

**ASSESSING DEVELOPMENT OUTCOMES WHEN
WEATHER, LAND, AND PEOPLE DIFFER**

By

Jarrad Godwin Farris

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Agricultural, Food, and Resource Economics - Doctor of Philosophy

2020

ABSTRACT

ASSESSING DEVELOPMENT OUTCOMES WHEN WEATHER, LAND, AND PEOPLE DIFFER

By

Jarrad Godwin Farris

Agricultural development and policy design rely on careful analysis of the factors that influence the welfare and decision making of agricultural households. This dissertation leverages diverse data types, cross-disciplinary knowledge, and applied econometrics to assess the underlying factors that influence child welfare and agricultural production decisions. In doing so, it reveals the role of observed and unobserved differences in shaping our understanding of decision making in agricultural production systems.

The first chapter evaluates the impacts of *in utero* rainfall on child growth in rural Rwanda and assesses whether estimates based on aggregate *in utero* rainfall are attenuated by intra-seasonal *in utero* rainfall effects. The *in utero* period of a child's life is critical for his or her development. For families relying on rain-fed agricultural production, such development can be severely impacted by the timing of rainfall shocks. Yet, evidence of *in utero* rainfall effects has been mixed. My results suggest that this mixed evidence may be driven by a focus on aggregate rainfall measures, which ignore cropping period specific heterogeneity in rainfall effects on human health.

The second chapter assesses the separability hypothesis which posits that agricultural households make their production decisions separately from their consumption decisions. This theory relates closely to the completeness of markets and provides an important avenue for understanding how agricultural households are likely to respond to new policies and programs. The current standard identification strategy for testing whether this separability hypothesis holds is to estimate reduced form regressions of household labor demand on household demographic characteristics, using household fixed effects to address unobserved

heterogeneity. Using plot panel data from Rwanda, I apply an alternative test that controls for unobserved heterogeneity in land quality. Using simulations, I then show that the standard approach based on household fixed effects is prone to omitted variable bias from the endogeneity of household demographic characteristics with unobserved land characteristics. Simulations indicate that this bias is exacerbated as the land market becomes more active.

The third chapter examines the role of farmer personality in the effectiveness of a community-based extension program for promoting improved bean varieties in Tanzania. I develop a conceptual framework which shows that the information gained from community-based extension activities may be heterogeneous by farmer personality. I then examine this potential heterogeneity empirically using a unique dataset of the Big Six personality traits measured using the Midlife Development Inventory (MIDI). My findings suggest that personality characteristics influenced which farmers benefited from the extension program. In particular, more extraverted farmers appear to have benefited more from residing in villages that received trial packs of improved bean seed relative to less extraverted farmers. This is consistent with their increased sociability and has implications for the types of farmers likely to gain from community-based extension programs.

ACKNOWLEDGMENTS

A huge thank you to my advisor, Mywish Maredia, for always encouraging me to ask questions, seek out diverse research, and find ways to improve. Thank you for being such a wonderful role model.

I am also forever grateful for the other members of my committee: Maria Porter, for instilling a problem solving tenacity; Songqing Jin, for believing in a plot panel and making sure it happened; and Robby Richardson, for encouraging the integration of cross-disciplinary ideas. I am very fortunate for all of the support and guidance from my committee throughout this process.

I would also like to thank Nicky Mason, for her unwavering support and for making difficult development topics accessible; David Ortega, for introducing me to the idea of pursuing the intersection of psychology and economics; Roy Black, for teaching me about tart cherry production; and Trey Malone, for not giving up on a student's research. A heartfelt thank you also goes to the many, many faculty, staff, and students in the Department of Agricultural, Food, and Resource Economics who shared knowledge, feedback, advice, and kindness throughout this process.

I am also grateful for the fieldwork and data collection support from Incisive Africa, the Agricultural Research Institute-Uyole, CIP, PIM, Imbaraga, and YWCA. Thanks also go to all the farmers and their family members who took the time to share their knowledge and make the surveys possible as well as to the SurveyCTO staff for answering technical CAPI programming questions at all hours of the night.

To my amazing family, you are the support system that kept me going. A huge hug to my parents, for showing me how to appreciate the wonders of the backyard and globe alike; Leigh, for your open ear, kind heart, and helping me de-stress; Fern, Khalil, Gwen, Andy, and Whitney, for teaching me how to look up with an open mind and balance critical

thinking with deep breaths; and Sarah, for all the emotional support, fresh perspectives, good vibes, and laughs—thanks for putting up with me and making me go outside.

TABLE OF CONTENTS

LIST OF TABLES	viii
LIST OF FIGURES	x
Chapter 1 Growing Pains: Cropping Period-Specific In Utero Rainfall Shocks Impact Child Growth in Rural Rwanda	1
I. Introduction	2
II. Crop Yield Response to Rainfall	5
III. Data	7
III.1 Household Survey and Anthropometrics Data	7
III.2 Rainfall Data and In Utero Rainfall Shocks	8
IV. Empirical Strategy	9
V. Empirical Results	11
V.1 Main Findings	11
V.2 Robustness Checks	14
VI. Conclusion	15
APPENDIX	17
REFERENCES	26
Chapter 2 Does Unobserved Land Quality Bias Separability Tests?	30
I. Introduction	31
II. Theoretical Model	33
III. Empirical Strategy	36
IV. Data	38
V. Results	40
VI. Simulation	42
VI.1 Simulation Methods	42
VI.2 Simulation Results	46
VII. Conclusion	48
APPENDICES	50
APPENDIX A: Tables and Figures	51
APPENDIX B: Supplementary Tables	60
REFERENCES	65
Chapter 3 Farmer Personality and Community-Based Extension Effectiveness in Tanzania	68
I. Introduction	69
II. Background and Experimental Design	71
III. Conceptual Framework	73
IV. Data	80

V. Empirical Strategy	85
VI. Results	87
VII. Conclusions	91
APPENDICES	93
APPENDIX A: Tables and Figures	94
APPENDIX B: Supplementary Tables	109
REFERENCES	112

LIST OF TABLES

Table 1.1: Summary Statistics	18
Table 1.2: Annual In Utero Rainfall Effects on Child Growth (Height-for-Age Z-Score)	19
Table 1.3: Cropping Period-Specific In Utero Rainfall Effects on Child Growth (Height-for-Age Z-Score)	20
Table 1.4: In Utero Rainfall Effects on Child Growth (Height-for-Age Z-Score) by Demographic and Household Characteristics	21
Table 1.5: In Utero Rainfall Effects on Height-for-Age Z-Score: Incorporating Minor Season and Rainfall in Other Periods	22
Table 2.1: Household Summary Statistics	52
Table 2.2: Plot Summary Statistics	53
Table 2.3: Farm and Plot Labor Demand (Log of Person Days per Season) on Household Characteristics	54
Table 2.4: Sum of Labor Demand Across Plot Panel Plots (Log of Person Days per Season) on Household Characteristics	55
Table 2.5: Simulated Second Period Land Changes	56
Table 2.6: Number of Observed Type I Errors in 1,000 Replications	57
Table 2.7: Labor Demand (Log of Person Days per Season) on Household Characteristics: Accounting for Potential Child Productivity Differences	61
Table 2.8: Farm Labor Demand (Log of Person Days per Season) on Household Characteristics: Controlling for Land Quality Proxy	62
Table 2.9: Farm Labor Demand (Log of Person Days per Season) on Household Characteristics: Incorporating All Households	63
Table 2.10: Farm and Plot Labor Demand (Log of Person Days per Season) on Household Characteristics: Excluding Migrant Households	64

Table 3.1: VBAA Balance Tests	95
Table 3.2: Farmer Survey Balance Tests	96
Table 3.3: Household Attrition Tests	97
Table 3.4: Key Characteristics by Agricultural Year and Treatment Group	98
Table 3.5: Midlife Development Inventory (MIDI) Personality Traits	99
Table 3.6: Personality Trait Cronbach’s Alpha Values	100
Table 3.7: Farmer-VBAA Personality Traits Comparison	101
Table 3.8: Physical Distance Summary Statistics	102
Table 3.9: Farmer Adoption of Improved Bean Varieties in 2018 Ag. Year	103
Table 3.10: Heterogeneous Trial Pack Effects on Farmer Adoption of Improved Bean Varieties in 2018 Ag. Year	104
Table 3.11: Farmer Interactions with VBAA	105
Table 3.12: Farmer Interactions with Other Bean Farmers	106
Table 3.13: Heterogeneous Trial Pack Effects on the Proportion of Bean Farming House- holds that the Farmer has Sought Farming Advice From	107
Table 3.14: Heterogeneous Trial Pack Effects on the Proportion of Bean Farming House- holds that the Farmer has Discussed Bean Farming With	108
Table 3.15: Farmer Adoption of Improved Bean Varieties in 2018 Ag. Year: Controlling for Pre-Intervention Outcome	110
Table 3.16: Heterogeneous Trial Pack Effects on Farmer Adoption of Improved Bean Varieties in 2018 Ag. Year: Controlling for Pre-Intervention Outcome	111

LIST OF FIGURES

Figure 1.1: Rwanda’s Cropping Period Calendar Around Two Rainy Seasons – Includes Minor Season	23
Figure 1.2: Rwanda’s Cropping Period Calendar – Includes Minor Season and Restricts Major Seasons	24
Figure 1.3: Histogram of Child Birth Months	25
Figure 2.1: Empirical Distribution via Household Fixed Effects	58
Figure 2.2: Empirical Distribution via Plot Fixed Effects	59

Chapter 1

Growing Pains: Cropping Period-Specific In Utero Rainfall Shocks Impact Child Growth in Rural Rwanda

I. Introduction

Throughout the developing world, households reliant on rain-fed agricultural production face increasing risk of food insecurity due to increased variability and uncertainty in rainfall (e.g. Funk et al., 2019). Climate change has increased both the risk of prolonged drought, as well as heavy rainfall (Lehmann et al., 2018). The increasing prevalence of such weather shocks is likely to have very dire long-term consequences for those who are exposed to such shocks in the most critical early stages of human development.

The *in utero* period of a child’s life is critical for his or her development, with long-term implications for human capital accumulation (Almond et al., 2017; Barker, 1998). Exposure to poor environmental conditions during this period has been linked to a wide array of negative consequences. For example, negative rainfall shocks in the year before birth result in lower birth weights and increased infant mortality (Rocha & Soares, 2015). The negative consequences of *in utero* shocks need not be limited to infancy, nor do they require extreme events. In their recent meta-analysis, Almond et al. (2017) emphasize that even moderate early life shocks can have long-term negative consequences.

This article contributes to previous research studying the effects of *in utero* rainfall shocks by differentiating between cropping-period rainfall shocks that are specific to planting, mid-season, and harvest seasons. My identification strategy captures potential differences in the marginal effects of rainfall depending on the crop development stage in which the rainfall shocks occur. Such differences would be subsumed in aggregate measures of rainfall (e.g. deviations from annual averages) that have been used in many previous studies.

I argue that the direction of the effect of rainfall shocks on child growth differs according to the cropping period at the time of the rainfall shock, which implies that aggregating rainfall outcomes into a single measure will attenuate the estimated effects on child growth. By not taking such seasonal differences into account, rainfall measures that are homogeneous across different cropping periods in a given year may not be capturing important differences

in cropping period-specific shocks to child growth. Such seasonal differences are particularly important in countries where agricultural production is primarily rain-fed, where irrigation is largely infeasible, and where there are multiple production seasons, implying relatively brief windows for planting, growing and harvesting seasons.

Such agricultural production is exemplified in Rwanda, the focus of this study. In rural Rwanda, as in many rural developing country contexts, agriculture is the principal source of livelihood. As rural Rwandan households rely primarily on rainfed agriculture, rainfall outcomes are a key source of exogenous variability in household incomes and hence nutrition availability.

The literature measuring rainfall impacts on health has led to contradictory conclusions. For example, two meta-analyses combining DHS data on African countries find both positive and negative effects of rainfall on mortality (Comfort, 2016; Kim, 2010). In Mexico, rainfall that was one standard deviation below average during the wet season reduced height-for-age z-scores (HAZ) for a subsample of children in the north region, but increased HAZ for a subsample of children in the center and south regions (Skoufias et al., 2011). Alongside these mixed findings, other studies that have measured the effects of either annual rainfall shocks (e.g. Abiona, 2017; Burgess et al., 2017; Comfort, 2016; Kim, 2010; Maccini & Yang, 2009; Rocha & Soares, 2015) or rainfall shocks during the growing season (e.g. Kudamatsu et al., 2012; Shively, 2017; Skoufias et al., 2011) have found a positive relationship between rainfall and health outcomes.

One response to the mixed findings on the rainfall-health relationship is a strand of the literature that seeks to assess this relationship using non-rainfall measures. Early work in Rwanda analyzed a crop failure in 1988-1989 that occurred for a particular cohort of children in one region of the country (Akresh et al., 2011). Girls born in the region during the famine had reduced growth rates compared to girls born in the same region at a different time, or girls born at the same time but in a different region. No effect was found for boys. Similarly,

in a panel analysis of children in Zimbabwe, children who were 12-24 months old during a 1994-1995 food shortage were 1.5 to 2 cm shorter on average than children who reached that age in a non-food shortage year (Hoddinott & Kinsey, 2001). Among communities in Ethiopia who were asked to directly report the months when food was relatively scarce, children exposed to more food-scarce months *in utero* had significantly lower heights at age eight (Miller, 2017).

Rainfall-based measures have several advantages over non-rainfall-based indicators. Rainfall data can assess the impact of small deviations without extreme events or self-reporting of food scarcity. The rainfall data used in this study is at the 0.05 degree resolution, allowing daily rainfall to be estimated at the household level. As rainfall data and rainfall-based measures improve, rainfall-based measures become increasingly useful in teasing out complex relationships between nutrition availability and health.

This study contributes in several ways to an emerging literature exploring the rainfall-health relationship using intra-season rainfall measures. First, I allow for cropping period-specific rainfall shocks to differ in whether they improve or harm children's health. Prior studies which have incorporated intra-season heterogeneity have forced all positive or all negative deviations to have the same-signed effect, regardless of the cropping period during which the rainfall occurred (Cornwell & Inder, 2015; Tiwari et al., 2017). Grouping all deviations in the same direction into a single category ignores potentially important differences in the timing of rainfall. Second, I examine rainfall effects on health from three different major cropping periods, as opposed to examining effects from rainfall in only one cropping period (Aguilar & Vicarelli, 2011).

I show that aggregate measures of rainfall shocks conceal acute differences in the relationship between rainfall and child growth. As the intra-year relationship between rainfall and agricultural productivity need not be constant, the relationship between rainfall experienced *in utero* and child growth may also vary according to the cropping period. I find

that during planting and in mid-season, increases in rainfall have a positive relationship with child growth. Maize yields can be particularly sensitive to water deficits during the mid-season flowering and grain-filling stages. Less rainfall during the harvest period, however, can improve yields, as this reduces grain water content (Barron et al., 2003). Indeed, I find that increases in rainfall during harvest reduce child growth.

The remainder of the paper proceeds as follows. Section two discusses the crop yield response to rainfall as it relates to rural Rwanda and describes the cropping period specific rainfall measures. The third and fourth sections describe the data sources and the empirical strategy respectively. The fifth section outlines the empirical results and the final section concludes.

II. Crop Yield Response to Rainfall

Two seasons with the same total rainfall can have very different yield outcomes depending on how rainfall was distributed within the season (Brown, 2008; HarvestChoice, 2010). The crop-yield response to rainfall varies over the course of the growing season (Brown, 2008; Doorenbos & Kassam, 1979; Steduto et al., 2012). For most crops, yield response to water requirements is relatively low in early and late growth periods and relatively high during periods of flowering and crop formation. For example, for one of Rwanda’s most important food crops – beans – the yield response to water deficit is four to five times larger during flowering and pod filling phases compared to during its vegetative and ripening periods (Doorenbos & Kassam, 1979).

The relationship between rainfall and yield is not always positive. Extreme rainfall can lead to waterlogging and aeration stress. As in crop stress from insufficient rainfall, aeration stress from excess rainfall can lead to crop yield reductions (Steduto et al., 2012). For example, maize – a major staple in Rwanda – is relatively sensitive to water stress during the flowering period (Doorenbos & Kassam, 1979).

Even in the absence of extreme rainfall events, there may be periods when less rainfall is beneficial. A period of little rainfall at harvest can improve maize yield by reducing grain water content (Barron et al., 2003). Similarly, a period of no rain for 20 to 25 days before dry bean harvest is ideal (Doorenbos & Kassam, 1979).

To capture key differences in crop yield response, I distinguish between three distinct periods for each agricultural season. The land preparation and planting period is the period when most crops are beginning their growth. The mid-season period is the period when most crops are in their critical growth stages and are relatively sensitive to water stress. Finally, the harvest period is the period when most crops are mature and may benefit from a tapering of rainfall.

As shown in Figure 1.1, Rwanda’s first major agricultural season, which extends from February to July, is split into: land preparation and planting (February and March), mid-season (April and May), and harvest (June and July). Similarly, Rwanda’s other major agricultural season, which extends from August to January, is split into land preparation and planting (August and September), mid-season (October and November), and harvest (December and January).¹

For each of these individual cropping periods, I estimate a distinct marginal effect of *in utero* rainfall shocks on child height-for-age z-scores, holding all other potential determinants constant (i.e. parental characteristics, and other environmental factors).² I define the *in utero* period of a child’s life as the entire year before birth.³ I denote rainfall shocks (specific to child i in household h) in the land preparation and planting, mid-season, and harvest periods while *in utero* ($t = 0$) by the following: R_{ih0}^1 , R_{ih0}^2 , and R_{ih0}^3 . These rainfall shocks determine the height (i.e., length) at birth of child i in household h (H_{ih0}), which in turn

¹I also estimate robustness specifications incorporating alternative season and cropping period definitions.

