



PhD thesis

Anna Folke Larsen

Agricultural technology adoption, food security, poverty and child health

Assessments of an agricultural intervention in Tanzania



Academic advisor: Mette Ejrnæs

Submitted: 07/05/15

Acknowledgments

This PhD dissertation could not have been written if I had not entered into a fruitful collaboration with the Rockwool Foundation and its Research Unit. The Rockwool Foundation has funded the RIPAT project which I study in this dissertation along with the data collection upon which I build all empirical analyses. I am very grateful for their financial support for the data collection. I would like to thank all staff at RE-CODA who worked hard to ensure a high quality data collection, Maria Fibæk for supporting me during field work, Cathrine Søgaard Hansen for research assistance and Jens Vesterager for agricultural technical guidance.

In particular, I owe a great thanks to Helene Bie Lilleør, Head of Evaluation at the Rockwool Foundation Research Unit, who is coauthoring two of my PhD chapters. She has guided me through many difficult decisions in the process of data collection, data cleaning and the empirical analyses. Besides offering good support, she has also been great fun to work with.

I owe a special thanks to my supervisor, Mette Ejrnæs, whose door has always been open to me whether I needed to discuss the greater picture of my dissertation or I was struggling with the calculation of sampling weights and such details. She has trained me in understanding what constitutes a good identification strategy which is maybe the most important thing I take with me from my time as a PhD student.

I would like to thank Abhijit Banerjee for hosting me at Massachusetts Institute of Technology during the fall 2012. The visit provided me with a great insight into the research frontier of Development Microeconomics and fostered new research ideas. I acknowledge travel grants from the Augustinus Foundation, the Oticon Foundation and the Foundation of Knud Højgaard.

Writing this dissertation would have been a lot more boring had I not enjoyed the great company of my fellow PhD colleagues on many Coffee Buses over the years. Thanks guys!

Stepping seven years back in time, I would like to thank Claus Thustrup Kreiner for his role as a mentor very early in my academic career as his research assistant. He was

probably the first to articulate that I should pursue a PhD. I thank Finn Tarp for the time as his research assistant where I had the opportunity to write my first academic article together with Nina Torm and John Rand. This experience gave me a first insight in how to conduct academic work. I would also like to thank Bertel Schjerning for further encouraging me to apply for a PhD scholarship which gave me confidence that this was the right path for me.

Last but not least: Thank you so much, Jesper Neumann, for your great love and support throughout the entire process and for our continuing methodological discussions which keep me on my toes.

Summary

The three self-contained chapters of this dissertation evolves around different aspects of a Farmer Field School intervention taking place in northern Tanzania, studying the diffusion of agricultural technologies and how agriculture links to food security, poverty and child health. The intervention is called RIPAT (Rural Initiatives for Participatory Agricultural Transformation) and was funded by the Rockwool Foundation. As a consultant for the Rockwool Foundation I administered a large scale data collection for an impact evaluation of RIPAT.

All three chapters of my dissertation build on these data. The first two chapters are coauthored with Helene Bie Lilleør. A recurrent theme in this dissertation is the identification of causal effects. Since participation in RIPAT is voluntary, the data does not offer direct experimental variation which I can exploit for identification, and there exists no baseline data collected before the implementation of RIPAT I to control for selection. I pursue different identification strategies which are detailed in the three chapters. In the following, I provide a preview of the findings.

Food security and poverty: The first chapter provides a broad impact evaluation of RIPAT with a focus on food security and poverty which were the development objectives stated in the original project documentation. Despite the strong potential of agricultural interventions to affect food security and the poverty status of small scale farmers, existing studies of Farmer Field School interventions focus on short term outcomes such as knowledge, take-up, and agricultural yields. We employ four different estimators to identify the impact of RIPAT on food security and poverty, and they all yield consistent results: The RIPAT households become more food secure in particular in the hungry season, but we do not detect any impacts on our poverty indicators. These finding can be explained by a reallocation of labor resources toward own agricultural production and improved production smoothing which may have improved food security while leaving poverty unaffected.

Child health: The second chapter studies the impact of RIPAT on the health of children below five years of age living in the RIPAT households. Agricultural production

is an underlying determinant of child nutrition, and the improvement we found in food security among the RIPAT households has the potential to materialize in better child nutrition and thereby taller children. We find that young children conceived after the phase-in of RIPAT have become 0.8 standard deviations taller and are 17.6 percentage points less likely to be stunted (i.e. severely reduced height-for-age). The large impacts may be explained by the fact that the area was hit by a drought in 2009. The RIPAT project is designed to make the agricultural production more resistant to droughts, and the results show that the children in RIPAT households were indeed better shielded against the adverse consequences of the drought than children from comparison households regardless of whether we compare to comparison households in nearby villages or to non-RIPAT households within RIPAT villages. We investigate if the results could be driven by self-selection into RIPAT based on better drought coping capabilities, but we find no evidence in support of this.

Adoption of technologies: The main component of the RIPAT intervention was the introduction of improved banana cultivation. In the third chapter I assess how the adoption of improved banana cultivation among non-RIPAT farmers in RIPAT villages depends on their links to RIPAT participants who grow improved bananas. In the existing literature on networks and technology adoption, network effects are interpreted as social learning. I show that a farmer's network can affect the adoption of a new crop not only through social learning, but also by providing necessary inputs for adoption. I set up a simple model for adoption of a new crop where the farmer's network can provide both information about the expected yield of the new crop and necessary inputs for adoption. I derive the same model implications for how the network affects adoption regardless of whether the network provides inputs or information. Empirically, I find that a farmer is 39 percentage points more likely to adopt banana cultivation if there is at least one farmer growing improved bananas in the farmer's network. The data suggests that the provision of inputs (banana seedlings) through networks plays an important role for the strong network effects found. This is particularly important for the diffusion of new agricultural technologies in areas that suffer from poor infras-

structure which impedes the distribution of agricultural inputs.

Resumé (Danish summary)

De tre selvstændige kapitler i denne afhandling behandler forskellige aspekter af en Farmer Field School intervention i det nordlige Tanzania. Jeg studerer hvordan en ny landbrugsteknologi spredt sig fra bonde til bonde, og hvordan landbrugsprojektet påvirker fødevarsikkerhed, fattigdom og børns sundhed. Interventionen hedder RIPAT (Rural Initiatives for Participatory Agricultural Transformation) og blev finansieret af Rockwool Fonden. Som konsulent for Rockwool Fonden har jeg administreret en stor dataindsamling, som havde det formål at evaluere RIPAT.

Alle tre kapitler i min afhandling bygger på disse data. De to første kapitler har jeg skrevet sammen med Helene Bie Lilleør. Et gennemgående tema i denne afhandling er identifikation af kausale effekter. Da deltagelsen i RIPAT er frivillig, indeholder dataene ikke nogen direkte eksperimentel variation, som jeg kan udnytte til identifikation, og der er ikke blevet indsamlet baseline data før gennemførelsen af RIPAT, som ellers kunne bruges til at kontrollere for selektionen ind i projektet. Jeg udnytter forskellige strategier til at identificere kausale effekter, som jeg beskriver i detaljer i de tre kapitler. Her giver jeg et samlet overblik over resultaterne.

Fødevarsikkerhed og fattigdom: Det første kapitel indeholder en bred effektevaluering af RIPAT med fokus på fødevarsikkerhed og fattigdom, som var de udviklingsmål, der var anført i de oprindelige projektdokumenter. På trods af det store potentiale landbrugsinterventioner har til at påvirke sult og fattigdom blandt små landmænd, fokuserer eksisterende studier af Farmer Field School interventioner på kortsigtede resultater, såsom viden, anvendelse af nye teknologier, samt landbrugsudbytte. Vi anvender fire forskellige estimators for at identificere effekten af RIPAT på sult og fattigdom, og de giver alle ensartede resultater: RIPAT-husholdningerne sulter mindre især i sulte-sæsonen umiddelbart før næste høst, men vi kan ikke påvise nogen effekter på vores fattigdomsindikatorer. Disse umiddelbart overraskende resultater kan potentielt

forklares ud fra en ændring i hvordan husholdningerne anvender deres arbejdskraft, samt en udjævning af landbrugsproduktion hen over landbrugssæsonen.

Børns sundhed: I det andet kapitel undersøger vi effekten af RIPAT på ernæring for de børn som er under fem år gamle og bor i en RIPAT-husholdning. Landbrugsproduktionen er en underliggende faktor som påvirker børneernæring, og den forbedring vi fandt i fødevarerikthed blandt RIPAT-husholdningerne har potentiale til at materialisere sig i bedre børneernæring og dermed højere børn. Vi finder, at børn, der er undfanget efter indfasningen af RIPAT er blevet 0,8 standardafvigelser højere og er 17,6 procentpoint mindre tilbøjelige til at blive for lave i forhold til deres alder. De store effekter kan forklares ved, at området blev ramt af en tørke i 2009. RIPAT-projektet er designet til at gøre landbrugsproduktionen mere modstandsdygtig over for tørke, og resultaterne viser, at børnene fra RIPAT-husholdninger faktisk var bedre beskyttet mod de negative konsekvenser af tørken end børn fra sammenlignings-husholdninger. Vi undersøger, om resultaterne kunne være drevet af, at husholdninger, der i forvejen var bedre rustet til at klare sig gennem en tørke, valgte at deltage i RIPAT, men vi finder ikke noget tegn på, at det skulle være tilfældet.

Anvendelse af en ny teknologi: Det vigtigste element i RIPAT var introduktionen af en ny og bedre måde at dyrke bananer på. I det tredje kapitel anvender jeg data fra ikke-RIPAT bønder i RIPAT landsbyer for at undersøge, hvem der vælger at dyrke bananer afhængigt af, om de kender RIPAT deltagere som dyrker bananer. I den eksisterende litteratur om netværk og teknologianvendelse fortolkes netværkseffekter som tegn på, at der videregives information i netværket. Jeg viser, at en bondes netværk kan påvirke beslutningen om at dyrke bananer ikke kun gennem videregivelse af information, men også ved at videregive et nødvendigt input til at dyrke bananer, nemlig forbedrede bananstiklinger. Jeg opstiller en simpel model for beslutningen om at dyrke en ny afgrøde, hvor bondens netværk både kan videregive et nødvendigt input for at dyrke den nye afgrøde og oplysninger om det forventede udbytte af den nye afgrøde. Uanset om jeg antager at netværket videregiver input eller information, kan jeg udlede de samme forudsigelser fra modellen om, hvordan netværket påvirker

beslutningen om at dyrke den nye afgrøde. Min dataanalyse viser, at en bonde er 39 procentpoint mere tilbøjelig til at dyrke bananer, hvis vedkomne kender mindst én bonde som dyrker bananer. Det fremgår desuden af data at videregivelse af bananstiklinger gennem bøndernes netværk spiller en vigtig rolle for de stærke netværkseffekter jeg finder. Dette er især vigtigt for udbredelsen af nye landbrugsteknologier i områder med dårlig infrastruktur, hvor distribution af nye landbrugsinput enten er meget dyr eller ikke-eksisterende.

Contents

I Beyond the field: The impact of Farmer Field Schools on food security and poverty alleviation	1
II Can agricultural interventions improve child health? Evidence from Tanzania	45
III The network at work: Diffusion of banana cultivation in Tanzania	98

Beyond the field: The impact of Farmer Field Schools on food security and poverty alleviation

Anna Folke Larsen and Helene Bie Lilleør*

Abstract

We estimate the impact of a Farmer Field School intervention among small-scale farmers in northern Tanzania on two main development objectives: food security and poverty. We employ a series of evaluation methodologies, including a Quasi-Difference-in-Difference setup, to account for potential selection into the project, despite lack of baseline data. We find strong positive effects on food security, but no effect on poverty. Investigating possible mechanisms for this result shows that reallocation of labor resources toward own agricultural production and improved production smoothing may have led to improved food security while poverty remained unaffected.

1 Introduction

The majority of poor households in developing countries rely on subsistence agriculture for their own food production and as a source of income. Over the past few decades, various initiatives have been taken aimed at increasing food production by

*Anna Folke Larsen: Department of Economics, University of Copenhagen, afl@econ.ku.dk. Helene Bie Lilleør: Rockwool Foundation Research Unit, hbl@rff.dk. Published in 2014 in *World Development*, 64: 843-859. Final revision accepted: July 9, 2014. All the views expressed in this paper are those of the authors and do not necessarily reflect the views of the Rockwool Foundation. The Rockwool Foundation funded the RIPAT implementation and the data collection for the evaluation. We are grateful for comments and advice from RECODA, Jens Vesterager, Mette Ejrnæs, Torben Tranæs, Cathrine Søgaaard Hansen, Maria Fibæk, as well as from participants in the CAM Christmas workshop 2011. All remaining errors are of course our own.

closing the technology gap faced by subsistence farmers. Such initiatives have worked either directly, through the supply of new technologies such as fertilizer, seeds of improved plant varieties or new animal breeds, or more indirectly, through agricultural extension and advisory services, or both (Anderson & Feder 2007, Lunduka et al. 2013, Rawlins et al. 2014) .

Agricultural extension has long been seen as a key element in improving agricultural development. However, the effectiveness of two dominant approaches to agricultural extension services in particular - Training and Visit(T&V)¹ and Farmer Field Schools (FFS)² - has been widely debated. The T&V approach relies on the “top-down” extension of technical information, with specialists and field staff transferring knowledge to “contact farmers” in villages, who in turn are responsible for diffusing knowledge into the local community. As a response to this top-down approach, FFS were developed as a “bottom-up” approach to extension with a focus on participatory, experiential, and reflective learning to improve the problem-solving capacity of farmers through highly trained facilitators working with farmer groups (Anderson & Feder, 2007). In this paper, we assess the impact on food security and poverty of an intervention which seeks to combine both the top-down and bottom-up approaches and which has been implemented among smallholders in northern Tanzania. The intervention, locally known as RIPAT (Rural Initiatives for Participatory Agricultural Transformation), is designed as a modified FFS approach taking its starting point in farmer groups and experiential learning, but with a strong element of traditional technology transfer through the introduction of a “basket” of new technology options. We find that RIPAT has had a large impact on food security, but no detectable impact on poverty.

FFS have been implemented and adopted worldwide (Braun et al., 2006). Nonetheless, the ability of the approach to ensure both sustained technology adoption and increased productivity is still subject to an ongoing debate about appropriate evalua-

¹This was primarily promoted by the World Bank in the 1970s and 1980s and developed to tackle some of the inefficiencies present at the time in traditional public extension services.

²The Farmer Field School concept was originally developed by the FAO to promote integrated pest management among Indonesian rice farmers in the late 1980s, but since then has spread to many countries and over the years has been so widely adopted and locally adapted that there is no longer a single model for either its technical content or the educational format(van den Berg & Jiggins, 2007).

tion methodologies, when to evaluate, and choice of outcome measures (Feder et al., 2008; van den Berg & Jiggins, 2008; Davis & Nkonya, 2008; Mancini & Jiggins, 2008; Feder et al., 2010; Braun & Duveskog, 2011). More recently, a thorough survey of FFS impact studies was provided by Davis et al. (2012: Table 1), highlighting the fact that the outcomes selected for examination are very mixed, as are the findings. While some papers find positive impacts on adoption, agricultural yields, productivity and agricultural income, others do not. Most papers studying the impact on various aspects of empowerment find that empowerment increases, which has led to an argument being advanced that FFS is more a model of adult learning than of agricultural extension (Van den Berg & Jiggins 2007, Friis-Hansen & Duveskog 2012) .

The debate in the FFS evaluation literature was initially sparked by Feder et al. (2004) criticizing earlier FFS evaluation methodologies for not taking the potential positive bias of non-random program placement and selection of participants into account in their assessments of impact. This led to discussions of evaluation timing and problems of spillover effects. Measuring outcomes using a relatively long time horizon, as Feder et al. (2004) do, allows for an assessment of impact sustainability - unless the estimated impact is confounded by spillovers from FFS graduates to control farmers living nearby, as suggested by Van den Berg & Jiggins (2007, 2008), but proven by Yamazaki & Resosudarmo (2008) not to be the case using the same data as Feder et al (2004).

The best way to obtain an unbiased estimate of impact would be to conduct a randomized controlled trial, but to our knowledge, this has not been done for FFS yet. Given non-random program placement, a few papers, including Godtland et al (2004), Rejesus et al. (2012), Davis et al. (2012) and Todo & Takahashi (2013), do attempt to take this selection factor into account in a careful manner. However, all four of these studies suffer from having relatively small sample sizes (ranging from 142 to 486 within each country), which may have resulted in no significant impact being found simply due to lack of statistical power, and from operating with a very short time horizon (one to two years since project start). They therefore have to assess the impact on outcomes

that are very closely related to project activities, such as knowledge transfer, technology adoption, yields or agricultural income.³ Again findings are mixed, though with some indications of improved technology knowledge transfer and adoption leading to higher yields and thus to increased agricultural income.

While it is of value to assess the impact of FFS on farmers' knowledge, technology transfer, take-up and agricultural production, it should be kept in mind that households may simply divert resources away from other activities towards the new project-related activities. It is therefore also important to analyze the impact on broader welfare indicators for the participating households. Although it has become popular to assess empowerment, it is not in itself a welfare measure; rather, it can be a channel through which people may obtain improved welfare. We have not found any studies within the conventional peer-reviewed literature that analyze the impact of FFS on broader welfare factors such as food security or poverty.

This paper is intended to contribute to filling this gap in the literature by presenting a rigorous impact evaluation of RIPAT FFS to examine whether the program improved food security and reduced poverty among participating households. In our evaluation, we have sought to address the main points raised in the FFS evaluation debate summarized above.

We let the original project documentation guide us in the choice of outcome measures. It was explicitly stated that the overall development objectives of the intervention were to increase food security and alleviate poverty among participating households. Any effect on these outcome measures can only be expected to be observable in the medium or long term, as participating households have to first adopt and then implement the new technologies throughout a full agricultural cycle before impacts on food security and poverty can occur. By developing our evaluation strategy and the associated survey instrument accordingly, we have effectively tied the analysis - and our hands - to these outcome measures, and thereby reduced the possibilities of

³Although Davis et al. (2012) state in the title of their paper that they analyze the impact of FFS on agricultural productivity and poverty, they in fact analyze the impact on crop income and livestock income, which they sum as agricultural income.

“cherry-picking” convenient results. However, we did not have a full pre-analysis plan laid out, as suggested by Casey et al. (2012).

In explaining our choice of impact assessment methodologies, we discuss the extent to which we can overcome the potential endogeneity issues noted by Feder et al. (2004) and Godtland et al. (2004) that stem from non-random program placement and self-selection of participants. To address these issues we collected household data from two different areas: Arumeru district, where RIPAT I was implemented, and Karatu district, where RIPAT II was started two years later. In both areas we collected data from virtually all RIPAT households and from a sample of control households in nearby villages. In addition, we also collected data from non-RIPAT households in RIPAT I villages. We employ four different methodologies to assess the impact of RIPAT I: a simple cross-sectional comparison of RIPAT I and control households in a multivariate setting to control for observable characteristics; an intention to treat estimation, in which we include non-RIPAT households within RIPAT I villages, to circumvent the problem of self-selection at the household level; a matching estimation to increase comparability of observable characteristics between RIPAT I and control households and villages;⁴ and finally a Quasi Difference-in-Difference estimation exploiting data from the later RIPAT II households and their controls to account for selection. Under the assumption that the household- and village-level selection mechanisms in the two districts were the same, the Quasi Difference-in-Difference takes selection for both observable and unobservable characteristics into account, i.e. we circumvent the endogeneity problems of non-random program placement and self-selection of participants. To the extent that there was already some initial impact among RIPAT II farmers on food security and poverty indicators at the time of the data collection in 2011, which was more than one year after RIPAT I completion and half way through the RIPAT II project period, our impact assessment will be a conservative estimate of the true impact. We thereby avoid the problem of positive selection bias. Throughout the paper, the impact assessment is an assessment of RIPAT I only, unless explicitly stated otherwise.

⁴We thank an anonymous referee for suggesting intention to treat and matching estimation.

To address the potential problem of timing and spillover to control farmers diluting the impact of the intervention, as described by Van den Berg & Jiggins (2008), we use control farmers living at a sufficient distance from the RIPAT intervention villages. Although there had been spillover within RIPAT I villages at the time of data collection, qualitative findings confirm that we do not have to worry about any potential spillover in food security and poverty from RIPAT I at the distances used.⁵ In addition, by assessing the impact of RIPAT I almost five years after project start and more than one year after completion, we are also able to address issues of sustainability, at least in the medium term.

Our analyses are based on interviews with 2,041 farming households using a highly-structured closed-form questionnaire administered in 36 villages, of which 16 were intervention villages. We thus have a large sample size compared to previous FFS impact evaluations.⁶

The vast majority of participants in RIPAT Farmer Field Schools were involved in the project throughout the full project period. We see that half-way through the project period in RIPAT II and one year after project completion in RIPAT I the participating households were more likely to have adopted virtually all the key technologies promoted through the basket of options than farmers in the control villages. This indicates both immediate and sustained adoption of the new technologies. We find that the participating households were more likely to be cultivating improved varieties of banana, to have a larger degree of crop diversification, to be keeping improved breeds of livestock and to be members of savings groups.

Most importantly, we find that these high levels of engagement and technology take-up resulted in considerable improvements in food security levels, suggesting an increase in overall household welfare. In this medium term, i.e. five years after project start, we find that RIPAT I households were up to 24 percentage points less likely to experience hunger, that their diet contained more animal proteins, and that their children

⁵In RIPAT II there were no reports of spillover within intervention villages in 2011, let alone to the control villages.

⁶We only know of one study, by Davis et al. (op.cit.), with a similar sample size (1,126 households). However, this is spread across 8-10 districts and three countries.

were more likely to have at least three meals per day. These are substantial impacts, which we believe will be sustainable in the longer term, given the timing of the evaluation.

We do not find any significant impact of RIPAT I on any of our poverty indicators occurring by 2011. This suggests that RIPAT I households might have had a more urgent need to overcome food insecurity than to invest in the more material goods that are typically used as poverty indicators, e.g. good floors or mobile phones. We analyze two possible mechanisms that might have led to our results: reallocation of labor resources towards own agricultural production and production smoothing over the agricultural cycle. We find indications of both.

We have organized the paper as follows. Section 2 describes the RIPAT intervention, and Section 3 presents the data: summary statistics for household and village characteristics, participants, adoption of technologies, and choice of outcome measures. In Section 4 we explain our evaluation strategy, while we turn to the results in Section 5. We analyze the role of labor reallocation and production smoothing in the findings in Section 6. Section 7 concludes.

2 The intervention

In this paper we evaluate the agricultural project RIPAT I, but for some elements of the evaluation strategy we also exploit data from a later project, RIPAT II. Both projects were implemented by a local Tanzanian NGO, RECODA. They targeted small and medium-sized farmers in rural villages with at least one acre and in principle no more than five acres of land. Village leaders were asked to form two groups of 30-35 farmers in each village and to assist the groups in getting access to a joint group field. Membership was voluntary. RECODA explained to village leaders that members should not be rich in terms of the wealth ranking of the village, had to be committed to active participation (attendance records were kept and strict rules enforced), had to be willing to share their knowledge with and demonstrate agricultural techniques to their fellow

villagers and should therefore be of good standing in the community, and had to live within the village administrative area. Furthermore, the leaders were told that each group should have an equal number of men and women and only one member per household (Maguzu et al., 2013). The two RIPAT projects were each implemented in eight villages; these were selected by district officials as being the poorest villages in the given district.

There were thus two sources of endogenous selection into the project. One was endogenous village selection: since program placement was not random and if district officials followed the guidelines given to them, RIPAT villages were less wealthy than the other villages in the district at the outset of the project, i.e. there was a negative selection effect. Secondly, since participation was voluntary, households self-selected into the project (provided they met the targeting criteria) and hence we would expect participating RIPAT households to have been more motivated than other households, resulting in a positive selection effect. The sign of the *net* selection effect thus cannot be assumed *a priori*.

The RIPAT Farmer Field Schools draw on a bottom-up experiential and reflective approach to learning and practical demonstrations of farming techniques, as do most FFS. However, they are described as less participatory and more top-down than other FFS approaches (Aben et al., 2013). A key difference is a strong element of traditional technology transfer through training in a predetermined but locally adapted “basket of technology options”, rather than in just one technology. These agricultural technology options are chosen by the implementing NGO on the basis of their strong agricultural expertise and in prior consultation with the villages in question. By equipping the farmers with the necessary information, knowledge and hands-on experience in the use of different relevant and efficient technologies, the program provides each farmer with the means to choose which technologies to adopt in his or her own agricultural production. Each group meets weekly at its demonstration plot or group field. At these meetings, progress is followed and discussed throughout the agricultural cycles. The crops and technologies introduced in the “basket of options” are very diverse and

cannot all be fully introduced in a single agricultural cycle as in typical FFS programs; the implementing NGO therefore works with the RIPAT FFS for a three-year period, after which the farmers “graduate”. The standard “basket” includes improved varieties of banana with new cultivation techniques, conservation agriculture and crop diversification, improved animal husbandry, fruit and multipurpose trees, soil and water conservation, post-harvesting technologies, and encouragement to participate in savings groups. However, the basket is always adapted to suit local conditions, taking into account, for example, soil, water and climate.⁷

The two RIPAT FFS projects commenced two years and four months apart in two districts in the Arusha Region. The implementation of RIPAT I in Arumeru District was from May 2006 until the end of 2009, while RIPAT II was implemented in Karatu District from September 2008 until August 2012 (see Figure 1). The implementation strategies for the two projects were the same except for minor adjustments to the content of the basket of options.⁸ We exploit this gradual roll-out in one of our empirical strategies below to address the potential problems caused by self-selection of participating farmers and non-random program placement at village level due to unobservable factors.

3 Data and summary statistics

In January 2011 we conducted a large-scale quantitative household survey in both RIPAT and control villages in the two intervention districts. This was one year after completion of RIPAT I and around halfway through implementation of RIPAT II. We used a highly structured closed-form pilot-tested questionnaire to capture the extent to which participating farmers had adopted the technologies introduced through the RIPAT farmer groups and to discover whether this in turn had had an impact on their

⁷For more detailed descriptions, see (Maguzu et al., 2013) and (Vesterager et al., 2013) for shorter and longer accounts, respectively.

