all operated plots during the February-to-June agricultural season of a given year. The omitted demographic share in (2) and (4) is male members 61 years or older. All models control for age, education, and gender of the household head, the natural log of farmed land area in a given year, household fixed effects, and community-time fixed effects.
* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 10: Farm and Plot Labor Demand (Log of Person Days per Season) on Household Characteristics: Excluding Migrant Households

	Household Fixed Effects				Plot Fixed Effects	
	Excl. Migrant HHs		Excl. Migrant & Non-Plot Panel HHs		Excl. Migrant HHs	
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Household size)	-0.492*	-0.349	-0.389	-0.051	-0.814***	-0.566
	(0.251)	(0.376)	(0.275)	(0.414)	(0.282)	(0.433)
Share of female members 0 to 14 years old		0.312		-0.751		0.541
		(1.659)		(1.466)		(1.501)
Share of male members 0 to 14 years old		0.667		-0.348		0.358
		(1.642)		(1.454)		(1.483)
Share of female members 15 to 60 years old		0.484		-0.364		0.692
		(1.844)		(1.709)		(1.758)
Share of male members 15 to 60 years old		0.873		0.176		0.921
		(1.557)		(1.378)		(1.465)
Share of female members 61 years or older		1.515		0.748		1.242
		(1.897)		(1.739)		(1.897)
F-test p-value for demogr. vars. joint signif.	0.051	0.255	0.158	0.345	0.004	0.073
Number of observations	2,208	2,208	1,864	1,864	2,962	2,962
Number of FE groups	1,104	1,104	932	932	1,481	1,481

Note: Coefficient estimates with village-level cluster-robust standard errors in parentheses. All columns are restricted to the subsample of households where no person left or joined the household except for children born between survey waves. Columns (3) and (4) further restrict the household sample to exclude households without a plot in the plot panel. Dependent variable in (1)-(4) is the natural log of the sum of pre-harvest person days of labor used by a household across all operated plots during the February-to-June agricultural season of a given year. Dependent variable in (5) and (6) is the natural log of pre-harvest person days of labor used on a plot during the February to June agricultural season of a given year. The omitted demographic share in (2), (4), and (6) is male members 61 years or older. All models control for the age, education, and gender of the household head and community-time fixed effects. (1)-(4) include household fixed effects and the natural log of farmed land area in a given year. (5) and (6) include plot fixed effects. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

REFERENCES

REFERENCES

Benjamin, D. (1992). Household Composition Labor Markets and Labor Demand, Testing for Separation in Ag Household Models. *Econometrica*, *60*(2), 287–322.

Conley, T. G., & Udry, C. R. (2010). Learning about a New Technology: Pineapple in Ghana. *American Economic Review*, *100*(1), 35–69.

Dillon, B., & Barrett, C. B. (2017). Agricultural factor markets in Sub-Saharan Africa: An updated view with formal tests for market failure. *Food Policy*, *67*, 64–77.

Dillon, B., Brummund, P., & Mwabu, G. (2019). Asymmetric non-separation and rural labor markets. *Journal of Development Economics*, *139*, 78–96.

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.

Kopper, S. A. (2018). *Agricultural labor markets and fertilizer demand: Intensification is not a single factor problem for non-separable households.*

LaFave, D., & Thomas, D. (2016). Farms, Families, and Markets: New Evidence on Completeness of Markets in Agricultural Settings. *Econometrica*, *84*(5), 1917–1960.

Pitt, M., & Rosenzweig, M. (1986). Agricultural Prices, Food, Consumption and the Health and Productivity of Indonesian Farmers. In I. Singh, L. Squire, & J. Strauss (Eds.), *Agricultural Household Models: Extensions, Applications, and Policy.* Baltimore: The Johns Hopkins University Press.

Sheahan, M., Black, R., & Jayne, T. (2013). Are Kenyan farmers under-utilizing fertilizer? Implications for input intensification strategies and research. *Food Policy*, *41*, 39–52.

Singh, I., Squire, L., & Strauss, J. (1986). *Agricultural Household Models: Extensions, Applications, and Policy.* Baltimore: The Johns Hopkins University Press.

Suri, T. (2011). Selection and Comparative Advantage in Technology Adoption. *Econometrica*, *79*(1), 159–209.

Udry, C. R. (1999). Efficiency and Market Structure: Testing for Profit Maximization in African Agriculture. In G. Ranis & L. Raut (Eds.), *Trade, Growth and Development: Essays in Honor of T.N. Srinivasan.* Amsterdam: Elsevier Science.

Wooldridge, J. (2010). *Econometric Analysis of Cross Section and Panel Data* (2nd ed.). Cambridge, MA: MIT Press.