²As height is a long-term measure of health, there are natural dynamics in the production process for child height. Strauss and Thomas (1998) and Maccini and Yang (2009) provide useful simple reduced forms for dynamic health production functions.

³The additional three months preceding pregnancy are included as a child’s health endowment can be influenced by the mother’s health pre-conception (Rocha & Soares, 2015).

affects subsequent height-for-age z-scores in year t (HAZ_{iht}).

III. Data

III.1 Household Survey and Anthropometrics Data

This study uses household survey data collected in 2014 and 2017 from households who had children under the age of five.⁴ This survey was a two round panel survey of households from eight districts across three provinces in Rwanda: Northern, Southern, and Eastern. In the enumeration areas, there are significant concerns regarding food security and child malnutrition.

In 2014, the survey used a three-stage cluster sampling method where the primary, secondary, and tertiary sampling units were the sector, village, and household respectively. A census was first conducted in each of the 252 selected villages to collect basic information of all households. A random sample of households with either a pregnant woman or child under the age of five were then selected (Peters et al., 2015). The survey then followed up with these same households in 2017 for the second round of the survey.

The survey instrument was comprehensive, covering a wide array of household and individual characteristics. It collected detailed child-specific information on nutrition and health, as well as weight, height, age, and gender.

This study uses the anthropometric information collected for all children less than the age of five. If a child was under the age of five and was measured in both survey rounds, the most recent measurement from the second round of surveying is used.

The rate of stunting in the sample is 39%, well above the expected rate of 2.3% in a healthy population (World Health Organization [WHO], 1997). This is not surprising given the rural, developing country context. The mean and median HAZ in the sample are -1.57

⁴The survey was conducted as part of an impact evaluation of the International Potato Center's Scaling Up Sweet Potato Through Agriculture and Nutrition (CIP-SUSTAIN) project.

and -1.64 respectively, indicating that most of the children are below average height for age (see Table 1.1). These averages also highlight the relevance of this research for the study area.

III.2 Rainfall Data and In Utero Rainfall Shocks

The most reliable source of rainfall data is considered to be the Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS), according to the Famine Early Warning Systems Network (Famine Early Warning Systems Network [FEWS NET], 2018). This publicly available dataset combines 0.05 degree satellite data with in-station data to provide daily rainfall estimates from January 1st 1981 to near present day (Funk et al., 2015).

I estimate average rainfall within five kilometers of each household, with GPS coordinates for each day from January 1st, 1981 to December 31st, 2017. A five-kilometer radius captures the majority of household plots, as more than 85% of all plots were within an hour's walk from households' dwellings. For households missing GPS coordinates, I estimate the average rainfall among all other households in the same village.

Rainfall shocks are commonly measured as the deviation of the log of rainfall in a given year from the log of historical annual average rainfall (Maccini & Yang, 2009; Rocha & Soares, 2015). However, a focus on annual rainfall shock measures may mask cropping period-specific heterogeneity in the effects of *in utero* rainfall shocks. I address this limitation by applying the same log-deviation, *in utero* shock measure as used by Rocha and Soares (2015), but take a cropping period approach. I estimate *in utero* rainfall shocks (R_{imh0}^k) as the deviation of the log of total rainfall in cropping period k that child i born in month m in household h experienced while *in utero*, from the log of the historical average for that period. The equation for R_{imh0}^k is:

$$R_{imh0}^k = \ln \left(\sum_{t=m-12}^{m-1} z_{kt} * r_{ht} \right) - \ln(\bar{r}_{kh}) \quad \forall k = 1, 2, 3 \quad (1.1)$$

where m is the birth month of child i in household h ;⁵ z_{kt} is an indicator variable equal to one if month t is included in cropping period k and zero otherwise; r_{ht} is the monthly rainfall (measured in millimeters) that occurred within a 5 kilometer radius of household h in month t ; and \bar{r}_{hj} is household h 's long-run yearly average rainfall for cropping period k over the 1981-2017 time period.

I illustrate R_{imh0}^k with a particular example of a child born in October 2014. The harvest period rainfall would be the sum of the rainfall experienced by the child's household in the months of December 2013 and January, June, and July 2014. This summation represents the total rainfall the child experienced while *in utero* during this cropping period. The *in utero* shock measure is then the natural log of the average harvest period rainfall for the child's household in all years from 1981-2017, subtracted from the natural log of the *in utero* summation. This same process is repeated for the land preparation and planting and mid-season periods.

The measures of *in utero* rainfall shocks vary considerably across cropping periods (see Table 1.1). During land preparation and planting, rainfall shocks are about 6% above historical averages. But there is less overall variability during the mid-season period, as average *in utero* rainfall shocks are about 1.6% below historical averages. During the harvest period, *in utero* rainfall shocks are 2.7% below historical averages.

IV. Empirical Strategy

My identification strategy is to differentiate any *in utero* rainfall effects from other spatial and temporal child growth determinants that may be correlated with rainfall outcomes. By including a broad set of controls in the econometric specification, I isolate the effects of *in utero* rainfall deviations on child growth from the average conditions of children in similar

⁵Each month in the rainfall dataset is assigned a running integer value (e.g. January 1981 = 1, February 1981 = 2, etc.).

circumstances and localities. The main empirical specification is given by the following:

$$HAZ_{imhvd} = \mathbf{R}_{imh0}\boldsymbol{\theta} + age_i\gamma + age_i * \mathbf{R}_{imh0}\boldsymbol{\delta} + \mathbf{X}_i\boldsymbol{\beta} + \omega_i + \sigma_m + c_v + \tau_d t + u_{imhvd} \quad (1.2)$$

The outcome variable is HAZ_{imhvd} , the height-for-age z-score (HAZ) for child i born in month m in household h in village v and district d . The main coefficients of interest, $\boldsymbol{\theta}$, are the coefficients on \mathbf{R}_{imh0} - the log-deviations of *in utero* rainfall for the three cropping periods (based on each child's birthday and household GPS location). I interact \mathbf{R}_{imh0} with child age in months, to account for the fact that rainfall shocks during the *in utero* period may impact more recently born children differently than older children. All regressions also include: two child-specific controls, gender and age squared (X_i); and a survey round indicator (ω_i).

Birth month may influence a child's development path and would be correlated with season-specific *in utero* rainfall. For instance, children born during lean months are likely to have different child growth outcomes compared to children born in post-harvest months. To control for such differences, I include birth month fixed effects, σ_m . The distribution of birth months is relatively even throughout the year, with slight jumps in January, May, and December (see Figure 1.3 for a histogram of children's birth months).

Village location and other time constant village-level unobservables may influence child growth and be correlated with rainfall deviations. For example, certain villages may be located in more mountainous terrain with little access to healthcare. To address this potential issue, I include village fixed effects, c_v .

The economies of Rwanda's districts develop at different rates, and the average growth path of children will vary according to local economic development (Maccini and Yang 2009). To control for differences in economic development paths across districts, I include a time trend t which is allowed to vary by the district-specific coefficient, τ_d .

To take account of village-level spatial correlation in the idiosyncratic error term, u_{imhvd} ,

I cluster standard errors at the village level.

A remaining concern with my identification strategy is due to potential serial correlation in log deviations in rainfall. If rainfall outcomes are serially correlated, then the estimated marginal effects could be driven by the effect of rainfall outcomes in other periods (Maccini & Yang, 2009; Rocha & Soares, 2015). I rule out this possibility, as the estimates of *in utero* period effects are robust to including rainfall shocks prior to a child’s *in utero* period.

Finally, I obtain similar estimates when I control for unobserved household-level heterogeneity in child growth outcomes, such as family heritage and the household environment. To do so, I estimate a robustness regression with household fixed effects. As this specification is restricted to the subsample of households with anthropometric measurements for at least two children under age five, the sample size is too small to obtain precise estimates.

V. Empirical Results

V.1 Main Findings

In Table 1.2, I present results from estimating equation 1.2, where I use only one annual rainfall measure. These estimates are based on the assumptions that the *in utero* rainfall effect on HAZ is homogenous throughout the year. In column (1), I control for demographic characteristics only. In column (2), I add birth month fixed effects to control for general temporal differences in child growth outcomes that may be correlated with rainfall. In column (3), I add village fixed effects to control for spatial differences in birth outcomes that could be correlated with rainfall. In column (4), I add further controls for unobserved district-specific time trends that could influence children’s development paths and be correlated with *in utero* rainfall outcomes.

As estimates are not statistically significantly different from zero, I cannot reject the null hypothesis that the impact is in fact zero. The point estimates suggest that increases in *in utero* rainfall may raise HAZ. Estimates are consistent across these four specifications, all

with wide confidence intervals.

The wide confidence interval in estimates could be due to one of two reasons. First, the relatively high fluctuation in annual rainfall may lead to high standard errors. Second, this aggregate, annual measure of rainfall may be suppressing intra-year, seasonal heterogeneity in how rainfall impacts yield at various times of the year, and therefore also children’s growth outcomes.

Estimates of equation 1.2, using cropping-period specific rainfall deviations rather than a single annual deviation, suggest that the annual measure attenuates impact estimates due to seasonal differences in impacts. Increases in *in utero* rainfall in the land preparation and planting or mid-season periods significantly raise HAZ. In contrast, increases in *in utero* rainfall during the harvest period significantly lowers HAZ. These results are summarized in Table 1.3, where covariates in each specification mirror those in Table 1.2. Estimates are consistent and statistically significant across all specifications. Results in column (4) of Table 1.3 imply that a one standard deviation increase in *in utero* rainfall during the land preparation and planting or mid-season periods raises the HAZ of a one year old child by 0.28 or 0.26 standard deviations, respectively.⁶ These effects represent 18% and 16% increases at the sample mean HAZ of -1.57. In the harvest period, a one standard deviation increase in *in utero* rainfall reduces a one year old child’s expected HAZ by 0.24 standard deviations, a decrease of 15% from the sample mean.

The findings also suggest that households may be able to compensate for some of the growth effects of *in utero* rainfall. *In utero* rainfall effects are largest for infants, declining as children age. These differences are statistically significant. This declining impact over

⁶The calculation, using land preparation and planting as an example, is as follows:

$$0.143 * (2.560 - 0.049 * 12) \approx 0.28$$

where 0.143 is a one standard deviation or 14.3% increase in *in utero* land preparation and planting period rainfall, 2.560 is the coefficient on the log deviation in *in utero* land preparation and planting rainfall, -0.049 is the coefficient on the interaction between the rainfall term and child age in months, and 12 is the age in months of a one year old child.

ages is distinct from the general decline in HAZ with age that is typical of the developing country context (Akresh et al., 2011; Groppo & Kraehnert, 2016).

I find further heterogeneity in rainfall impacts on HAZ, with estimated effects being concentrated among girls. As the sample is evenly split between boys and girls, I estimate separate regressions for each of them (see columns (1) and (2) in Table 1.4). There are several potential drivers for this concentration of *in utero* rainfall effects among girls. For instance, this difference could be a result of weaker girls being more likely to survive pregnancy and early infancy than boys. Boys have a lower *in utero* survival rate than girls (Bruckner and Catalano 2018). Boy infants also have the highest death rate of any age-sex group; only the strongest boys survive and I do not observe the boys who die *in utero*. Alternatively, the observed difference could be due to parents' behaviors in response to rainfall shocks. Gender-preferences in parenting could result in the sampled boys being less susceptible to long-term growth effects from *in utero* rainfall shocks.

Another potential source of heterogeneity in the rainfall impacts on HAZ is households' degrees of exposure to rainfall shocks. Some households are likely to be better positioned to react to changes in rainfall than others. As the dataset was collected after each child's *in utero* period, I cannot directly address this potential heterogeneity in this study and it remains a promising area for future work. I do, however, have data on household wellbeing at the time of the anthropometric measurements which can serve as an imperfect proxy of household wellbeing while the child was *in utero*. I rank the likelihood of each sampled household being below Rwanda's national poverty line at the time of the child's anthropometric measurements using a Rwanda-specific poverty index developed in Schreiner (2010). This poverty index uses household characteristics and assets to predict the likelihood of a given household falling below Rwanda's national poverty line. Column (3) of Table 1.4 summarizes estimates of equation 1.2 for children in relatively poor households (bottom 50th percentile of poverty index). In Column (4), I present the same estimates for relatively wealthy house-

holds (top 50th percentile). Overall, this subsample analysis by current household poverty status suggests that even the relatively wealthy households in the rural sample of agricultural households are unable to substantially mitigate the effects of rainfall shocks on their children’s growth outcomes.

V.2 Robustness Checks

The overall findings are robust to a number of different specifications. First, results are nearly identical when I add birth order and sibling controls, as shown in column (5) of Table 1.4. In column (6) of Table 1.4, I estimate a robustness regression with household fixed effects in lieu of village fixed effects. This specification is restricted to the subsample of households with anthropometric measurements for at least two children under age five. For comparison, in column (7) I also estimate the village fixed effects specification from Table 1.3 column (4) on this subsample of multiple child households. Unlike village fixed effects, household fixed effects controls for unobserved household-level heterogeneity in child growth outcomes, such as family heritage and the household environment. The results are qualitatively similar, with increases in early and mid-season rainfall associated with increases in expected HAZ and increases in harvest period rainfall associated with decrease in expected HAZ.

The main findings are robust to accounting for Rwanda’s minor agricultural season, which I have ignored thus far. Figure 1.1 shows that a minor agricultural season occurs simultaneously with that of the two major seasons. The minor harvest occurs just before the onset of the lean season in October and November. Approximately 70% of the households in the sample farm in this season. To test whether the findings may be sensitive to the incorporation of this minor season, I first estimate regressions for the subsample of households who farmed in the minor season. The results, presented in column (2) of Table 1.5, show nearly identical results to those found with the broad sample. I further assess the sensitivity of the results to the season and cropping period specifications via two alternative season

definitions. First, I incorporate the minor season by assigning the minor season months of July, August, and September to multiple cropping periods (see Figure 1.1). For example, *in utero* September rainfall is included simultaneously in both the harvest period (to reflect minor season activity), and the land preparation and planting period (to reflect second major season activity). Alternatively, as shown in Figure 1.2, I restrict the first major season harvest and second major season land preparation and planting to June and September respectively. The results of these alternative definitions are presented in columns (4) and (5) of Table 1.5. The findings are robust across all definitions of cropping periods.

Finally, results are not driven by the possibility of serial correlation in rainfall outcomes. I test for this possible serial correlation by controlling for the rainfall measures and their interactions for the 13-24 month period before a child’s birth. As shown in column (6) of Table 1.5, when controlling for the same rainfall shocks in another year, all season-specific rainfall effects remain very similar to original estimates. In addition, coefficient estimates are not statistically significant or meaningful in magnitude for the season-specific rainfall shocks during the 13-24 month period prior to birth, or their interactions with child age.

VI. Conclusion

In this article, I have shown that aggregate rainfall measures often ignore intra-seasonal heterogeneity in rainfall effects and attenuate estimated impacts of *in utero* rainfall shocks on children’s health. My analysis contributes to prior studies analyzing the relationship between rainfall and child health, which have primarily used annual or growing season rainfall measures. I identify intra-seasonal heterogeneity in rainfall effects by using high resolution rainfall data and a unique dataset of at-risk children in rural Rwanda. I have shown that intra-year or intra-season rainfall effects differ in sign.

The results are consistent with agronomic research finding that low yield can be due to both insufficient rainfall in the mid-season flowering and yield formation periods, as well as

excess rainfall during the harvest period (Barron et al., 2003; Doorenbos & Kassam, 1979). I find that increases in *in utero* rainfall in the land preparation and planting or mid-season periods increases expected child height-for-age z-scores. But during the harvest period, increases in *in utero* rainfall decreases expected child HAZ. I have shown these effects to be robust to alternative definitions of the cropping periods, as well as to estimates where I have controlled for potential serial correlation in rainfall. The findings also suggest that households may be compensating for some of the growth effects of *in utero* rainfall.

These estimates provide a lower bound, because of potential attenuation bias due to classical measurement error (Wooldridge, 2010). The measurement of *in utero* rainfall approximates the actual rainfall experienced by a child when *in utero*. In addition, to ensure that cropping period effects are not endogenous to household's planting decisions, I have assumed cropping periods to be constant across space and time. In fact, they vary with household cropping decisions and with the onset of rains.

The relationship between weather shocks and individual health is an issue that is becoming increasingly salient as extreme weather events - particularly drought and flooding - become more frequent over time. Understanding the complexity of this weather-health relationship can help inform policies and programs which aim to identify and distribute resources towards at risk pregnant women and their children. I have shown that identifying or predicting the timing of rainfall such events in the context of local agricultural production seasons is important in anticipating negative consequences for children's long-term health.