⁸Savings group participation was encouraged but not facilitated during the RIPAT I project. Furthermore, during RIPAT I it became clear that a more efficient distribution system for the improved breeds of goats would be needed in future projects. Finally, in Karatu there was an additional demand for an improved breed of pigs, which was then also included in the basket of options.

food security and poverty levels relative to farmers in control villages. A selection of non-RIPAT households were also surveyed in the RIPAT I villages in order to gather data for separate diffusion analyses of the more popular technologies.⁹ We interviewed a total of 2,374 households in 36 villages; of these, 2,041 households are included in the analyses in this paper.¹⁰ The bottom row of Table 1 shows how these households are distributed across RIPAT I & II and their respective control households, as well as the stratified random sample of non-RIPAT households surveyed in RIPAT I villages. We aimed at interviewing all the farmers who originally signed up for the RIPAT Farmer Field Schools, including those who later dropped out - provided they had remained in the village. In Arumeru district, 90 percent of the original RIPAT I farmers were interviewed, and 96 percent of the RIPAT II farmers were interviewed in Karatu district.

In each household, an interview was conducted with the person mainly responsible for agricultural decisions, often the head of the household. However, in RIPAT households, the person interviewed was always the RIPAT group member, who typically was the head or spouse of the head. The project aimed at achieving gender balance in the RIPAT farmer groups, which resulted in a larger share of female-headed households among the RIPAT farmers than otherwise in the village. The same degree of overrepresentation of female-headed households was sought among the control households. A village-level questionnaire was administered to representatives of each village government as a supplement to the household interviews. We thus have household- and village-level information in the data.

Table 1 lists means (and standard deviations) for the household- and village-level variables. Column (1) presents the averages for all households in the data, while the

⁹The data collection and data entry were closely supervised by us in cooperation with a survey management team from the Economic Development Initiative (a Tanzanian survey company). RECODA assisted in the hiring of a team of local interviewers and data entry clerks.

¹⁰We excluded from the dataset all farmers with more than eight acres of land and less than one acre of land in 2006 (for RIPAT, non-RIPAT and control farmers), as these did not comply with the original target criteria for RIPAT participation (174 households). We capped the acres at eight rather than five, as the data show that 17 percent of the RIPAT farmers did in fact have more than five acres of land, but only six percent had more than eight acres in 2006. Excluding households with more than five acres from the analysis below does not change the overall conclusions. We also excluded all newcomers to the villages (48 households). Finally, we excluded households with missing observations for any of our variables (111 households).

remaining columns represent subsets of the data used for different analyses. Columns (2) through (4) provide data for Arumeru district and Columns (5) and (6) for Karatu district. Column (2) shows data for households in RIPAT I, Column (3) has non-RIPAT households in RIPAT I villages, and Column (3) shows the averages for the households used as control households for RIPAT I. Columns (5) and (6) present data for RIPAT II households and their control households respectively.

It can be seen from Column (1) that the households included in the analysis generally had around 3 acres of land, that the majority of household heads had completed seven years of primary school, and that heads were typically middle-aged males with between one and two children living at home. We tested the farmers' math skills with two simple math problems;¹¹ 36 percent answered correctly. 16 percent of the households had participated in other development projects in the past. We also included the average historical rainfall level at 1:1 km resolution based on the household's GPS coordinates from secondary data,¹² since these households mainly rely on rain-fed agriculture. There is a large difference between the two districts, with Karatu receiving almost 200 mm more rainfall than Arumeru.

At the village level, we see that the average distance from each village to its most important market for agricultural output was eight kilometres, that two-thirds of the villages had secondary schools, and that half of the villages had hosted a development project in the past.

In the main part of the analyses below, we will be comparing RIPAT I households to control households from Arumeru district. It is therefore important that these are indeed comparable in terms of observable characteristics. We find that the two groups are well balanced; the only characteristic in Table 1 that differs significantly between RIPAT I and control households is whether the household had participated in another development project in the past, tested at the five percent level with cluster standard errors.

¹¹The farmer was considered "Good at math" if s/he correctly answered both questions, $29-13=?$ and $50/10=?$

¹²We used interpolated data for yearly precipitation measured in mm from the period 1950-2000 available from <http://www.worldclim.org/>.

In general, we cluster standard errors at the village level, as this is the most conservative approach when testing against the null hypothesis of no impact. However, this implies that we do not have enough degrees of freedom to include all the household and village characteristics shown. For consistency, we control for the log of acres, education, age, age squared and gender, and for all village characteristics, in all specifications.¹³ In Appendix B, we present regression results corresponding to the analyses below, where all characteristics are included and standard errors are clustered at the sub-village level instead.¹⁴ The inclusion of all household characteristics does not alter the results markedly.

3.1 Who Participated?

We know that the RIPAT project was not randomly allocated and it is interesting to take a closer look at the sources of selection: self-selection of households within villages, and non-random program placement across villages. Table 2 presents estimates from a logit regression of whether or not a household participated in RIPAT I on household and village characteristics. In Column (1) we compare RIPAT I households with non-RIPAT households in RIPAT I villages, which isolates self-selection of households. We can see that RIPAT I households were typically older, better at math and more likely to have participated in other projects in the past. The last two points suggest that RIPAT I households were more entrepreneurial than non-RIPAT households, as we expected, which could lead to a positive bias in the impact assessment. In Column (2), we compare the household characteristics of RIPAT I and control households and find that the same differences persist. In addition, RIPAT I households were more likely to be female-headed and received more precipitation than the control households. We include all households ever enrolled in RIPAT I, even if they later dropped out. This is done to ameliorate the issue of household selection, and we consider this to be the most conservative approach.

¹³When including observations from Karatu, we allow village characteristics to have district-specific coefficients.

¹⁴Clustering at the sub-village level leads to 52 clusters in regressions with Arumeru data only, and 130 clusters when all villages are included.

In the last column of Table 2, we add village characteristics to the regression. We see that RIPAT I villages were further away from their main markets, less likely to have a secondary school and more likely to have hosted a development project in the past. These differences all point to RIPAT I villages being less wealthy than the control villages, confirming the accounts from the original project documentation. This suggests that we might underestimate an impact when comparing the two. With respect to household characteristics, the difference in the likelihood of having participated in a development project is absorbed by the corresponding village-level difference. In addition, when we control for village-level characteristics, we find that heads of RIPAT I households were better educated than heads of control households, supporting the suggestion of positive self-selection of households into RIPAT I.

In Section 4 we present the four evaluation methodologies we employ to address household and village selection on the basis of observable and unobservable characteristics.

Finally, we note that among the RIPAT I participants who initially enrolled in the RIPAT FFS, there was a high level of engagement: the vast majority stayed with the project throughout the three year project period. In RIPAT I, more than 80 percent of the participants graduated, and the picture is similar for RIPAT II. It should be noted that participating in RIPAT FFS is rather time-consuming. Attendance rules were strictly enforced, and the need to attend is given as the main reason for dropping out by those who left the RIPAT farmer groups (Lilleør & Pedersen, 2013).¹⁵

3.2 Technology Adoption

The next obvious question is to examine whether or not RIPAT farmers also adopted on their own farms the technologies introduced through the RIPAT farmer groups. Farmers' engagement in project activities and the decision to allocate household resources (labor and land) towards adopting the proposed crops, livestock and new agricultural

¹⁵In RIPAT I, 77 households dropped out of their farmer groups before the end of implementation, while in RIPAT II, 96 households dropped out. All these drop-outs are included in the analyses throughout the paper and still considered to be RIPAT farmers or RIPAT participants regardless of when they dropped out.

practices are in themselves indicators of project implementation success, but they also represent a prior and necessary condition for finding any impact on broader welfare indicators as an outcome of the intervention.

We examine technology adoption among participating farmers in both RIPAT I and RIPAT II, since participants in the latter project had also been exposed to the full set of technologies examined here by the time of the survey in 2011. Because the basket of options entails a myriad of technologies and other elements (Maguzu et al., 2013), we have focused the analysis on six of the main components. We use simple means to indicate whether, relative to their control households, RIPAT I & II households were more likely to have adopted improved banana cultivation, to use more crop diversification, to grow fruit trees, to keep improved breeds of small livestock, to practice zero-grazing in their livestock husbandry, and to participate in savings groups (which was encouraged by RECODA).

In Table 3, we list the means (and standard deviations) for these key adoption measures for RIPAT I and II households and for their respective control households in the two districts, Arumeru and Karatu. Around two-thirds of the RIPAT households were found to be growing an improved banana variety. On average, they were growing around six different types of crop. About half had fruit trees, a quarter of them kept improved poultry breeds, 20-40 percent kept improved breeds of milking goats, and non-negligible fractions practiced zero grazing and were members of local savings groups.

To see whether there were significant differences between RIPAT and control households in the two districts, we carried out a series of cluster-robust t-tests for the difference being zero. A quick glance at the associated p-values shows that both the RIPAT I graduate households and the RIPAT II households had adopted all the analyzed components of the basket of options to a significantly greater extent than the households surveyed in the control villages. Only zero-grazing restrictions and the use of fruit trees seem not to have caught on in any significant way in RIPAT I compared to the

control villages in this simple bivariate setting.¹⁶ It should be noted that improved pig breeds were introduced only in RIPAT II.

This suggests that there was a considerable degree of take-up of the proposed technology options among the RIPAT farmers. It is unlikely that all of these significant differences in take-up could be driven by selection into the project, especially because the improved varieties of crops and breeds of livestock did not exist in the area prior to RIPAT.

Furthermore, these take-up rates indicate both that there was a high level of immediate take-up among RIPAT II farmers, who were only half-way through the project cycle, and high rates of sustained take-up among RIPAT I farmers, who at the point of data collection were more than one year beyond graduation and project closure.

When we analyze the overall degree of take-up, we find that all of the components in the basket of options were adopted by some farmers. No single element was adopted by all farmers, although all farmers were growing or keeping at least one of the promoted crops or animals breeds. This suggests that the element of choice built into the basket of options was indeed used by farmers to pick and choose according to their specific needs and resources.

3.3 Choice of Outcome Measures

We evaluate the impact of RIPAT on the basis of the development objectives that it was intended to improve, as stated in the original project documentation: namely, better food security and reduced poverty among the participating households.

3.3.1 Food security measures

To assess the food security situation among the respondent households, we employed a household level measure capturing access to food: the “Household Hunger Scale”

¹⁶When we control for household and village-level characteristics in the comparison of technology adoption between RIPAT I households and their control households, we find that all the listed adoption measures were in fact used to a greater extent among RIPAT I households, with significance levels of $p < 0.01$ or $p < 0.05$.

(HHS).¹⁷ It is based on three questions asking whether, due to lack of resources, anyone in the household 1) went to sleep at night hungry; 2) had no food to eat of any kind in the household; and 3) went a whole day and night without eating. The response codes are 0: never; 1: rarely or sometimes; 2: often. The HHS is simply the sum of the responses to the three questions resulting in an index from zero to six where zero corresponds to “no hunger” and six corresponds to “severe hunger”.

Due to considerable seasonal variations in the food security status of households, we take three different reference periods into account - the self-assessed best and worst months in terms of food security during the previous year, and the last four weeks prior to the survey.¹⁸ Since this area of Tanzania is not subject to severe and prolonged periods of starvation, we would expect to find most variation in the measure when the period of reference is the self-assessed worst month in terms of food security within the previous year. As it is difficult to interpret the magnitude of an impact on HHS because it is an ordinal measure, we also consider the simple binary variable “No hunger”, which is one if the household did not suffer from hunger at any point during the past year according to HHS and zero otherwise. To see whether children benefitted from RIPAT I, we measured their food consumption by looking at the prevalence of households where children had at least three meals per day during each of the three periods.

Finally, we aim to capture the nutritional quality of the overall household diet by analyzing whether household members had meat, eggs or dairy products to eat during the previous week.

From the raw averages in Table 4 we see that households in this region did not suffer from food insecurity throughout the year, but that food insecurity was rather pronounced during the worst periods of the year, typically the lean season immediately before harvest. Only 30-40 percent of households did not experience any hunger during the worst period of the year. Similarly, virtually all children had at least three

¹⁷The HHS is a modern food security instrument developed by US Aid to ensure cross-cultural comparability. It has been validated in five sub-Saharan African countries (Ballard et al., 2011).

¹⁸Households were interviewed in January, which is neither immediately after harvest nor in the worst hungry period, so we expected the hunger situation in the previous four weeks to have been somewhere in between the best and the worst months.

meals per day during the best part of the year, while on average about a quarter of the households with children served two meals or fewer per day during the lean season. Households in Arumeru seemed to report higher levels of food insecurity than households in Karatu, but in terms of reported nutritional quality, the weekly consumption of meat, eggs and dairy products was generally lower in the latter district.

A glance at the p-values for the cluster-robust t-tests of whether there were significant differences between RIPAT and control households in the two districts reveals that the raw means of the food security outcome variables are rather similar when we do not control for selection, household or village characteristics.

3.3.2 Poverty measures

Poverty is a complex outcome to measure. It is a relative measure, and it depends on local circumstances. Tanzania operates with a national poverty line of TZS 492 per adult equivalent per day (or roughly USD 1 per day using Purchasing Power Parity), representing the monetary cost of fulfilling basic needs (Schreiner, 2012).

Household income and consumption levels are notoriously difficult and time-consuming measures to capture, especially if this is to be done using a reasonably short survey instrument (Beegle et al. 2012a, 2012b). We therefore use an asset-based indicator of poverty as a short-cut. The “Progress out of Poverty Index” (PPI), as developed by Schreiner (2012), captures the probability that a household falls below the national poverty line. The PPI is country-specific and is based on ten simple questions that together provide a statistically strong and simple predictor of whether a household’s consumption level is likely to be below the national poverty line as established in the 2007 Household Budget Survey of 10,466 representative households from all over Tanzania.¹⁹ The PPI score ranges from 0 (most likely to be below a poverty line) to 100 (least likely to be below a poverty line).

We have taken the Progress out of Poverty Index as our key poverty indicator because it is a widely-used measure for identifying poverty *levels* and the only one avail-

¹⁹See Figure A.1 in Appendix A for the list of questions used in the latest PPI measure for Tanzania. Summing the points gives the overall PPI score.

able for Tanzania at the time of data collection. Schreiner (2012) notes that the PPI scorecard also aims to measure *changes* in poverty through time, and therefore in selecting indicators and holding other considerations constant, preference should be given to more sensitive indicators, e.g. ownership of a lantern. However, it places a lot of weight on more static measures, here fertility and female literacy.²⁰ We have therefore also considered the two best single predictors of poverty, according to Schreiner (2012), in isolation; namely the quality of the floor in the main dwelling and whether or not the household owns a (mobile) phone. In table 4, the raw averages for the poverty measures show clearly that households in Karatu are on average poorer than households in Arumeru.²¹ There are no significant poverty level differences between RIPAT households and their respective control households.

Finally, we also examine the supply of casual labor, as this is often an important source of income for poor households, but also something that is associated with stigma. It is a possible channel for RIPAT households to adjust their allocation of resources, if they can afford to do so; we return to this below. We see that among the control households, 15-20 percent relied on casual labor as one of the most important sources of income; but also that RIPAT households in both districts relied significantly less on supplying casual labor than the control households, and were also more likely to hire labor to work on their farms.

4 Evaluation strategy

In order to estimate the impact of RIPAT on participating households, we need a good estimate of the counterfactual situation - of what would have happened to the RIPAT households had they *not* participated in the project. We approach the counterfactual from four different angles, which in different ways and to different degrees take into account the participant self-selection and the non-random project placement at village

²⁰Such measures are often not helpful in analyses of poverty change; for example, we would not expect RIPAT to affect literacy adult females.

²¹District means for PPI, good quality floor and mobile phone are all significantly different at the one percent level.

level.

First, we undertake a *simple cross-sectional impact assessment* comparing outcomes of RIPAT I households to outcomes of households in control villages in a multivariate regression analysis using Ordinary Least Squares (OLS). To the extent that the household and village-level characteristics included in this multivariate setting do not fully account for the endogenous selection at household- and village-level, this simple cross-sectional estimation of the impact may be biased. It will be upward biased if the farmers that chose to participate and thus self-selected into the project were more motivated and entrepreneurial than the average farming household in a control village, *ceteris paribus*. It will be downward biased if the RIPAT I villages were indeed less wealthy than the control villages prior to project implementation, as suggested by the project documentation, and if this difference is not captured by the village characteristics included in the regressions.

Second, to take household self-selection into account, we estimate the impact at village rather than household level. That is, we explore the fact that we have surveyed non-RIPAT households in RIPAT I villages and estimate the *intention to treat* (ITT) effect, which pools both RIPAT and non-RIPAT households in RIPAT villages, since they were all intended for treatment. This does not give us an estimate of the average treatment effect among those who initially signed up for the project, but rather an average village-level effect among all those who could have signed up. The ITT estimator is free from self-selection bias and is only biased to the extent that the village level characteristics included do not fully account for the non-random project placement.

Third, to increase the comparability between RIPAT I households and their control households, we employ a *matching estimator*. This allows us to match more closely each RIPAT I household with a control household that has similar household and village-level characteristics. More specifically, we employ Mahalanobis matching with one nearest neighbor, which implies that a higher weight is given to control observations that are similar to RIPAT observations compared to OLS.²² In this way, we address the

²²We have also employed a propensity score matching estimation and get very similar results. However, we choose to present Mahalanobis matching estimates to obtain valid confidence intervals (Abadie

potential bias in the simple cross-sectional comparison due to unbalanced observables, but we still rely on the assumption of no selection on unobservable characteristics.

Finally, we propose a *Quasi-Difference-in-Difference* (QDiD) approach exploiting the gradual roll-out of the project. RIPAT II started more than two years after RIPAT I. RIPAT II participants were at the time of data collection still one and a half years away from graduation. Assuming that the selection mechanisms into RIPAT were the same in the two districts at both household and village levels, we can adjust for this selection in the simple cross-sectional impact assessment of RIPAT I in Arumeru District by subtracting the differences found between RIPAT II and control households in Karatu District. Doing this in a multivariate regression framework results in the QDiD estimator, which does not suffer from selection bias. The central assumption here is that the differences in outcomes due to household and village selection between treated and control households should be the same in the two districts in absence of treatment. Examining the observable characteristics in Table 1 above, we find indications that the RIPAT - control differences in the two districts are very similar. Out of the 12 characteristics listed, only two are significantly different at the five percent level (age and gender of head). This QDiD approach is similar to the evaluation strategy initially employed by Coleman (1999; 2006).

Ideally, for the perfect QDiD estimation, our data collection should have taken place exactly at project start-up of RIPAT II. The fact that the data collection took place two and a half years after project start of RIPAT II may result in QDiD *underestimating* the average treatment effect, since the high level of take-up of the different components in the basket of options could already have resulted in a beginning impact on the broader development outcomes at the time of the survey. There are three reasons why we are not very worried about this. First, it is always better to under-estimate than to overestimate, making any significant effect found more credible. Second, during the first year of RIPAT II, a severe drought hit the entire area (both Karatu and Arumeru districts) and caused the virtually complete failure of all agricultural activities in the

& Imbens, 2006). All observations are within the common support of the propensity score.

area. The project was therefore in effect postponed by one year, and project activities were resumed in the following agricultural season. Third, there is a natural time lag in both agricultural production and livestock breeding from the adoption of a new technology until its yields can be harvested, and in any case most households adopt additional new technologies gradually.²³

5 Results

To assess the food security impact of RIPAT, we consider whether households experienced any hunger, their HHS, whether the children in the households had at least three meals per day, and whether the households had eaten meat, eggs or dairy products during the previous week.

Table 5 is a compilation of the estimated effects. Each column represents an estimation method, while the rows refer to different outcome measures. Columns (1) and (2) present estimated coefficients for the RIPAT indicator variable from cross-sectional comparison regressions and for the RIPAT village indicator variable from ITT regressions respectively. Column (3) shows the differences between RIPAT and control households from Mahalanobis matching, while Column (4) gives the estimated impact from the QDiD specification.²⁴ All regressions include village characteristics and the restricted set of household characteristics, and standard errors are clustered at the village level. The same variables are used for the matching procedure. In Appendix C, we show the full set of regressors for the simple cross-sectional comparison and the QDiD regression with the HHS in worst month as the outcome variable.

In the first row of Panel A of Table 5, we show the estimated impact on the *No hunger* indicator. Reading across the columns, we see that RIPAT I increased the probability of being free from hunger by 17-24 percentage points, depending on the evaluation

²³Most RIPAT participants spend the first agricultural season learning about the new agricultural practices at a demonstration plot before they then in a later agricultural season choose which ones to adopt on their own farms.

²⁴This corresponds to the regression coefficient for the interaction term between the RIPAT indicator variable and the Arumeru indicator variable in a regression where both indicators are also included separately.

methodology, with the village level ITT impact being the lowest, as expected, but still large and statistically significant.

Having taken self-selection into account in the ITT estimation, the remaining worry is whether the impact is driven by pre-existing village differences, as the project was not randomly placed. In Column (3) we match on village and household characteristics, and thereby aim at a better balance of observables between RIPAT and the control villages and households. The impact on hunger persists in magnitude but is not statistically significant.

However, there might still be remaining *unobserved* differences between villages that we have not fully accounted for, and we therefore employ the QDiD approach using differences between RIPAT II and control households in Karatu to account for potential selection in Arumeru. Assuming that the selection mechanisms were the same in the two districts, this regression provides unbiased estimates of the impact. The result is reconfirmed: RIPAT I households are 24 percentage points less likely than their controls to have suffered from hunger when selection is accounted for. The fact that the QDiD estimate is so close to the estimated impacts from the other specifications suggests that selection on the basis of unobservables did not play a major role.²⁵

We also analyze the impact on the HHS for three different reference periods, recalling that higher values on the HHS correspond to more severe hunger. Consistently across all four specifications, we find that RIPAT I significantly reduced hunger in the worst period of the year. We do not see any impact in the best period or the four weeks immediately prior to the time of the interview. From Table 4 we note that there was only a little room for improvement, especially in the best month, as control households in Arumeru had an average HHS value of 0.04.

The reduction in hunger is associated with an increase in the number of meals for the children.²⁶ We see a consistent impact on the likelihood of having at least three meals in the best period of the year. This is a significant and substantive impact of

²⁵To the extent that RIPAT II had already had a (positive) impact on food security or poverty, we underestimate the impact of RIPAT I.

²⁶Because some households did not have any resident children, we lose 91 observations in Columns (1) and (3), 129 observations in Column (2) and 191 observations in Column (4).

seven to ten percentage points of improvement, depending on the specification. For the worst period we estimate an impact almost double that in magnitude for most specifications, but statistically less significant. With respect to the previous four weeks, the picture is more blurred. Taken together, however, these figures suggest that participating in RIPAT not only affected the food security status of households in the lean period as measured by the HHS, but it also improved children's intake of food at other times of the year.

Regarding the nutritional quality of the diet, we find that in general RIPAT I households were significantly more likely than controls to have eaten meat or eggs in the week before the interview, although the ITT results are weak for meat. We do not find a consistent increase in the intake of dairy products.

Based on these findings, we conclude that overall RIPAT I had a clear impact on food security in terms of reducing hunger, increasing the number of meals provided to children and improving the intake of animal protein.

The next question is then whether RIPAT also succeeded in improving the situation with regard to the other development objective of poverty alleviation. Turning to Panel B of Table 5, the first row shows that we do not find any significant impact of RIPAT on poverty as measured by the PPI. Estimates for the two additional time-variant indicators which have proven to be strong individual predictors of poverty status in Tanzania, quality of the floor and ownership of a mobile phone, are also insignificant. We have also checked for various degrees of heterogeneity in these results, but the conclusion remains the same: RIPAT has not had any significant impact on any of these poverty indicators and thus we believe the overall level of wealth of the participating households to have remained virtually unchanged.

In order to address potential gender differences, we split the results by gender of household head, and a few interesting findings emerge.²⁷ The female-headed RIPAT I households were more food secure than the female-headed control households during the *best* period of the year (suggesting that there was room for improvement among

²⁷The results are not shown, but are available upon request.

this subset of households), and they were more likely to have eggs as part of their daily diet. However, they were less likely to consume dairy products, which could be linked to the fact that they were also less likely to have adopted the improved breeds of milking goats.²⁸

6 Possible mechanisms

The fact that we find significant improvements in food security among RIPAT households, but no improvement in their poverty status, has led us to wonder why this should be so.

One explanation could be that resources were scarce for RIPAT households at the outset; when they experienced an improvement in their level of resources, they simply prioritized more secure and improved food consumption over higher non-food consumption. We cannot empirically test this any further, but it would explain the above finding.

A second explanation could be that households reallocated their use of labor resources within the household, e.g. shifted from cash income activities towards own agricultural production. This would have meant that the households produced more food, but earned less cash income, which again could have resulted in better food security (from own production) at the expense of lost income. This would make it unlikely that there would be a positive impact on poverty indicators.

Finally, a third explanation could be that the agricultural technologies introduced did not increase the total annual agricultural production, but only smoothed production over the agricultural cycle, thereby increasing food security in what typically would have been the lean period. We analyze the two last explanations empirically below.

²⁸This is consistent with the qualitative gender research among these women, which highlights the fact that the improved milking goat breeds introduced in the RIPAT groups were zero-grazing goats, which had to be fed. Whereas grazing goats is typically a male task in the local context, collecting fodder and firewood is a female task. Some female RIPAT farmers were therefore against keeping the milking goats, as this would increase the burden of collecting fodder. Later, specific fodder plants, e.g. elephant grass, were introduced to reduce this burden for the women (Mogensen & Pedersen, 2013)

6.1 Casual Labor

During qualitative interviews, it became clear that in this local setting casual labor is considered a “last resort”, an income source turned to when all other options are exhausted and hence, is greatly stigmatized. However, a reduction in the supply of casual labor could result in a substantial decrease in income, since casual labor can be a remunerative income source for Tanzanian smallholders.²⁹ We see from Table 4 that casual labor was indeed relatively widespread among the control households in Arumeru and Karatu districts; 15 and 20 percent respectively of these households supplied casual labor as a primary source of income. However, it was significantly less prevalent among RIPAT I and II farmers; only 5 and 11 percent respectively in these districts relied primarily on casual labor.³⁰ This suggests that RIPAT households might have chosen to cut back on casual labor because they had experienced an increase in their agricultural income. Such a cut-back would offset partially or completely any increase in income from agriculture, but still result in a welfare increase, because the household would avoid the stigma of supplying casual labor and at the same time become more food secure.

As to whether households *hired* labor on their own farms during 2010, we see from Table 4 that RIPAT households were 13 and 12 percentage points more likely than the control households to have hired labor on their own farms in Arumeru and Karatu districts respectively.

These impacts may not be causal and could be fully driven by selection, though controlling for household and village characteristics in a simple cross-sectional comparison only increases the differences found and the statistical significance. The ITT estimates are not significantly different from zero.³¹ This is not surprising, however, as we would expect labor markets to be local; if RIPAT farmers increased their demand for casual labor, non-RIPAT farmers in RIPAT villages might start to rely more on income

²⁹For example, weeding one acre of land pays TZS 2,000, which is four times the daily national poverty line.

³⁰Controlling for household and village characteristics only increases the estimated differences and the statistical significance.

³¹Results available upon request.

from casual labor, which would then even out the village average.

Nevertheless, taken together, the two results are suggestive in providing a possible explanation for why we find a profound impact on food security but no impact from RIPAT on the poverty measures used. The RIPAT households seem to have re-optimized the allocation of labor within their households and begun to invest in their own agricultural production.

6.2 Production Smoothing

The agricultural technologies introduced in RIPAT Farmer Field Schools were selected to enhance production smoothing over the agricultural cycle. Households generally experience large seasonal variation in food security, and they do not seem able to smooth consumption. In the lean period, 70 percent of the households in Arumeru control villages experienced some kind of hunger, while only two percent experienced any hunger just after harvest. Limited access to proper storage facilities and financial markets inhibit the ability of households to smooth consumption.³² Several elements in the basket of options introduced by RIPAT are production-smoothing technologies that provide the households with food even in the lean period. Banana plants fruit outside of the main harvest season as long as they receive some water, improved breeds of poultry lay more eggs, and improved breeds of goat produce more milk all year round than their traditional counterparts. It is therefore important to consider whether the impact of RIPAT on food security was mainly driven by the adoption of these three production-smoothing technologies that ease the smoothing of food consumption over the year and thus increase food security in the typical lean period.

The first two columns of Table 6 show that participation in RIPAT I and RIPAT II increased the probability of adopting at least one of the three production-smoothing technologies by 60-65 percentage points, controlling for household and village characteristics. As discussed in Section 3(b) above, participation in RIPAT is significantly correlated with adopting either banana cultivation or the keeping of improved breeds

³²Our results are not driven by access to savings, as they are robust to controlling for membership of a savings group.

of poultry or goats. The next two columns of Table 6 present regressions of the HHS in the worst month on a smoothing technology dummy that equals one if the household adopted any of the production-smoothing technologies and zero otherwise, for RIPAT and control households respectively. In both groups, households using production-smoothing technologies also experienced significantly less hunger than households that did not use any of the three technologies. This is not necessarily a causal relationship, as the decision to adopt the smoothing technologies was endogenous.

In Column (5) we limit the sample to those households that adopted any of the production-smoothing technologies and run the QDiD regression on this sub-sample in order to analyze whether RIPAT FFS participation brought about any additional degree of food security. The estimates suggest that RIPAT households adopting the production-smoothing technologies achieved the same level of food security as the selected sample of control households that adopted the smoothing technologies. However, 82 and 74 percent respectively of RIPAT I and RIPAT II households employed such technologies, whereas this was the case for only 25 and 8 percent of the control households in Arumeru and Karatu respectively. In Column (6) we see the QDiD regression results for the sub-sample of households *not* adopting any of the smoothing technologies. Among these, RIPAT I households experienced less hunger than controls at the five percent significance level after taking the selection into account, suggesting that the impact of RIPAT on food security was not purely driven by the production-smoothing technologies; other elements of the basket of options also improved the food security of households in the lean period.³³

We can thus conclude that although the use of smoothing technologies is associated with greater food security, the overall impact of RIPAT I on food security was not driven solely by these, as the basket of technology options seems to contain other elements that are also relevant for the food security of households not applying the main smoothing technologies.

³³We reach the same conclusions from a QDiD regression on the full sample where the RIPAT dummy, the district dummy and their interaction term are all interacted with the smoothing dummy.

7 Conclusion

This is, to the best of our knowledge, the first paper which rigorously analyzes the impact of a locally adapted Farmer Field School project on broader welfare indicators and development objectives, namely food security and poverty alleviation, and not just on intermediate and very project-related agricultural outcomes, such as technology knowledge transfer, technology adoption, or agricultural yields from the technologies promoted.

We find that there were strong and sustained positive effects on food security among the participating households more than one year after end of project, in terms of access to food, food consumption and quality of diet. Participating households experienced less hunger in the lean period, were more likely to have animal protein in their weekly diet, and were more likely to give the children in the household at least three meals per day. We find no impact of the RIPAT project on poverty indicators. There is suggestive evidence that the positive impacts on food security measures, but lack of impact on poverty indicators, could be caused by RIPAT households having prioritized food over non-food consumption, reallocated their labor resources towards improving their own agricultural production, and reduced seasonal peaks and troughs in food production.

Taken together - and when compared with earlier FFS evaluations - these results point to the importance of allowing the passage of time for assessing outcomes. Although the RIPAT II farmers, who had completed the project more recently than the RIPAT I farmers at the time of the survey, were also more likely than their control farmers to have adopted the full range of technologies examined, the impacts on food security can only be expected where the technology adoption has had sufficient time to raise food security levels, which in this case was among RIPAT I farmers. Timing is thus an important factor both when considering the length of the project (a typical RIPAT Farmer Field School runs for at least three years, as opposed to the one agricultural cycle of standard FFS projects) and when considering the timing of the impact evaluation and the outcomes selected for examination, allowing impacts from a change in agricultural systems to materialize. For instance, although we do not find any im-

impact on poverty indicators, it could be that such an impact will materialize in an even longer time horizon, when food security is no longer a concern for RIPAT households. Only time can tell.

A final question which may spring to mind concerns the costs of producing the food security impacts found. The total cost per participating RIPAT household per year was USD 200,³⁴ which is from three to 20 times as high as the various FFS cost estimates listed in Van den Berg & Jiggings (2007). However, it should be borne in mind that RIPAT projects differ from the typical FFS in that they offer a full *basket* of technology options, combine top-down teaching with participatory learning methods, have very close follow-up during the phasing-in period for the new technologies and are implemented over a substantially longer time horizon. Although these key differences clearly increase the cost per farmer, we also believe that they are vital to the impacts found above. None of the existing FFS evaluations have documented improved food security, so potentially the extra money was well spent.

Furthermore, apart from the objectives of improved food security and poverty alleviation among participating households, RIPAT also has the aim of ensuring that the participants are willing to share their knowledge with and demonstrate agricultural techniques to their fellow villagers, thus increasing the probability of diffusion of the improved techniques within RIPAT villages. A study by Gausset (2013) highlights the fact that a reasonably high degree of diffusion of the various RIPAT technologies has taken place. In particular, the improved banana variety has been popular, and by 2011 it had been adopted by one in eight non-RIPAT farmers in RIPAT I villages (Larsen, 2012). With this focus of RIPAT FFS on diffusion as in conventional agricultural extension programs, one can argue that the relevant cost-benefit analysis should be carried out at village rather than household level. The average cost per household is then only USD 30, and this expenditure led to an overall outcome of a 17-percentage-point increase in the probability of households being free from hunger in RIPAT I villages.³⁵ In

³⁴It should be noted that since the RIPAT interventions described here were the first out of a series of such projects, some piloting costs are also included.

³⁵The potential impact on non-RIPAT households need not only come through increased technology adoption. The analysis above of the demand for hired labor suggests that RIPAT also brought about in-

comparison, a large nation-wide community-based child nutrition program in Ethiopia resulted in an improvement of only seven percentage points measured on the same household hunger scale (White & Mason, 2012).

References

- Abadie, A. & Imbens, G. W. (2006). Large sample properties of matching estimators for average treatment effects. *Econometrica*, 74(1), 235–267.
- Aben, C., Duveskog, D., & Friis-Hansen, E. (2013). Evaluation of the RIPAT concept. In H. B. Lilleør & U. Lund-Sørensen (Eds.), *Farmers' Choice. Evaluating an Approach to Technology Adoption in Tanzania*. Practical Action Publishing.
- Anderson, J. R. & Feder, G. (2007). Agricultural extension. In *Handbook of Agricultural Economics, Agricultural Development: Farmers, Farm Production and Farm Markets*, volume 3 (pp. 2343–2378). Elsevier.
- Ballard, T., Coates, J., Swindale, A., & Deitchler, M. (2011). *Household Hunger Scale: Indicator Definition and Measurement Guide*. Technical report, FANTA-2 Bridge, FHI 360, Washington, DC.
- Beegle, K., Carletto, C., & Himelein, K. (2012a). Reliability of recall in agricultural data. *Journal of Development Economics*, 98(1), 34–41.
- Beegle, K., De Weerd, J., Friedman, J., & Gibson, J. (2012b). Methods of household consumption measurement through surveys: Experimental results from Tanzania. *Journal of Development Economics*, 98(1), 3–18.
- Braun, A. & Duveskog, D. (2011). The farmer field school approach: History, global assessment and success stories. *Background paper for the IFAD Rural poverty report*.
-
- creased economic activity on the local markets, which could have beneficial effects on the local economy in general.

- Braun, A., Jiggins, J., Röling, N., Van den Berg, H., & Snijders, P. (2006). A global survey and review of farmer field school experiences. In *Report prepared for the International Livestock Research Institute*. ILRI.
- Casey, K., Glennerster, R., & Miguel, E. (2012). Reshaping institutions: Evidence on aid impacts using a preanalysis plan. *The Quarterly Journal of Economics*, 127(4), 1755–1812.
- Coleman, B. E. (1999). The impact of group lending in Northeast Thailand. *Journal of Development Economics*, 60(1), 105–141.
- Coleman, B. E. (2006). Microfinance in Northeast Thailand: Who benefits and how much? *World Development*, 34(9), 1612–1638.
- Davis, K. & Nkonya, E. (2008). Developing a methodology for assessing the impact of farmer field schools in East Africa. In *Proceedings of the 24th Annual Meeting of the Association for International Agricultural and Extension Education (AIAEE)* (pp. 93–99).
- Davis, K., Nkonya, E., Kato, E., Mekonnen, D., Odendo, M., Miiro, R., & Nkuba, J. (2012). Impact of farmer field schools on agricultural productivity and poverty in East Africa. *World Development*, 40(2), 402–413.
- Feder, G., Anderson, J. R., Birner, R., & Deininger, K. W. (2010). *Promises and realities of community-based agricultural extension*. IFPRI discussion paper 959, International Food Policy Research Institute (IFPRI).
- Feder, G., Murgai, R., & Quizon, J. (2008). “Investing in farmers—the impacts of farmer field schools in relation to integrated pest management”—a comment. *World Development*, 36(10), 2103–2106.
- Feder, G., Murgai, R., & Quizon, J. B. (2004). Sending farmers back to school: The impact of farmer field schools in Indonesia. *Applied Economic Perspectives and Policy*, 26(1), 45–62.

- Friis-Hansen, E. & Duveskog, D. (2012). The empowerment route to well-being: An analysis of farmer field schools in East Africa. *World Development*, 40(2), 414–427.
- Gausset, Q. (2013). Local adoption of social and agricultural technologies. In H. B. Lilleør & U. Lund-Sørensen (Eds.), *Farmers' Choice. Evaluating an Approach to Technology Adoption in Tanzania*. Practical Action Publishing.
- Godtland, E. M., Sadoulet, E., Janvry, A. d., Murgai, R., & Ortiz, O. (2004). The impact of farmer field schools on knowledge and productivity: A study of potato farmers in the Peruvian Andes. *Economic Development and Cultural Change*, 53(1), 63–92.
- Larsen, A. F. (2012). *Adoption of Banana Cultivation and Information Networks: An Empirical Study of Northern Rural Tanzania*. Master thesis, University of Copenhagen, Copenhagen.
- Lilleør, H. B. & Pedersen, E. K. (2013). The RIPAT groups. In H. B. Lilleør & U. Lund-Sørensen (Eds.), *Farmers' Choice. Evaluating an Approach to Technology Adoption in Tanzania*. Practical Action Publishing.
- Lunduka, R., Ricker-Gilbert, J., & Fisher, M. (2013). What are the farm-level impacts of Malawi's farm input subsidy program? A critical review. *Agricultural Economics*, 44(6), 563–579.
- Maguzu, C. W., Ringo, D., & Vesterager, J. M. (2013). Presentation of RIPAT: core components and project implementation. In H. B. Lilleør & U. Lund-Sørensen (Eds.), *Farmers' Choice. Evaluating an Approach to Technology Adoption in Tanzania*. Practical Action Publishing.
- Mancini, F. & Jiggins, J. (2008). Appraisal of methods to evaluate farmer field schools. *Development in Practice*, 18(4-5), 539–550.
- Mogensen, H. O. & Pedersen, E. K. (2013). Household dynamics and gender politics: female farmers in RIPAT 1. In H. B. Lilleør & U. Lund-Sørensen (Eds.), *Farmers' Choice. Evaluating an Approach to Technology Adoption in Tanzania*. Practical Action Publishing.

- Rawlins, R., Pimkina, S., Barrett, C. B., Pedersen, S., & Wydick, B. (2014). Got milk? the impact of Heifer International's livestock donation programs in Rwanda on nutritional outcomes. *Food Policy*, 44, 202–213.
- Rejesus, R., Mutuc, M., Yasar, M., Lapitan, A., Palis, F., & Chi, T. (2012). Sending vietnamese rice farmers back to school: Further evidence on the impacts of farmer field schools. *Canadian Journal of Agricultural Economics*, 60(3), 407–426.
- Schreiner, M. (2012). *A Simple Poverty Scorecard for Tanzania*. <http://www.microfinance.com/>.
- Todo, Y. & Takahashi, R. (2013). Impact of farmer field schools on agricultural income and skills: Evidence from an aid-funded project in rural Ethiopia. *Journal of International Development*, 25(3), 362–381.
- van den Berg, H. & Jiggins, J. (2007). Investing in farmers—the impacts of farmer field schools in relation to integrated pest management. *World Development*, 35(4), 663–686.
- van den Berg, H. & Jiggins, J. (2008). “ Investing in farmers—the impacts of farmer field schools in relation to integrated pest management”—a reply. *World Development*, 36(10), 2107–2108.
- Vesterager, J., Ringo, D., & Maguzu, C. (2013). *RIPAT Manual*. www.ripat.org: Rock-wool Foundation.
- White, J. & Mason, J. (2012). *Assessing the Impact on Child Nutrition of the Ethiopia Community-based Nutrition Program*. Report to UNICEF of an evaluation study, Tulane University School of Public Health and Tropical Medicine, New Orleans.
- Yamazaki, S. & Resosudarmo, B. P. (2008). Does sending farmers back to school have an impact? Revisiting the issue. *The Developing Economies*, 46(2), 135–150.

Figures




Figure 1: Time line of RIPAT projects and data collection

Tables

Table 1: Summary statistics for background characteristics

	All (1)	Arumeru			Karatu	
		RIPAT I (2)	Non-RIPAT (3)	Control (4)	RIPAT II (5)	Control (6)
Acres 2006	3.04 (1.69)	3.29 (1.78)	2.95 (1.74)	3.02 (1.72)	3.03 (1.61)	2.90 (1.61)
Head less than 7 yrs educ.	0.32 (0.47)	0.30 (0.46)	0.30 (0.46)	0.32 (0.47)	0.28 (0.45)	0.41 (0.49)
Head more than 7 yrs educ.	0.05 (0.22)	0.07 (0.26)	0.06 (0.24)	0.06 (0.24)	0.04 (0.20)	0.04 (0.19)
Age of head	46.82 (14.45)	48.19 (13.57)	44.97 (16.10)	46.32 (16.06)	45.72 (11.53)	48.58 (15.21)
Head is female	0.15 (0.35)	0.19 (0.40)	0.15 (0.35)	0.19 (0.39)	0.07 (0.26)	0.16 (0.37)
Number of children of head	1.73 (1.55)	1.49 (1.31)	1.32 (1.34)	1.35 (1.33)	2.37 (1.68)	1.86 (1.70)
Good at math	0.36 (0.48)	0.41 (0.49)	0.36 (0.48)	0.38 (0.49)	0.36 (0.48)	0.31 (0.46)
Participation in other projects	0.16 (0.37)	0.27 (0.45)	0.13 (0.33)	0.16 (0.37)	0.15 (0.35)	0.09 (0.29)
Household rainfall, mm/year	818.59 (106.97)	751.26 (53.99)	749.53 (54.38)	703.84 (41.29)	930.23 (50.83)	905.28 (61.30)
Village distance to market, km	8.53 (4.83)	9.59 (3.68)	10.09 (3.71)	5.43 (4.86)	8.42 (5.96)	8.98 (3.93)
Village has secondary school	0.68 (0.47)	0.60 (0.49)	0.65 (0.48)	0.88 (0.33)	0.63 (0.48)	0.67 (0.47)
Village hosted devel. project	0.51 (0.50)	0.63 (0.48)	0.70 (0.46)	0.39 (0.49)	0.36 (0.48)	0.49 (0.50)
Observations	2,041	420	335	359	491	436

Notes: Means (and standard deviations) of household and village characteristics for all households in the sample are shown in Column (1) and for subsets of the sample in Columns (2)-(6). The means are unweighted. Since non-RIPAT households are overrepresented in some villages, the village level means differ slightly between Column (2) and (3).

Table 2: Who participated in RIPAT I?

Comparison households:	Non-RIPAT (1)	Control (2)	Control (3)
Log acres 2006	0.218 (0.17)	0.091 (0.22)	-0.091 (0.30)
Head less than 7 yrs educ.	-0.238 (0.25)	-0.219 (0.29)	-0.663** (0.34)
Head more than 7 yrs educ.	0.216 (0.33)	0.038 (0.36)	0.322 (0.36)
Age of head	0.163*** (0.05)	0.112*** (0.04)	0.177*** (0.05)
Age of head, squared /100	-0.138*** (0.04)	-0.100** (0.04)	-0.162*** (0.04)
Head is female	0.335 (0.24)	0.335* (0.20)	0.257 (0.22)
Number of children of head	0.030 (0.07)	0.107 (0.08)	-0.035 (0.11)
Good at math	0.267* (0.15)	0.066 (0.21)	0.135 (0.27)
Participation in other projects	0.830*** (0.21)	0.494** (0.24)	0.236 (0.37)
Household rainfall, mm/year	-0.000 (0.00)	0.024** (0.01)	0.021*** (0.01)
Village distance to market			0.335*** (0.12)
Village has secondary school			-3.191** (1.32)
Village hosted devel. project			2.359* (1.33)
Constant	-4.450** (2.09)	-20.185*** (7.53)	-20.706*** (5.47)
N	755	779	779

Notes: Logit estimates with a RIPAT I household indicator as the outcome variable. The samples consist of households within RIPAT I villages in Column (1), and of RIPAT I and their control households in Columns (2) and (3). Standard errors in parentheses are clustered at the sub-village level. Significance levels are denoted by * 0.1, ** 0.05 and *** 0.01.

Table 3: Summary statistics for adoption measures

	Arumeru			Karatu		
	RIPAT I	Control	(p-value)	RIPAT II	Control	(p-value)
Improved banana	0.69 (0.46)	0.12 (0.33)	0.00	0.64 (0.48)	0.01 (0.08)	0.00
Number of crops in 2010	5.62 (2.30)	4.76 (2.22)	0.02	6.65 (2.71)	4.69 (2.12)	0.00
Fruit tree(s)	0.66 (0.48)	0.56 (0.50)	0.46	0.49 (0.50)	0.28 (0.45)	0.02
Improved poultry breeds	0.27 (0.44)	0.02 (0.14)	0.00	0.25 (0.44)	0.01 (0.10)	0.00
Improved goat breeds	0.40 (0.49)	0.15 (0.36)	0.00	0.19 (0.40)	0.05 (0.22)	0.00
Improved pig breeds	0.00 (0.00)	0.00 (0.00)		0.18 (0.38)	0.00 (0.05)	0.00
Zero grazing	0.30 (0.46)	0.29 (0.45)	0.93	0.21 (0.41)	0.09 (0.29)	0.02
Savings	0.23 (0.42)	0.03 (0.18)	0.00	0.30 (0.46)	0.11 (0.31)	0.01
Observations	420	359		491	436	

Notes: The table shows the means (standard deviations) as well as the p-values of a cluster-robust t-test of the differences in means being equal to zero, clustering at the village level.

Table 4: Summary statistics for development outcomes

	Arumeru			Karatu		
	RIPAT I	Control	(p-value)	RIPAT II	Control	(p-value)
No hunger	0.40 (0.49)	0.29 (0.46)	0.19	0.39 (0.49)	0.39 (0.49)	0.90
HHS worst month	1.43 (1.47)	1.65 (1.46)	0.43	1.23 (1.25)	1.23 (1.26)	0.98
HHS best month	0.07 (0.35)	0.04 (0.27)	0.38	0.03 (0.24)	0.03 (0.26)	0.97
HHS previous four weeks	0.25 (0.66)	0.32 (0.73)	0.57	0.19 (0.53)	0.32 (0.74)	0.01
At least 3 meals, worst month	0.63 (0.48)	0.62 (0.49)	0.89	0.82 (0.38)	0.82 (0.39)	0.92
At least 3 meals, best month	0.94 (0.24)	0.91 (0.29)	0.39	0.99 (0.11)	0.98 (0.13)	0.42
At least 3 meals, previous 4 weeks	0.87 (0.34)	0.84 (0.37)	0.43	0.96 (0.19)	0.95 (0.21)	0.62
Meat	0.74 (0.44)	0.69 (0.46)	0.48	0.40 (0.49)	0.39 (0.49)	0.70
Eggs	0.56 (0.50)	0.36 (0.48)	0.00	0.45 (0.50)	0.38 (0.49)	0.29
Dairy products	0.87 (0.34)	0.83 (0.38)	0.44	0.63 (0.48)	0.60 (0.49)	0.62
PPI	44.29 (14.81)	44.68 (14.04)	0.89	32.00 (16.41)	33.49 (14.84)	0.56
Good quality floor (not earth)	0.26 (0.44)	0.31 (0.46)	0.55	0.13 (0.33)	0.11 (0.32)	0.81
Mobile phone	0.68 (0.47)	0.67 (0.47)	0.83	0.61 (0.49)	0.56 (0.50)	0.27
Rely on casual labour	0.05 (0.22)	0.15 (0.36)	0.02	0.11 (0.31)	0.20 (0.40)	0.02
Hired labour	0.62 (0.49)	0.49 (0.50)	0.03	0.45 (0.50)	0.33 (0.47)	0.05
Observations	420	359		491	436	

Notes: The table shows the means (standard deviations) as well as the p-values of a cluster robust t-test of the differences in means being equal to zero, clustering at the village level.

Table 5: Impact of RIPAT on development outcomes

	(1) Simple CS	(2) ITT	(3) Matching	(4) QDiD
PANEL A: Food security outcomes				
No hunger	0.208 *** (0.052)	0.172 ** (0.062)	0.189 (0.204)	0.238 *** (0.063)
HHS, worst month	-0.714 *** (0.193)	-0.699 ** (0.277)	-0.723 *** (0.204)	-0.809 *** (0.226)
HHS, best month	-0.004 (0.031)	0.004 (0.037)	-0.043 (0.052)	-0.023 (0.034)
HHS, previous 4 weeks	-0.146 (0.127)	-0.153 (0.106)	-0.146 (0.126)	-0.046 (0.133)
At least 3 meals, worst month	0.158 * (0.085)	0.156 * (0.089)	0.083 (0.062)	0.170 * (0.098)
At least 3 meals, best month	0.069 * (0.038)	0.100 ** (0.045)	0.076 ** (0.038)	0.065 * (0.037)
At least 3 meals, previous 4 weeks	0.106 ** (0.037)	0.080 (0.057)	0.070 (0.049)	0.101 ** (0.040)
Had meat previous week	0.132 * (0.064)	0.044 (0.069)	0.143 *** (0.042)	0.148 * (0.076)
Had eggs previous week	0.223 *** (0.044)	0.145 ** (0.057)	0.189 *** (0.061)	0.163 ** (0.065)
Had dairy products prev. week	0.068 (0.072)	0.005 (0.083)	0.120 ** (0.047)	0.050 (0.092)
PANEL B: Poverty outcomes				
PPI	3.472 (2.086)	0.829 (3.075)	0.077 (0.050)	4.047 (3.105)
Has good quality floor (not earth)	-0.002 (0.076)	-0.062 (0.087)	1.351 (1.481)	-0.006 (0.081)
Has (mobile) phone	0.055 (0.033)	-0.063 (0.037)	-0.031 (0.060)	0.064 (0.047)
Observations	779	1,114	779	1,706

Notes: Each row represents a dependent variable. Columns (1), (2) and (4) show OLS regression coefficients: Column (1) gives the coefficient to the RIPAT I indicator in a simple cross-sectional (CS) comparison using data from Arumeru district only; Column (2) gives the coefficient for the RIPAT I village indicator in ITT regressions including non-RIPAT households from RIPAT I villages, applying inverse sampling probability weights, and using data from Arumeru district only; Column (3) gives the Mahalanobis matching estimates yielded when RIPAT I households are matched to controls in Arumeru district; and Column (4) gives the coefficients for the interaction term between the RIPAT dummy and the Arumeru district dummy in the QDiD specification, i.e. these estimations include both RIPAT I and RIPAT II households and their respective control households in the two districts. Village characteristics and household characteristics as described in text are controlled for in all specifications. Standard errors in parentheses are clustered at the village level. Significance levels are denoted by * 0.1, ** 0.05 and *** 0.01. The numbers of observations are reduced for the "Less than 3 meals" outcomes to 688, 985, 688 and 1515 respectively for the four columns.