APPENDIX

Table 1.1: Summary Statistics

Variables	Mean	Std. Dev.	Min	Max	Median
Height-for-Age Z-Score (HAZ)	-1.572	1.598	-5.980	5.610	-1.640
Stunted (1=yes)	0.390	0.488	0	1	
<i>In utero</i> annual rainfall	0.017	0.106	-0.277	0.374	0.004
<i>In utero</i> land preparation and planting period rainfall	0.060	0.143	-0.346	0.435	0.047
<i>In utero</i> mid-season rainfall	-0.016	0.171	-0.395	0.531	-0.025
<i>In utero</i> harvest period rainfall	-0.027	0.237	-0.738	0.473	0.010
Child's age in months	33.692	15.550	0	59	36
Child is female (1=yes)	0.499	0.500	0	1	
First born (1=yes)	0.248	0.432	0	1	
Second born (1=yes)	0.277	0.448	0	1	
Number of siblings	2.022	1.646	0	9	2
Poverty likelihood	47.6	19.6	0	100	51.8
Survey round indicator (1=data from 2017)	0.551	0.497	0	1	

Note: 3,093 child observations. Stunting is defined as a HAZ below -2. Rainfall variables are the natural log deviations of rainfall for a given cropping period in the 12 months before birth from the historical average for that period. Cropping period definitions are defined as in Figure 1a. Birth order and number of siblings are based on siblings currently living in the household; data on these two variables are missing for 47 children. Poverty likelihood is the approximate percent chance that a child's household is below the Rwanda national poverty line based on the Rwanda poverty index defined in Schreiner (2010).

Table 1.2: Annual In Utero Rainfall Effects on Child Growth (Height-for-Age Z-Score)

Variables	(1)	(2)	(3)	(4)
<i>In utero</i> annual rainfall	1.220 (1.004)	0.997 (1.034)	1.038 (1.098)	1.546 (1.168)
<i>In utero</i> annual rainfall X child age in months	−0.010 (0.026)	−0.003 (0.026)	−0.006 (0.028)	−0.016 (0.029)
Child is female (1=yes)	0.185*** (0.056)	0.185*** (0.056)	0.181*** (0.058)	0.180*** (0.058)
Child's age in months	−0.083*** (0.009)	−0.083*** (0.009)	−0.088*** (0.010)	−0.136*** (0.041)
Child's age in months squared	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)
Birth month fixed effects?	No	Yes	Yes	Yes
Village fixed effects?	No	No	Yes	Yes
District-specific time trends?	No	No	No	Yes
Observations	3,093	3,093	3,093	3,093
R-squared	0.070	0.077	0.081	0.086
Number of villages	251	251	251	251

Note: Robust standard errors clustered at the village level in parentheses. The dependent variable is HAZ. Rainfall variables are the natural log deviations of rainfall 12 months before birth from the historical annual average. All specifications include a survey round indicator and an overall constant. *** $p < 0.01$.

Table 1.3: Cropping Period-Specific In Utero Rainfall Effects
on Child Growth (Height-for-Age Z-Score)

Variables	(1)	(2)	(3)	(4)
<i>In utero</i> land preparation and planting period rainfall	1.234* (0.655)	1.405** (0.671)	1.473** (0.696)	2.560*** (0.818)
<i>In utero</i> mid-season rainfall	1.597*** (0.603)	1.478** (0.616)	1.611** (0.660)	2.033*** (0.690)
<i>In utero</i> harvest period rainfall	-1.099** (0.444)	-1.191*** (0.445)	-1.322*** (0.481)	-1.556*** (0.500)
<i>In utero</i> land prep. and planting pd. rainfall X child age in months	-0.021 (0.017)	-0.024 (0.017)	-0.026 (0.017)	-0.049** (0.020)
<i>In utero</i> mid-season rainfall X child age in months	-0.035** (0.015)	-0.031** (0.015)	-0.036** (0.017)	-0.044*** (0.017)
<i>In utero</i> harvest period rainfall X child age in months	0.035*** (0.012)	0.037*** (0.012)	0.040*** (0.013)	0.047*** (0.013)
Child is female (1=yes)	0.185*** (0.055)	0.186*** (0.056)	0.185*** (0.057)	0.185*** (0.058)
Child's age in months	-0.072*** (0.010)	-0.072*** (0.010)	-0.075*** (0.011)	-0.128*** (0.041)
Child's age in months squared	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)
Birth month fixed effects?	No	Yes	Yes	Yes
Village fixed effects?	No	No	Yes	Yes
District-specific time trends?	No	No	No	Yes
Observations	3,093	3,093	3,093	3,093
R-squared	0.076	0.083	0.088	0.096
Number of villages	251	251	251	251

Note: Robust standard errors clustered at the village level in parentheses. The dependent variable is HAZ. Rainfall variables are the natural log deviations of rainfall for a given cropping period in the 12 months before birth from the historical average for that period. All specifications include a survey round indicator and an overall constant. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 1.4: In Utero Rainfall Effects on Child Growth (Height-for-Age Z-Score)
by Demographic and Household Characteristics

Variables	Female (1)	Male (2)	Relatively Poor (3)	Relatively Wealthy (4)	Sibling Controls (5)	Multiple Child Subsample	
						Household Fixed Effects (6)	Village Fixed Effects (7)
<i>In utero</i> land preparation and planting period rainfall	3.194*** (1.080)	1.751 (1.219)	3.067*** (1.093)	2.080* (1.162)	2.490*** (0.811)	1.505 (1.151)	2.092** (0.997)
<i>In utero</i> mid-season rainfall	1.848* (0.977)	2.307** (1.046)	0.901 (1.007)	2.680*** (0.884)	2.264*** (0.690)	0.477 (1.053)	1.405 (0.933)
<i>In utero</i> harvest period rainfall	-2.070*** (0.658)	-0.984 (0.760)	-1.403** (0.655)	-1.605** (0.647)	-1.493*** (0.500)	-1.792*** (0.630)	-1.676*** (0.586)
<i>In utero</i> land preparation and planting period rainfall X child age in months	-0.056** (0.026)	-0.038 (0.029)	-0.050* (0.026)	-0.047* (0.028)	-0.049** (0.019)	-0.020 (0.027)	-0.028 (0.024)
<i>In utero</i> mid-season rainfall X child age in months	-0.038 (0.025)	-0.056** (0.026)	-0.012 (0.024)	-0.064*** (0.022)	-0.051*** (0.017)	-0.005 (0.027)	-0.031 (0.023)
<i>In utero</i> harvest period rainfall X child age in months	0.056*** (0.018)	0.032 (0.020)	0.048*** (0.017)	0.047*** (0.018)	0.045*** (0.013)	0.052*** (0.018)	0.057*** (0.016)
Observations	1,539	1,546	1,570	1,503	3,046	1,943	1,943
R-squared	0.119	0.084	0.132	0.105	0.100	0.162	0.137
Number of villages	245	246	243	237	251	242	242

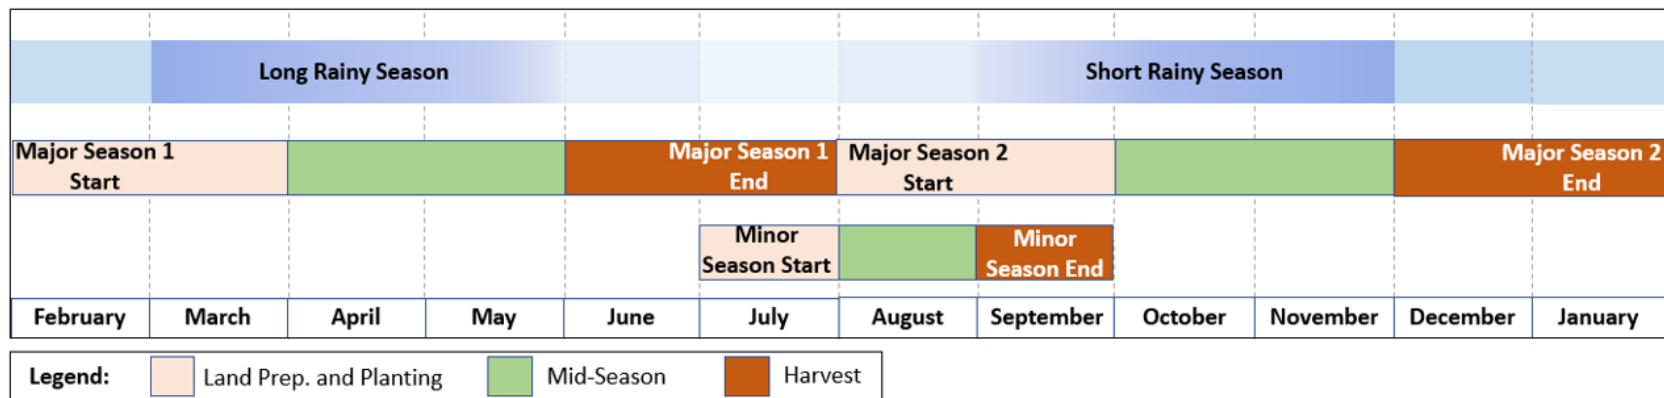
Note: Robust standard errors clustered at the village level in parentheses. Dependent variable is HAZ. Rainfall variables are the natural log deviations of rainfall for a given cropping period in the 12 months before birth from the historical average for that period. Cols. (1) and (2) split the sample by the child's sex. Cols. (3) and (4) split the sample by households in the bottom and top 50th percentile, respectively, using the Rwanda poverty index defined in Schreiner (2010). Col. (5) includes controls for first born, second born, and number of siblings. Cols. (6) and (7) are restricted to the subsample of 883 households with two or more child observations. All specifications include age in months, age in months squared, birth month fixed effects, district specific time trends, a survey round indicator, and an overall constant. All cols. except (1) and (2) include a female indicator. All cols. except (6) include village fixed effects. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 1.5: In Utero Rainfall Effects on Height-for-Age Z-Score:
Incorporating Minor Season and Rainfall in Other Periods

Variables	Main (1)	Minor Season Prod. Subsample (2)	Season Definitions Include Minor with: Overlap (3)	Limited Overlap (4)	Controlling for Rainfall 12-24 Months Before Birth (5)
<i>In utero</i> land preparation and planting period rainfall	2.560*** (0.818)	2.794*** (1.020)	3.056*** (0.960)	2.433*** (0.935)	1.808* (0.963)
<i>In utero</i> mid-season rainfall	2.033*** (0.690)	1.901** (0.839)	2.376*** (0.685)	2.507*** (0.715)	1.900** (0.744)
<i>In utero</i> harvest period rainfall	-1.556*** (0.500)	-1.931*** (0.577)	-2.329*** (0.772)	-2.107*** (0.740)	-1.122* (0.586)
<i>In utero</i> land preparation and planting period rainfall X child age in months	-0.049** (0.020)	-0.050** (0.023)	-0.062*** (0.022)	-0.053** (0.021)	-0.031 (0.023)
<i>In utero</i> mid-season rainfall X child age in months	-0.044*** (0.017)	-0.041** (0.020)	-0.052*** (0.017)	-0.056*** (0.018)	-0.041** (0.018)
<i>In utero</i> harvest period rainfall X child age in months	0.047*** (0.013)	0.057*** (0.015)	0.069*** (0.018)	0.064*** (0.018)	0.035** (0.015)
Observations	3,093	2,150	3,093	3,093	3,093
R-squared	0.096	0.105	0.099	0.097	0.098
Number of villages	251	249	251	251	251

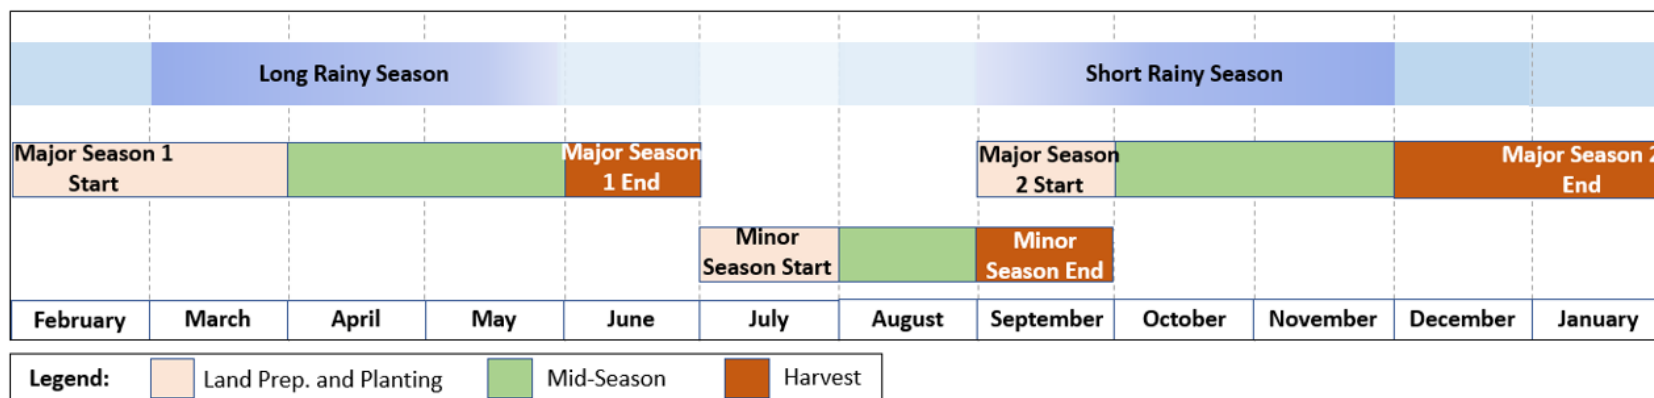
Note: Robust standard errors clustered at the village level in parentheses. Dependent variable is HAZ. Rainfall variables are the natural log deviations of rainfall for a given cropping period in the 12 months before birth from the historical average for that period. Col. (1) repeats the preferred specification from Table 3 Col. (4). Col. (2) restricts the sample to children in households that farmed in the minor season. Col. (3) assigns *in utero* rainfall in the minor season months of July, August, and September to multiple cropping periods as shown in Figure 1a. Col. (4) assigns *in utero* rainfall in September to multiple cropping periods as shown in Figure 1b. Col. (5) controls for the cropping period rainfall variables for the period 13-24 months before birth (none of which are significant). All specifications include a female indicator, age in months, age in months squared, birth month fixed effects, village fixed effects, district specific time trends, a survey round indicator, and an overall constant. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Figure 1.1: Rwanda's Cropping Period Calendar Around Two Rainy Seasons – Includes Minor Season




Note: Figure adapted from Famine Early Warning Systems Network (2017).

Figure 1.2: Rwanda's Cropping Period Calendar – Includes Minor Season and Restricts Major Seasons



Note: Figure adapted from Famine Early Warning Systems Network (2017).

Figure 1.3: Histogram of Child Birth Months



REFERENCES

REFERENCES

- Abiona, O. (2017). Adverse Effects of Early Life Extreme Precipitation Shocks on Short-term Health and Adulthood Welfare Outcomes. *Review of Development Economics*, 21(4), 1229–1254.
- Aguilar, A., & Vicarelli, M. (2011). *El Nino and Mexican Children: Medium-Term Effects of Early-Life Weather Shocks on Cognitive and Health Outcomes* (Working Paper).
- Akresh, R., Verwimp, P., & Bundervoet, T. (2011). Civil War, Crop Failure, and Child Stunting in Rwanda. *Economic Development and Cultural Change*, 59(4), 777–810.
- Almond, D., Currie, J., & Duque, V. (2017, January). *Childhood Circumstances and Adult Outcomes: Act II* (NBER Working Paper No. 23017). National Bureau of Economic Research. Cambridge, MA.
- Barker, D. J. P. (1998). In Utero Programming of Chronic Disease. *Clinical Science*, 95(2), 115–128.
- Barron, J., Rockström, J., Gichuki, F., & Hatibu, N. (2003). Dry Spell Analysis and Maize Yields for Two Semi-Arid Locations in East Africa. *Agricultural and Forest Meteorology*, 117(1-2), 23–37.
- Brown, M. E. (2008). *Famine Early Warning Systems and Remote Sensing Data*. Berlin, Springer.
- Burgess, R., Deschenes, O., Donaldson, D., & Greenstone, M. (2017, April 20). *Weather, Climate Change, and Death in India*.
- Comfort, A. B. (2016). Long-Term Effect of In Utero Conditions on Maternal Survival Later in Life: Evidence from Sub-Saharan Africa. *Journal of Population Economics*, 29(2), 493–527.
- Cornwell, K., & Inder, B. (2015). Child Health and Rainfall in Early Life. *The Journal of Development Studies*, 51(7), 865–880.
- Doorenbos, J., & Kassam, A. (1979). *Yield Response to Water* (Agriculture Organization of the United Nations (FAO) Irrigation and Drainage Paper No. 33). Rome, Italy.
- Famine Early Warning Systems Network. (2018). *Building Rainfall Assumptions for Scenario Development: Guidance Document Number 2*. Washington, D.C.