Table 6: Smoothing mechanisms

Outcome variable	Adoption		HHS, worst month			
	(1) RIPAT I	(2) RIPAT II	(3) RIPAT	(4) Control	(5) Smooth	(6) Nonsmooth
RIPAT	0.601*** (0.07)	0.648*** (0.04)			0.167 (0.20)	0.344*** (0.12)
Smooth			-0.357* (0.18)	-0.369*** (0.12)		
District			0.654 (0.56)	0.893*** (0.29)	1.199** (0.55)	0.745** (0.27)
RIPAT*District					-0.397 (0.32)	-0.826** (0.33)
N	779	927	911	795	828	878

Notes: OLS estimates. The dependent variable in Columns (1) and (2) is an adoption indicator equal to one if the household had adopted any of the smoothing technologies; in Columns (3)-(6) it is HHS in worst month. Village characteristics and household characteristics as described in text are controlled for in all specifications. Standard errors in parentheses are clustered at the village level. Significance levels are denoted by * 0.1, ** 0.5 and *** 0.01.

Appendix

A Progress out of Poverty Indicator (PPI)

The PPI is constructed by Schreiner (2012) based on ten simple questions listed in what he refers to as a scorecard; see the example from Tanzania below.

<u>Entity</u>	<u>Name</u>	<u>ID</u>	<u>Date</u> (DD/MM/YY)
Member:	_____	_____	Joined: _____
Field agent:	_____	_____	Today: _____
Service point:	_____	_____	Household size: _____

Indicator	Value	Points	Score
1. How many household members are 17-years-old or younger?	A. Four or more	0	
	B. Three	8	
	C. Two	15	
	D. One	23	
	E. None	30	
2. Do all children ages 6 to 17 attend school?	A. No	0	
	B. Yes, or no children ages 6 to 17	1	
3. Can the female head/spouse read and write?	A. No	0	
	B. Yes, but not in Kiswahili nor English	0	
	C. No female head/spouse	0	
	D. Yes, only in Kiswahili	5	
	E. Yes, in English (regardless of others)	12	
4. What is the main building material of the floor of the main dwelling?	A. Earth	0	
	B. Concrete, cement, tiles, timber, or other	11	
5. What is the main building material of the roof of the main dwelling?	A. Mud and grass	0	
	B. Grass, leaves, bamboo	8	
	C. Concrete, cement, metal sheets (GCI), asbestos sheets, tiles, or other	11	
6. How many bicycles, mopeds, motorcycles, tractors, or motor vehicles does your household own?	A. None	0	
	B. One	1	
	C. Two or more	11	
7. Does your household own any radios or radio cassettes?	A. No	0	
	B. Yes	5	
8. Does your household own any lanterns?	A. No	0	
	B. Yes	6	
9. Does your household own any irons (charcoal or electric)?	A. No	0	
	B. Yes	6	
10. How many tables does your household own?	A. None	0	
	B. One	2	
	C. Two	4	
	D. Three or more	7	

Microfinance Risk Management, L.L.C., <http://www.microfinance.com> **Total score:**

Figure A.1: A simple poverty scorecard for Tanzania

B Robustness results

Table A.1: Impact of RIPAT on development outcomes

	(1)	(2)	(3)	(4)
	Simple CS	ITT	Matching	QDiD
PANEL A: Food security outcomes				
No hunger	0.203 *** (0.056)	0.155 *** (0.055)	0.224 (0.164)	0.225 *** (0.075)
HHS, worst month	-0.590 *** (0.195)	-0.524 ** (0.211)	-0.935 *** (0.164)	-0.649 *** (0.231)
HHS, best month	0.012 (0.050)	-0.024 (0.045)	-0.013 (0.036)	-0.008 (0.050)
HHS, previous 4 weeks	-0.062 (0.147)	-0.082 (0.118)	-0.172 ** (0.080)	0.035 (0.152)
At least 3 meals, worst month	0.137 ** (0.066)	0.143 ** (0.068)	0.225 *** (0.071)	0.143 * (0.081)
At least 3 meals, best month	0.088 ** (0.035)	0.133 *** (0.036)	0.050 (0.044)	0.085 ** (0.034)
At least 3 meals, previous 4 weeks	0.124 *** (0.041)	0.093 * (0.049)	0.103 * (0.056)	0.123 *** (0.042)
Had meat previous week	0.173 ** (0.066)	0.099 (0.082)	0.156 *** (0.059)	0.174 ** (0.077)
Had eggs previous week	0.212 *** (0.057)	0.179 *** (0.063)	0.221 *** (0.051)	0.146 ** (0.071)
Had dairy products prev. week	0.066 (0.067)	0.019 (0.072)	0.084 ** (0.042)	0.033 (0.082)
PANEL B: Poverty outcomes				
PPI	0.197 (2.473)	-1.087 (2.767)	0.097 (0.064)	0.562 (2.918)
Has good quality floor (not earth)	-0.098 (0.062)	-0.131 * (0.068)	2.968 * (1.592)	-0.105 (0.068)
Has (mobile) phone	0.027 (0.043)	-0.057 (0.048)	-0.036 (0.048)	0.039 (0.052)
Observations	779	1,114	779	1,706

Notes: Each row represents a dependent variable. Columns (1), (2) and (4) show OLS regression coefficients: Column (1) gives the coefficient for the RIPAT I indicator in a simple cross-sectional (CS) comparison using data from Arumeru district only; Column (2) gives the coefficient for the RIPAT I village indicator in ITT regressions, including non-RIPAT households from RIPAT I villages, applying inverse sampling probability weights, and using data from Arumeru district only; Column (3) gives the Mahalanobis matching estimates yielded when RIPAT I households are matched to controls in Arumeru district; and Column (4) gives the coefficient for the interaction term between the RIPAT dummy and the Arumeru district dummy in the QDiD specification, i.e. this estimation includes both RIPAT I and RIPAT II households and their respective control households in the two districts. Village characteristics (Distance to market, Has secondary school, Hosted development project in 2006-2010, and in Column (4) all three interacted with Arumeru district dummy) and household characteristics (Log acres in 2006; Household head's gender, education, math skills, age and age squared; Number of children of head; Whether household has participated in other project in the past; Historical rainfall (interacted with the Arumeru district dummy in Column (4))) are controlled for in all specifications. Standard errors in parentheses are clustered at the sub-village level. Significance levels are denoted by * 0.1, ** 0.5 and *** 0.01. The numbers of observations are reduced for the "Less than 3 meals" outcomes to 688, 985, 688 and 1,515 respectively for the four columns.

Table A.2: Smoothing mechanisms

Outcome variable	Adoption		HHS, worst month			
	(1) RIPAT I	(2) RIPAT II	(3) RIPAT	(4) Control	(5) Smooth	(6) Nonsmooth
RIPAT	0.567*** (0.09)	0.650*** (0.04)			0.051 (0.27)	0.320*** (0.10)
Smooth			-0.348* (0.18)	-0.313** (0.12)		
District			6.403** (2.45)	9.356* (5.09)	3.244 (3.35)	5.194* (2.90)
RIPAT*District					-0.108 (0.40)	-0.851** (0.35)
N	779	927	911	795	828	878

Notes: OLS estimates. The dependent variable in Columns (1) and (2) is an adoption indicator equal to one if the household has adopted any of the smoothing technologies; in Columns (3)-(6) the dependent variable is HHS in worst month. Village characteristics (Distance to market, Has secondary school, Hosted development project in 2006-2010, and all three interacted with the Arumeru district dummy) and household characteristics (Log acres in 2006; Household head's gender, education, math skills, age and age squared; Number of children of head; Whether household has participated in other project in the past; Historical rainfall, and the last interacted with the Arumeru district dummy) are controlled for in all specifications. Standard errors in parentheses are clustered at the sub-village level. Significance levels are denoted by * 0.1, ** 0.5 and *** 0.01.

C HHS in worst month, all regression coefficients

Table A.3: Impact of RIPAT on HHS in worst month; all regression coefficients shown

	Simple CS			QDiD		
	(1)	(2)	(3)	(4)	(5)	(6)
RIPAT dummy	-0.226 (0.28)	-0.737*** (0.19)	-0.714*** (0.19)	0.003 (0.15)	0.044 (0.13)	0.096 (0.13)
Arumeru district dummy				0.422* (0.22)	0.991*** (0.34)	1.054*** (0.34)
RIPAT*District				-0.228 (0.31)	-0.782*** (0.23)	-0.809*** (0.23)
Village distance to market		0.044** (0.02)	0.043** (0.02)		0.023* (0.01)	0.023** (0.01)
Village has secondary school		-0.741*** (0.23)	-0.726*** (0.23)		0.075 (0.12)	0.059 (0.11)
Village had devel. project		0.506** (0.22)	0.521** (0.23)		0.197* (0.11)	0.244** (0.11)
Village distance to market*District					0.021 (0.02)	0.019 (0.02)
Village has secondary school*District					-0.816*** (0.25)	-0.784*** (0.25)
Village had devel. project*District					0.309 (0.24)	0.280 (0.25)
Log acres 2006			-0.349*** (0.10)			-0.365*** (0.07)
Head less than 7 yrs educ.			0.123 (0.13)			0.130 (0.09)
Head more than 7 yrs educ.			-0.529*** (0.17)			-0.594*** (0.12)
Age of head			0.019 (0.03)			0.022 (0.01)
Age of head, squared			-0.009 (0.02)			-0.014 (0.01)
Head is female			-0.115 (0.12)			-0.038 (0.11)
Constant	1.652*** (0.20)	1.871*** (0.31)	1.543** (0.70)	1.229*** (0.10)	0.880*** (0.15)	0.472 (0.39)
N	779	779	779	1706	1706	1706

Notes: OLS estimates; Dependent variable is Household Hunger Scale in worst month. Standard errors in parentheses are clustered at the village level. Significance levels are denoted by * 0.1, ** 0.5 and *** 0.01. Columns (1)-(3) are based on data from RIPAT I and control households in Arumeru district, while Columns (4)-(6) also include data from Karatu district. 'District' refers to an interaction term with the Arumeru district dummy.

Can agricultural interventions improve child health? Evidence from Tanzania

Anna Folke Larsen* and Helene Bie Lilleør†

Abstract

Severely reduced height-for-age due to undernutrition is widespread in young African children, with serious implications for their health and later economic productivity. It is primarily caused by growth faltering due to hunger spells in critical periods of early child development. We assess the impact on child health, measured as height-for-age, of an agricultural intervention that improved food security among smallholder farmers by providing them with a “basket” of new technology options. We find that height-for-age measures among children from participating households increased by about 0.8 standard deviations and the incidence of stunting among them decreased by about 17 percentage points.

Acknowledgements All views are those of the authors and do not necessarily reflect the views of the Rockwool Foundation. The Rockwool Foundation funded the intervention and the associated data collection used in this paper. We are grateful for comments and advice from RECODA, Jens Vesterager, Mette Ejrnæs, Torben Tranæs, Henrik Friis and Pernille Kæstel, as well as from participants at the Nordic Econometrics Meeting 2013, the Nordic Conference in Development Economics 2013, the Summer School on Socioeconomic Inequality 2013 and seminars in the Development Economics Research Group and the Centre for Applied Microeconometrics, University of Copenhagen. All remaining errors are of course our own.

*University of Copenhagen, afl@econ.ku.dk

†Rockwool Foundation Research Unit, hbl@rff.dk

1 Introduction

Undernutrition is a key reason for poor child health in many developing countries. In Sub-Saharan Africa, around 40 per cent of children under the age of five suffer from stunted growth, i.e. severely reduced height-for-age relative to their growth potential (de Onis, Blössner and Borghi, 2011). Stunting is a result of periods of undernutrition in early childhood, and it has been found to have a series of adverse long-term effects in those who survive childhood. It is negatively associated with mental development (Martorell, 1999), with human capital accumulation (Jamison, 1986; Glewwe, Jacoby and King, 2001; Maluccio et al., 2009), with adult health (Victora et al., 2008; Adair et al., 2013), and with economic productivity and income levels in adulthood (Hoddinott et al., 2008, 2013).¹

It is by now well established that height-for-age can be seen as a “summary indicator” of the health and development of children during the first 1,000 days of their lives, from conception to two years of age (Hoddinott et al., 2013). During this period, children have very high growth rates; and consequently, when subject to spells of growth faltering, children quickly fall behind the height-for-age growth curves of their peers, with limited chances of catching up subsequently (Victora et al., 2010).²

In this paper, we assess the impact on early childhood health, measured as height-for-age, of an agricultural intervention that improved food security in the lean season among smallholder farmers by providing them with a “basket” of new technology options. The intervention targeted smallholder farmers in Northern Tanzania by organizing farmer groups similar to those used in the widespread Farmer Field Schools approach. On a common group plot, each group was trained in and given the opportunity to experiment with a basket of agricultural and animal husbandry technology options based on locally available resources over a time horizon of three and a half years. Each farmer then adopted his or her preferred technologies in accordance with

¹Although Vogl (2014) shows that a sizeable fraction of higher adult wages may be mediated by occupational choice and better education.

²Although an opportunity window for catch-up may exist in the later puberty period, as recently shown by Hirvonen (2013).

his or her own needs and resources.

Roughly half of the participating households had children under the age of five years. To identify the impact of the agricultural intervention on early childhood health in terms of height-for-age, we employ a difference-in-differences comparison of cohorts conceived before and after the phase-in of the project, where only the latter cohort lived all of their first 1,000 days under full project implementation. The height-for-age data for the older cohort allow us to control for systematic differences in nutritional levels between children in treatment and comparison households prior to the onset of intervention activities.³

Because stunting is widespread in developing countries and has serious long-term implications, its causes and potential prevention strategies have been subject to careful scrutiny. The prevention strategies focus on the nutrition of pregnant women, infants and young children. They include disease prevention strategies, breastfeeding practices, micronutrient supplements, food fortification, and food security strategies (Allen and Gillespie, 2001; Bhutta et al., 2008; Schroeder, 2008). The various authors all note that the evidence of the effectiveness of these strategies in preventing undernutrition is as mixed as the range of strategies itself. Although breastfeeding promotion and providing micronutrient supplements are effective strategies for reducing stunting, they cannot fully prevent stunting in food-insecure environments where the mother is undernourished or there are numerous deficiencies in micronutrients (Schroeder, 2008).

There is a general agreement in the nutritional literature that there is no magic bullet for solving the undernutrition problem. Rather, it is believed that “to eliminate stunting in the longer term, these [nutritional] interventions should be supplemented by improvements in the underlying determinants of undernutrition” (Bhutta et al., 2008). It is argued that it is necessary to combine nutrition programs with income growth (Alderman, Hoogeveen and Rossi, 2006) and with broader food systems (Miller and Welch, 2013), and to focus more on overall dietary quality, bringing local needs, cultural conditions and resource constraints into play (Schroeder, 2008) to achieve sustain-

³This follows closely the identification strategy of Duflo (2003).

able solutions to undernutrition. Recently, this has been stressed even more strongly by Ruel and Alderman (2013). They argue that nutrition-sensitive interventions or programs have enormous potential to improve nutrition without being nutrition-specific.⁴ They review the nutritional potential of interventions and programs in four different sectors, including agriculture. They conclude that within agriculture in particular, the potential for positive nutritional impacts is great, because agricultural interventions can support livelihoods, increase food production and enhance access to diverse diets. They note, however that the empirical evidence is very scanty, largely due to the poor quality of evaluations.

A recent systematic review by Masset et al. (2012) focuses specifically on whether agricultural interventions, such as home gardens, animal husbandry, and the production of bio-fortified crops, all aimed at improving the nutritional status of children, actually succeeded in doing so. They find that although there is a positive effect on the production and consumption of the agricultural goods promoted, the impact on the overall diet is unclear, and very little positive evidence was found of an effect on the nutritional status of young children. However, Masset et al. stress that weak evaluation methodologies and lack of sufficient statistical power cast serious doubt on the validity of an overall and somewhat counterintuitive conclusion that there was a limited impact of the agricultural interventions on nutrition. They therefore call for more rigorous research on the subject in order to be able to answer the question of whether agricultural interventions can reduce undernutrition and therefore should play a more prominent role in the prevention of growth faltering among young children.

We contribute to this literature by providing a careful and rigorous impact assessment on height-for-age and stunting among young children of one agricultural intervention. To use the terminology of Ruel and Alderman (2013) above, the intervention was nutrition-sensitive in that it targeted food security broadly, but it was not

⁴Ruel and Alderman (2013) define nutrition sensitive interventions to be interventions or programs that address the underlying determinants of fetal and child nutrition and development, such as food security, whereas nutrition-specific interventions are interventions or programs that address the immediate determinants of fetal and child nutrition and development, such as adequate food and nutrient intake by children, feeding, caregiving and parenting practices, and low burden of infectious diseases.

nutrition-specific. It promoted a more constant level of food security throughout the year by introducing perennial crops and improved breeds of livestock to help increase food availability during the lean season.

Using post-treatment data, we analyze whether the three-and-a-half-year-long intervention led to an improvement in the height-for-age measures among children young enough to have lived all their lives under the intervention. We measure this one year after the completion of the intervention. To identify the impact, we follow the identification strategy in Duflo (2003) and exploit the fact that the height-for-age measure is a strong biological marker of undernutrition in a well-defined age window, from conception to 24 months of life. The intuitive reasoning is as follows. If the intervention indeed reduced spells of undernutrition or hunger among participating households, children conceived after the phase-in of the project should be taller for their age than their older peers who lived (part of) their first two years of life before the project could have had any impact on food security. However, there may have been a general change in the food security status of children during the project period. To control for this, we employ a cohort difference-in-differences strategy and compare the relative height differential between young children in participating and comparison households to the height differential between their older peers.

We find that young children from participating households on average experienced a health improvement, in that their standardized height-for-age measures increased by about 0.8 standard deviations. In addition, we not only find improvements on average, but also in the lower tail of the height-for-age distribution. Looking at the prevalence rates of stunting, which is defined as having a height more than two standard deviations below the mean of a global reference distribution,⁵ we find indications that prevalence rates dropped by 17.6 percentage points. Compared to the literature, these are sizable impacts and larger than most from nutrition interventions, but comparable to nutritional impacts of cash-transfer programs. We show that improved food security in (severe) hunger periods is a probable mechanism behind this result. Furthermore,

⁵We use the international WHO growth standards (WHO, 2006).

we examine our identifying common trend assumption and test our results against various alternative specifications and explanations, and find that they are highly robust.

Overall, our findings suggest that agricultural interventions can in fact influence the underlying determinants of undernutrition to such an extent that they translate directly into children coming closer to their full growth potential. Although this is only *one* impact assessment of *one* agricultural intervention, and more rigorous impact assessments are needed to shape policy recommendations, our findings show that in the context studied it is possible to reduce early childhood stunting considerably through a broad nutrition-sensitive agricultural intervention.

The remainder of the paper is organized as follows. In section 2, we describe the characteristics of the agricultural intervention in more detail, while the data and summary statistics are described in section 3. In section 4, we present our identification strategy, and in section 5, our main results and their robustness. In section 6, we examine our identifying assumption further, and in section 7 we conclude with a discussion of the relative magnitude of the impact and the project costs.

2 The agricultural intervention

The agricultural intervention is called “Rural Initiatives for Participatory Agricultural Transformation”, or RIPAT.⁶ The specific instance of this intervention that we evaluate was the first RIPAT program (RIPAT I), implemented by a local NGO, RECODA, in eight villages in Arumeru District in the Arusha Region of Northern Tanzania between 2006 and 2009 (see figure 2.1). Subsequently, another three similar RIPAT interventions have been implemented in nearby districts. The stated overall development goal of RIPAT is to reduce poverty and improve food security among smallholder farmers by facilitating high and sustainable levels of adoption of improved agricultural and live-

⁶See www.ripat.org or Lilleør and Lund-Sørensen (2013) for a thorough description and discussion of the intervention.

stock technologies disseminated through local farmer groups. The intervention has strong similarities with the Farmer Field Schools (FFS) approach as outlined in Van den Berg and Jiggins (2007), the main differences being that RIPAT offers a variety of technology options as opposed to one technology in FFS, combines both top-down teaching and participatory learning methods, and runs for three years with close follow-up as opposed to one agricultural season in FFS (Aben, Duveskog and Friis-Hansen, 2013).

Participation in RIPAT is not randomly allocated, which makes perfect sense from an implementation perspective, but which poses a challenge for the evaluation of the project. Poor villages with suitable agricultural conditions are selected at the district level. In the chosen villages, interested farmers (typically up to 70 in a village) are organized in farmer groups of 30-35 voluntary participants selected by the village council. In finding target participants, the village council is asked to select individuals who will be committed to the project (strict attendance records are kept), who are willing to share their new knowledge with fellow villagers, and who are not rich in terms of the internal village wealth ranking. However, to facilitate individual technology adoption, participants must own at least one acre (and no more than five acres) of farm land.

Once the groups have been organized, the facilitators from the implementing NGO meet with the group on a weekly basis during the phase-in period. The first tasks of the group are to agree on a group constitution and elect group leaders. Each farmer group then has to rent an appropriate group field of around one acre of land which can function as a demonstration plot, typically renting land belonging to a fellow farmer or the village community. All group meetings are subsequently held at the group field.

The group is offered training in a full basket of technology options, which covers a broad range of local needs. The technology options include new banana cultivation techniques; new improved banana and other perennial and annual crop varieties; conservation agriculture for improved land utilization (such as minimum soil disturbance, cover crops, intercropping, rotation and diversification of crops); post-harvesting technologies; improved animal husbandry; multipurpose trees for fodder, fruit, or firewood; soil and water conservation, including rain water harvesting; and

savings groups. During the phase-in period of one year, the facilitators from the implementing NGO (typically agronomists) train the group members gradually in each of the technology options according to the agricultural seasons. After this period, the main role of the facilitators is to monitor and provide guidance on a bi-monthly or monthly basis.

The farmer groups are exposed to the full basket of options at the group demonstration plot, where the new techniques are implemented and compared with traditional methods under the guidance of the skilled facilitator. This reduces the individual risks involved in trying out or learning new technologies. Each farmer is free to choose which technologies to adopt on his/her own farm according to his/her own needs, constraints and resources. Groups are given an initial set of inputs for free for the training in, demonstration of and testing of technologies on the group field, including roosters of improved breeds to cross-breed with local hens. However, individual farmers wanting to adopt the new technologies must purchase inputs from the implementing NGO at cost prices. In the case of improved varieties of banana seedlings and goats, solidarity chains are implemented to promote local diffusion.⁷ While some technologies may be generally more popular than others, adoption varies considerably from farmer to farmer, and often takes place after a time lag.

In the area of the implementation of RIPAT in the present study, food insecurity is most pronounced during the lean season of the year, i.e. during the months leading up to the annual harvest of the main staple crop, maize. The project implementation started in the beginning of the growing season in 2006, and hence we would expect the earliest impact on food insecurity to have taken place in the lean season of 2007. Children who were fully exposed to the potential benefits of RIPAT are therefore defined to be those conceived in January 2007 or thereafter (see below).

Although the intervention was not nutrition-specific, it was nutrition-sensitive in

⁷After the phase-in period and once banana seedlings are available from the group plot, the farmers can obtain free seedlings in exchange for agreeing to pass on three times the number of the seedlings received to other farmers within or outside the farmer group. The farmer tending a she-goat of an improved breed can keep the goat after passing on the first female offspring to another farmer on the same condition.




Figure 2.1: Timeline

its strong focus on achieving food security by promoting agricultural and livestock technologies which were more drought resistant, more varied, and led to a smoother rate of food production throughout the year. In Larsen and Lilleør (2014) we show that the intervention did in fact lead to improved food security levels among the full sample of participating households⁸ in terms of reduced hunger during the lean season, higher intake of animal protein in terms of meat and eggs, and more meals per day.⁹ Based on the nutrition literature, we expect children who were exposed to RIPAT *in utero* and during the first two years of their lives to have benefited from this, as children’s physical growth is particularly sensitive to insufficient nutrition. In this paper, we therefore contribute to the nutrition literature, when we investigate whether households in the subset with young children were in fact able to shield those children from nutrition-related setbacks in their growth.

3 Data and summary statistics

Our main outcome variable in this paper is the height-for-age z-score for children, which is a very powerful indicator of severe early childhood or *in-utero* undernutrition, as described above. We construct height-for-age z-scores (HAZ) by subtracting the means and dividing by the standard deviations of the age- and gender-specific lengths or heights from the reference distribution established in the WHO Multicentre

⁸That sample also included households without young children, as opposed to the sample of this paper.

⁹We also examined whether the intervention succeeded in alleviating poverty. Based on a broad range of single indicators and one composite poverty indicator, we did not detect any impact of RIPAT on poverty levels.

Growth Reference Study, which was based on healthy children from Brazil, Ghana, India, Norway, Oman and USA (de Onis et al., 2004). Though children below 24 months of age were measured recumbent, and hence we measured length rather than height, we henceforth refer to both length and height measurements as height.

We also look at the prevalence of stunting, using an indicator variable which equals one for those children whose height is less than two standard deviations below the age- and gender-specific mean.

3.1 Data

As indicated in the timeline in figure 2.1, we collected household-level data more than one year after the project was completed.¹⁰ We interviewed 506 of the 561 original RIPAT households from the eight intervention villages and 395 households from eight comparable non-intervention control villages in the same district.¹¹ The comparison households were sampled at random among farming households with one to eight acres of land.¹² Out of these 901 households, 469 of them had children aged five years or less, in total 645 children. We are able to construct height-for-age z-scores for 482 children from 382 households. The main reason for attrition is that enumerators were not obliged to measure all children if some children were not present at the time of the interview.¹³ The second most important reason for attrition is that not all parents knew the month of birth of their child, which is a requisite for finding the relevant height from the reference distribution to construct the HAZ. We disregard 14 child

¹⁰In January 2011, we conducted a large scale quantitative household survey using a closed-form highly structured pilot-tested questionnaire to capture the impact of RIPAT on technology adoption, food security and poverty. The data collection and data entry were closely supervised by us in cooperation with a survey management team from the Economic Development Initiative (a Tanzanian survey company). RECODA assisted in the hiring of a team of local interviewers and data entry clerks. Both the project implementation and the data collection were financed by the Rockwool Foundation.