- Funk, C., Hoell, A., Nicholson, S., Korecha, D., Galu, G., Artan, G., Teshome, F., Hailer-mariam, K., Segele, Z., Harrison, L., Tadege, A., Atheru, Z., Pomposi, C., & Pedreros, D. (2019). Examining the Potential Contributions of Extreme “Western V” Sea Surface Temperatures to the March-June East African Drought. *Bulletin of the American Meteorological Society*, 100(1), S55–S60.
- Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., Husak, G., Rowland, J., Harrison, L., Hoell, A., & Michaelsen, J. (2015). The Climate Hazards Infrared Precipitation with Stations—A New Environmental Record for Monitoring Extremes. *Scientific Data*, 2, 150066.
- Grosso, V., & Kraehnert, K. (2016). Extreme Weather Events and Child Height: Evidence from Mongolia. *World Development*, 86, 59–78.
- HarvestChoice. (2010). *Rainfall Variability and Crop Yield Potential*. International Food Policy Research Institute. Washington, D.C. and University of Minnesota, St. Paul, MN. Retrieved January 16, 2019, from <http://harvestchoice.org/node/2240>
- Hoddinott, J., & Kinsey, B. (2001). Child Growth in the Time of Drought. *Oxford Bulletin of Economics and Statistics*, 63(4), 409–436.
- Kim, Y. S. (2010). *The Impact of Rainfall on Early Child Health* (Working Paper).
- Kudamatsu, M., Persson, T., & Stromberg, D. (2012). *Weather and Infant Mortality in Africa* (Working Paper).
- Lehmann, J., Mempel, F., & Coumou, D. (2018). Increased Occurrence of Record-Wet and Record-Dry Months Reflect Changes in Mean Rainfall. *Geophysical Research Letters*, 45(24), 13, 468–13, 476.
- Maccini, S., & Yang, D. (2009). Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall. *American Economic Review*, 99(3), 1006–1026.
- Miller, R. (2017). Childhood Health and Prenatal Exposure to Seasonal Food Scarcity in Ethiopia. *World Development*, 99, 350–376.
- Peters, C., Farris, J., Porter, M., Maredia, M. K., & Jin, S. (2015). *Impact Evaluation of Scaling up Sweet Potato Through Agriculture and Nutrition (SUSTAIN) Project in Rwanda: Baseline Report* (Unpublished research report to SUSTAIN project, International Potato Center, Nairobi, Kenya).
- Rocha, R., & Soares, R. R. (2015). Water scarcity and birth outcomes in the Brazilian semiarid. *Journal of Development Economics*, 112, 72–91.

- Schreiner, M. (2010). Simple Poverty Scorecard Poverty-Assessment Tool Rwanda, 118. http://www.simplepovertyscorecard.com/RWA_2005_ENG.pdf
- Shively, G. E. (2017). Infrastructure mitigates the sensitivity of child growth to local agriculture and rainfall in Nepal and Uganda. *Proceedings of the National Academy of Sciences*, 114(5), 903–908.
- Skoufias, E., Vinha, K., & Conroy, H. V. (2011, February). *The impacts of climate variability on welfare in rural Mexico* (World Bank Working Paper No. 5555). The World Bank.
- Steduto, P., Hsiao, T., Fereres, E., & Raes, D. (Eds.). (2012). *Crop Yield Response to Water*. Rome, Italy, Food and Agriculture Organization of the United Nations.
- Strauss, J., & Thomas, D. (1998). Health, Nutrition, and Economic Development. *Journal of Economic Literature*, 36(2), 766–817.
- Tiwari, S., Jacoby, H. G., & Skoufias, E. (2017). Monsoon Babies: Rainfall Shocks and Child Nutrition in Nepal. *Economic Development and Cultural Change*, 65(2), 167–188. <https://doi.org/10.1086/689308>
- Wooldridge, J. (2010). *Econometric Analysis of Cross Section and Panel Data* (2nd ed.). Cambridge, MA, MIT Press.
- World Health Organization. (1997). *WHO Global Database on Child Growth and Malnutrition [Database]*. World Health Organization. Geneva, Switzerland.

Chapter 2

Does Unobserved Land Quality Bias Separability Tests?

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 agricultural households make their production decisions separately from their consumption characteristics and preferences. The existence of separability affects households' production responses to new opportunities and shocks and provides an indication of the completeness of markets (Benjamin, 1992; Singh et al., 1986).

Numerous studies implicitly or explicitly assume separability of agricultural production decisions (e.g. Conley & Udry, 2010; Foster & Rosenzweig, 1995; Sheahan et al., 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 member consumption preferences. 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 households' 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.

¹This 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 et al., 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 an active land market as this 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 future research to incorporate more robust means of controlling for unobserved land quality, such as plot panel data which enables 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 drive biased inference 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 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:

$$\max E \left[\sum_{t=1}^{\infty} \beta^{t-1} U(\mathbf{x}_{\mathbf{mt}}, \mathbf{x}_{\mathbf{at}}, x_{lt}; \mathbf{D}_{\mathbf{t}}, \mu_t) \right] \quad (2.1)$$

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, $\mathbf{x}_{\mathbf{mt}}$, a vector of agricultural goods, $\mathbf{x}_{\mathbf{at}}$, and leisure, x_{lt} . The utility derived from these goods differs according to household consumption preferences observed by the analyst (e.g. demographic characteristics), $\mathbf{D}_{\mathbf{t}}$, 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:

$$\mathbf{p}_{\mathbf{mt}} \cdot \mathbf{x}_{\mathbf{mt}} + \mathbf{p}_{\mathbf{at}} \cdot \mathbf{x}_{\mathbf{at}} + w_t x_{lt} + \frac{1}{1 + \tau_t} W_{t+1} = w_t E_t + \pi_t + W_t \quad (2.2)$$

where the prices of the market goods, agricultural goods, and leisure are $\mathbf{p}_{\mathbf{mt}}$, $\mathbf{p}_{\mathbf{at}}$, 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:

$$\pi_t = \sum_{i=1}^{N_t} p_{qt} f(L_t^i, \tilde{A}^i, \mathbf{Z}_{\mathbf{t}}^i; v_t) - w_t L_t^i - r_t \tilde{A}_t^i - \mathbf{p}_{\mathbf{zt}} \cdot \mathbf{Z}_{\mathbf{t}}^i \quad (2.3)$$

where N_t is the number of plots the household farms 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, \tilde{A}^i , a vector of other inputs, $\mathbf{Z}_{\mathbf{t}}^i$, 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, \tilde{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:

$$\tilde{A}^i = \theta^i A^i \quad \forall i \in \{1, 2, \dots, N_t\} \quad (2.4)$$

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

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:

$$p_{qt} \frac{\partial f(L_t^i, \tilde{A}^i, \mathbf{Z}_{\mathbf{t}}^i; v_t)}{\partial L_t^i} = w_t \quad (2.5)$$

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 2.5 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, $\mathbf{D}_{\mathbf{t}}$, 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.

²For 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:

$$\ln L_{cht} = \kappa + \delta \mathbf{D}_{\mathbf{cht}} + \beta \mathbf{X}_{\mathbf{cht}} + \eta_{ct} + \eta_h + \epsilon_{cht} \quad (2.6)$$

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, $\mathbf{D}_{\mathbf{cht}}$ is a vector of household demographic characteristics, $\mathbf{X}_{\mathbf{cht}}$ is a vector of other observed characteristics which affect labor demand and are potentially correlated with $\mathbf{D}_{\mathbf{cht}}$, 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 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 arbitrarily correlated 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 2.6, however, does not control for plot-level unobservables which may be correlated with household demographics within a particular community. For

³ $\mathbf{X}_{\mathbf{cht}}$ 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 adequately control 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 2.6, the idiosyncratic error can be approximated as follows:

$$\epsilon_{cht} = \sum_{i=1}^{N_t} \eta_i + v_{cht} \quad (2.7)$$

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 2.6 will control for unobserved plot characteristics.

If some portion of the households in the sample alter their 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

$$\widetilde{\ln L_{chit}} = \kappa + \delta \mathbf{D}_{\text{cht}} + \beta \mathbf{X}_{\text{cht}} + \eta_{ct} + \eta_i + \epsilon_{chit} \quad (2.8)$$

⁴This 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 of their farmed plots over the study time period, 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 2.8; 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 Rwanda 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 pregnant woman

⁵Separable 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.

⁶Let $\mathbf{s}_i = [s_{i1}, s_{i2}, \dots, s_{iT}]$ 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 2.8 as $\mathbf{Z}_{it} = [\mathbf{D}_{\text{cht}}, \mathbf{X}_{\text{cht}}, \eta_{ct}]$ and $\mathbf{Z}_i = [\mathbf{Z}_{i1}, \mathbf{Z}_{i2}, \dots, \mathbf{Z}_{iT}]$. 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 child under five. The survey collected detailed plot-level agricultural information as well as household demographic information.

This analysis focuses on total plot-specific labor demand during the major February to June agricultural season. 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 as harvest labor requirements are typically proportional to yield 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, approximately 51% of plots observed in 2017 were successfully linked back to 2014 plot observations.⁷

Table 2.1 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

⁷Respondent 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 2.2 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 2.6 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 2.3 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 robust to taking into account potential productivity differences in child household members (Appendix Table 2.7 columns 1 and 2), controlling for a land quality proxy (Appendix Table 2.8 columns 3 and 4), and to including the households without a plot in the plot panel sample (Appendix Table 2.9).

Separability tests based on the plot fixed effects specification in equation 2.8 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 2.3 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 taking into account potential productivity differences in child household members (Appendix Table 2.7 columns 3 and 4). The findings suggest that this sample of small-scale, agricultural households integrate 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 2.4, 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 2.10). 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 a child birth between

⁸LaFave 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 pregnant woman or child under five.

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

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

periods by the following process:

$$Y_{chit} = \exp(\kappa + \eta_c + \eta_h + \eta_i + \epsilon_{chit}) \quad (2.9)$$

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

Data is simulated for 250 communities and 25,000 households, 100 households per community, using this data generating process. 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:

$$\eta_i = aX_{c,h,t=1} + bZ_{chi} \quad (2.10)$$

where $X_{c,h,t=1} = \tilde{X}_{c,h,t=1} + 1$, $\tilde{X}_{c,h,t=1} \sim \text{Poisson}(4)$, 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:

$$a = \frac{\text{Cov}(X_{c,h,t=1}, \eta_i)}{\text{Var}(X_{c,h,t=1})} = \frac{\text{Cov}(X_{c,h,t=1}, \eta_i)}{4} \quad (2.11)$$

¹⁰This 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:

$$b = \frac{Cov(Z_{chi}, \eta_i)}{Var(Z_{chi})} = \frac{Cov(Z_{chi}, \eta_i)}{4} \quad (2.12)$$

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

$$Corr(X_{c,h,t=1}, \eta_i) = \frac{Cov(X_{c,h,t=1}, \eta_i)}{\sigma_X \sigma_\eta} = \frac{Cov(X_{c,h,t=1}, \eta_i)}{4\sqrt{a^2 + b^2}} \quad (2.13)$$

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 of its component parts as shown in equation 2.9.

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 $\delta_h = \tilde{\delta}_h + 1$ where $\tilde{\delta}_h \sim Poisson(2)$. 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 $N_{ht} = \tilde{N}_{ht} + 1$ where $\tilde{N}_{ht} \sim Poisson(2)$.

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 2.5, 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 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 2.6 reports the number of Type I errors under different levels of land market activity for each estimator over the 1,000 replications. As this is simulated data, 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,

¹²These 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 2.1 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 2.2); 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 households' production responses to new opportunities and shocks; its existence also provides an indication of 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 Rwanda 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 plot selection into the plot panel will lead 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 plot panel. The inclusion or exclusion of households not in the plot panel subsample, however, does not impact 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 households' operated plots. If the land market is volatile and few plots are maintained by the same household overtime, 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, this paper's 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 mis-characterization of households' 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.

APPENDICES

APPENDIX A:

Tables and Figures

Table 2.1: Household Summary Statistics

	Plot Panel Households	All Households
Household size	5.05 (1.78)	5.05 (1.78)
Share of female members 0 to 14 years old	0.24 (0.18)	0.24 (0.18)
Share of male members 0 to 14 years old	0.24 (0.18)	0.24 (0.18)
Share of female members 15 to 60 years old	0.26 (0.12)	0.26 (0.12)
Share of male members 15 to 60 years old	0.22 (0.13)	0.22 (0.14)
Share of female members 61 years or older	0.02 (0.09)	0.02 (0.09)
Share of male members 61 years or older	0.02 (0.06)	0.01 (0.06)
Farmed land area across all plots (ha)	1.08 (12.92)	1.27 (13.08)
Number of farmed plots	2.44 (1.40)	2.34 (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 2.2: Plot Summary Statistics

	Panel Subsample	All Plots
Plot size (ha)	0.22 (2.96)	0.44 (8.09)
Total labor demand (labor-days)	24.81 (36.22)	24.21 (35.29)
Family labor demand (labor-days)	21.30 (25.23)	20.76 (25.38)
Hired labor demand (labor-days)	3.50 (24.16)	3.45 (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 2.3: 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** (0.111)	0.356*** (0.134)	-0.012 (0.112)	-0.021 (0.140)
Share of female members 0 to 14 years old		0.222 (0.687)		1.842*** (0.609)
Share of male members 0 to 14 years old		0.221 (0.675)		2.057*** (0.621)
Share of female members 15 to 60 years old		0.909 (0.650)		2.298*** (0.612)
Share of male members 15 to 60 years old		0.683 (0.659)		2.221*** (0.596)
Share of female members 61 years or older		0.940 (0.672)		2.096*** (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 2.4: Sum of Labor Demand Across Plot Panel Plots (Log of Person Days per Season) on Household Characteristics

	(1)	(2)
ln(Household size)	-0.052 (0.110)	-0.069 (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 2.5: 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 2.6: Number of Observed Type I Errors in 1,000 Replications

Estimator	Percentage of Simulated Households with Land Change							
	0%		16%		30%		37%	
	Significance Level		Significance Level		Significance Level		Significance Level	
	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 2.1: Empirical Distribution via Household Fixed Effects





Figure 2.2: Empirical Distribution via Plot Fixed Effects



APPENDIX B:

Supplementary Tables

Table 2.7: 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*** (0.123)	0.355*** (0.133)	0.055 (0.128)	-0.019 (0.139)
Share of female members 0 to 14 years old		0.442 (0.665)		1.825*** (0.582)
Share of male members 0 to 14 years old		0.442 (0.656)		2.041*** (0.592)
Share of female members 15 to 60 years old		0.900 (0.652)		2.295*** (0.613)
Share of male members 15 to 60 years old		0.679 (0.660)		2.219*** (0.596)
Share of female members 61 years or older		0.937 (0.671)		2.096*** (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 2.8: 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** (0.111)	0.356*** (0.134)	0.240** (0.111)	0.358*** (0.134)
Share of female members 0 to 14 years old		0.222 (0.687)		0.211 (0.686)
Share of male members 0 to 14 years old		0.221 (0.675)		0.213 (0.675)
Share of female members 15 to 60 years old		0.909 (0.650)		0.906 (0.650)
Share of male members 15 to 60 years old		0.683 (0.659)		0.664 (0.660)
Share of female members 61 years or older		0.940 (0.672)		0.991 (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 2.3. (3) and (4) control for household's reported high quality farmed land area. 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 2.9: 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** (0.111)	0.356*** (0.134)	0.190** (0.092)	0.332*** (0.112)
Share of female members 0 to 14 years old		0.222 (0.687)		0.063 (0.632)
Share of male members 0 to 14 years old		0.221 (0.675)		0.102 (0.600)
Share of female members 15 to 60 years old		0.909 (0.650)		0.875 (0.595)
Share of male members 15 to 60 years old		0.683 (0.659)		0.687 (0.576)
Share of female members 61 years or older		0.940 (0.672)		0.820 (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 2.3. 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 2.10: 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, July 30). *Agricultural labor markets and fertilizer demand: Intensification is not a single factor problem for non-separable households* (Working Paper).
- 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.

Chapter 3

Farmer Personality and Community-Based Extension Effectiveness in Tanzania

I. Introduction

Diffusion of new agricultural technologies is critical to agricultural transformation in developing countries. The development literature abounds with examples of new agricultural technologies that have low adoption rates despite experimental studies demonstrating their superior profitability relative to the prevailing technology. Low adoption of a new agricultural technology may be driven by farmers' rational assessments of its lack of suitability to their local context or cultural practices. For example, the relative benefits of a new agricultural technology may require complementary inputs that a farmer cannot access due to the lack of well-functioning credit or input markets (Foster & Rosenzweig, 2010; Jack, 2013). In other cases, however, a new agricultural technology may have low adoption despite being well-suited to the local context. Lack of take-up in these latter cases can be characterized as an asymmetric information problem whereby the technology 'buyer' (e.g. the farmer) is not aware of or does not trust the credibility of available information signals on the relative benefits of the new agricultural technology (Macho-Stadler & Pérez-Castrillo, 2001).

One means of providing a credible signal is extension programs which provide a crucial link between pure research and the general public. Extension programs are particularly important in the developing country context where agricultural innovations have the potential to lift millions out of poverty, but only if those innovations are known and adopted by farmers. Evidence on the effectiveness of extension programs has been mixed, particularly in Sub-Saharan Africa (Anderson & Feder, 2004; Krishnan & Patnam, 2014; Norton & Alwang, 2020). Lead-farmer extension, which relies on local farmers to act as community-based extensionists in their villages, provides one potential means of increasing the trust in and effectiveness of extension efforts (BenYishay & Mobarak, 2019). This form of community-based extension has been shown to be effective in the study area of Tanzania (Nakano et al., 2018). Even within a community-based extension program, however, there may be heterogeneity in a farmer's likelihood of receiving and trusting the extension 'signal.'

Farmer personality may be an important source of such heterogeneity in community-based extension effectiveness.¹ Extension, and particularly community-based extension, has historically relied on the ability of extensionists to establish and nurture relationships with farmers. The psychology literature provides ample evidence for the role of personality in such interpersonal outcomes (Brandstätter et al., 2018; Ozer & Benet-Martínez, 2006; Weidmann et al., 2017; Wilson et al., 2016). There is also increasing evidence for the role of personality in a variety of consumer, health, labor, and social outcomes (Bazzani et al., 2017; Lin et al., 2019; Ozer & Benet-Martínez, 2006; Soto, 2019). A recent meta-replication study of 78 published personality trait-outcome associations found a vast majority to be replicable (Soto, 2019).

Despite this, few studies have explored the potential role of personality in extension and technology adoption in the agricultural context (Ali et al., 2017). There is also a dearth of personality studies in developing country contexts (Van der Linden et al., 2018). Building on a randomized control trial (RCT), this study addresses these gaps by assessing whether farmer personality influences the effectiveness of extension approaches for promoting improved bean varieties in Tanzania. I motivate this analysis with a conceptual framework for recipient farmers' responses to extension activities under asymmetric uncertainty and assess the potential heterogeneity empirically using a unique dataset of the Big Six personality traits measured using the Midlife Development Inventory (MIDI) (Lachman & Weaver, 1997, 2005). My analysis provides insights on the types of farmers most likely to benefit from community-based extension initiatives as reflected in their adoption behavior and social interactions.

The remainder of this paper is organized as follows. First, I describe in more detail the extension program and RCT which encompass this study's contextual background and experimental design. Second, I develop a conceptual model of adoption under asymmetric information which informs my hypotheses. I describe the data, including the measures of

¹Personality traits can be broadly defined as "...relatively enduring patterns of thoughts, feelings, and behaviors" (Roberts, 2009, pg. 140)

personality, in the third section. The fourth and fifth sections detail the empirical model and results respectively. I conclude in the final section with a summary of key findings and implications.

II. Background and Experimental Design

My empirical analysis relies on a community-based extension program in the Southern Highlands of Tanzania developed by Farm Input Promotions Africa Ltd. (FIPS-Africa). FIPS-Africa is a non-governmental organization that aims to improve small farmers' welfare by increasing awareness of and access to improved agricultural inputs. As part of the extension program, one lead-farmer, which FIPS-Africa refers to as a Village-Based Agricultural Advisor (VBAA), is selected for each village in the program (Melkani and Mason 2018). These VBAAAs are selected by the village members themselves based on a variety of factors, including farming experience, communication skills, and willingness to coordinate with FIPS-Africa (Morgan, 2018). The VBAAAs are volunteers, but receive training from FIPS-Africa on good agronomic practices and small business development as well as support for village-based extension services. FIPS-Africa has traditionally focused on maize and provided funding for all VBAAAs to conduct a demonstration plot comparison of varieties as well as distribute small trial packs of improved varieties to farmers in the village for personal trials (Melkani & Mason, 2018).

FIPS-Africa, in conjunction with Agricultural Research Institute Uyole (ARI-Uyole), the International Center for Tropical Agriculture (CIAT), and Michigan State University, conducted an RCT in 2017 to assess the marginal value-added of these small trial pack distributions. As FIPS-Africa had focused on improved maize varieties in its past activities, the RCT was based around improved bean varieties. In particular, the RCT focused on Njano Uyole and Uyole 96, two high-yielding disease tolerant bean varieties developed by ARI-Uyole and CIAT in the Mbeya Region (Ibid).

While the full RCT encompassed all 230 active VBAAAs in Tanzania’s Southern Highlands, this study focuses on a subsample of 32 VBAAAs in the Mbeya and Mbozi Districts for which a detailed farmer level survey was conducted. Each VBAA in the subsample led a demonstration plot comparison in the 2017 Major Season (March 2017-July 2017) of Njano Uyole, Uyole 96, Uyole 03, and a prevailing local bean variety.² Each demonstration plot was divided into 16 equal-sized sub-plots (four for each variety). This allowed for a demonstration plot comparison of the four varieties when untreated, treated with a fungicide/insecticide (Apron Star), treated with fertilizer (YaraMila CEREAL), and treated with both the fungicide/insecticide and fertilizer. While each VBAA conducted a demonstration plot comparison of bean varieties, half were randomly assigned to provide free trial packs of Uyole 96 and Njano Uyole seed for farmers to conduct their own personal trials. The VBAAAs that conducted only the demonstration plot are not a “true control.” Rather, the RCT assesses the marginal value of VBAA’s providing trial packs along with a demonstration plot relative to a demonstration plot alone (Ibid).

To reduce the probability of pre-treatment differences between the demonstration plot only (demo-only) and demonstration plot plus trial pack (demo-trial) subsamples, the VBAAAs were stratified by district and, within each district, paired according to a Mahalanobis greedy pairwise matching index.³ One VBAA from each pair was randomly assigned to the demo-trial group. In addition to the common demonstration plot, the VBAAAs randomly selected to the demo-trial group received 150 trial packs of improved bean seed to give out to the farmers in their village at the start of the 2017 Major Season (March 2017-July 2017). VBAAAs followed FIPS-Africa’s usual practice of distributing trial packs, which is to give them to participants at the demonstration plot planting, keeping to one trial pack per household.⁴

²While all 32 villages in this subsample used Uyole 03 as the third improved bean variety in the demonstration plot, different improved varieties were used in its place in the other villages in the larger RCT. This analysis focuses on Njano Uyole and Uyole 96 as these were the two common improved varieties with the most widespread applicability (Melkani & Mason, 2018).

³The characteristics used are described in Melkani and Mason (2018).

⁴If any extra trial packs remained following the demonstration plot planting, distribution was up to the

Each trial pack contained 400 grams (g) of bean seed equally split between the local bean variety and one of the three improved bean varieties used on the demonstration plot. Of the 200g of the bean seed for a given variety contained in each pack, 100g was pre-treated with the insecticide/fungicide while 100g was left untreated. The VBAs instructed each of the farmers that received the trial pack on how to setup a demo plot on their own farm split into four sub-plots. Thus, the farmers that received trial packs had an opportunity to mimic a portion of the larger village-level demonstration plot on their own farms (Ibid).

The stratified, pairwise matching treatment randomization was designed to ensure balance on VBA characteristics. Melkani and Mason (2018) conducted a baseline survey of all 230 VBAs in early 2017 prior to the start of the RCT. Table 3.1 provides the balance test results across key characteristics for the subsample of VBAs in this study.⁵ The results indicate that VBAs in the demo-only and demo-trial groups are balanced across a variety of demographic characteristics as well as past FIPS-Africa activities.

This previously implemented community-based extension RCT generated experimental variation in the potential “signals” recipient farmers received about the new, improved bean varieties from their VBAs. I exploit this variability in this paper. In the next section I develop a conceptual model to illustrate the likely value of the added trial pack information and how it may be heterogeneous according to the personality traits of recipient farmers.

III. Conceptual Framework

In this section, I develop a conceptual model to explore the role of community-based extension activities and farmer personality traits in addressing the information asymmetries inherent in the introduction of a new seed variety. I show how community-based extension activities that prioritize information signals from varied contexts can better enable farmers to update their beliefs about a new variety. Furthermore, I show that farmer personality traits can

VBAs’ discretion, keeping to the one per household requirement.

⁵The full set of balance tests across all 230 VBAs is available in Melkani and Mason (2018).

influence the likelihood of receiving these information signals which has implications for the types of farmers likely to benefit from community-based extension activities.

Prior to its release, a new seed variety k undergoes rigorous performance evaluation under researcher or farmer-managed trials conducted at experiment stations, central experimental plots in communities, or to a limited extent on farmers' fields. It is released as a new and improved variety if, on average across these trials, its performance was deemed better than existing popular varieties, at least in one trait (e.g. in profitability, yield, nutrition quality, processing quality, resistance to drought, pests and disease, etc.). In the conceptual model of varietal development, I assume that despite being considered an improved variety, under a farmer's growing conditions and given his or her varietal trait preferences, this newly developed variety k can be one of two types relative to the prevailing variety, a high performance type (θ_H) or a low performance type (θ_L).⁶ When a new variety is promoted, farmers do not know whether a variety is θ_H or θ_L on traits they consider important and must learn about it from the extensionist, agricultural input dealers, other farmers, and their own experience. The performance of the prevailing variety (θ_0) is given by π_0 and is assumed constant.⁷ If adopted, a new variety can either have a good performance outcome (π_G) or bad performance outcome (π_B), where $\pi_G > \pi_0 > \pi_B$. While a newly developed variety has a nonzero chance of either outcome regardless of type, the probability of a good (bad) outcome is higher (lower) for newly developed high type varieties than newly developed low type varieties:

$$\begin{aligned} \rho_{GH} &= P(\pi_G | \theta = \theta_H) > P(\pi_G | \theta = \theta_L) = \rho_{GL} \\ \rho_{Bk} &= (1 - \rho_{Gk}) \quad \forall k = H, L \end{aligned} \tag{3.1}$$

This framework reflects the idea that farmers are very familiar with the prevailing local

⁶Performance includes, but is not limited to, profitability. Farmers may also value non-monetary varietal characteristics.

⁷This is not as strong an assumption as it first appears. π_0 can be thought of as the long-run certainty equivalent performance of the prevailing technology which serves as the benchmark for comparing the relative benefits of new technologies.

variety's long-term performance potential while newly developed varieties must be evaluated on the basis of shorter term, noisy outcomes. Furthermore, it reflects the idea that some new varieties are better than others in terms of their performance potential relative to the prevailing technology.

I assume farmers are risk averse expected utility maximizers and that the relative probability distributions of good and bad outcomes imply that if a farmer knew that the new variety was θ_H with certainty, then he or she would adopt the new variety. In contrast, I assume that if a farmer knew that a new variety was θ_L with certainty, then he or she would not adopt the new variety. These conditions imply the following:

$$\begin{aligned} EU(\theta_H) &> EU(\pi_0) > EU(\theta_L) \\ \rho_{GH}U(\pi_G) + \rho_{BH}U(\pi_B) &> U(\pi_0) > \rho_{GL}U(\pi_G) + \rho_{BL}U(\pi_B) \end{aligned} \tag{3.2}$$

This setup is developed to illustrate a simple case where, given full information, a farmer would rationally choose to adopt a new variety that is type θ_H , while rationally choosing not to adopt a new variety that is type θ_L despite both having the potential for higher performance under “good” conditions. As a farmer does not know the type of a newly developed variety, he or she must develop a prior of the probability of a new variety's type. Without any additional information on the type of a new variety, this prior is given by γ_H and γ_L for θ_H and θ_L , respectively, where $\gamma_L = 1 - \gamma_H$. Given no means of updating this prior about the type of a new variety, I assume that a vast majority of farmers will not adopt a newly developed seed variety. That is:

$$\gamma_H EU(\theta_H) + \gamma_L EU(\theta_L) \leq U(\pi_0) \tag{3.3}$$

If this were not the case, then most farmers would adopt a newly developed variety as soon as it became available without any additional information. Thus, this condition provides

the opportunity to explore the role of extension efforts to facilitate technology diffusion by helping farmers update their beliefs about the relative benefits of a new technology.

As a farmer does not know the performance characteristics of a new variety and extensionists may promote varieties of either performance type, this context describes an asymmetric information problem.⁸ A farmer must make probabilistic inferences about the type of a new variety based on the information signals he or she receives from extension activities. The quantity and quality of the information available will depend on the promotional activities of the extension program. For example, if an extensionist promotes a new variety by establishing a demonstration plot that compares the new variety side-by-side with the prevailing local variety, then this activity provides one outcome for a farmer to assess. If, however, an extensionist distributes free trial packs within the village for farmers to test a new variety on their own farms, then this creates many more potential outcomes for a farmer to evaluate. I denote N as the number of potential outcomes that extension activities create in order for farmers to assess the underlying type of a new variety.

Suppose $N = 1$ and a farmer must base his or her belief about the underlying type of the variety based on a single outcome (e.g. the extensionist provides a demonstration plot, but does not distribute trial packs). If a farmer observes this outcome (either good or bad), I assume he or she will update his or her beliefs about the variety's type via Bayes' Rule:

$$P(\theta_H|\pi_i^1) = \frac{\rho_{iH}\gamma_H}{\rho_{iH}\gamma_H + \rho_{iL}\gamma_L} \text{ for } i \in \{G, B\} \quad (3.4)$$

where π_i^1 is the observed outcome, ρ_{iH} and ρ_{iL} are the conditional probabilities of observing that outcome i given that the new variety is of the high type or low type respectively, and γ_H and γ_L are a farmer's prior beliefs about the probability of a high or low type respectively. As defined in equation 3.1, $\rho_{GH} > \rho_{GL}$, which implies that $P(\theta_H|\pi_G) > P(\theta_L|\pi_G)$. That is, if a

⁸Extensionists are informed technology 'sellers' given the task of promoting varieties (of either type) by outside institutions (e.g. non-governmental organizations or public research institutes).

good outcome is observed, then a farmer will adjust his or her belief about the underlying type of the new variety towards θ_H (and vice versa if a poor outcome is observed). Furthermore, as $\rho_{iL}\gamma_L > 0$, a farmer will not know the type of the new variety with certainty even after observing the outcome. A good outcome might be a “lucky” low type result while a poor outcome could be an “unlucky” high type result.

This posterior updating, however, assumes that the outcome π_i^1 will be costlessly observed and incorporated by each farmer. More realistically, whether a farmer observes and incorporates π_i^1 is likely to depend on the quantity and quality of a farmer’s interactions with the extensionist and other farmers. If a farmer does not hear of π_i^1 or is distrustful of the information source, then he or she will not use this outcome to update his or her prior beliefs (γ_H and γ_L).

Let $\delta(\pi_i^1|\mathbf{X})$ be the probability that a farmer learns of and incorporates π_i^1 into his or her beliefs about the new variety where \mathbf{X} are individual characteristics (e.g. farmer personality) that influence the quantity and quality of a farmer’s interactions. The affect of outcome π_i^1 on a farmer’s belief about the type of the underlying variety is then:

$$P(\theta_H|\pi_i^1, \mathbf{X}) = \begin{cases} [\text{Equation 3.4}] & \text{w.p. } \delta(\pi_i^1|\mathbf{X}) \\ \gamma_H & \text{w.p. } 1 - \delta(\pi_i^1|\mathbf{X}) \end{cases} \quad \text{for } i \in \{G, B\} \quad (3.5)$$

That is, a farmer learns of and pays attention to the information inherent in π_i^1 with probability $\delta(\pi_i^1|\mathbf{X})$ and disregards it (either due to lack of awareness or distrust) with probability $1 - \delta(\pi_i^1|\mathbf{X})$.

If $N > 1$ and the farmer is aware of and incorporates a second outcome, π_i^2 , into his or her decision making, then he or she will further update his or her beliefs through an extension

of Bayes' rule as follows:

$$P(\theta_H|\pi_i^1, \pi_j^2) = \frac{P(\pi_j^2|\theta_H, \pi_i^1)P(\theta_H|\pi_i^1, \mathbf{X})}{P(\pi_j^2|\theta_H, \pi_i^1)P(\theta_H|\pi_i^1, \mathbf{X}) + P(\pi_j^2|\theta_L, \pi_i^1)P(\theta_L|\pi_i^1, \mathbf{X})} \text{ for } i, j \in \{G, B\} \quad (3.6)$$

where $P(\theta_H|\pi_i^1, \mathbf{X})$ is as defined in equation 3.5. This delineation demonstrates that the additional value of the second outcome is inversely related to the absolute value of the covariance between the two outcomes. To see this, note that if the two outcomes are perfectly correlated (i.e. $\pi_j^2 = \pi_i^1$ with certainty) then $P(\pi_j^2|\theta_H, \pi_i^1) = 1$ and equation 3.6 collapses to equation 3.5. That is, the second outcome provides no new information over that of the first.⁹ In contrast, suppose $P(\pi_j^2|\theta_H, \pi_i^1) = P(\pi_j^2|\theta_H)$, so knowing the first outcome (π_i^1) provides no information about the likelihood of the second outcome (π_j^1). Then, the value of the second outcome is “as good as” that of the first outcome in terms of updating a farmer’s probabilistic beliefs. Thus, while a second outcome helps a farmer update his or her beliefs about the underlying type of a new variety, the marginal contribution will be larger if the second outcome comes from a more dissimilar context than that of the first outcome.¹⁰

Therefore, receiving information signals from more diverse sources will, on average, better enable farmers to correctly pin down the underlying type of the new variety.¹¹ These results highlight the potential role of the trial pack distribution and farmer personality traits; trial pack distribution creates a diverse set of information sources to learn from and farmer personality traits may affect how farmers react to and process available information.

The updating in equation 3.6, however, is contingent on the farmer’s receipt of the information signal. As in π_i^1 , the existence of π_i^2 alone does not guarantee that a farmer will

⁹In other words, any potential value of a perfectly correlated second outcome is encapsulated in equation 3.5. After observing the first outcome, observing a perfectly correlated second outcome does not help the farmer update his or her beliefs beyond what he or she already knew.

¹⁰Marginal contribution refers to the magnitude of the shift of a farmer’s belief about the type of a variety in equation 3.6 relative to equation 3.5.

¹¹Although observing outcomes from similar sources reduces the likelihood of a farmer receiving conflicting information, this also increases the likelihood of honing in on an incorrect assessment.

benefit from the additional information it provides. The affect of outcome π_i^2 on a farmer's belief about the type of the underlying variety depends on the farmer's awareness of the information source. Thus, the more realistic effect of an additional outcome is:

$$P(\theta_H|\pi_i^1, \pi_i^2, \mathbf{X}) = \begin{cases} [\text{Equation 3.6}] & \text{w.p. } \delta(\pi_i^2|\mathbf{X}) \\ [\text{Equation 3.5}] & \text{w.p. } 1 - \delta(\pi_i^2|\mathbf{X}) \end{cases} \quad \text{for } i \in \{G, B\} \quad (3.7)$$

where $\delta(\pi_i^2|\mathbf{X})$ is the probability of the farmer observing and incorporating π_i^2 into his or her decision making given \mathbf{X} , the characteristics (e.g. farmer personality traits) that affect the quantity and quality of a farmer's interactions.