¹¹The initial target was 12 comparison villages, but only eight villages in the district were comparable to RIPAT with respect to relevant characteristics, e.g. agriculture being the most important economic activity for them.

¹²During pilot testing of the survey, we became aware that some RIPAT participant did in fact hold more than five acres of land in 2011. To increase comparability, we therefore allowed households in control villages to have up to eight acres of land. We control for land area in all the conditional estimations below and impose a restriction on the number of acres in the robustness section.

¹³They were required to measure at least one child per household where there were children below six years of age.

observations with missing values in the household characteristics and 11 child observations with an absolute HAZ larger than five standard deviations in order to avoid extreme outliers. Furthermore, following the convention in the literature (e.g. Bhutta et al., 2008; de Onis, Blössner and Borghi, 2011; Masset et al., 2012), we focus the analysis on children up to 60 months old, in order to avoid the influence of environmental factors on the heights of the children. This results in a final sample of 335 households with 396 children.

In addition, we interviewed 427 non-participating households in RIPAT villages for a study of diffusion of improved banana cultivation using a stratified random sample (Larsen, 2012).¹⁴ From the households with young children we have HAZ measurements of 195 children, which we use in section 6.1 as an alternative comparison group. We apply sampling weights to account for stratification. See table 11 in the appendix for an overview of the sample composition and the different reasons for attrition. In section 5.1 below, we address the attrition in various different ways. We show that results are robust to the use of a Heckman selection correction model to account for the fact that the probability of being measured may not have been random. Furthermore, we show that results are also robust to the inclusion of outliers in HAZ and of children aged 61-71 months.

3.2 Summary statistics

In table 1 we list the mean values of key child, parent, household and village characteristics for the RIPAT households in column (1), and the corresponding values for the comparison households in column (2). In column (3) we present wild cluster bootstrap p-values from two-sided t-tests of whether the means differ between RIPAT and comparison households, clustered at the village level.¹⁵

Looking at the characteristics of children in the sample, we see that the overall HAZ

¹⁴Non-participating households were therefore oversampled in villages with a larger degree of diffusion, and households growing improved bananas were sampled with a slightly higher probability than other households (see Larsen (2012) for details of the sampling scheme).

¹⁵We use wild cluster bootstrap-t p-values for all inferences in the paper because we only have 16 clusters (villages), and with few clusters the usual asymptotic theory does not apply (Cameron, Gelbach and Miller, 2008).

is about one standard deviation below the WHO reference population mean, indicating that they suffer from undernutrition in general. One in four children are stunted, and although this might appear to be a high level of prevalence, it is well below the regional stunting prevalence rate of 44 percent as found in the 2010 Demographic and Health Survey (DHS, 2010). This indicates that the children in our sample are somewhat better off than the regional average, possibly reflecting better socio-economic conditions, as the area is reasonably fertile and in close proximity to Arusha town.

Slightly more than half of our sample were girls, and most were children of the household head. Their fathers were typically in their late 30s, while their mothers were around 30 years old. Both parents had between six and seven years of schooling, corresponding to having almost completed primary education. However, there is a tendency for the parents in RIPAT households, especially the mothers, to be older and slightly more educated than the parents in comparison households.¹⁶

The children lived in households with, on average, five other household members, these being fairly evenly distributed across the four age groups shown. In 2006, prior to the commencement of the RIPAT project, the households owned on average three to four acres of land. The math skills of the farmers interviewed were tested through two simple math questions; less than half answered both of them correctly. We have also included the average historical rainfall level at the household level,¹⁷ since the households mainly rely on rain-fed agriculture. In accordance with the village selection criteria of suitable agricultural conditions, RIPAT villages received more rain than the comparison villages. Both RIPAT households and RIPAT villages were more likely to have participated in a development project in the past than their comparison equivalents. However, these differences are not statistically significant. The RIPAT villages were situated further away from the main local market and they were less likely to have a secondary school, and although these differences are insignificant they suggest

¹⁶When we have not been able to identify the parents, we have imputed the sample mean following Duflo (2003).

¹⁷We used interpolated data on yearly precipitation on a one-by-one kilometer grid measured in mm from the period 1950-2000 and available from <http://www.worldclim.org/>. The rainfall data were matched to households using GPS coordinates.

that the program allocation procedure targeted wetter and more remote villages.

From table 1 it is thus clear that there are some differences in observables between participating and comparison households, although only few of these are significant at a conventional level. We return to these below. It is, however, still important to account for these characteristics in the analyses below in order to increase comparability.

4 The identification strategy

The participation selection process at both village and individual levels suggests that more motivated farmers from poorer villages were likely to become project participants. Furthermore, no baseline data were collected prior to the intervention, and therefore we cannot rely on standard difference-in-differences estimates to establish counterfactual outcomes. We want to identify the impact of household participation in RIPAT on the nutritional status of children measured by their height-for-age z-scores (HAZ). To find an unbiased estimate of the average treatment effect, we therefore need to account for project placement and self-selection. We do so by employing the identification strategy of Duflo (2003), which exploits the fact that height is a stock variable reflecting accumulated nutrition and infections since conception.

This identification strategy relies on the findings in the medical literature that the *in-utero* period and the first two years of life are critical periods for childhood development. The length of new-born infants and the height of young children is considered to be more sensitive to the nutritional intake than the height of older children (Martorell and Habicht, 1986; Ruel, 2001), and stunting at birth or in early childhood is found to be a strong predictor of later childhood stunting (Adair, 1999; Saleemi et al., 2001). Thus, because stunting is persistent, the HAZ of older children represents reliable recall data, as it is a biological marker of their past nutrition in early childhood (Hoddinott et al., 2013; Victora et al., 2010). We exploit this fact to identify the impact of RIPAT with a difference-in-differences estimator: the HAZ difference between young RIPAT and comparison children conceived after the phase-in of the project, net of the difference

for the older children. The difference in height-for-age of the older children captures any systematic differences in nutritional status between RIPAT and comparison children before a potential impact of the project. That is, it captures nutritional-level differences due to the non-random selection and thereby accounts for the selection into the project.

In other words, the idea of the identification strategy is to estimate whether children who were conceived after project phase-in were taller for their age than their older peers who were conceived earlier, relative to a similar cohort difference between younger and older children from comparison households. The young RIPAT children would have been fully exposed to potential benefits of the project during the first critical 1,000 days of their lives, while the older RIPAT children would only have been partly exposed, or not at all. The difference in their HAZ can be assigned to RIPAT after accounting for general time variation in nutrition and infections by deducting the HAZ difference between young and older comparison children. The identifying assumption is that—in absence of treatment—the height-for-age of treated and comparison children would follow a common growth profile.¹⁸ We capture a growth profile curvature by controlling for age in months quadratically. Our results could be misleading if the growth profiles differ between treated and comparison children in absence of treatment. Below we therefore investigate whether there were any confounding time-varying differences between participating and comparison households, such as changes in fertility patterns or different coping abilities in times of drought (see section 6).

We estimate the average treatment effect of RIPAT with ordinary least squares (OLS) using the specification in equation (1).

$$Y_i = \beta_1 RIPAT_h + \beta_2 young_i + \beta_3 RIPAT_h \cdot young_i + C_i \delta + P_i \phi + X_h \eta + W_v \gamma + \varepsilon_i \quad (1)$$

Y_i is the outcome for child i in household h in village v . The variable $RIPAT_h$ indicates whether household h had ever participated (i.e. including those that dropped

¹⁸This corresponds to the common trends assumption in a classical difference-in-differences set-up.

out) in a RIPAT farmers' group; $young_i$ indicates whether child i was younger than a certain threshold described below; and $RIPAT_h \cdot young_i$ gives the interaction between the two last variables. Thus, β_3 will give the estimate of the average treatment effect of RIPAT on the nutritional status of young children, net of selection. We control for child characteristics, denoted as C_i , parent characteristics, P_i , household characteristics, X_h , and village characteristics, W_v , all of which are listed in table 1. Age in months is included quadratically. We take the logarithm of acres of land owned in 2006. Finally, we allow for errors to be correlated within villages, $\varepsilon_{i,v}$.

We have a small subsample of households with measurements of both young and older siblings. This allows us to also provide estimates with household fixed effects instead of parent, household and village characteristics as a simple robustness check.¹⁹

There is some flexibility in how we define the relevant threshold for the *young* dummy, as it depends on when we can expect an impact of RIPAT on food security to have taken place in the households. Food insecurity in this area is highly seasonal, and is only pronounced in the lean seasons (January to May).²⁰ This implies that the earliest time we can expect an impact on nutrition of pregnant women and young children is in the first lean season after project start, January-May in 2007. Hence, we define the *young* dummy such that it is equal to one for children conceived in January 2007 or later (henceforth referred to as “young” children). Regardless of the choice of threshold, some children classified as *old* may also be affected by the improved nutrition. If there is any such catch-up growth, it will lead to an underestimation of the impact. We return to the choice of threshold in section 5.1.

5 Results

Before turning to the estimation results we compare the distributions of HAZ presented in figure 5.1 for the *old* and *young* children separately. We have conditioned

¹⁹We do not include parent characteristics in the fixed effects regressions, as there is naturally very little variation within households.

²⁰We define the span of the lean season according to self-assessment by the households in the sample. The majority of households mentioned the months January-May as part of the “worst period in terms of having enough food for everyone in your household [during 2010]”.




Figure 5.1: Distributions of the HAZ

on child, parent, household and village characteristics to reduce noise. We see that the conditional distribution of HAZ for the *old* RIPAT children is closely aligned to that of *old* comparison children, suggesting that these children are indeed highly comparable. For the *young* children the RIPAT distribution is clearly shifted to the right of the comparison distribution. Obviously, this graphical inspection does not constitute a formal test; however, it does suggest that not only were the *young* RIPAT children taller for their age than the comparison children *on average*, but it appears that the intervention has affected the entire HAZ distribution of *young* RIPAT children, in particular the lower tail.

Table 2 shows OLS estimation results for the average treatment effect of RIPAT using the econometric specification given in equation 1. Columns (1) to (3) present estimated impacts on the height-for-age z-score (HAZ) of *young* children in participating households, hence the impact on the mean value of the HAZ distribution. To analyze whether RIPAT affects the lower part of the distribution and decreases the prevalence of severe undernutrition, we also provide estimates for the impact on the likelihood of children being stunted in column (4) using the linear probability model. The coefficient to the *RIPAT* and *young* interaction term gives an estimate of the average treatment effect of RIPAT on the HAZ or the probability of being stunted among the younger children who grew up under the influence of RIPAT. In column (1) we show the unconditional estimates, in column (2) we control for child, parent, household, and village characteristics, and in column (3) we allow for household fixed effects.

The unconditional estimate of the impact of RIPAT on HAZ is an average improvement of 0.57 standard deviations (SD) of the WHO reference distribution. When we control for child, parent, household, and village characteristics, the estimate of the impact increases to 0.88 SD. This means that young children in RIPAT households were 0.88 SD taller than their peers in comparison households, controlling for any pre-project differences among the older children. When we include household fixed effects to account for unobserved household characteristics, the point estimate further increases to 1.38 SD. The fact that we still find a positive impact after the introduction of household fixed effects suggest that the results are not driven by unobserved differences in the selection into the project between households with young and older children. However, fixed effect estimation relies on variation in a relatively small subset of the sample, as only 21 RIPAT households and 19 comparison households had both young and older children in the sample, and we therefore only include it as a robustness check of the conditional estimates.

Because RIPAT is a village intervention, we cluster standard errors at the village level, and the corresponding significance levels are reported with the customary use of asterisks. However, since we only have 16 villages and thus 16 clusters, the standard asymptotic theory cannot be applied for inference. We therefore also report p-values in square brackets based on wild cluster bootstrapped t-statistics for the impact coefficients, as suggested by Cameron, Gelbach and Miller (2008).

Turning to the impact on stunting in column (4), we see that the average impact on height-for-age also translates into an impact in the lower part of the HAZ distribution, as suggested by figure 5.1. Compared to children in control villages, we find that *young* RIPAT children experienced a reduction in the prevalence of stunting of 17.6 percentage points, significant at the ten percent level. We have less statistical power compared to our results for HAZ, since we discard information by reducing the continuous HAZ to a binary variable. However, this does suggest that the nutritional improvements also reach children in the lower parts of the height-for-age distribution who suffer from severe undernutrition.

When we measure the impact of RIPAT on HAZ, we measure the impact on a nutritional stock (height). We expect RIPAT to affect the stock through improvements in the nutritional flows generated from the ongoing agricultural production.²¹ This suggests that the effect of RIPAT on height-for-age should increase with the duration of exposure to RIPAT. The longer children were exposed to improved nutrition, the more the impact accumulates in their stock, i.e. their height. On the other hand, we are analyzing the impact of an agricultural project that was gradually phased in, and for which there was a lag from the onset of new project activities to a tangible nutritional outcome from the fields or the livestock. Children born early in the project period therefore received a weaker nutritional improvement during their first 1,000 days than children born later, when new agricultural technologies could potentially have been adopted, and this factor works in the opposite direction.

In table 3 we present estimates from a model that allow for cohort-specific impacts: instead of a *young* indicator we include age indicators for the years zero to three, along with the RIPAT indicator and their interaction terms. Four-year-old comparison children then form the reference group. Overall, the impact is driven by the one- and two-year-olds. Both groups have estimated impacts on the HAZ of 1 SD (see column 2), which suggests that the accumulation of impact we expected to see in the two-year-old children is offset by the gradual impact of the project. We see no significant impact on the youngest children, possibly because the differences in nutritional intakes between RIPAT and control children (and their mothers while they were *in utero*) are not yet detectable in the height measurement. The three-year-old children belong to the group of *old* children and, as expected, there is no significant difference between the three- and four-year-old RIPAT cohorts relative to the comparison cohorts.²² The latter result thus supports our common growth profile assumption for treated and comparison groups

²¹Instead of measuring the impact of RIPAT on the height-for-age of the children, it would be more direct to measure the impact on the nutritional intake of the children in every period. However, it is difficult to collect diary data with precise measurements of calorie and micronutrient intakes, so it is convenient to use the height-for-age as a simple summary measurement of the nutritional status of the child.

²²Actually, the *young* threshold is 39 months, i.e. three years and three months, so 22 of the 90 three-year-old children are considered *young* in the main analysis.

prior to any impact.

Mechanisms

Although we cannot pin down the exact channel through which RIPAT has influenced the nutritional status of young children, we can examine the most likely chain of events, namely whether our sub-sample of RIPAT households are more likely to have adopted the variety of technologies provided through the basket of options and whether they are also more likely to be food secure.

Elsewhere we have examined the impact of RIPAT on food security and poverty for the full sample of RIPAT households using a gradual roll-out of the project into a nearby district to account for selection, (Larsen and Lilleør, 2014). We found robust impacts on a variety of food security measures, but no impact on poverty. In this paper, we focus on the sub-sample of households with young children, and assess the impact on their height-for-age with an identification strategy that is only valid for this particular outcome. To confirm that high levels of adoption and food security are also more pronounced for the sub-sample of RIPAT households with young children, we therefore resort to a simple cross-sectional comparison.²³

In table 4, columns (1) and (2), we show the average adoption rates (and standard deviations) for RIPAT and comparison households for six of the central technology options and the number of different crops cultivated in 2010, in order to capture crop diversification. To test the difference in adoption between RIPAT and comparison households we regress the technologies on a RIPAT indicator and household and village characteristics.²⁴ Column (3) presents estimates for the RIPAT indicator with cluster standard errors in parentheses and wild cluster bootstrap-t p-values in square brackets. We see that there are high rates of technology adoption. Our RIPAT house-

²³In Larsen and Lilleør (2014), we estimate the average treatment effect using simple cross-sectional comparisons between treatment and control groups, matching estimators, and a difference-in-differences estimator exploiting the gradual roll-out. The findings are reasonably robust across estimation methods, suggesting that selection into the project is not a major driver of results. We are therefore confident that when we employ simple cross-sectional comparisons to this subsample, it will give a good indication of whether there was also increased adoption and improved food security levels in the sub-sample of RIPAT households with young children.

²⁴We include age and education of the household head instead of that of the parents.

holds are significantly more likely to have been growing improved banana varieties, to have been keeping improved breeds of chickens and goats, to have been practicing zero-grazing with their livestock by keeping their animals in smaller enclosures and feeding them, to have been participating in savings groups,²⁵ and to have had a larger degree of crop diversification. We find no significant difference in the adoption of fruit trees. It is worth noting that both perennial crops (like banana) and improved livestock technologies (poultry providing eggs and meat, and milking-goats providing milk) enhanced production smoothing over the agricultural cycle and thereby also helped to facilitate the smoothing of food consumption over the year.

The fact that the RIPAT farmers practiced zero-grazing among their livestock by keeping the animals in smaller enclosures may have reduced the exposure of young children to disease, since they would have been less exposed to animal excrement. Similarly, in our data we can see that RIPAT households were more likely to have a roof over their pit-latrines as recommended by RIPAT facilitators (along with village and government officials), and this would have reduced the spread of bacteria through flies. A reduced exposure to disease could therefore be another channel through which children's growth and thus height-for-age is positively affected (Bhutta et al., 2008; Adair et al., 2013).

Increased adoption of new agricultural and animal husbandry technologies should lead in turn to higher levels of food security. In table 5, columns (1) and (2), we list RIPAT and comparison household means (and standard deviations) for eleven different outcome measures of food security, and in column (3) we show RIPAT regression coefficients from regressions of the food security measurements on a RIPAT indicator and household and village characteristics, as in table 4.

We find that RIPAT households experienced a significantly shorter hunger season than comparison households. When asked about the worst period in terms of having enough food during the previous 12 months, RIPAT households reported an 11 per-

²⁵Although later RIPAT projects (RIPAT 2-4) actively used Village Savings and Loans Associations as one of the basket options, membership in external savings groups was simply encouraged in RIPAT 1, the project which we study here.

cent shorter period than comparison households, *ceteris paribus*. Similarly, we see that RIPAT households are 16 percentage points less likely to have experienced any hunger during the 12 months before the interview. We measure hunger using the Household Hunger Scale (HHS)²⁶ using three different reference periods: the previous four weeks, and the worst and the best months during the previous 12 months. We see that hunger is reduced by 0.6 on the HHS (corresponding to 32 percent) during the lean season (worst month), while there is no significant impact on the level of food security in the best period of the year or the four weeks preceding the interview, where the prevalence of hunger was relatively low. Next, we look at whether the children in the household had at least three meals per day in the best and worst periods of the year as well as in the previous four months. The coefficients are all positive, and in particular the estimated difference for the worst month is large, this being where we also see the greatest room for improvement. However, once the small number of clusters are taken into account, there is not enough power to yield statistically significant results. This is also the case for the analysis of whether the households consume meat, eggs or dairy products, all sources of animal protein.²⁷

All in all, this suggests that the positive impact on the height-for-age of young RIPAT children is likely to come about through higher levels of technology adoption promoting higher levels of food security in the lean season of the year. Not being exposed to hunger spells seems to have long-lasting consequences for the growth curves of these young children. The effect may be reinforced by less exposure to animal- and excrement-related bacteria. We also examined whether RIPAT households had lower poverty levels than the comparison households, but find no clear evidence of such differences (see appendix table 12). This suggests that RIPAT did not bring a large income

²⁶The HHS is a modern food security instrument developed by US Aid to ensure cross-cultural comparability. It has been validated in five sub-Saharan African countries. It is based on three questions asking whether, due to lack of resources, anyone in the household 1) went to sleep at night hungry; 2) had no food to eat of any kind in the household; and 3) went a whole day and night without eating. The response codes are 0: never; 1: rarely or sometimes; 2: often. The HHS is simply the sum of the responses to the three questions resulting in an index from zero to six where zero corresponds to "no hunger" and six corresponds to "severe hunger". See Ballard et al. (2011).

²⁷In the full sample we find statistically significant impacts on almost all of these measures (Larsen and Lilleør, 2014).

effect with it, but rather that the main impact came about through better smoothing mechanisms shielding households and more specifically children against hunger in the lean season.

5.1 Robustness of the results

In analyses like the one in this paper, one worries whether the results might be driven by systematic errors or decisions concerning data. In this section we therefore analyze whether our results are robust to accounting for attrition, to different thresholds of the *young* indicator, and to changing the sample selection with respect to children's age, number of acres owned, outliers and data quality considerations.

5.1.1 Attrition

Not all children living in the surveyed households were measured. If there are systematic differences in which children were measured across RIPAT and comparison households, this could potentially affect our results. We address this issue with a Heckman selection model (Heckman, 1979) where we exploit the variation in enumerator meticulousness as an instrument for the probability of a child being in the sample. The 25 enumerators were instructed to measure at least one child in each household of zero to five years of age, and preferably all available children. The instrument is constructed as the average share of children that the enumerator measured in other households, not including the household in question.²⁸ In this way, the instrument is unaffected by household-specific characteristics that determine whether a child is measured or not, and as we can see from panel B in Table 6 it is highly correlated with the probability of being measured. Estimation results are presented in table 6.²⁹

The estimates for the impact of RIPAT on HAZ when correction is made for selection in measurement are shown in panel A. We see that the results are robust and not

²⁸This share varied between 0.40 and 0.95.

²⁹Alternatively, we can also just include enumerator dummy variables as instruments. In that case, we cannot obtain convergence of the maximum likelihood estimator, but we obtain similar results to the ones presented in table 6 if we apply a two-step estimator instead with these alternative instruments. Results available upon request.

driven by selection in who was measured, as we find a large and significant impact of RIPAT just 0.1 standard deviations lower than the corresponding results presented in table 2.

Panel B shows the estimates from the selection equation, and here we see no significant differences between RIPAT and comparison households in the likelihood of being measured, either for young or for older children. Young children are more likely to have been measured than older children, probably because they were more likely to be around at the time of the interview.

5.1.2 Threshold for the *young* indicator

Next, we turn to the choice of threshold for the *young* indicator. We expect some lag from the introduction of new agricultural methods on the common demonstration plot to a change in the agricultural practices of the households and a subsequent improvement in the food security of the household. Following this reasoning, the threshold of the *young* indicator should be later in time than January 2007. On the other hand, with a conception threshold in January 2007, all children born before October 2007 are classified as *old* even though they lived the main part of their first two critical years of life after the implementation of RIPAT. This would speak in favor of an earlier threshold. In table 7 we show results where we move the threshold between May 2006, the start of RIPAT, and January 2008, the second lean season after the start of RIPAT. All estimated impacts are within the confidence bounds of their counterparts in table 2, and apart from column (5) they are all statistically significant at the ten percent level. The latest threshold (January 2008) results in the lowest and least significant impact. This estimate will be downward biased if children born before October 2008 were affected by the project, which may very well have been the case. It is reassuring that the positive impact found is not specific to a certain choice of threshold for the *young* indicator.

5.1.3 Alternative sample selections

In the main analysis we consider children up to 60 months of age, which is common practice in the nutrition literature. As we have height measures for children up to 71 months of age, we investigate whether the results are robust to the inclusion of these older children. The results are shown in table 8 column (1), and we see in panel A that the estimated impact on HAZ is reduced to 0.6 standard deviations and is significantly different from zero at the ten percent level. Similarly, the estimated impact on the stunting indicator (shown in panel B) is reduced but the reduction is statistically insignificant; however, it is still economically significant, with an estimated 11 percentage points reduction in the prevalence of stunting.

Though RIPAT participants are required to have between one and five acres of farm land, our data show that this requirement was violated in many cases in RIPAT I. Comparison households were therefore chosen among households with one to eight acres to mirror the actual distribution of farm land holdings among RIPAT households. However, 19 RIPAT households and four comparison households reported having less than one or more than eight acres of land. In column (2) of table 8 we exclude the 28 children from households with areas of farm land that lay outside the range of one to eight acres. In panel A, this results in a stronger estimated impact on HAZ of 0.9 standard deviations, significant at the one percent level. This result indicates that the impact of RIPAT was higher among households targeted for the intervention, which is in line with intuition.³⁰ The estimated impact on the prevalence of stunting is unaltered.

When we calculate the height-for-age z-scores, a few observations have very extreme values. In the main sample we have disregarded children with a HAZ larger than five in absolute value (11 observations). Column (3) shows that we obtain even stronger results if we include outlier observations. It is particularly noteworthy that the estimated impact on the prevalence of stunting is higher and more significant than the main result. The stunting indicator is not affected by large outliers in HAZ, so for this specification it can indeed be argued that the outlier observations should be

³⁰Results are very similar if we further restrict the sample to only include the 268 households which had between one and five acres of land.

included.

In the last two columns we consider the data quality of the HAZ. Children below 24 months of age should be measured recumbent, while children above 24 months of age should be measured standing. Not all enumerators followed these guidelines,³¹ so in column (4) of table 8 we present regression results excluding children measured in the incorrect position for their age. We obtain fairly similar results to those in table 2; however, the wild bootstrap p-values suggest that the estimated impacts on HAZ and the prevalence of stunting cannot be distinguished from zero.

The calculation of the HAZ is based on the age of the child, and there might be uncertainty about parents' recall of their children's birth dates. For a sub-population of children we have their birth dates confirmed by official clinic cards, and regression results in column (4) show that the results are still significant and within one standard error of the estimates in table 2 when we consider this sub-population.

In general, the impact of RIPAT on height-for-age is fairly robust to the selection of the sample, with results ranging from 0.6 to 1.2 standard deviations, all but one of them being significant at the ten percent level. Though we have less power when we consider the prevalence of stunting, we also consistently find large impacts of participation in RIPAT.

6 Possible alternative explanations

Our identification strategy relies on the standard assumption of treatment and control groups sharing a common trend in the absence of treatment. In our setting this translates into an assumption that children from RIPAT and comparison households would share a common growth profile in the absence of treatment. That is, our difference-in-differences set-up allows for differences in child nutrition *levels* between RIPAT and comparison households, but not for differences in *trends* or in time-varying differences not caused by the intervention. If such differences exist, our results could be mislead-

³¹If a child was measured recumbent though older than 24 months or *vice versa*, we adjusted the measurement by 0.7 cm, in accordance with WHO guidelines (WHO, 2006).

ing. We study three potential factors which could lead to differences in the growth profiles, namely time-varying differences between RIPAT and control villages, differences in fertility patterns between RIPAT and comparison households, and differences in households' coping capabilities in times of drought.

6.1 Village differences

In our main analysis above we compare children in RIPAT households with children in comparison households in control villages. If the two groups of villages were differentially exposed to shocks, e.g. there was a serious drought in 2009 which could have hit the comparison villages harder than the treated villages, or vice versa, our impact estimates may be confounded. We address this issue by comparing the RIPAT children to other children *within* the RIPAT villages who did not live in participating households. The data we have from a stratified random sample of non-RIPAT households within RIPAT villages allow us to examine whether the estimated impacts on HAZ and stunting found above are in fact driven by time-varying village level differences. If so, we should expect to see no difference in nutritional levels between children from RIPAT and non-RIPAT households within the RIPAT villages. Nevertheless, we have to keep in mind that there has been considerable diffusion of technologies within the RIPAT villages (e.g. 13 percent of non-RIPAT farmers in RIPAT villages grow improved banana varieties; see Gausset and Larsen (2013)). A comparison between households within the RIPAT villages may therefore underestimate an impact.

Table 9 corresponds to table 2 above, only using children from non-RIPAT households in RIPAT villages as a comparison, rather than children from comparison villages. Standard errors are again clustered at village level (note that now there are only eight villages) and p-values based on wild cluster bootstrapped t-statistics are shown in square brackets. Furthermore, we have added an additional column allowing for village fixed effects, column (4). The estimated impact on HAZ is much in the same order of magnitude as in table 2 and appears slightly more robust across specifications, but the small number of clusters affects the bootstrapped p-values and we have

less power. The estimated impact on the stunting indicator increases to a 26.7 percentage point reduction in stunting. It is reassuring that we find the same positive impact regardless of comparison group. This rules out the possibility that the estimates are driven purely by differences in village-level shocks.

6.2 Fertility patterns

Could project participation itself lead to endogenous changes in fertility patterns and thus in cohort composition among the participating households relative to comparison households, such that the estimated impact found above is a result of this phenomenon?

First, if RIPAT induces households to have fewer children, that would imply that households would have more resources per child, which could have led to an improvement in the nutritional status of the children born. However, since we control for the number of household members between zero and five years of age, this cannot be the mechanism for the impact we find.

Second, if participation in the project changed the timing of fertility, this could potentially affect the group composition of *old* and *young* RIPAT children *vis-a-vis* the comparison children. Table 1 shows that the group of RIPAT children were on average slightly older (three months) than the group of comparison children. As the HAZ trends downwards for undernourished children, the difference-in-differences estimate will be upward biased if *young* RIPAT children were on average younger than *young* comparison children or if *old* RIPAT children were older than *old* comparison children. We have tested whether the average ages within the *old* and *young* groups are correlated with RIPAT, and we find no significant correlations. We further test the composition of the age cohorts by regressing age indicators on a RIPAT indicator, while controlling for household and village characteristics. Coefficients and wild cluster bootstrap confidence bounds for the RIPAT indicator are presented in figure 6.1 for each age indicator. None of the age groups are significantly under- or overrepresented

among the children in RIPAT households.³²

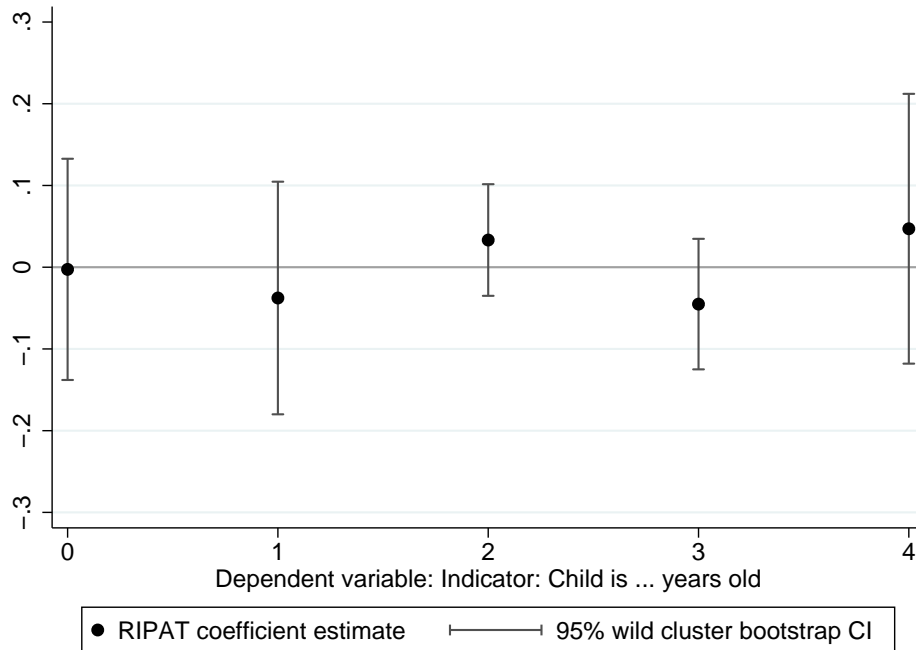


Figure 6.1: Cohort composition

Notes: OLS estimates and wild cluster bootstrap-t confidence intervals for the coefficient to the RIPAT indicator in regressions with age indicators for ages 0, 1, 2, 3, and 4 years as dependent variables. The regressions also control for household and village characteristics. The wild bootstrap-t procedure is clustered at the village level, following Cameron, Gelbach and Miller (2008), and confidence intervals are constructed by finding the highest possible null hypothesis (from below) that is rejected at the five percent level and imposing symmetry.

Third, if the project affected timing of conception over the year, RIPAT children might have been differently exposed to the lean season relative to the comparison children, which again could affect our results. Hence, we run twelve regressions with month of birth indicators as dependent variables using the same specification as in equation 1.³³ With this difference-in-differences specification we test whether potential differences in the seasonal timing of fertility between *old* RIPAT and *old* comparison children persist among the *young* children. If the seasonal pattern changed remarkably, it could be driving the results. As can be seen in figure 6.2, the only significant difference we find is that *young* RIPAT children are less likely than *young* comparison children to be born in November relative to any differences among their older peers.

³²This also holds with OLS confidence intervals if clusters are not taken into account.

³³All children, household and village characteristics are included except the child's age.

For this difference to be driving our results it would need to be very unfavorable to be born in November as compared to other months of the year.³⁴ Our results are robust to excluding children born in November (results available upon request) and hence, we do not expect this small difference to be driving the large impact that we find.

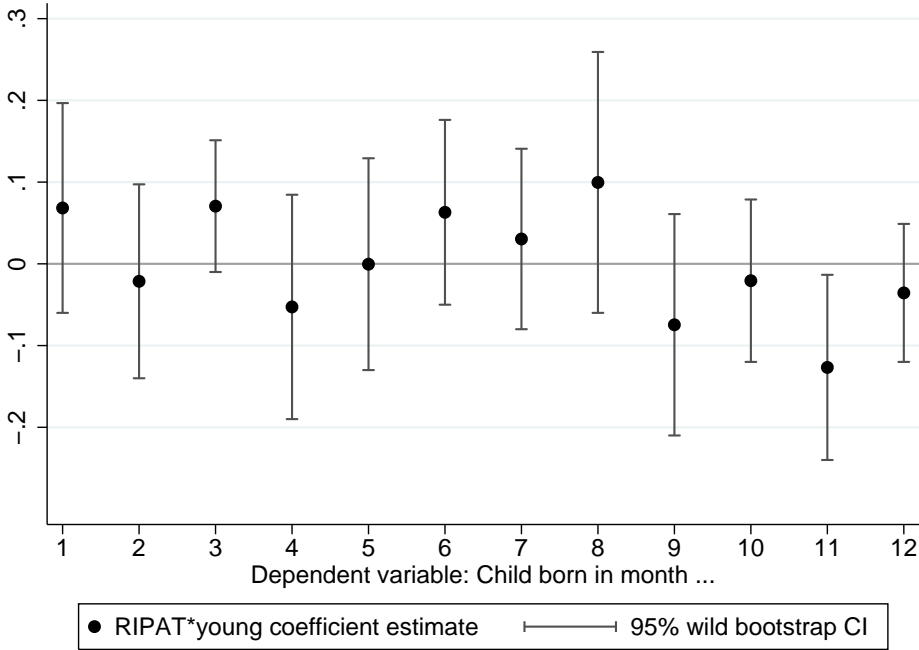


Figure 6.2: Seasonality in fertility

Notes: OLS estimates and wild cluster bootstrap confidence intervals for the coefficients for the RIPAT and young interaction terms in regressions where the dependent variables are indicators for the month when the child was born. 1 corresponds to January, 2 is February, and so forth. The regressions also control for a RIPAT indicator, a *young* indicator, and child, parent, household, and village characteristics, excluding the age of the child. Wild cluster bootstrap-t confidence intervals are constructed as in figure 6.1. We do not correct for multiple hypothesis testing.

Taken together, this suggests that the positive impacts found on height-for-age using the cohort difference-in-differences estimator are not driven by changes in fertility patterns or cohort composition during the project period.

³⁴It is difficult to hypothesize whether a child is better off being born in November compared to June, say. In the former case the child is exposed to the lean season during the first trimester *in utero*, while in the latter case the child is exposed during second and third trimesters. The timing of the weaning period will also be different, and this might also play a role.

6.3 Capabilities for coping with drought

Finally, the common growth profile assumption would also be violated if the RIPAT and comparison households had been subject to different shocks or had coped with a common shock in different ways, regardless of project participation. Above we showed that we can reject the possibility that our main results are driven by a difference in *village*-level shocks, since we obtain similar results when using comparison children from RIPAT villages as opposed to control villages.

With respect to coping with shocks at the *household* level, we should keep in mind that RIPAT aims at reducing vulnerability to drought shocks by introducing drought-resistant crops and production-smoothing technologies, so we should in fact expect that RIPAT households would have become better at coping with drought shocks. But we need to address the concern that households who selected into RIPAT may *initially* have had different coping strategies than the comparison households. Coupled with the drought in 2009, this could have driven the impacts that we find.

We address this potential selection bias in three ways. First, we investigate whether the impact is driven by any of the observed differences in parent and household characteristics between RIPAT and comparison households in control villages. Table 1 shows that parent characteristics differ significantly between RIPAT and comparison households in terms of father's and mother's age. Furthermore, mother's education, which is often a strong predictor of children's health, is also marginally different, with children in RIPAT households having more educated mothers. If, say, older or better-educated mothers were better at nourishing their children during the 2009 drought, we would overestimate the impact, since RIPAT mothers were on average better educated.

We demean these key parental variables and interact them with the *young* indicator, the *RIPAT* indicator and their interaction term respectively, to allow for the treatment effect to depend on, for example, mother's age. This results in the following specification, where Q_h represents one of the demeaned parent or household characteristics, $Z_{i,h,v}$ comprises all child, parent, household and village characteristics, and $\zeta_{i,v}$ is an

error term with intra-village correlation:

$$\begin{aligned}
 Y_i = & \mu_1 RIPAT_h + \mu_2 young_i + \mu_3 RIPAT_h \cdot young_i + \mu_4 RIPAT \cdot Q_h \\
 & + \mu_5 young \cdot Q_h + \mu_6 RIPAT \cdot young \cdot Q_h + Z_{i,h,v} \vartheta + \zeta_{i,v}
 \end{aligned} \tag{2}$$

When we allow for the relative difference between *young* and *old* children to depend on the age of the mother, the average impact estimate captured by μ_3 is unaffected by any possible influence of RIPAT mothers' age on the nutrition of their young children during the drought spell in 2009.

Table 10 shows estimates of equation (2) with interactions with parental variables in columns (1)-(3). The estimate of the mean impact of RIPAT relative to comparison children in control villages is remarkably stable across these columns, confirming that the impact found above is not driven by any of the differences in observed parental characteristics. However, we see that the impact is negatively correlated with father's age, indicating that younger households benefited more from RIPAT than older households did.

Second, using the same method, we examine whether three of the selection criteria could be driving the results. Households self-selected into the project, but had to fulfill a land ownership criteria. Villages were partly chosen based on suitable agricultural conditions, including sufficient rainfall. We interact the difference-in-differences variables with three variables capturing this selection: historical rainfall, log of farm land acres in 2006, and a proxy for self-selection using participation in other projects in the past. From table 10 we see that differences in land ownership (column (4)) or prior participation in other projects do not alter the estimated impact of RIPAT on the HAZ of *young* children. However, turning to rainfall in column (6), we see that part of the impact of RIPAT on HAZ is driven by a positive interaction with rainfall, reducing the average effect of RIPAT on the HAZ to 0.77 SD. The impact of RIPAT is increased by 0.01 SD of the HAZ per additional millimeter of historical rainfall. Given that the majority of the technology options also rely on adequate rainfall, especially in the phase-in period, this is not a surprising finding.


Third, selection into the project could still have been based on intrinsic *unobserved* differences in strategies for coping with shocks between participating and comparison households. Due to the drought in 2009, such differences might have led to the improvements in height-for-age that we find for the young RIPAT children. However, intrinsic differences in strategies for coping with drought between RIPAT and comparison households should then also be detectable when comparing the HAZ of children exposed to an earlier drought spell. To measure weather shocks, we follow Harari and La Ferrara (2013) and examine monthly Standardized Precipitation and Evapotranspiration Indices (SPEIs) for the geographical area under study, using the average of the four preceding months and considering values of the SPEI below one SD as negative climate shocks. We consider March to June to be the main growing season based on the Food and Agriculture Organization crop calendar.³⁵

Figure 6.3 shows SPEIs for the period 2004 to 2011 with three data points per year: the four-month average SPEIs for the growing season March-June and for the non-growing seasons July-October and November-February.³⁶ It can clearly be seen from figure 6.3 that the growing season in 2009 was particularly dry. But we also see that the area was hit by a drought during the growing season in 2006.

This implies that if RIPAT and comparison households initially had different coping strategies, we should expect to see differences in the HAZ of children conceived just before or during 2006. These are precisely the children we define as *old*, and where we find no significant difference in their height-for-age between RIPAT and comparison children. Thus, we argue that the improved nutrition among the young RIPAT children cannot be driven by differences in drought coping strategies across treated and comparison households *a priori*. On the contrary, we propose that RIPAT farmers had improved their ability to cope with the 2009 drought through the adoption of drought-resistant crops and production-smoothing technologies. The magnitude of our estimated average treatment effect on HAZ might therefore have been consider-

³⁵<http://www.fao.org/agriculture/seed/cropcalendar/welcome.do>

³⁶The graph is from a grid covering half of the villages in our sample; the graph from the neighboring grid covering the remaining villages is very similar and is available from the authors. The global SPEI database can be found at <http://sac.csic.es/spei/database.html>.



Source: The global SPEI database

Figure 6.3: Standardized Precipitation and Evapotranspiration Index

ably smaller if the area had experienced years of bumper harvest and thus little food insecurity and no hunger spells. In that sense, the 2009 drought has increased the degree of variation in our data, enabling us to identify a larger impact on nutrition.

7 Discussion

In this paper, we have estimated the impact on early childhood nutrition of an holistic agricultural intervention aimed at improving food security and poverty among small-holder farmers. Given the widespread prevalence of severe undernutrition resulting in stunted growth and the relatively recent acknowledgment of its many long-term adverse implications, combating undernutrition of unborn and infant children has become a very important subject that attracts attention from both researchers and policy-makers; see for example the recent Lancet reviews by Bhutta et al. (2008); Victora et al. (2008); Ruel and Alderman (2013) and the Cost Of Hunger in Africa report by African Union Commission et al. (2014).

The RIPAT program studied here is an agricultural intervention. The specific intervention that we examined did not have a direct nutritional aim, but rather an overall aim of improving food security: a nutrition-sensitive intervention in the terminology of Ruel and Alderman (2013). RIPAT is a broad intervention with a strong focus on im-

proving drought resilience through a basket of technology options including crop diversification, perennial crops, conservation agriculture, improved animal husbandry, and land use management. This holistic approach may have been key in improving the nutritional status of young children in the participating households, as it is argued elsewhere that these components help to improve the nutritional quality of farming output (Miller and Welch, 2013). We find that the RIPAT intervention had a significant positive impact of about 0.8 SD on the height-for-age z-scores of young children who had been fully exposed to the project in their early life. Similarly, we see a reduction in stunting prevalence among the *young* group of RIPAT children of around 17 percentage points.

There are two important points to note concerning these impacts. First, they were measured almost five years after the start of the project, which lasted three and a half years, suggesting that these are sustainable impacts, but not necessarily quick impacts. Second, towards the end of the project implementation period, a serious drought hit the area, worsening and lengthening the annual hunger period. This has possibly increased the difference in undernutrition levels found between participating and comparison households, since the intervention was designed to increase the drought resistance of farmers and shield their food production, rather than to boost agricultural output during bumper years.

According to Masset et al. (2012) and Ruel and Alderman (2013), there has been no rigorous empirical investigation showing a significant nutritional impact of an agricultural intervention among young children. Compared to impacts found of more narrow non-agricultural nutritional interventions in Bhutta et al. (2008) and Caulfield, Huffman and Piwoz (1999), the impacts of the RIPAT intervention on HAZ and stunting prevalence are sizable.³⁷ More recent papers by Linnemayr and Alderman (2011) and Powell-Jackson et al. (2014) find no overall effect of a randomized nutrition program in Senegal or a randomized free health care program in Ghana, although the former

³⁷Bhutta et al. (2008) report that the provision of food supplements in populations with insufficient food can increase the HAZ by 0.41 SD, while Caulfield, Huffman and Piwoz (1999) review efficacy trials to improve infant dietary intakes and find improvements in HAZ of 0.04-0.46 SD.

do detect a positive impact among the youngest on weight-for-age z-scores of 0.27 SD (and similar impacts on HAZ).

Based on the limited impacts found in the nutrition literature, one could suspect that height-for-age in itself is a rigid measure that is hard to influence. However, there is ample evidence from conflict-prone areas in Africa which shows that this is by no means the case. Two different African conflicts have been shown to have a negative impact on HAZ of about 0.4 SD among young children exposed to those conflicts (Akresh, Lucchetti and Thirumurthy, 2012; Minoiu and Shemyakina, 2014). In addition, Baez (2011) shows that children who are not directly exposed to conflict can also be negatively affected (their HAZ drops by 0.6 SD) by a large and sudden influx of poor refugees into the local communities.

In fact, the magnitude of our results is comparable in size to large-scale cash transfer programs. The nutritional impacts of RIPAT on young children are comparable to those found by (Duflo, 2003) in assessing the impact on young grand-daughters of extending a generous public old-age pension scheme to low-income families in South Africa. Duflo finds that this increased the HAZ of young girls in the household by more than one standard deviation if the recipient was the grandmother. Similarly, when analyzing the impact on childhood nutrition of a large-scale conditional cash transfer program aimed at increasing both health and education among Mexican children, PROGRESA, Behrman and Hoddinott (2005) find that the prevalence of stunting drops to a third of the level among comparison children.

Such an impact on HAZ does of course not come without a cost. Indeed, the cost per household of the RIPAT intervention is also comparable to PROGRESA, which has an annual cost of approximately USD 300³⁸. The total cost per household for the 3.5-year RIPAT intervention studied here was USD 700.³⁹ Although roughly USD 200 per year per family may be a relatively high cost—and considerably higher than the

³⁸The annual budget was USD 777 million for a program covering 2.6 million families (Behrman and Hoddinott, 2005)

³⁹This was the first in a series of interventions, which have gradually become more cost-effective. The cost per household is now USD 625 for a full three-year project.

average Farmer Field Schools cost⁴⁰—it must be judged against the benefits found here in terms of improved food security and taller children. Based on the findings in the nutrition literature regarding the adverse impact of stunting on health, economic and social outcomes in adulthood, this positive impact is likely to follow the young RIPAT children throughout life. In addition, we expect the adopted technologies to sustain the improvement in food security into the future, also positively affecting children to come.

All in all, this shows that a broad and highly-sustainable agricultural intervention such as the one studied here, building on local resources, needs and constraints and offering a basket of technology options for farmers to choose from, can result not only in sustainable technology adoption and increased food security among farmers, but also in substantial long-term impacts on the lives of the young children in participating households. Indeed, there are reasons to believe that precisely because of the holistic nature of the intervention and its focus on shielding farmers' food production against adverse impacts of drought, the nutritional and thus growth impacts on young children are sizable and larger than those typically found in more narrow nutrition interventions as reviewed in Bhutta et al. (2008) and Caulfield, Huffman and Piwoz (1999). As hypothesized by both Masset et al. (2012) and Ruel and Alderman (2013), our study confirms that there is scope for agricultural interventions in alleviating undernutrition and that they can indeed be very effective.

References

- Aben, Charles, Deborah Duveskog, and Esbern Friis-Hansen.** 2013. "Evaluation of the RIPAT concept." In *Farmers' Choice*, ed. Helene Bie Lilleør and Ulrik Lund-Sørensen. Practical Action Publishing.
- Adair, Linda S.** 1999. "Filipino Children Exhibit Catch-Up Growth from Age 2 to 12 Years." *The Journal of Nutrition*, 129(6): 1140–1148.

⁴⁰Van den Berg and Jiggins (2007) and Waddington, White and Anderson (2014) report most FFS costs to be between USD 20-40 per farmer.

- Adair, Linda S, Caroline HD Fall, Clive Osmond, Aryeh D Stein, Reynaldo Martorell, Manuel Ramirez-Zea, Harshpal Singh Sachdev, Darren L Dahly, Isabelita Bas, Shane A Norris, Lisa Micklesfield, Pedro Hallal, and Cesar G Victora.** 2013. "Associations of linear growth and relative weight gain during early life with adult health and human capital in countries of low and middle income: findings from five birth cohort studies." *The Lancet*, 382(9891): 525–534.
- African Union Commission, NEPAD Planning and Coordinating Agency, UN Economic Commission for Africa, and UN World Food Programme.** 2014. "The Cost of Hunger in Africa: Social and Economic Impact of Child Undernutrition in Egypt, Ethiopia, Swaziland and Uganda." Addis Ababa: UNECA.
- Akresh, Richard, Leonardo Lucchetti, and Harsha Thirumurthy.** 2012. "Wars and child health: Evidence from the Eritrean–Ethiopian conflict." *Journal of Development Economics*, 99(2): 330–340.
- Alderman, Harold, Hans Hoogeveen, and Mariacristina Rossi.** 2006. "Reducing child malnutrition in Tanzania: Combined effects of income growth and program interventions." *Economics & Human Biology*, 4(1): 1–23.
- Allen, L., and S. R. Gillespie.** 2001. *What works?: A review of the efficacy and effectiveness of nutrition interventions.* United Nations, Administrative Committee on Coordination, Sub-Committee on Nutrition.
- Baez, Javier E.** 2011. "Civil wars beyond their borders: The human capital and health consequences of hosting refugees." *Journal of Development Economics*, 96(2): 391–408.
- Ballard, Terri, Jennifer Coates, Anne Swindale, and Megan Deitchler.** 2011. "Household Hunger Scale: Indicator Definition and Measurement Guide." FANTA-2 Bridge, FHI 360, Washington, DC.
- Behrman, Jere R., and John Hoddinott.** 2005. "Programme Evaluation with Unobserved Heterogeneity and Selective Implementation: The Mexican PROGRESA Impact on Child Nutrition." *Oxford Bulletin of Economics and Statistics*, 67(4): 547–569.

- Bhutta, Zulfiqar A, Tahmeed Ahmed, Robert E Black, Simon Cousens, Kathryn Dewey, Elsa Giugliani, Batool A Haider, Betty Kirkwood, Saul S Morris, Hps Sachdev, and Meera Shekar.** 2008. "What works? Interventions for maternal and child undernutrition and survival." *The Lancet*, 371(9610): 417–440.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller.** 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics*, 90(3): 414–427.
- Caulfield, Laura E, Sandra L Huffman, and Ellen G Piwoz.** 1999. "Interventions to improve intake of complementary foods by infants 6 to 12 months of age in developing countries: impact on growth and on the prevalence of malnutrition and potential contribution to child survival." *Food & Nutrition Bulletin*, 20(2): 183–200.
- de Onis, Mercedes, Cutberto Garza, Cesar G. Victora, Adelheid W. Onyango, Edward A. Frongillo, and Jose Martines.** 2004. "The WHO Multicentre Growth Reference Study: planning, study design, and methodology." *Food & Nutrition Bulletin*, 25(Supplement 1): 15S–26S.
- de Onis, Mercedes, Monika Blössner, and Elaine Borghi.** 2011. "Prevalence and trends of stunting among pre-school children, 1990-2020." *Public Health Nutrition*, 1(1): 1–7.
- DHS.** 2010. "Tanzania Demographic and Health Survey." National Bureau of Statistics, Dar es Salaam, ICF Macro, Calverton.
- Duflo, Esther.** 2003. "Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa." *The World Bank Economic Review*, 17(1): 1–25.
- Gausset, Quentin, and Anna Folke Larsen.** 2013. "Social constraints on the adoption of improved banana varieties in Arumeru District." In *Farmers' Choice. Evaluating an approach to technology adoption in Tanzania.* , ed. Helene Bie Lilleør and Ulrik Lund-Sørensen. Practical Action Publishing.