These results illustrate the potential role of extension activities and farmer personality traits that encourage exposure and learning from varied contexts. As extension activities increase the number of information signals, N , the extensionist increases the likelihood of a farmer updating his beliefs about the new variety as well as increases the accuracy of a farmer's assessment; however, extension activities which encourage experimentation in varied contexts, like the distribution of trial packs, may be better suited to relieving the new variety adoption asymmetric information problem than additional trials on a single demonstration plot. Furthermore, farmer personality characteristics, which are likely to affect the quality and quantity of a farmer's interactions, may play a vital role in addressing this asymmetry. The effect of this additional extension effort is likely to be heterogeneous according to a farmer's awareness of and reaction to new information. Therefore, I hypothesize that the impact of community based extension will differ by farmer personality characteristics. In the next section, I describe the data used to test this hypothesis.

IV. Data

I utilize data from 740 bean farmers and 32 VBAs (one per village) from 32 villages which participated in the previously implemented RCT on VBA extension activities (Melkani & Mason, 2018). These 32 villages represent a subsample of those in the Mbozi and Mbeya Districts that participated in the larger RCT. In each of the 32 villages, a random sample of approximately 25 bean-growing households were selected for two rounds of a detailed survey conducted immediately following the 2017 and 2018 Major Seasons. As the population of interest is bean farmers in each village, non-bean growing households were excluded from the sampling frame.

The 2017 farmer survey collected a range of detailed household and respondent level information including variety and plot specific bean production and varietal use in the minor and major bean growing seasons in 2016 (pre-intervention), minor bean growing season in 2017 (pre-intervention), and in major season 2017 (during intervention) as well as household and respondent sociodemographic characteristics and household GPS location. Table 3.2 provides covariate balance tests based on the farmer sample.¹² Across most farmer characteristics, I fail to reject the null hypothesis of equality between the demo-only and demo-trial VBA villages. In particular, farmer adoption of Njano Uyole and Uyole 96 before the start of the intervention was low and not statistically different between demo-only and demo-trial villages. Households in demo-only villages, however, are slightly larger on average than those in demo-trial villages (mean household size was 5.67 and 5.19 in the demo-only and demo-trial villages respectively). Additionally, demo-only households were more likely to be related to the village chairman than those in demo-trial villages. I control for these unbalanced covariates in my analysis.

To assess the impact of the RCT intervention on farmer adoption of promoted bean

¹²The detailed 2017 farmer survey occurred after the incidence of the treatment in the 2017 Major Season. Thus, farmer characteristics which may have been influenced by the treatment (e.g. knowledge of improved bean varieties and VBA-activities) are excluded from the balancing tests.

varieties, a follow up survey of the same households was conducted immediately following the 2018 Major Season. Analysis of this second round of data from post-intervention can provide evidence of any short-term adoption effects and which subgroups of farmers were most likely to benefit (or not) from the trial packs treatment after the program’s completion. Of the 791 households surveyed in the 2017 round, 740 households were successfully resurveyed in the 2018 round (an attrition rate of 6%). To assess whether this household attrition was correlated with village treatment status, I estimated logit regressions of the household attrition indicator variable on village treatment status controlling for the VBAA-pair used in the random treatment assignment. As shown in Table 3.3, I do not find evidence that household attrition was related to village treatment assignment even after controlling for 2017 household size and relationship to the village chairmen.

Table 3.4 provides a comparison of key characteristics across three agricultural years by demo-only and demo-trial villages. In the year before the intervention, only 7.8% and 6.6% of sampled households had adopted Uyole 96 or Njano Uyole in demo-only and demo-trial villages respectively.¹³ When the RCT was implemented in the 2017 agricultural year, adoption in demo-only villages remained relatively flat at 8.5%. As expected, the reported adoption in demo-trial villages was higher (13%) in 2017, reflecting an increased planting of these varieties due to the distribution of trial packs of improved bean seed. By the 2018 agricultural year (i.e. the post-intervention year), the bean varieties had diffused more widely in both demo-only and demo-trial villages with 21% of sampled households growing at least one of the two varieties. This increase in the use of these varieties is suggestive of households receiving information signals in 2017 that induced take-up. While the adoption rate was more than 50 percent higher among the subsample of households residing in demo-trial villages relative to the demo-only subsample at the time of the intervention, the adoption rates between the two subsamples were similar in 2018 (20-21%), but represent an approxi-

¹³I define adoption as the respondent reporting that his or her household grew either variety in the given agricultural year.

mately 2.5 and 3 times increases from pre-intervention to post-intervention for the demo-only subsample and demo-trial subsamples respectively. In my subsequent empirical analysis, I explore the farmer characteristics which may have led to heterogeneity in the value of this additional trial pack signal.

In addition to the respondent and household characteristics collected in the 2017 survey round, farmer social interactions were also measured in the 2018 survey round (although not specific to that year). Each farmer was asked: “Have you ever met [Name]?” where [Name] was replaced by the name of the VBAA working in his or her village. If the farmer replied yes, he or she was asked several related follow-up questions including: (1) Was your VBAA named [Name]? (2) Have you ever gone to [Name] for advice about farming? and (3) Have you ever discussed bean farming with [Name]? Along with these VBAA-specific questions, each farmer was also asked whether he or she had ever met, ever sought farming advice from, and ever discussed bean farming with each of the approximately 24 other bean farming households that were sampled in that farmer’s village. Summary statistics for these social interaction outcomes are presented in Table 3.4.

In addition to social interaction outcomes, personality traits were also measured in the 2018 survey round via the Lachman and Weaver (1997) Midlife Development Inventory (MIDI). This scale is based on the widely studied Five Factor Model of personality which was expanded to include a sixth trait (e.g. Bazzani et al., 2017; Grebitus et al., 2013; Lin et al., 2019; Van der Linden et al., 2018). As described in Table 3.5, MIDI captures personality by grouping 31 adjectives into six major traits of agency, agreeableness, openness to experience, neuroticism, extraversion, and conscientiousness. Relative to other personality inventories based on the Five Factor Model, MIDI has the advantage of being relatively easy to elicit as well as measuring an additional sixth trait, agency. Respondents rated how well each adjective described them on a scale of: one (not at all), two (a little), three (some), and four (a lot). Respondents’ personality traits are measured as the average scores of the associated

adjectives.

The MIDI scale was translated into Swahili through multiple rounds of translation and backward translation. First, two experts were asked to translate each word independently. One was a Tanzanian professor at a U.S. university who teaches Swahili to English speaking students, and the other was a collaborator based in Tanzania. Another professor from a U.S. university who is based in Tanzania (and is Tanzanian) was then asked to backward translate the Swahili translated words received from the local collaborator. In the last round, the survey enumerators and two supervisors, who are all fluent in both Swahili and English, were presented with the two sets of translation and one set of backward translation and as a group went through a rigorous consultation process to come up with the final Swahili translation of each word. This final round with the enumerators and supervisors also ensured that a common understanding was reached among all on each personality trait being measured.

This rigorous approach to translating and localizing the MIDI scale is particularly important given the rural, developing country sample. Capturing personality traits in developing country contexts accurately is challenging. Laajaj et al. (2019) compare the validity of personality trait measurements across a variety of contexts, finding that personality trait measurements from developing country surveys tended to be less reliable. Although the data collection for this study preceded the publication of Laajaj et al. (2019), the extensive translation and enumerator training for the personality module is a main approach recommended by Laajaj et al. (2019) for improving the measurement of personality in developing country contexts.

Furthermore, the Cronbach's alpha values of the personality traits in this study's sample, the main measure of internal consistency and reliability for personality scales, compare favorably to the developing country in-person surveys analyzed in Laajaj et al. (2019). Cronbach's alpha increases (i.e. improves) as the within-group correlation between grouped items

increases and as the number of items increases (Laajaj et al., 2019). As shown in Table 3.6, the majority of the Cronbach's alpha values for the personality traits in this study's sample are 0.7 or greater with an average of 0.69. The six developing country surveys analyzed in Laajaj et al. (2019) with 44 item personality scales (i.e. more than the 31 item MIDI scale used in this study) had an average Cronbach's alpha value of 0.62 (ibid).

As discussed previously, each village self-selected their VBAA based on a variety of farming characteristics and other skills. As village-selected lead-farmers, I would expect VBAA's to differ from the average farmer on skill-related variables such as farming experience. Whether these differences include personality, however, is uncertain *a priori*. Farmers might prefer to have a VBAA that is of similar personality temperament or one that stands out from the average. As shown in Table 3.7, in this study's sample I find the latter. That is, VBAA's tend to score higher than the average farmer across a variety of personality dimensions. Of the six personality traits, the three traits in which VBAA's differ the most from the farmer sample are in openness to experience, extraversion, and agency. Individuals high in openness to experience are more willing to experiment with new ideas, making them good candidates for a lead-farmer position. Similarly, being high in extraversion, which is a measure of sociability and tendency to be outgoing, may make a lead-farmer position more attractive. High agency, which measures self-confidence, dominance, and outspokenness, may also help a farmer win the position of VBAA and succeed in marketing new ideas once doing so. Along with these major traits, VBAA's also stood out from the farmer sample in terms of agreeableness and conscientiousness. Agreeableness measures kindness and likeability, characteristics which may make a candidate more likely to be chosen for a lead farmer position. Conscientiousness, which is a measure of organization, responsibility, and work ethic, is a trait which may be particularly important for a VBAA's effectiveness in his or her role. All of these differences are indicative of lead-farmer personality being markedly different than the average farmer. The trait where VBAA's are most similar to the

farmer sample is neuroticism, which is a negative personality trait associated with anxiety and worry (Greibitus et al., 2013). In my empirical analysis, I control for these VBAA personality traits.

Along with controlling for VBAA personality in my analysis, I also control for a more common dimension of similarity, physical distance. Farmers that interact more frequently with their VBAA are more likely to be exposed to extension information and this frequency of interaction is likely correlated with physical distance. I measure physical distance by capturing the GPS location of each VBAA's and sampled farmer's homestead as well as that of the VBAA-led demonstration plot. A farmer's distances to the VBAA's homestead and demonstration plot are measured via the GPS-based linear distances in kilometers. Table 3.8 reports summary statistics for these physical distances.

V. Empirical Strategy

My first outcome of interest is farmer adoption of the improved bean varieties post-intervention (i.e. in the 2018 agricultural year). This adoption analysis provides an indication of the marginal value (in terms of farmer adoption probability) of the additional information signals from the distribution of trial packs of bean seed relative to the information signals from a demonstration plot alone. It also enables an assessment of whether farmers with certain personality traits were more likely to benefit (in terms of adoption) from these additional information signals.

Along with adoption, I also investigate the relationships between the trial pack distribution and farmer personality traits on farmers' social interactions with their VBAA and other farmers. Analysis of these intermediate outcomes serves two main purposes. First, it evaluates the increase in the bilateral exchange of information between farmers and VBAA's—a key mechanism through which community-based extension is intended to encourage farmer take-up of the improved bean varieties. Second, it assesses whether farmer personality char-

acteristics predict differences in these key social mechanisms for community-based extension effectiveness.¹⁴

I analyze three farmer-VBAA interaction outcomes: whether the farmer (1) has ever met his or her VBAA and identified him/her as the VBAA, (2) has ever sought farming advice from his/her VBAA, and (3) has ever discussed bean farming with his/her VBAA. Similarly, I analyze three other peer interaction outcome variables: the proportion of a random sample of ~ 24 bean farming households in the farmer's village that he/she (1) has ever met, (2) has ever sought farming advice from, and (3) has ever discussed bean farming with. The analysis of these six outcomes provides evidence on the extent to which farmer personality plays a role in the quantity and quality of VBAA and peer farmer interactions.

For each outcome, I specify the following equation:

$$E(Y_{ijk}|Trial_{jk}, \mathbf{P}_{ijk}, \mathbf{X}_{ijk}) = G(\alpha + \beta Trial_{jk} + \mathbf{P}_{ijk}\delta + \mathbf{X}_{ijk}\zeta) \quad (3.8)$$

where Y_{ijk} is the binary or fractional outcome variable for farmer i in village j of VBAA-pair k , $G(\cdot)$ is a logistic function, $Trial_{jk}$ is an indicator variable equal to one if village j of VBAA-pair k was randomly assigned to the demo-trial group and zero otherwise, \mathbf{P}_{ijk} is a vector of farmer personality characteristics, \mathbf{X}_{ijk} is a vector of controls including farmer and VBAA sociodemographics, VBAA personality characteristics, and VBAA-pair indicator variables. For the binary response and fractional response outcome variables, I report average marginal effects from the logit regression and fractional logit regression respectively (Papke & Wooldridge, 1996; Wooldridge, 2010). This specification evaluates the effect of the trial pack treatment as well as whether farmer personality types are associated with different adoption and social outcomes.

¹⁴With only 32 villages and VBAA's in the sample (one VBAA per village), I am unable to assess the potential effect of VBAA personality characteristics on community adoption of the improved varieties. This remains a promising area for future work with an expanded sample.

I also estimate the following supplementary linear model:

$$E(Y_{ijk}|Trial_{jk}, \mathbf{P}_{ijk}, \mathbf{X}_{ijk}) = \alpha + \beta Trial_{jk} + Trial_{jk} * \mathbf{P}_{ijk}\gamma + \mathbf{P}_{ijk}\delta + \mathbf{X}_{ijk}\zeta \quad (3.9)$$

where farmer personality characteristics, \mathbf{P}_{ijk} , are interacted with the trial indicator variable. This specification enables an assessment of potential heterogeneity in the effect of the trial pack treatment by farmer personality traits.

In addition to these main specifications, I also specify adoption models for equations 3.8 and 3.9 which control for the pre-intervention outcome (i.e. adoption in the 2016 agricultural year). The results for these specifications, based on the ANCOVA framework outlined in McKenzie (2012), are nearly identical to that of the main specifications (see Appendix Tables 3.15 and 3.16). As in the balancing tests, these robustness specifications suggest that residing in a village that received trial packs of bean seed is exogenous and that my empirical strategy controls for unobserved, pre-intervention differences between farmers in demo-trial and demo-only villages.

VI. Results

A key outcome in terms of the diffusion of these improved bean varieties is whether the trial pack treatment or personality traits are associated with changes in the likelihood of farmer adoption of the new varieties. The results based on farmer adoption of Uyole 96 or Njano Uyole in the 2018 agricultural year are presented in Table 3.9. As discussed previously, the overall adoption rate among the sample of households increased dramatically from 10.7% in 2017 to 21% in 2018. The adoption gap between demo-only and demo-trial villages, however, declined. Given this, it is not surprising that I find limited evidence that residing in a village that received trial packs increased farmer adoption of these improved varieties on average. Across most specifications, I cannot reject the null that, for the average farmer

in the sample, residing in a village with additional information signals from the trial pack distribution did not increase the likelihood of adopting the improved varieties relative to residing in a village where only the VBAA-led demonstration plot took place. I also do not find evidence of differences in the likelihood of farmer adoption by personality type. That is, farmer personality traits alone do not predict improved variety adoption in the 2018 agricultural year.

There may, however, be important heterogeneity that this initial specification cannot characterize. As shown in the conceptual framework, personality traits may influence the likelihood that a farmer seeks out additional information signals as well as how he or she reacts to the information received; therefore, the effect of the trial pack distribution may be heterogeneous by farmer personality types. Table 3.10 presents the results for equation 3.9 which allows the effect of residing in a trial pack village on farmer adoption of Njano Uyole or Uyole 96 to vary by farmer personality traits. While the impact of trial packs on the diffusion of these improved bean varieties was attenuated for the average farmer, extraverted farmers residing in trial pack villages were more likely to adopt these improved bean varieties than their peers. This suggests that farmer personality may play a vital role in the likelihood of benefiting from community-based extension efforts which aim to increase the exchange of information between farmers.

I assess this potential social mechanism of the community-based extension program via outcome variables that measure farmers' social interactions with their VBAA and other farmers. Table 3.11 columns 1-3 present the results for equation 3.8 based on whether a farmer has ever met the VBAA in his or her village and correctly identified him or her as the VBAA. The findings show that farmers residing in trial pack villages were more likely to be able to identify their VBAA relative to farmers in demonstration plot only villages. This suggests that the trial pack distribution may have boosted farmer participation in and knowledge of extension activities. Furthermore, farmers scoring higher in agreeableness,

which is a measure of friendliness, and openness to experience were more likely to correctly identify their VBAA. Farmers scoring high in neuroticism, which is a measure of anxiety and worry, however, were less likely to be able to identify their VBAA. These results are generally robust to additional demographic controls and suggest that farmer personality traits are associated with differences in farmers' basic awareness of the existence of a VBAA in their village.

The results based on differences in a farmer's likelihood of ever seeking out farming advice from his or her VBAA, a stronger measure of information exchange, were more mixed (Table 3.11 columns 4-6). There is some evidence that the trial pack distribution may have increased the likelihood of farmers seeking advice from their VBAA, but this result is not robust. Similarly, there is some evidence that farmers scoring high in agency, which is a measure of self-confidence, were less likely to seek out farming advice. This is consistent with these farmers being more sure of their own practices and less willing to seek out others' opinions. Scoring high in openness to experience, however, may offset this effect.

A more middle ground means of information exchange relative to having met a VBAA or sought farming advice from him or her, is whether farmers indicate that they have ever discussed bean farming with their VBAA; results based on this outcome, which are presented in Table 3.11 columns 7-9, show more robust evidence for trial pack and personality differences. In particular, residing in a village that received trial packs increases the likelihood that a farmer has discussed bean farming with his or her VBAA by approximately seven percentage points. Higher agency (i.e. more self-confident) farmers are also much less likely to have ever discussed bean farming with their VBAA; the predicted reduction from scoring a single point higher in agency would more than offset the gain in likelihood from residing in a village that received trial packs.

Along with a farmer's interactions with his or her VBAA, other important social outcomes are his or her interactions with other peer farmers in the village. The results based on

the proportion of a sample of ~ 24 bean farming households in a farmer's village that he or she has ever met (Table 3.12 columns 1-3), ever sought farming advice from (Table 3.12 columns 4-6), and ever discussed bean farming with (Table 3.12 columns 7-9), are similar to those based on farmer-VBAA interaction. The findings suggest that the trial pack distribution increased the proportion of bean farming households that farmers have met and have discussed bean farming with. Similar to that of the VBAA outcomes, however, there is no evidence that it increased the proportion of bean farming households that the farmer has sought advice from. This suggests that the trial pack distribution increased the bi-lateral exchange of information between farmers in the village, but not to an extent that farmers were more willing to seek out farming advice.

The role of personality traits in farmer interactions with their peers are also very similar to that of the VBAA social outcomes. In particular, higher agency and neuroticism are associated with declines in peer social interaction while openness to experience is associated with an increase in peer social interaction. There is also some evidence that increases in agreeableness (friendliness) reduced the willingness of farmers to seek advice from their peers, but this result is not robust to the inclusion of demographic controls.

Finally, a heterogenous effects analysis on the proportion of bean farming households that a farmer has sought advice from (Table 3.13) and discussed bean farming with (Table 3.14) provides evidence for why extraverted farmers may have benefited more from the trial pack distribution. The results show that more extroverted farmers residing in trial pack villages were more likely to have sought out relevant social interactions with peer farmers. This is consistent with extraversion playing an important role in the quantity and quality of trial pack information signals a farmer receives and leading to an increase in take-up of the improved bean varieties. This exploratory social outcome analysis suggests that farmer personality may raise farmer awareness of extension information and also provides evidence that the trial pack treatment increased the exchange of information.

VII. Conclusions

In this study, I examine the role of community-based extension and farmer personality on adoption of improved bean varieties and social interactions. I develop a conceptual framework demonstrating how farmer adoption of new varieties can be modeled as an asymmetric information problem whereby the farmer seeks to discover the underlying, unobserved relative benefits of a new bean variety based on extension activities. I show that farmer personality can influence the likelihood of receiving a benefits signal. I then examine this heterogeneity empirically using a unique dataset of the Big Six personality traits measured using the Midlife Development Inventory (MIDI) for bean farmers in the Mbeya Region of Tanzania (Lachman & Weaver, 1997).

In terms of adoption impacts, although the trial pack distribution increased immediate planting of the improved varieties in the year of the intervention, diffusion rates of the improved bean varieties converged between demo-only and demo-trial villages in the one year post-intervention. Correspondingly, I find limited evidence that providing additional extension resources to VBAAAs in the form of trial packs of bean seed increased the average farmer's post-intervention likelihood of adoption of the improved bean varieties relative to only providing VBAAAs resources for a bean demonstration plot. This suggests that, on average, the demonstration plot alone provided enough of a meaningful signal to farmers to kickstart diffusion.

Although the benefits of the trial pack treatment were attenuated for the average farmer, I do find evidence of heterogeneous treatment effects by personality traits. In particular, extraverted farmers residing in trial pack villages were more likely to adopt Njano Uyole or Uyole 96. These farmers were also more likely to seek out farming advice and discuss bean farming with their neighbors, suggesting that the benefits of the trial pack treatment may have been greater (smaller) for more (less) sociable farmers.

I also find that farmers residing in villages randomly selected to receive trial packs were

more likely to be able to identify their VBAA as well as have discussed bean farming with him or her. Furthermore, farmer personality traits were also associated with differences in social outcomes related to extension effectiveness. In particular, farmers with higher openness to experience were more likely to know their VBAA as well as have discussed bean farming with him or her. Similarly, farmers that scored higher on the openness to experience personality trait also discussed bean farming with a greater proportion of farmers in their village on average. These findings are consistent with the trial pack treatment and farmer personality traits affecting farmer awareness of information on bean farming.

Although these findings should be interpreted as predictors and not necessarily causal, they have important implications for the design of community-based extension programs and warrant future research. Community-based extension programs are inherently social, relying on already established informal institutional connections to facilitate information flows. Personality influences interpersonal relationships, yet it has received relatively little attention in the technology adoption and extension literature. My findings suggest that, much like in social relationships, personality characteristics may influence who benefits from community-based extension programs. In this study's context, the marginal gains from providing an experience-based information treatment (like the trial pack) appear to have been higher for extroverted recipient farmers. This suggests that personality influences the potential beneficiaries of community-based extension programs and that community-based extension effectiveness may be increased by considering the role of personality in the quality and quantity of farmers interactions within their communities.

APPENDICES

APPENDIX A:

Tables and Figures

Table 3.1: VBAA Balance Tests

Characteristics (as of baseline survey)	Mean Value		P-value from t-test
	Demo-only (n=16)	Demo-trial (n=16)	
Age (years)	44.44	46.75	0.508
Is female	0.19	0.19	1.000
Level of education is primary or less	0.75	0.88	0.382
Farming experience (years)	23.56	23.06	0.889
Maize farming experience	23.56	23.06	0.889
Bean farming experience	21.44	20.00	0.729
Land area owned in acres (2016 ag. year)	14.28	10.97	0.476
Household size	8.31	6.56	0.156
Indicators of activities			
Distributed free maize seed	0.94	0.94	1.000
Distributed free bean seed	0.56	0.56	1.000
Setup maize demonstration plot	1.00	0.81	0.083
Setup bean demonstration plot	0.56	0.38	0.303
Sold commercial maize seed	0.50	0.50	1.000
Sold commercial bean seed	0.06	0.00	0.333
Sold commercial fertilizer	0.06	0.06	1.000
Sold pesticides or seed treatments	0.13	0.06	0.560
No. free maize seed packs allocated by FIPS-Africa	198.44	214.63	0.743
No. of farmers given free maize seed packs	225.38	211.06	0.782
No. free bean seed packs allocated by FIPS-Africa	131.94	123.44	0.888
No. of farmers given free bean seed packs	131.94	104.81	0.641
No. of maize demonstration plots setup	1.19	1.00	0.480
No. of people involved in maize demonstration plot	12.31	8.44	0.235
No. of bean demonstration plots setup	0.75	0.44	0.250
No. of people involved in bean demonstration plot	4.31	5.31	0.688
Personality traits			
Agency	3.35	3.29	0.590
Agreeableness	3.71	3.60	0.429
Openness to experience	3.65	3.55	0.476
Neuroticism	2.09	2.14	0.820
Extraversion	3.67	3.54	0.357
Conscientiousness	3.67	3.67	1.000

Note: Tests of equality of means. t-test assumes unequal variances. Unless otherwise noted, VBAA characteristics are based on current status at the time of the baseline VBAA survey in early 2017. Indicators of VBAA performance are for the 2016 agricultural year. Personality traits were measured in 2018 on a scale of one (not at all), two (a little), three (some), and four (a lot).

Table 3.2: Farmer Survey Balance Tests

Characteristics	Mean Value		P-value from t-test
	Demo-only (n=399)	Demo-trial (n=392)	
Grew Njano Uyole or Uyole 96 in 2016 ag. year	0.08	0.07	0.537
Received bean seed from FIPS-Africa VBAA in 2016 ag. year	0.01	0.01	0.688
Bean production area in 2016 ag. year (acres)	0.94	0.85	0.245
Land owned (acres)	3.94	4.11	0.545
Household head is female	0.18	0.22	0.205
Household head's age (years)	47.88	48.47	0.757
Household's highest level of education is primary or less	0.64	0.67	0.311
Household size	5.67	5.19	0.016
Respondent or spouse related to the village chairman	0.31	0.19	0.000
Respondent or spouse related to the village extension officer	0.02	0.02	0.590
Physical distance (km) between farmer's and VBAA's homestead	1.33	1.49	0.087
Physical distance (km) between farmer's homestead and VBAA demo. plot	1.78	1.67	0.307
Personality traits			
Agency	2.88	2.91	0.309
Agreeableness	3.31	3.36	0.103
Openness to experience	2.90	2.91	0.876
Neuroticism	1.95	1.92	0.497
Extraversion	3.12	3.17	0.257
Conscientiousness	3.34	3.35	0.862

Note: Tests of equality of means. t-test assumes unequal variances. Unless noted below, characteristics were collected in the 2017 farmer survey. Household head gender and age were collected in 2018. Physical distance is based on GPS data collected in 2018. Personality traits were collected in 2018 on a scale of one (not at all), two (a little), three (some), and four (a lot).

Table 3.3: Household Attrition Tests

Variable	(1)	(2)
Village treatment status (1=demo-trial; 0=demo only)	0.313 (0.294)	0.158 (0.302)
Household size in 2017		−0.260*** (0.073)
Respondent or spouse related to village chairman in 2017		−0.763* (0.460)
Observations	791	791
Pseudo-R2	0.0277	0.0776

Note: Logit regressions. Dependent variable equals one if the household is only observed in the first survey round (2017), zero otherwise. All regressions include indicator variables for VBAA-pair used in the random assignment. Standard errors in parentheses.

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 3.4: Key Characteristics by Agricultural Year and Treatment Group

Variable	2016 Ag. Year		2017 Ag. Year		2018 Ag. Year	
	Pre-intervention		Intervention year		Post-intervention	
	Demo-only (n=399)	Demo-trial (n=392)	Demo-only (n=399)	Demo-trial (n=392)	Demo-only (n=377)	Demo-trial (n=363)
Information specific to the given ag. year:						
Farmer's household grew Njano Uyole or Uyole 96	0.08 (0.27)	0.07 (0.25)	0.09 (0.28)	0.13 (0.34)	0.20 (0.40)	0.21 (0.41)
Farmer's household has ever received training on Uyole 96 or Njano Uyole	-	-	0.05 (0.22)	0.05 (0.23)	0.16 (0.37)	0.17 (0.37)
Farmer's household received a trial pack of bean seed	0.01 (0.09)	0.01 (0.10)	0.05 (0.22)	0.16 (0.37)	0.03 (0.18)	0.08 (0.28)
Bean production area (acres)	0.94 (1.11)	0.85 (1.05)	1.35 (1.07)	1.22 (1.09)	1.13 (1.07)	1.04 (1.08)
Information collected in 2018, but not specific to that ag. year:						
Farmer has ever met VBAA and identified him/her as the VBAA	-	-	-	-	0.552 (0.498)	0.634 (0.482)
Farmer has ever sought farming advice from VBAA	-	-	-	-	0.475 (0.500)	0.474 (0.500)
Farmer has ever discussed bean farming with VBAA	-	-	-	-	0.491 (0.501)	0.521 (0.500)
Proportion of a random sample of ~ 24 bean farming households in the farmer's village that he/she... ...has ever met	-	-	-	-	0.786 (0.225)	0.847 (0.192)
...has ever sought farming advice from	-	-	-	-	0.130 (0.223)	0.122 (0.209)
...has ever discussed bean farming with	-	-	-	-	0.151 (0.228)	0.158 (0.235)

Note: Mean values by agricultural year and treatment group. Standard deviations in parentheses. Dash indicates outcome not measured for given year. Characteristics for 2016 agricultural year were collected retroactively in 2017. The trial pack variable for 2016 is for any bean seed received from a FIPS-Africa VBAA (more broad than the Njano Uyole or Uyole 96 bean seed definition used for 2017 and 2018).

Table 3.5: Midlife Development Inventory (MIDI) Personality Traits

Trait	Corresponding adjectives
Agency	Self-confident, Forceful, Assertive, Outspoken, Dominant
Agreeableness	Helpful, Warm, Caring, Softhearted, Sympathetic
Openness to experience	Creative, Imaginative, Intelligent, Curious, Broadminded, Sophisticated, Adventurous
Neuroticism	Moody, Worrying, Nervous, Calm(-)
Extraversion	Outgoing, Friendly, Lively, Active, Talkative
Conscientiousness	Organized, Responsible, Hardworking, Careless(-), Thorough

Note: Adapted from Lachman and Weaver (2005). Respondents rated how well each adjective described them on a scale of: one (not at all), two (a little), three (some), and four (a lot). Traits are the average scores of the associated adjectives. (-) indicates that the adjective was reverse coded when scoring.

Table 3.6: Personality Trait Cronbach's Alpha Values

Trait	Alpha (N=772)
Agency	0.60
Agreeableness	0.78
Openness to Experience	0.74
Neuroticism	0.59
Extraversion	0.72
Conscientiousness	0.70

Table 3.7: Farmer-VBAA Personality Traits Comparison

Trait	Farmers (N=740)	VBAAAs (N=32)	P-value from t-test
Agency	2.89 (0.52)	3.32 (0.32)	0.000
Agreeableness	3.33 (0.48)	3.66 (0.39)	0.000
Openness	2.91 (0.47)	3.60 (0.38)	0.000
Neuroticism	1.94 (0.59)	2.12 (0.57)	0.088
Extraversion	3.14 (0.5)	3.61 (0.41)	0.000
Conscientiousness	3.35 (0.47)	3.67 (0.4)	0.000

Note: t-tests of equality of means assuming unequal variances. Personality traits were measured on a scale of one (not at all), two (a little), three (some), and four (a lot).

Table 3.8: Physical Distance Summary Statistics

Variable	Mean	Std. Dev.	Min	Max
Physical distance (km) between farmer's and VBAA's homestead	1.41	1.34	0.02	7.11
Physical distance (km) between farmer's homestead and VBAA demonstration plot	1.73	1.48	0.05	7.24

Note: N=740. Physical distances are GPS-based linear distances.

Table 3.9: Farmer Adoption of Improved Bean Varieties in 2018 Ag. Year

Variable	(1)	(2)	(3)
Trial	0.007 (0.032)	0.047 (0.033)	0.097** (0.046)
Agency		0.028 (0.039)	0.032 (0.038)
Agreeableness		0.019 (0.038)	0.021 (0.040)
Openness to experience		-0.042 (0.051)	-0.048 (0.054)
Neuroticism		0.011 (0.021)	0.012 (0.019)
Extraversion		0.062 (0.041)	0.062 (0.041)
Conscientiousness		-0.019 (0.042)	-0.023 (0.044)
VBAA-pair indicators	Yes	Yes	Yes
VBAA personality and distance controls \a	No	Yes	Yes
Farmer and VBAA demographic controls \b	No	No	Yes
Observations	740	740	740
Pseudo R-squared	0.171	0.211	0.220
P-value of test that farmer personality average partial effects are jointly zero		0.158	0.104

Note: Logit average partial effects with standard errors clustered at the village level in parentheses. Dependent variable is equal to one if the farmer's household adopted Njano Uyole or Uyole 96 in the 2018 agricultural year. Trial is an indicator variable equal to one if the farmer resides in a village where the VBAA received trial packs. Personality traits are measured on a scale of one (not at all) to four (a lot).

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

\a VBAA personality controls are the same as that of the farmer. Distance controls are the GPS-based physical distances (km) from the farmer's homestead to the VBAA's homestead and VBAA demonstration plot.

\b Demographic controls (2017 ag. year) are: farmer household's size; indicators for farmer or spouse related to village chairmen; farmer's and VBAA's gender and age; indicator for farmer household's and VBAA's highest education achieved is primary or less.

Table 3.10: Heterogeneous Trial Pack Effects on Farmer Adoption of Improved Bean Varieties in 2018 Ag. Year

Variable	(1)	(2)
Trial	0.234 (0.219)	0.312 (0.235)
Trial*Agency	-0.048 (0.073)	-0.049 (0.075)
Trial*Agreeableness	-0.071 (0.082)	-0.068 (0.087)
Trial*Openness to experience	-0.002 (0.107)	-0.007 (0.112)
Trial*Neuroticism	0.008 (0.041)	0.000 (0.042)
Trial*Extraversion	0.152** (0.072)	0.149** (0.072)
Trial*Conscientiousness	-0.090 (0.078)	-0.092 (0.078)
VBAA-pair indicators	Yes	Yes
Farmer and VBAA personality and distance controls \a	Yes	Yes
Farmer and VBAA demographic controls \b	No	Yes
Observations	740	740
R-squared	0.219	0.227

Note: Linear probability model coefficients with standard errors clustered at the village level in parentheses. Dependent variable is equal to one if the farmer's household adopted Njano Uyole or Uyole 96 in the 2018 agricultural year. Trial is an indicator variable equal to one if the farmer resides in a village where the VBAA received trial packs.

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

\a Farmer and VBAA personality controls are their six personality traits measured on a scale of one (not at all) to four (a lot). Distance controls are the GPS-based physical distances (km) from the farmer's homestead to the VBAA's homestead and VBAA demonstration plot.

\b Demographic controls (2017 ag. year) are: farmer household's size; indicators for farmer or spouse related to village chairmen; farmer's and VBAA's gender and age; indicator for farmer household's and VBAA's highest education achieved is primary or less.