- Glewwe, Paul, Hanan G Jacoby, and Elizabeth M King.** 2001. "Early childhood nutrition and academic achievement: a longitudinal analysis." *Journal of Public Economics*, 81(3): 345–368.
- Harari, Mariaflavia, and Eliana La Ferrara.** 2013. "Conflict, Climate and Cells: A disaggregated analysis." Centre for Economic Policy Research Discussion Paper 9277.
- Heckman, James J.** 1979. "Sample Selection Bias as a Specification Error." *Econometrica*, 47(1): 153–161.
- Hirvonen, Kalle.** 2013. "Measuring catch-up growth in malnourished populations." *Annals of Human Biology*, 1–10.
- Hoddinott, John, Jere R. Behrman, John A. Maluccio, Paul Melgar, Agnes R. Quisumbing, Manuel Ramirez-Zea, Aryeh D. Stein, Kathryn M. Yount, and Reynaldo Martorell.** 2013. "Adult consequences of growth failure in early childhood." *The American Journal of Clinical Nutrition*.
- Hoddinott, John, John A Maluccio, Jere R Behrman, Rafael Flores, and Reynaldo Martorell.** 2008. "Effect of a nutrition intervention during early childhood on economic productivity in Guatemalan adults." *The Lancet*, 371(9610): 411–416.
- Jamison, Dean T.** 1986. "Child malnutrition and school performance in China." *Journal of Development Economics*, 20(2): 299–309.
- Larsen, Anna Folke.** 2012. "Adoption of Banana Cultivation and Information Networks: An Empirical Study of Northern Rural Tanzania." Master's diss. University of Copenhagen, Copenhagen.
- Larsen, Anna Folke, and Helene Bie Lilleør.** 2014. "Beyond the Field: The Impact of Farmer Field Schools on Food Security and Poverty Alleviation." *World Development*, 64: 843–859.
- Lilleør, Helene Bie, and Ulrik Lund-Sørensen,** ed. 2013. *Farmers' Choice: Evaluating an Approach to Agricultural Technology Adoption in Tanzania*. Practical Action Publishing.

- Linnemayr, Sebastian, and Harold Alderman.** 2011. "Almost random: Evaluating a large-scale randomized nutrition program in the presence of crossover." *Journal of Development Economics*, 96(1): 106–114.
- Maluccio, John A., John Hoddinott, Jere R. Behrman, Reynaldo Martorell, Agnes R. Quisumbing, and Aryeh D. Stein.** 2009. "The Impact of Improving Nutrition During Early Childhood on Education among Guatemalan Adults." *The Economic Journal*, 119(537): 734–763.
- Martorell, R., and J.-P. Habicht.** 1986. "Growth in early childhood in developing countries." In *Human growth: a comprehensive treatise*. Vol. 3. 2 ed., , ed. F. Falkner and J. Tanner. Plenum Press, New York.
- Martorell, Reynaldo.** 1999. "The nature of child malnutrition and its long-term implications." *Food & Nutrition Bulletin*, 20(3): 288–292.
- Masset, Edoardo, Lawrence Haddad, Alexander Cornelius, and Jairo Isaza-Castro.** 2012. "Effectiveness of agricultural interventions that aim to improve nutritional status of children: systematic review." *British Medical Journal*, 344.
- Miller, Dennis D., and Ross M. Welch.** 2013. "Food system strategies for preventing micronutrient malnutrition." *Food Policy*, 42: 115–128.
- Minoiu, Camelia, and Olga N. Shemyakina.** 2014. "Armed conflict, household victimization, and child health in Côte d'Ivoire." *Journal of Development Economics*, 108: 237–255.
- Powell-Jackson, Timothy, Kara Hanson, Christopher J. M. Whitty, and Evelyn K. Ansah.** 2014. "Who benefits from free healthcare? Evidence from a randomized experiment in Ghana." *Journal of Development Economics*, 107: 305–319.
- Ruel, Marie T, and Harold Alderman.** 2013. "Nutrition-sensitive interventions and programmes: how can they help to accelerate progress in improving maternal and child nutrition?" *The Lancet*, 382(9891): 536–551.

- Ruel, MT.** 2001. "The natural history of growth failure: importance of intrauterine and postnatal periods." Vol. 47, 123–158, Philadelphia; Lippincott-Raven; 1999.
- Saleemi, MA, RN Ashraf, L Mellander, and S Zaman.** 2001. "Determinants of stunting at 6, 12, 24 and 60 months and postnatal linear growth in Pakistani children." *Acta Paediatrica*, 90(11): 1304–1308.
- Schreiner, Mark.** 2012. "A Simple Poverty Scorecard for Tanzania."
- Schroeder, Dirk G.** 2008. "Malnutrition." In *Nutrition and health in developing countries*. 341–376. Springer.
- Van den Berg, Henk, and Janice Jiggins.** 2007. "Investing in Farmers—The Impacts of Farmer Field Schools in Relation to Integrated Pest Management." *World Development*, 35(4): 663–686.
- Victora, Cesar G, Linda Adair, Caroline Fall, Pedro C Hallal, Reynaldo Martorell, Linda Richter, Harshpal Singh Sachdev, et al.** 2008. "Maternal and child undernutrition: consequences for adult health and human capital." *Lancet*, 371(9609): 340–357.
- Victora, Cesar Gomes, Mercedes de Onis, Pedro Curi Hallal, Monika Blössner, and Roger Shrimpton.** 2010. "Worldwide timing of growth faltering: revisiting implications for interventions." *Pediatrics*, 125(3): 473–e480.
- Vogl, Tom S.** 2014. "Height, skills, and labor market outcomes in Mexico." *Journal of Development Economics*, 107: 84–96.
- Waddington, H, H White, and J Anderson.** 2014. "Farmer Field Schools: From Agricultural Extension to Adult Education." *Systematic Review Summary*, 1.
- WHO, Multicentre Growth Reference Study Group.** 2006. *WHO Child Growth Standards: Length/height-for-age, weight-for-age, weight-for-length, weight-for-height and body mass index-for-age: Methods and development*. Geneva:World Health Organization.

Tables

Table 1: Summary statistics

		RIPAT	Comparison	P-value
Outcome variables	Height-for-Age Z-score	-0.94 (1.66)	-1.05 (1.66)	0.59
	Stunting indicator	0.25 (0.44)	0.27 (0.45)	0.65
Child characteristics	Young indicator	0.61 (0.49)	0.65 (0.48)	0.26
	Age in months	34.11 (15.36)	31.20 (15.52)	0.11
	Girl	0.57 (0.50)	0.52 (0.50)	0.19
	Child of head	0.83 (0.37)	0.87 (0.33)	0.45
Parent characteristics	Father's education	6.78 (1.68)	6.53 (1.67)	0.25
	Father's age	39.12 (8.10)	36.99 (8.25)	0.02
	Mother's education	6.70 (1.50)	6.08 (2.66)	0.12
	Mother's age	31.85 (7.17)	28.67 (6.70)	0.00
Household characteristics	Household size	6.20 (2.01)	5.95 (1.99)	0.40
	HH members age 0-5	1.58 (0.78)	1.60 (0.66)	0.90
	HH members age 6-14	1.61 (1.20)	1.66 (1.25)	0.80
	HH members age 15-24	0.98 (1.03)	0.84 (1.00)	0.34
	HH members age 25-49	1.63 (0.66)	1.58 (0.67)	0.49
	Head is widow(er)	0.06 (0.24)	0.03 (0.18)	0.14
	Acres 2006	4.07 (5.32)	3.11 (1.79)	0.19
	Good in math	0.41 (0.49)	0.42 (0.50)	0.86
	Participation in other projects	0.27 (0.44)	0.16 (0.37)	0.14
	Household rain	738.67 (47.86)	706.91 (45.64)	0.21
Village characteristics	Village distance to market	9.88 (3.90)	5.76 (5.00)	0.14
	Village has secondary school	0.57 (0.50)	0.86 (0.35)	0.29
	Village had devel. project	0.60 (0.49)	0.41 (0.49)	0.52
Number of children		214	182	
Number of households		182	153	
Number of villages		8	8	

Notes: Variable means in samples of RIPAT and comparison children. Standard deviations in parentheses. Column 3 gives wild cluster bootstrap-t p-values from two-sided t-tests of equal means of the RIPAT and comparison children, calculated as suggested by Cameron, Gelbach and Miller (2008). Clustering is at the village level.

Table 2: Impact of RIPAT on HAZ

	HAZ			Stunting
	(1)	(2)	(3)	(4)
RIPAT and young	0.569*	0.879***	1.377***	-0.176*
	(0.29)	(0.29)	(0.43)	(0.09)
	[0.062]	[0.012]	[0.004]	[0.094]
RIPAT	-0.240	-0.215		0.090
	(0.20)	(0.24)		(0.06)
Young	-0.025	-0.133	-0.302	0.060
	(0.11)	(0.30)	(0.71)	(0.09)
Child characteristics	No	Yes	Yes	Yes
Other characteristics	No	Yes	No	Yes
Household fixed effects	No	No	Yes	No
Clusters (villages)	16	16	16	16
Observations	396	396	396	396

Notes: OLS estimates with HAZ as dependent variable, cluster standard errors in parentheses, and wild cluster bootstrap-t p-values in square brackets. "Other characteristics" include parent, household, and village characteristics as described in the text. Statistical significance based on standard inference is indicated by ***, **, and * for the 1, 5, and 10 percent levels respectively.

Table 3: Cohort specific impacts on HAZ

	(1)	(2)
RIPAT and age 0	-0.237 (0.65) [0.666]	0.274 (0.67) [0.700]
RIPAT and age 1	0.666 (0.49) [0.198]	1.097* (0.53) [0.064]
RIPAT and age 2	0.473 (0.46) [0.282]	1.012** (0.40) [0.018]
RIPAT and age 3	-0.332 (0.34) [0.388]	0.106 (0.43) [0.786]
RIPAT	-0.042 (0.29)	-0.179 (0.35)
Age 0	0.681 (0.52)	-0.366 (1.37)
Age 1	-0.096 (0.30)	-0.751 (0.88)
Age 2	-0.200 (0.30)	-0.663 (0.54)
Age 3	0.107 (0.31)	-0.102 (0.44)
All characteristics	No	Yes
Clusters	16	16
Observations	396	396

Notes: OLS estimates with HAZ as dependent variable, cluster standard errors in parentheses, and wild cluster bootstrap-t p-values in square brackets. "All characteristics" includes child, parent, household, and village characteristics as described in the text. Statistical significance based on standard inference is indicated by ***, **, and * for the 1, 5, and 10 percent levels respectively.

Table 4: Adoption of technologies

	(1) RIPAT	(2) Comparison	(3) Conditional difference
Improved banana cultivation	0.657 (0.476)	0.121 (0.327)	0.523*** (0.103) [0.030]
Fruit tree(s)	0.590 (0.493)	0.497 (0.502)	0.232* (0.120) [0.322]
Improved breed of poultry	0.309 (0.463)	0.013 (0.115)	0.243*** (0.055) [0.032]
Improved breed of goats	0.354 (0.480)	0.128 (0.335)	0.227*** (0.044) [0.006]
Zerograzing	0.275 (0.448)	0.242 (0.430)	0.230** (0.083) [0.080]
Savings scheme	0.191 (0.394)	0.040 (0.197)	0.149*** (0.027) [0.024]
Number of crops in 2010	5.551 (2.538)	4.852 (2.126)	0.791* (0.390) [0.072]
Number of households	178	149	327

Notes: Variable means in samples of RIPAT and comparison children and standard deviations in parentheses in columns (1) and (2). Column (3) presents OLS estimates from regressions of the technology on a RIPAT indicator, cluster standard errors are in parentheses, and wild cluster bootstrap-t p-values are in square brackets. Regressions also control for education and age of the household head and household and village characteristics as described in the text. Statistical significance based on standard inference is indicated by ***, **, and * for the 1, 5, and 10 percent levels respectively.

Table 5: Food security

	(1)	(2)	(3)
	RIPAT	Comparison	Conditional difference
Number of worst months	3.831 (1.338)	4.150 (1.445)	-0.438*** (0.135) [0.038]
No hunger	0.365 (0.483)	0.265 (0.443)	0.159*** (0.050) [0.036]
HHS, worst month	1.494 (1.427)	1.789 (1.415)	-0.566** (0.205) [0.070]
HHS, best month	0.062 (0.304)	0.048 (0.270)	-0.034 (0.034) [0.532]
HHS, last four weeks	0.281 (0.680)	0.306 (0.679)	-0.067 (0.164) [0.764]
At least three meals, worst month	0.708 (0.456)	0.667 (0.473)	0.163 (0.110) [0.390]
At least three meals, best month	0.955 (0.208)	0.932 (0.253)	0.061* (0.029) [0.236]
At least three meals, last four weeks	0.904 (0.295)	0.878 (0.329)	0.068* (0.038) [0.246]
Meat consumption last week	0.764 (0.426)	0.694 (0.462)	0.183** (0.085) [0.192]
Egg consumption last week	0.607 (0.490)	0.408 (0.493)	0.152** (0.068) [0.180]
Dairy consumption last week	0.843 (0.365)	0.810 (0.394)	0.086 (0.128) [0.664]
Number of households	178	147	325

Notes: Variable means in samples of RIPAT and comparison children and standard deviations in parentheses in columns (1) and (2). Column (3) presents OLS estimates from regressions of the food security variables on a RIPAT indicator, cluster standard errors are in parentheses, and wild cluster bootstrap-t p-values are in square brackets. Regressions also control for education and age of the household head and household and village characteristics as described in the text. Statistical significance based on standard inference is indicated by ***, **, and * for the 1, 5, and 10 percent levels respectively.

Table 6: Impact on HAZ, Heckman selection model

	(1)	(2)
Panel A: Regression equation		
RIPAT and young	0.492 (0.33)	0.784*** (0.22)
RIPAT	-0.072 (0.22)	-0.105 (0.26)
Young	-0.234 (0.23)	-0.354* (0.21)
Panel B: Selection equation		
RIPAT and young	-0.191 (0.25)	-0.094 (0.28)
RIPAT	0.182 (0.18)	0.119 (0.23)
Young	0.172 (0.17)	0.697** (0.32)
Share measured in other households by same enumerator	1.084** (0.46)	1.663*** (0.59)
All characteristics	No	Yes
Clusters (villages)	16	16
Observations	535	535

Notes: Maximum Likelihood estimates from a Heckman selection model, and cluster standard errors in parentheses. Panel A gives the estimates from the regression equation with HAZ as dependent variable when controlling for selection in measurement. Panel B gives the estimates from the selection equation. The instrument in the selection equation is the share of children measured in other households by the enumerator. The sample consists of children zero to four years of age, out of which 73 percent were measured. When indicated, "All characteristics" are included in both the regression and selection equation. They include child, parent, household, and village characteristics corresponding to the specification in table 2, with two modifications: 1) we include the age and education of the household head instead of the parents' characteristics, as we can only identify the parents of measured children; 2) we include age in months instead of age in years, as we do not have precise ages of unmeasured children. For the same reason, the threshold for the *young* indicator is adjusted to three years of age (36 months) instead of three years and three months (39 months). Statistical significance is indicated by ***, **, and * for the 1, 5, and 10 percent levels respectively.

Table 7: Changing threshold of the *young* indicator

	Young threshold, month of conception				
	(1)	(2)	(3)	(4)	(5)
	May 06	Sep 06	Jan 07	Jul 07	Jan 08
RIPAT and young	0.880** (0.402) [0.048]	0.788* (0.393) [0.078]	0.879*** (0.285) [0.012]	0.587** (0.275) [0.050]	0.478 (0.289) [0.144]
RIPAT	-0.336 (0.32)	-0.207 (0.31)	-0.215 (0.24)	0.042 (0.25)	0.163 (0.23)
Young	-0.700 (0.44)	-0.422 (0.44)	-0.133 (0.30)	-0.488 (0.30)	-0.668** (0.31)
All characteristics	Yes	Yes	Yes	Yes	Yes
Clusters (villages)	16	16	16	16	16
Observations	396	396	396	396	396

Notes: OLS estimates with HAZ as dependent variable, cluster standard errors in parentheses, and wild cluster bootstrap-t p-values in square brackets. The column headings refer to the threshold of the *young* indicator where children conceived in or after the month mentioned are coded as *young*. "All characteristics" includes child, parent, household, and village characteristics as described in the text. Statistical significance based on standard inference is indicated by ***, **, and * for the 1, 5, and 10 percent levels respectively.

Table 8: Alternative sample selections

	(1)	(2)	(3)	(4)	(5)
	Incl. older	Acres	Outliers	Position	Clinic card
Panel A: Outcome variable: HAZ					
RIPAT and young	0.602*	0.942***	1.171***	0.652*	0.604**
	(0.31)	(0.28)	(0.31)	(0.34)	(0.25)
	[0.090]	[0.008]	[0.004]	[0.112]	[0.024]
RIPAT	-0.047	-0.256	-0.103	0.009	-0.241
	(0.23)	(0.21)	(0.28)	(0.28)	(0.28)
Young	-0.150	-0.193	-0.357	0.100	0.135
	(0.29)	(0.35)	(0.31)	(0.37)	(0.29)
Panel B: Outcome variable: Stunting indicator					
RIPAT and young	-0.107	-0.171*	-0.219**	-0.140	-0.155*
	(0.09)	(0.08)	(0.09)	(0.13)	(0.07)
	[0.266]	[0.070]	[0.042]	[0.336]	[0.064]
RIPAT	0.037	0.095*	0.102*	0.069	0.149**
	(0.05)	(0.05)	(0.05)	(0.08)	(0.06)
Young	0.015	0.044	0.096	0.020	-0.027
	(0.08)	(0.10)	(0.10)	(0.13)	(0.13)
All characteristics	Yes	Yes	Yes	Yes	Yes
Clusters	16	16	16	16	16
Observations	457	368	406	328	307

Notes: OLS estimates with HAZ as dependent variable in panel A and the stunting indicator in panel B. In parentheses are cluster standard errors, and in square brackets are wild cluster bootstrap-t p-values. The columns represent different sample selections compared to the main sample: (1) includes children 61-71 months old; (2) excludes children from households with less than one or more than eight acres; (3) includes outliers in HAZ; (4) excludes children measured recumbent when older than 24 months and measured standing when younger than 24 months; and (5) excludes children whose month of birth could not be validated by a clinic card. "All characteristics" includes child, parent, household, and village characteristics as described in the text. Statistical significance based on standard inference is indicated by ***, **, and * for the 1, 5, and 10 percent levels respectively.

Table 9: Impact on HAZ and likelihood of stunting with weighted RIPAT village comparison sample

	HAZ				Stunting
	(1)	(2)	(3)	(4)	(5)
RIPAT and young	0.832* (0.43) [0.104]	0.788 (0.43) [0.138]	0.710* (0.36) [0.078]	0.834* (0.44) [0.474]	-0.267** (0.09) [0.056]
RIPAT	-0.129 (0.31)	-0.226 (0.28)		-0.257 (0.31)	0.083* (0.04)
Young	-0.287 (0.31)	-0.547 (0.35)	-0.588 (1.37)	-0.550 (0.35)	0.185 (0.16)
Child characteristics	No	Yes	Yes	Yes	Yes
Household characteristics	No	Yes	No	Yes	Yes
Village characteristics	No	Yes	No	No	Yes
Fixed effects	No	No	Household	Village	No
Clusters (villages)	8	8	8	8	8
Observations	409	409	409	409	409

Notes: OLS estimates using a comparison sample within RIPAT villages weighted with inverse sampling probabilities. Column headings refer to the dependent variable. In parentheses are cluster standard errors, and in square brackets are wild cluster bootstrap-t p-values. "Household characteristics" includes parental characteristics. Statistical significance based on standard inference is indicated by ***, **, and * for the 1, 5, and 10 percent levels respectively.

Table 10: Heterogeneous impacts on HAZ

	(1)	(2)	(3)	(4)	(5)	(6)
Q:	Father's age	Mother's education	Mother's age	Log acres 2006	Prior project participation	Historical rainfall
RIPAT and young	0.796** (0.279) [0.018]	0.892*** (0.273) [0.004]	0.781** (0.287) [0.032]	0.901*** (0.261) [0.006]	0.832** (0.296) [0.026]	0.773*** (0.240) [0.016]
RIPAT, young and Q	-0.078*** (0.025) [0.024]	0.015 (0.127) [0.902]	-0.024 (0.044) [0.544]	-0.379 (0.458) [0.394]	-0.749 (0.595) [0.286]	0.012*** (0.003) [0.014]
RIPAT	-0.196 (0.241)	-0.219 (0.239)	-0.177 (0.216)	-0.232 (0.235)	-0.170 (0.219)	-0.128 (0.194)
RIPAT and Q	0.024 (0.019)	-0.123 (0.080)	-0.026 (0.029)	0.163 (0.278)	0.958** (0.431)	-0.012** (0.005)
Young	-0.077 (0.318)	-0.173 (0.289)	-0.069 (0.311)	-0.133 (0.280)	-0.066 (0.290)	-0.163 (0.301)
Young and Q	0.065*** (0.014)	0.026 (0.049)	0.027 (0.035)	-0.004 (0.420)	0.253 (0.312)	-0.002 (0.002)
Q (not demeaned)	-0.051** (0.019)	0.004 (0.057)	0.034 (0.027)	0.003 (0.225)	-0.325* (0.179)	0.003 (0.003)
All characteristics	Yes	Yes	Yes	Yes	Yes	Yes
Clusters (villages)	16	16	16	16	16	16
Observations	396	396	396	396	396	396

Notes: OLS estimates, cluster standard errors in parentheses, and wild cluster bootstrap-t p-values in square brackets. Q refers to the variable stated in the column heading; the variable is demeaned when it enters an interaction term, but not when included in levels. "All characteristics" includes child, parent, household, and village characteristics as described in the text. Statistical significance based on standard inference is indicated by ***, **, and * for the 1, 5, and 10 percent levels respectively.

A Appendix

Table 11: Sample composition

		RIPAT	Compa- rison	Total	RIPAT village
Number of households	Total interviewed	506	395	901	427
	with children 0-5 years of age	254	215	469	239
	with at least one child with HAZ	208	174	382	195
	In final sample	182	153	335	171
Number of children	Total in interviewed households	344	301	645	329
	No HAZ	76	87	163	91
	Don't know month of birth	19	23	42	13
	Either parent or child refused	13	16	29	14
	Not all children measured in household	29	35	64	35
	Other reasons	15	13	28	29
	with HAZ	268	214	482	238
	with missing values in characteristics	9	5	14	6
	with $ HAZ > 5$	7	4	11	10
	older than 60 months	38	23	61	27
In final sample	214	182	396	195	

Notes: The table shows how the final sample of households and children used in the analysis is composed and the different reasons for attrition. The last column gives the numbers for households in RIPAT villages not participating in the RIPAT project (unweighted).

Table 12: Poverty

	(1)	(2)	(3)
	RIPAT	Comparison	Conditional difference
PPI	40.611 (13.733)	41.185 (12.710)	5.094* (2.733) [0.265]
Good quality floor	0.272 (0.446)	0.305 (0.462)	0.042 (0.099) [0.825]
(Mobile) phone	0.728 (0.446)	0.795 (0.405)	0.031 (0.107) [0.855]
Number of households	180	151	331

Notes: Variable means in samples of RIPAT and comparison children and standard deviations in parentheses in columns (1) and (2). Column (3) presents OLS estimates from regressions of the poverty variables on a RIPAT indicator, cluster standard errors are in parentheses, and wild cluster bootstrap-t p-values are in square brackets. The "Progress out of Poverty Index" (PPI), as developed by Schreiner (2012), captures the probability that a household falls below the national poverty line. The PPI is country-specific and is based on ten simple questions that together provide a statistically strong and simple predictor of whether a household's consumption level is likely to be below the national poverty line as established in the 2007 Tanzanian Household Budget Survey. The PPI score ranges from 0 (most likely to be below a poverty line) to 100 (least likely to be below a poverty line). The PPI regression also controls for age and gender of household head, log acres 2006 and village characteristics, while the floor quality and phone regressions control for education and age of the household head and household and village characteristics as described in the text. (The household characteristics included in the PPI regression differ because some of the household characteristics enter the PPI calculation and hence they cannot be used as covariates). Statistical significance based on standard inference is indicated by ***, **, and * for the 1, 5, and 10 percent levels respectively.

The network at work: Diffusion of banana cultivation in Tanzania

Anna Folke Larsen*

Abstract

This paper investigates the role of networks for diffusion of improved banana cultivation introduced by an agricultural project in Tanzania. In the existing literature on networks and technology adoption, network effects are interpreted as learning. I show that a farmer's network can affect the adoption of a new crop not only through social learning, but also by providing necessary inputs for adoption. I set up a simple model for adoption and derive similar model implications for the provision of either inputs or information through the network. Empirically, I find that a farmer is 39 percentage points more likely to adopt banana cultivation if there is at least one banana grower in the farmer's network. I use three falsification tests to support causal interpretation of the network effect on adoption. Provision of inputs (banana seedlings) through networks is found to play an important role for the network effects found.

1 Introduction

There are huge disparities in agricultural productivity across countries with agricultural output per worker being more than 100 times larger in the United States than in Sub-Saharan African countries (Gollin et al., 2014). As the majority of poor people in developing countries are employed in the agricultural sector, agricultural growth

*Department of Economics, University of Copenhagen, email: afl@econ.ku.dk, tel.: +4561661159. I would like to thank the Rockwool Foundation for financial support for the data collection, Helene Bie Lilleør, Mette Ejrnæs, Bo Honore, and Peter Norman Sørensen for invaluable comments and guidance, Maria Fibæk and Dominick Ringo, Catherine Maguzu and all other RECODA staff for great support during field work, Jens Vesterager for agricultural technical guidance and Cathrine Søgaard Hansen for research assistance. All errors remain my own.

has the strongest potential compared to other sectors to reduce poverty in developing countries, in particular among the poorest of the poor (Ligon & Sadoulet, 2011; de Janvry & Sadoulet, 2010; Christiaensen et al., 2011).¹ Though Africa's Green Revolution has been a long time coming, agricultural growth in Sub-Saharan African countries is still key to transforming their economies and reducing poverty. Indeed, population growth and declining farm sizes call for locally adapted technological change in the agricultural sector (Diao et al., 2010). Moreover, climate change increases the necessity of technological change in agriculture to adapt to the more erratic rainfall (Lybbert & Sumner, 2012).