Table 3.11: Farmer Interactions with VBAA

Variable	Met VBAA			Farmer has ever... Sought Farming Advice from VBAA			Discussed Bean Farming with VBAA		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Trial	0.081*	0.104***	0.059	−0.000	0.025	0.057**	0.031	0.068***	0.069**
	(0.044)	(0.036)	(0.037)	(0.035)	(0.021)	(0.028)	(0.038)	(0.023)	(0.029)
Agency		−0.054	−0.056		−0.069	−0.111**		−0.086*	−0.134***
		(0.039)	(0.038)		(0.053)	(0.051)		(0.047)	(0.046)
Agreeableness		0.073*	0.103**		−0.012	0.016		0.018	0.043
		(0.044)	(0.044)		(0.063)	(0.059)		(0.060)	(0.054)
Openness to experience		0.114**	0.057		0.097*	0.038		0.123**	0.076
		(0.049)	(0.055)		(0.056)	(0.059)		(0.054)	(0.059)
Neuroticism		−0.159***	−0.164***		0.018	0.015		0.028	0.025
		(0.039)	(0.037)		(0.044)	(0.041)		(0.045)	(0.043)
Extraversion		−0.078	−0.085		−0.094	−0.086		−0.047	−0.040
		(0.055)	(0.052)		(0.064)	(0.061)		(0.064)	(0.063)
Conscientiousness		0.011	−0.000		0.060	0.067		0.054	0.067
		(0.053)	(0.051)		(0.067)	(0.069)		(0.057)	(0.060)
VBAA-pair indicators	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
VBAA personality and distance \a	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Farmer and VBAA demographics \b	No	No	Yes	No	No	Yes	No	No	Yes
Observations	740	740	740	740	740	740	740	740	740
Pseudo R-squared	0.058	0.169	0.193	0.038	0.070	0.096	0.043	0.084	0.112
P-value of test that farmer personality average partial effects are jointly zero		0.000	0.000		0.064	0.002		0.011	0.001

Note: Logit average partial effects with standard errors clustered at the village level in parentheses. Trial is an indicator variable equal to one if the farmer resides in a village where the VBAA received trial packs. Personality traits are measured on a scale of one (not at all) to four (a lot). * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$. \a VBAA personality controls are the same as that of the farmer. Distance controls are the GPS-based physical distances (km) from the farmer's homestead to the VBAA's homestead and VBAA demonstration plot. \b Demographic controls (2017 ag. year) are: farmer household's size; indicators for farmer or spouse related to village chairmen; farmer's and VBAA's gender and age; indicator for farmer household's and VBAA's highest education achieved is primary or less.

Table 3.12: Farmer Interactions with Other Bean Farmers

	Proportion of Bean Farming Households that the Farmer has ever...								
	Met			Sought Farming Advice From			Discussed Bean Farming With		
Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Trial	0.417*** (0.144)	0.052*** (0.014)	0.080*** (0.021)	−0.086 (0.211)	0.009 (0.020)	0.026 (0.020)	0.049 (0.181)	0.013 (0.021)	0.035** (0.017)
Agency		0.006 (0.024)	−0.015 (0.024)		−0.066*** (0.019)	−0.092*** (0.017)		−0.061*** (0.021)	−0.086*** (0.020)
Agreeableness		0.008 (0.023)	0.013 (0.024)		−0.043** (0.022)	−0.026 (0.019)		−0.038 (0.025)	−0.026 (0.023)
Openness to experience		0.023 (0.029)	0.014 (0.027)		0.053*** (0.020)	0.010 (0.022)		0.079*** (0.019)	0.043* (0.023)
Neuroticism		0.039*** (0.012)	0.037*** (0.013)		−0.024 (0.016)	−0.024* (0.014)		0.002 (0.018)	0.003 (0.018)
Extraversion		0.034 (0.028)	0.033 (0.025)		0.000 (0.032)	0.011 (0.030)		0.033 (0.033)	0.047 (0.032)
Conscientiousness		−0.034 (0.029)	−0.022 (0.030)		0.031 (0.023)	0.036* (0.022)		0.039* (0.020)	0.047** (0.020)
VBAA-pair indicators	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
VBAA personality and distance \a	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Farmer and VBAA demographics \b	No	No	Yes	No	No	Yes	No	No	Yes
Observations	740	740	740	740	740	740	740	740	740
Pseudo R-squared	0.054	0.069	0.083	0.056	0.097	0.135	0.063	0.093	0.118
P-value of test that farmer personality average partial effects are jointly zero		0.003	0.004		0.001	0.000		0.000	0.000

Note: Fractional response logistic regression average partial effects with standard errors clustered at the village level in parentheses. Proportions are based on a random sample of ~ 25 bean farming households in the farmer's village. Trial is an indicator variable equal to one if the farmer resides in a village where the VBAA received trial packs. Personality traits are measured on a scale of one (not at all) to four (a lot).

* $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$. \a VBAA personality controls are the same as that of the farmer. Distance controls are the GPS-based physical distances (km) from the farmer's homestead to the VBAA's homestead and VBAA demonstration plot. \b Demographic controls (2017 ag. year) are: farmer household's size; indicators for farmer or spouse related to village chairmen; farmer's and VBAA's gender and age; indicator for farmer household's and VBAA's highest education achieved is primary or less.

Table 3.13: Heterogeneous Trial Pack Effects on the Proportion of Bean Farming Households that the Farmer has Sought Farming Advice From

Variable	(1)	(2)
Trial	-0.131 (0.140)	-0.130 (0.142)
Trial*Agency	-0.073* (0.039)	-0.069* (0.038)
Trial*Agreeableness	-0.004 (0.039)	-0.015 (0.038)
Trial*Openness to experience	-0.020 (0.048)	0.005 (0.045)
Trial*Neuroticism	0.060** (0.027)	0.060** (0.025)
Trial*Extraversion	0.122* (0.063)	0.110* (0.057)
Trial*Conscientiousness	-0.029 (0.044)	-0.030 (0.044)
VBAA-pair indicators	Yes	Yes
Farmer and VBAA personality and distance controls \a	Yes	Yes
Farmer and VBAA demographic controls \b	No	Yes
Observations	740	740
R-squared	0.195	0.254

Note: Linear probability model coefficients with standard errors clustered at the village level in parentheses. Dependent variable is the proportion of a random sample of ~ 25 bean farming households in the farmer's village that he/she has sought farming advice from. Trial is an indicator variable equal to one if the farmer resides in a village where the VBAA received trial packs. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

\a Farmer and VBAA personality controls are their six personality traits measured on a scale of one (not at all) to four (a lot). Distance controls are the GPS-based physical distances (km) from the farmer's homestead to the VBAA's homestead and VBAA demonstration plot.

\b Demographic controls (2017 ag. year) are: farmer household's size; indicators for farmer or spouse related to village chairmen; farmer's and VBAA's gender and age; indicator for farmer household's and VBAA's highest education achieved is primary or less.

Table 3.14: Heterogeneous Trial Pack Effects on the Proportion of Bean Farming Households that the Farmer has Discussed Bean Farming With

Variable	(1)	(2)
Trial	-0.206 (0.199)	-0.204 (0.205)
Trial*Agency	-0.061 (0.042)	-0.055 (0.043)
Trial*Agreeableness	-0.037 (0.044)	-0.047 (0.046)
Trial*Openness to experience	-0.021 (0.044)	-0.005 (0.042)
Trial*Neuroticism	0.054 (0.036)	0.056 (0.036)
Trial*Extraversion	0.157*** (0.057)	0.145** (0.054)
Trial*Conscientiousness	-0.011 (0.042)	-0.009 (0.043)
VBAA-pair indicators	Yes	Yes
Farmer and VBAA personality and distance controls \a	Yes	Yes
Farmer and VBAA demographic controls \b	No	Yes
Observations	740	740
R-squared	0.213	0.257

Note: Linear probability model coefficients with standard errors clustered at the village level in parentheses. Dependent variable is the proportion of a random sample of ~ 25 bean farming households in the farmer's village that he/she has discussed bean farming with. Trial is an indicator variable equal to one if the farmer resides in a village where the VBAA received trial packs. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

\a Farmer and VBAA personality controls are their six personality traits measured on a scale of one (not at all) to four (a lot). Distance controls are the GPS-based physical distances (km) from the farmer's homestead to the VBAA's homestead and VBAA demonstration plot.

\b Demographic controls (2017 ag. year) are: farmer household's size; indicators for farmer or spouse related to village chairmen; farmer's and VBAA's gender and age; indicator for farmer household's and VBAA's highest education achieved is primary or less.

APPENDIX B:

Supplementary Tables

Table 3.15: Farmer Adoption of Improved Bean Varieties in 2018 Ag. Year: Controlling for Pre-Intervention Outcome

Variable	(1)	(2)	(3)
Trial	0.007 (0.032)	0.048 (0.033)	0.098** (0.046)
Agency		0.029 (0.038)	0.033 (0.038)
Agreeableness		0.017 (0.038)	0.019 (0.040)
Openness to experience		-0.045 (0.051)	-0.050 (0.054)
Neuroticism		0.010 (0.021)	0.012 (0.019)
Extraversion		0.065 (0.041)	0.065 (0.040)
Conscientiousness		-0.018 (0.042)	-0.021 (0.044)
VBAA-pair indicators	Yes	Yes	Yes
VBAA personality and distance controls \a	No	Yes	Yes
Farmer and VBAA demographic controls \b	No	No	Yes
Observations	740	740	740
Pseudo R-squared	0.171	0.212	0.221
P-value of test that farmer personality average partial effects are jointly zero		0.132	0.0797

Note: Logit average partial effects with standard errors clustered at the village level in parentheses. Dependent variable is equal to one if the farmer's household adopted Njano Uyole or Uyole 96 in the 2018 agricultural year. All specifications control for adoption of Njano Uyole or Uyole 96 in the 2016 agricultural year. Trial is an indicator variable equal to one if the farmer resides in a village where the VBAA received trial packs. Personality traits are measured on a scale of one (not at all) to four (a lot). * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

\a VBAA personality controls are the same as that of the farmer. Distance controls are the GPS-based physical distances (km) from the farmer's homestead to the VBAA's homestead and VBAA demonstration plot.

\b Demographic controls (2017 ag. year) are: farmer household's size; indicators for farmer or spouse related to village chairmen; farmer's and VBAA's gender and age; indicator for farmer household's and VBAA's highest education achieved is primary or less.

Table 3.16: Heterogeneous Trial Pack Effects on Farmer Adoption of Improved Bean Varieties in 2018 Ag. Year: Controlling for Pre-Intervention Outcome

Variable	(1)	(2)
Trial	0.239 (0.219)	0.313 (0.236)
Trial*Agency	-0.047 (0.073)	-0.048 (0.075)
Trial*Agreeableness	-0.075 (0.080)	-0.071 (0.085)
Trial*Openness to experience	-0.002 (0.108)	-0.006 (0.113)
Trial*Neuroticism	0.008 (0.041)	0.000 (0.042)
Trial*Extraversion	0.154** (0.072)	0.151** (0.072)
Trial*Conscientiousness	-0.091 (0.077)	-0.093 (0.077)
VBAA-pair indicators	Yes	Yes
Farmer and VBAA personality and distance controls \a	Yes	Yes
Farmer and VBAA demographic controls \b	No	Yes
Observations	740	740
R-squared	0.220	0.228

Note: Linear probability model coefficients with standard errors clustered at the village level in parentheses. Dependent variable is equal to one if the farmer's household adopted Njano Uyole or Uyole 96 in the 2018 agricultural year. All specifications control for adoption of Njano Uyole or Uyole 96 in the 2016 agricultural year. Trial is an indicator variable equal to one if the farmer resides in a village where the VBAA received trial packs. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

\a Farmer and VBAA personality controls are their six personality traits measured on a scale of one (not at all) to four (a lot). Distance controls are the GPS-based physical distances (km) from the farmer's homestead to the VBAA's homestead and VBAA demonstration plot.

\b Demographic controls (2017 ag. year) are: farmer household's size; indicators for farmer or spouse related to village chairmen; farmer's and VBAA's gender and age; indicator for farmer household's and VBAA's highest education achieved is primary or less.

REFERENCES

REFERENCES

- Ali, D., Bowen, D., & Deininger, K. (2017). *Personality Traits, Technology Adoption, and Technical Efficiency: Evidence from Smallholder Rice Farms in Ghana* (Policy Research Working Paper No. 7959). World Bank.
- Anderson, J. R., & Feder, G. (2004). Agricultural Extension: Good Intentions and Hard Realities. *The World Bank Research Observer*, 19(1), 41–60.
- Bazzani, C., Caputo, V., Nayga, R. M., & Canavari, M. (2017). Revisiting consumers’ valuation for local versus organic food using a non-hypothetical choice experiment: Does personality matter? *Food Quality and Preference*, 62, 144–154.
- BenYishay, A., & Mobarak, A. M. (2019). Social Learning and Incentives for Experimentation and Communication. *The Review of Economic Studies*, 86(3), 976–1009.
- Brandstätter, H., Brandstätter, V., & Pelka, R. B. (2018). Similarity and Positivity of Personality Profiles Consistently Predict Relationship Satisfaction in Dyads. *Frontiers in Psychology*, 9.
- Foster, A. D., & Rosenzweig, M. R. (2010). Microeconomics of Technology Adoption. *Annual Review of Economics*, 2(1), 395–424.
- Grebitus, C., Lusk, J. L., & Nayga, R. M. (2013). Explaining differences in real and hypothetical experimental auctions and choice experiments with personality. *Journal of Economic Psychology*, 36, 11–26.
- Jack, K. (2013). *Market inefficiencies and the adoption of agricultural technologies in developing countries* (Working Paper).
- Krishnan, P., & Patnam, M. (2014). Neighbors and Extension Agents in Ethiopia: Who Matters More for Technology Adoption? *American Journal of Agricultural Economics*, 96(1), 308–327.
- Laaajaj, R., Macours, K., Pinzon Hernandez, D. A., Arias, O., Gosling, S. D., Potter, J., Rubio-Codina, M., & Vakis, R. (2019). Challenges to capture the big five personality traits in non-WEIRD populations. *Science Advances*, 5(7).
- Lachman, M., & Weaver, S. (1997). *The Midlife Development Inventory (MIDI) Personality Scales: Scale Construction and Scoring*. Brandeis University. Waltham, MA.

- Lachman, M., & Weaver, S. (2005). *Addendum for MIDI Personality Scales*. Brandeis University. Waltham, MA.
- Lin, W., Ortega, D. L., Caputo, V., & Lusk, J. L. (2019). Personality traits and consumer acceptance of controversial food technology: A cross-country investigation of genetically modified animal products. *Food Quality and Preference*, 76, 10–19.
- Macho-Stadler, I., & Pérez-Castrillo, J. (2001). *An Introduction to the Economics of Information: Incentives and Contracts* (2nd ed.). Oxford University Press.
- McKenzie, D. (2012). Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics*, 99(2), 210–221.
- Melkani, A., & Mason, N. M. (2018). *Tanzania Southern Highlands Farm Input Promotions-Africa Village-Based Agricultural Advisors Baseline Survey Report*. Michigan State University.
- Morgan, S. N. (2018). *The Experimental Science of Economic Behavior: Testing Theories of Participation, Valuation, and Innovation*. Michigan State University.
- Nakano, Y., Tsusaka, T. W., Aida, T., & Pede, V. O. (2018). Is farmer-to-farmer extension effective? The impact of training on technology adoption and rice farming productivity in Tanzania. *World Development*, 105, 336–351.
- Norton, G. W., & Alwang, J. (2020). Changes in Agricultural Extension and Implications for Farmer Adoption of New Practices. *Applied Economic Perspectives and Policy*, aepp.13008.
- Ozer, D. J., & Benet-Martínez, V. (2006). Personality and the Prediction of Consequential Outcomes. *Annual Review of Psychology*, 57(1), 401–421.
- Papke, L. E., & Wooldridge, J. M. (1996). Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics*, 11, 619–632.
- Roberts, B. W. (2009). Back to the future: Personality and Assessment and personality development. *Journal of Research in Personality*, 43(2), 137–145.
- Soto, C. J. (2019). How Replicable Are Links Between Personality Traits and Consequential Life Outcomes? The Life Outcomes of Personality Replication Project. *Psychological Science*, 1–17.

- Van der Linden, D., Dunkel, C. S., Figueredo, A. J., Gurven, M., Von Rueden, C., & Woodley of Menie, M. A. (2018). How Universal Is the General Factor of Personality? An Analysis of the Big Five in Forager Farmers of the Bolivian Amazon. *Journal of Cross-Cultural Psychology*, *49*(7), 1081–1097.
- Weidmann, R., Schönbrodt, F. D., Ledermann, T., & Grob, A. (2017). Concurrent and longitudinal dyadic polynomial regression analyses of Big Five traits and relationship satisfaction: Does similarity matter? *Journal of Research in Personality*, *70*, 6–15.
- Wilson, K. S., DeRue, D. S., Matta, F. K., Howe, M., & Conlon, D. E. (2016). Personality Similarity in Negotiations: Testing the Dyadic Effects of Similarity in Interpersonal Traits and the Use of Emotional Displays on Negotiation Outcomes. *Journal of Applied Psychology*, *101*(10), 1405–1421.
- Wooldridge, J. (2010). *Econometric Analysis of Cross Section and Panel Data* (2nd ed.). Cambridge, MA, MIT Press.