Hence, understanding barriers to adoption and diffusion of new agricultural technologies is key for agricultural development, poverty reduction and adaptation to climate change. The topic is not new; adoption of new agricultural technologies has been studied in a variety of countries and settings since the seminal work by Griliches (1957) (see reviews by Foster & Rosenzweig, 2010; Sunding & Zilberman, 2001; Evenson & Westphal, 1995; Feder et al., 1985).

In this paper I study how a farmer's network affects the decision to adopt a new agricultural technology in the context of African small-scale farming. The existing literature focus on the role of *social learning* through networks (Carter et al., 2014; Magnan et al., 2015; Conley & Udry, 2010; Bandiera & Rasul, 2006; Munshi, 2004; Foster & Rosenzweig, 1995; Krishnan & Patnam, 2014). These studies suggest that the network helps to relax an informational constraint faced by the farmer. I contribute to this literature by showing that a farmer's network can affect the adoption of a new crop not only through social learning, but also by providing necessary inputs for adoption. To my knowledge, the *provision of inputs* through networks has not been studied as an alternative or a complement to social learning.² This is an important distinction both because the role of information through networks may be exaggerated if it is confounded by

¹The contribution to poverty reduction from the agricultural sector stems not only from the size of the sector and the participation of poor people in the sector, but also from its indirect impact on growth in other sectors (Christiaensen et al., 2011).

²Emerick (2013) study the efficiency of input provision through networks as opposed to door-to-door visit, but he does not consider information provision through networks. Besides, he does not have data on networks but rely on sub-caste and last name as proxies for network connections.

input provision, but in particular because networks have the potential to mitigate not only imperfect information, but also input market imperfections. This can be used deliberately when designing future projects to increase diffusion of agricultural technologies, in particular in remote areas where input distribution is complicated by poor infrastructure.

Indeed, Spencer (1996) argues that the Green Revolution in Africa has been hindered by a low coverage of rural roads which impedes the distribution of inputs such as improved seeds and fertilizer. Road density in the low-income countries of Sub-Saharan Africa is less than half of that in low-income countries in the rest of the world (Carruthers et al., 2009). Malfunctioning input markets have been found to hamper the adoption of hybrid maize in Kenya (Suri, 2011) and improved pigeon pea varieties in Tanzania (Shiferaw et al., 2008). As many improved seed varieties can be multiplied locally, the local farmer network may work to relax the input constraint. In a recent paper by Shiferaw et al. (2015), constraints to information and access to seeds are considered jointly in the case of improved groundnut varieties in Uganda, but they do not explore how farmer networks can relax these constraints.

I study how networks affect the adoption of improved banana cultivation in the Arusha region in Tanzania. Improved banana variety seedlings and a new banana cultivation technique were introduced to participants of a Farmer Field School project called RIPAT in eight villages. I explore how the adoption among non-participants in the project villages depends on their informational links to project participants and to other farmers. The RIPAT project was designed to foster diffusion of banana cultivation to non-participants through a solidarity chain principle: Participants were obliged to pass on thrice as many seedlings as they received through the project to other farmers, free of charge. As the improved banana variety seedlings were not available through formal channels, input provision through networks becomes very important for adoption in this context.

To guide intuition for the adoption behavior among non-participating farmers, I set up a simple model of crop choice. Following the literature on social learning, I

first derive model implications under imperfect information and show that adopters in the network can affect the farmer's adoption decision through information about expected yields of the new crop. I then extend the model by allowing the network to provide inputs for the new crop when there is an imperfect input market. I can derive the exact same model implications with an imperfect input market as under imperfect information.

For the empirical analysis I use data on 509 non-participating farmers from households within the eight RIPAT villages collected for the purpose of this study. I find large network effects on adoption behavior: discussing farming issues with at least one farmer growing improved bananas increases the propensity to adopt by 39 percentage points. The data suggest that provision of inputs through the network contribute to this very strong network effect.

I further add to the literature by showing that network members who do not grow improved bananas have a negative effect on the propensity to adopt. The theoretical model provides the following intuition: Network members not growing bananas provide information or inputs that makes other crops more attractive, reducing the relative profitability of bananas. For a given amount of land, the farmer is then less likely to adopt banana cultivation. This finding points to the importance of controlling for network size when assessing the impact on adoption of adoption behavior in the network. Failing to do so (e.g. as in Bandiera & Rasul, 2006) potentially confounds the network estimate.

The network effects described are inherently difficult to identify (Manski, 1993, 2000; Brock & Durlauf, 2007; Bramoullé et al., 2009). Experimental variation in adoption in the network can facilitate identification (e.g. Magnan et al., 2015; Cai et al., 2015; Carter et al., 2014; Kremer & Miguel, 2007), but it is often not available. As participation in RIPAT is voluntary and hence subject to self-selection I must use a different approach to address causality.

First, I note that the network is captured *prior* to adoption (using a recall question), and there is a natural ordering in the timing of adoption as RIPAT participants are the

first to be introduced to improved banana cultivation. This mitigates the concern of a simultaneity problem. Next, I carefully investigate the different confounding factors and perform three falsification tests to address whether the estimated network effects are confounded by a) contextual effects; b) correlated effects; and c) self-selection into RIPAT, where a) and b) refers to the terminology of Manski (1993). *Contextual effects* cover the impact of the characteristics (rather than the behavior) of network members on individual behavior. I exploit detailed data on the RIPAT farmers in the network to test if the network effects are driven by the socioeconomic characteristics of network members. I do not find the characteristics to be driving the network effects found. *Correlated effects* capture the correlation in behavior within the network which is due to a common environment or a correlation in unobserved characteristics. I capture local growing conditions by the number of adopters within a radius of a half kilometer of the household and by subvillage fixed effects. I control for previous or current cultivation of traditional bananas to capture unobserved preferences for banana cultivation or prior knowledge of banana cultivation. Furthermore, I address the potential correlation of unobservables within networks in a placebo study. The network measures cannot predict adoption of three placebo crops, which leads me to conclude that the network effects found are not driven by a correlation in openness to new crops within networks. Finally, self-selection into RIPAT creates a concern for the interpretation of the results. My interpretation of the network effect is that the farmer is affected by the *adoption behavior* in his or her network either through the information or input channel. But as participation in RIPAT is voluntary the non-participants in my sample have implicitly self-selected out of RIPAT. They may have chosen to do so because they have network members who participate and they expect to receive information and inputs from them. In that case, they have decided to adopt regardless of the adoption behavior in the network and I would expect them to adopt as soon as possible. Hence, I explore the difference between early and late adopters to test if this behavior is driving the results. I find that the strong network effects persist among late adopters supporting my interpretation of the network estimates. Taken together, none of the evidence

suggests that the estimated network effects are confounded.

The remainder of the paper is structured as follows: Section 2 introduces the Farmer Field School project and the agricultural technology under study. In section 3 I set up a simple model of crop choice to illustrate how the adoption decision is affected by the network through either information or input provision and I derive testable implications of the model. I proceed with a description of the data and summary statistics in section 4, and subsequently, I present the empirical specification and estimation results in section 5. Section 6 discusses the identification of the network effects. In this section I address contextual and correlated effects and self-selection into RIPAT with three falsification tests and furthermore, I discuss how the provision of seedlings through networks could explain a part of the network effects found. Finally, section 7 concludes.

2 RIPAT and improved banana cultivation

The improved banana cultivation studied in this paper is introduced by a project called RIPAT (Rural Initiatives for Participatory Agricultural Transformation). RIPAT is a multifaceted agricultural and livestock project that aims to alleviate food insecurity and poverty among the participating households.³ A series of RIPAT projects have been implemented, and this study considers the first RIPAT project which took place in eight villages in Arumeru district in Northern Tanzania from 2006 to 2009. It was implemented by a local NGO RECODA and funded by the Rockwool Foundation.

Two Farmer Field School (FFS) groups are established in each village consisting of 30-35 farmers each. The farmers sign up voluntarily, but are only considered if they are dealing with agriculture already and if they have between one and five acres of land (however not rigorously abided).⁴ The FFS group cultivates a common plot, where RECODA facilitators demonstrate new agricultural techniques from a 'basket of options'. After learning about the new techniques and improved varieties the par-

³See thorough information on the project at www.RIPAT.org.

⁴The fact that RIPAT farmers self-select into the project also implies that the non-RIPAT farmers are a selected group. I discuss this issue in section 6.3.

ticipating farmers can choose to adopt on their own farm the components that best fit their soil, water accessibility, availability of household labor and land, preferences and taste.

The main component in the basket of options (and the most successful in terms of adoption) is a new technique of banana cultivation which is studied in this paper. It consists of special instructions for how to prepare the hard-pan soil and establish and tend a banana plantation, in conjunction with the introduction of five improved banana varieties which are more drought resistant than the traditional bananas grown in the area. The preparation of the soil consists of digging a one cubic meter hole which is then filled with a mixture of top soil and farm yard manure before planting the improved banana seedling. The soil around the plant can thereby contain more moisture which makes the plant more drought tolerant. The improved banana cultivation facilitates large scale plantations which is not possible with the traditional techniques in this area. The banana cultivation technique is indeed new in the area; when RIPAT was introduced at village meetings some people would laugh when banana plantations were mentioned because they knew it was not feasible—at least not with the existing techniques. The other components of the project include conservation agriculture, crop diversification, improved animal husbandry, fruit and multipurpose trees, soil and water conservation and post-harvesting technologies.

The project is designed to facilitate dissemination of the introduced technologies and varieties in several ways. A *solidarity chain* is established where participating farmers are obliged to pass on thrice as many improved banana seedlings to other farmers as they have received, free of charge. In addition, Super Farmers are chosen among the RIPAT farmers and educated to teach other farmers about banana cultivation. They are selected by the groups themselves among the best farmers to practice and teach the new methods. Furthermore, two criteria were set up for the formation of the Farmer Field School groups to foster dissemination of technologies: First, only farmers who were socially acceptable people and willing to share with others were admitted into the groups. Second, since each of the villages consist of two to five subvillages which

are not necessarily contiguous, it was ensured that all subvillages were represented in the FFS groups.

I study the role of networks in the local diffusion of improved banana cultivation from participating to non-participating farmers within the project villages. In particular, I study how the adoption of improved banana cultivation among non-RIPAT farmers residing in the project villages depend on whether they discuss farming issues with RIPAT farmers who have adopted the new technique. For expositional purposes I will only use the term 'improved' when it is important to distinguish between the existing 'traditional' banana cultivation and the new technique. Henceforth, 'banana cultivation' refers to the new technique.

The solidarity chain principle for improved banana seedlings was important for the diffusion of improved banana cultivation as the seedlings could not be purchased through formal channels in the area. Once the banana plant is established it produces seedlings which the farmer can only use if he or she wants to expand the banana plantation and hence the opportunity cost of giving them away is low. This is different from annual crops where the opportunity costs of the seed corn is to eat it or plant it on your own farm as you have to replant the crop every year.

The solidarity chain reduces the investment costs related to the establishment of a banana plantation. However, the opportunity costs of land and labor may still be considerable. The labor investment related to the establishment of the plantation is large as it is a very strenuous task to dig the big holes in the hard soil and some farmers may even choose to hire casual labor to dig the holes at a rate of around 2,000 Tanzanian Shillings (1.25 US dollars) per hole.⁵ Nevertheless, planting one or two banana plants is manageable and affordable for most farmers and a gradual expansion of the banana plantation can then be decided upon after testing the banana plantation on a small scale.

Figure 1 illustrates the adoption of banana cultivation among RIPAT farmers and the diffusion to non-RIPAT farmers over time. The maps are based on household GPS

⁵As noted by the anthropologist, Quentin Gausset, during field studies in the RIPAT villages.

location and adoption information from the data presented in section 4.1.⁶ Before project implementation in 2005 very few households in the sample had adopted banana cultivation. Already at the end of 2006, the first year of RIPAT, we see widespread adoption of banana cultivation among RIPAT households, and some few non-RIPAT households have followed suit at this early stage. By 2008 the number of adopting RIPAT households has almost doubled and the new technique is also catching on among non-RIPAT households. One year after the end of the project, 69 percent of the RIPAT households are growing improved bananas on their farm and the improved banana cultivation has spread to 20 percent of the non-RIPAT households. This is considered a very large degree of diffusion compared to the existing Farmer Field School literature, where only limited diffusion of the new technologies is documented (See reviews in Davis et al., 2012; Waddington et al., 2014). The high degree of adoption of improved banana cultivation among RIPAT participants and diffusion to non-RIPAT households suggests that banana cultivation indeed suits the local needs and preferences, that it is trialable at a smaller scale, and that it is profitable compared to existing crops and technologies.

3 A simple model of crop choice

To guide intuition for the empirical results, I set up a simple model that illustrates how the crop adoption decision is affected by the network of the farmer when the information is not perfect or when the input market is not functioning. The model allows me to derive testable implications for how the adoption decision is affected by the network of the farmer either through information or input provision.

I model how the adoption decision depends on the egocentric network including links to three different types of farmers: RIPAT farmers, non-RIPAT banana growers and other farmers who do not grow bananas. There are two main differences between the model I present and existing learning models (e.g. Foster & Rosenzweig, 1995; Bardhan & Udry, 1999; Munshi, 2004; Conley & Udry, 2010; Besley & Case, 1993;

⁶Data are collected in January 2011 and time of adoption is based on a recall question

Banerjee, 1992; Saha et al., 1994): 1) I show how social learning and provision of inputs through the network can lead to the exact same network effects on adoption behavior; and 2) I allow for network members growing other crops than the main crop studied to affect the adoption decision. To my knowledge, this has not been done before.

I will focus on model predictions for the extensive margin of the adoption decision (whether or not the farmer adopts) which corresponds to the empirical analysis. Initially, the farmer grows a traditional crop with a constant yield, y_a , and considers to adopt a new crop with a risky yield, $y_b = \mu + \varepsilon$, where the shock $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. As a benchmark I analyze the adoption decision under perfect information where the farmer knows the mean, μ , and variance, σ^2 , of the yield of the risky crop, and without imperfections on the input market. In this setting there will be no role for the network of the farmer.

3.1 Perfect information

The farmer can choose to adopt the new crop on some share, ω , of his or her land where the total farm area is normalized to one. The total farm yield will then be a weighted average of the yield from the traditional and the new crop:

$$y = \omega y_b + (1 - \omega) y_a, \quad 0 \leq \omega \leq 1$$

For simplicity I abstract from crop prices, but we could think of y_a and y_b as the value of the yields.⁷ If I assume that the input cost is linear in the yield and normalize the input price to zero for now,⁸ y_a and y_b represents the profits of the two crops. I return to the role of inputs in section 3.5. The crop choice of the farmer corresponds to a portfolio choice where the risk averse investor will trade-off mean and variance of the assets in his or her portfolio as exploited by Munshi (2004) in his model of acreage allocation and social learning. I follow Sargent (1979, pp.:150-151) and assume that the farmer values the total yield according to the utility function

⁷The analysis will be unaltered if I either assume constant prices or that the farmer only considers the current prices at the adoption decision.

⁸As long as input prices are constant, the analysis is unaffected.

$$U(y) = -e^{-\lambda y}, \lambda > 0$$

$U(y)$ is increasing and concave and λ captures the degree of risk aversion.⁹ This utility function is convenient because the expected utility can be rewritten to depend on the expected mean and variance of y , see Appendix A. The resulting expression for the expected utility is

$$E [U(y(\omega))] = -e^{-\lambda((1-\omega)y_a + \mu\omega - \frac{1}{2}\omega^2\lambda\sigma^2)}$$

The farmer maximizes the expected utility by choosing the optimal share of farm land, ω , to allocate to the risky crop. The interior solution is found by the first order condition:

$$\omega^* = \frac{\mu - y_a}{\lambda\sigma^2} \quad (3.1)$$

This result is quite intuitive: the optimal share of land allocated to the risky crop is increasing in the difference between the mean yield and the yield of the traditional crop. It is decreasing in the variance of the crop and in the risk aversion of the farmer. For a given increase in the variance of the yield, the more risk averse farmers will choose a larger reduction in the share of land allocated to the risky crop.

Assume for practical purposes that the share of land allocated to the risky crop cannot be infinitely small. I define a share ω_{min} which is the minimum feasible value of ω other than zero. This implies that I will consider ω^* as a latent variable and the observed adoption of the risky crop will be

$$\omega = \begin{cases} 0 & \text{if } \omega^* < \omega_{min} \\ \omega^* & \text{if } \omega_{min} \leq \omega^* < 1 \\ 1 & \text{if } 1 \leq \omega^* \end{cases}$$

⁹The Arrow-Pratt index of absolute risk aversion is $U''(y)/U'(y) = \lambda$ Pratt (1964).

Requiring a minimum share introduces the variance and the risk aversion in the extensive margin decision:

$$\omega = \begin{cases} 0 & \text{if } \mu - y_a < \lambda\sigma^2\omega_{min} \\ \omega^* & \text{if } \lambda\sigma^2\omega_{min} \leq \mu - y_a < \lambda\sigma^2 \\ 1 & \text{if } \lambda\sigma^2 \leq \mu - y_a \end{cases} \quad (3.2)$$

3.2 Imperfect information

Now I turn to the case where the expected yield of the risky crop, μ , is unknown to the farmer. This assumption is in line with models of Munshi (2004) and Besley & Case (1993), but in contrast to the target input type models where the subject of learning is the optimal input level (Foster & Rosenzweig, 1995; Conley & Udry, 2010). For simplicity I assume that the farmers know the dispersion of the yield, σ^2 , say because they are familiar with the dispersion of the rainfall. The farmer can discuss farming issues with other farmers who grow the risky crop (henceforth *informants*) to obtain information about the expected yield of the new crop. The farmer holds a belief about the expected yield:

$$\bar{\mu} \sim \mathcal{N}\left(\mu, \frac{1}{qN + k}\right)$$

I assume that the variance of the belief is inversely related to the number of informants, N , weighted by the quality of their information, q .¹⁰ When the farmer has no informants the variance of the belief is k^{-1} which is assumed to be very large (k is a very small positive number). As the farmer discusses farming issues with more people growing the new crop, his or her belief will approach the true expected yield of the new crop.

I can find the optimal share allocated to the risky crop following the same derivations as in section 3.1, but now replacing y_b by $\bar{y}_b = \bar{\mu} + \varepsilon$. Assuming that the belief about the expected yield and the yield shock are uncorrelated, \bar{y}_b will follow a nor-

¹⁰This assumption can be motivated by a Bayesian updating model where the variance of the signals from each informant is $1/q$ and the variance of the prior is $1/k$.

mal distribution with mean μ and variance $\sigma^2 + (qN + k)^{-1}$. Thus, due to uncertainty about the expected yield the farmer will consequently overestimate the variance of the yield. The optimal (latent) share which maximizes expected utility is then equal to:

$$\omega^* = \frac{\mu - y_a}{\lambda (\sigma^2 + (qN + k)^{-1})}$$

As expected, the optimal (latent) share of land allocated to the new crop is lower when uncertainty is introduced compared to the perfect information case (equation 3.1). The realized share will be

$$\omega = \begin{cases} 0 & \text{if } \mu - y_a < \lambda (\sigma^2 + (qN + k)^{-1}) \omega_{min} \\ \omega^* & \text{if } \lambda (\sigma^2 + (qN + k)^{-1}) \omega_{min} \leq \mu - y_a < \lambda (\sigma^2 + (qN + k)^{-1}) \\ 1 & \text{if } \lambda (\sigma^2 + (qN + k)^{-1}) \leq \mu - y_a \end{cases} \quad (3.3)$$

Equation 3.3 suggests the first testable empirical implication:

Model implication 1: Adopting the new crop is positively correlated with the number of informants growing the new crop.

When the farmer does not know anyone who grows the crop ($N = 0$) the variance of the belief about the risky yield is very large. For sufficiently small k ,¹¹ a risk averse farmer will not adopt a new crop which none of his or her informants grows. This is an alternative way of modeling that an information threshold has to be exceeded before adoption becomes feasible as in the model Saha et al. (1994). Discussing farming issues with just *one* farmer who grows the new crop will make the optimal latent share jump from (almost) zero to $(\mu - y_a) / (\lambda(\sigma^2 + (q + k)^{-1}))$. As long as $\mu > y_a$ the second order derivative of ω^* with respect to N is negative and hence, the extensive margin change is the unit change in the number of informants that leads to the largest increase in the propensity to adopt over the support of N .

Model implication 2: The change in the propensity to adopt is larger for ex-

¹¹ $k < ((\mu - y_a) / (\lambda \omega_{min}) - \sigma^2)^{-1}$

tensive than intensive margin changes in the number of informants.

The positive correlation between adoption and number of adopters in the network is also found in existing learning models, at least when there are few adopters in the network. Ambiguous effects for large networks are found in the target input model (Foster & Rosenzweig, 1995; Bardhan & Udry, 1999) where the subject of learning is the optimal amount of input rather than the expected yield. The farmer can learn about the optimal input both through learning by doing and learning from others which can create an incentive for strategic delay of adoption. When the farmer knows many adopters, s/he can free ride on the experimentation in the network and avoid costly experimentation on his or her own farm. This leads to an inverted U-shape relationship between the network and adoption which has been found empirically by Bandiera & Rasul (2006). In this model I do not specify how beliefs about the new crop are affected once the farmer has adopted the new crop because the empirical implications are not relevant in the context I consider.¹²

3.3 Information of different quality

Informants may not possess equally good information about the new crop. In the case of banana cultivation, RIPAT participants have received weekly training in the new cultivation technique for three years whereas non-RIPAT banana growers are likely to have less information about the new technique. When they pass on information on how to cultivate bananas I would expect information from RIPAT farmers to be of a higher quality than that of non-RIPAT farmers, $q_R > q_{nR}$. I can insert the sum of information from RIPAT and non-RIPAT informants, $qN = q_R N_R + q_{nR} N_{nR}$, into the expression for the optimal latent share of the new crop:

$$\omega^* = \frac{\mu - y_a}{\lambda \left(\sigma^2 + \frac{1}{q_R N_R + q_{nR} N_{nR} + k} \right)}$$

¹²In the sample 97 percent of the farmers discuss farming issues with no more than three banana growers. Hence, the number of informants growing bananas is not large enough to identify a non-linear relationship between the propensity to adopt and the banana network on the intensive margin. Bandiera & Rasul (2006) find the vertex of the inverted U to be at 10 adopters in the network.

High quality informants are better at reducing the variance of the expected yield than low quality informants. Hence, they also have a larger impact on the adoption decision. This leads to the third testable implication:

Model implication 3: The adoption decision is more affected by changes in the number of high quality than low quality informants.

3.4 Several risky crops

The model can be extended to include more than one risky crop. I consider the case where the farmer can choose to allocate land to two risky crops with yields y_b and y_c , which both outperform the traditional crop and hence in optimum no land is allocated to the traditional crop. The yields of the two crops are both assumed to be normally distributed and for simplicity, I assume that they are uncorrelated.

$$y_j \sim \mathcal{N}(\mu_j, \sigma_j^2), \mu_j > y_a, j \in \{b, c\}$$

Total yield is now a weighted average of the two risky crops, $y = \omega_b y_b + (1 - \omega_b) y_c$. Because the two yields are uncorrelated I can simply apply the same trick as in section 3.1 and expected utility under perfect information can be written as

$$E[U(y)] = -e^{-\lambda(\mu_b \omega_b + \mu_c(1-\omega_b) - \frac{1}{2}\lambda\omega_b^2\sigma_b^2 - \frac{1}{2}\lambda(1-\omega_b)^2\sigma_c^2)} \quad (3.4)$$

First I note that the expected disutility of risk is minimized when the farmer grows both crops ($\omega_b = \sigma_c^2 / (\sigma_b^2 + \sigma_c^2)$) because crop diversification reduces the variance of the total yield, in particular when the two yields are uncorrelated. Next, I maximize 3.4 with respect to ω_b and the first order condition gives

$$\omega_b^* = \frac{\mu_b - \mu_c + \lambda\sigma_c^2}{\lambda(\sigma_b^2 + \sigma_c^2)} \quad (3.5)$$

Again, I let the expected yields of the two crops be unknown to the farmer and let

the variance of the belief about the expected yields be the inverse of the number of informants growing the crop scaled by an information quality factor.¹³ Let the number of informants growing crop b and c be N_b and N_c respectively. Assuming that the beliefs for the two crops are independently distributed I can simply insert the inflated variances in equation 3.5

$$\omega_b^* = \frac{\mu_b - \mu_c + \lambda \left(\sigma_c^2 + \frac{1}{qN_c+k} \right)}{\lambda \left(\sigma_b^2 + \frac{1}{qN_b+k} + \sigma_c^2 + \frac{1}{qN_c+k} \right)}$$

and when I impose that the crops cannot be allocated a smaller non-zero share than ω_{min} , the realized share is

$$\omega_b = \begin{cases} 0 & \text{if } \omega_b^* < \omega_{min} \\ \omega_b^* & \text{if } \omega_{min} \leq \omega_b^* \leq 1 - \omega_{min} \\ 1 & \text{if } 1 < \omega_{min} + \omega_b^* \end{cases}$$

Discussing farming issues with farmers growing crop b still increases the propensity to adopt crop b . But the question is now whether the informants who grow crop c rather than crop b affect the choice to adopt crop b ? I differentiate ω_b^* with respect to N_c which yields

$$\frac{\partial \omega_b^*}{\partial N_c} = \frac{q \left[\mu_b - \mu_c - \lambda \left(\sigma_b^2 + \frac{1}{qN_b+k} \right) \right]}{\lambda (qN_c + k)^2 \left(\sigma_b^2 + \frac{1}{qN_b+k} + \sigma_c^2 + \frac{1}{qN_c+k} \right)^2} \quad (3.6)$$

$$\frac{\mu_b - \mu_c}{\lambda \left(\sigma_b^2 + \frac{1}{qN_b+k} \right)} < 1 \implies \frac{\partial \omega_b^*}{\partial N_c} < 0$$

Empirically, it would appear reasonable to assume that network size is positively correlated with the adoption of new crops even after controlling for the number of people in the network who grows the new crop, simply because the network size may correlate with unobserved characteristics such as entrepreneurship and openness. However, within this model framework I present how imperfect information may lead

¹³For simplicity, I let the information quality factor be equal for the two crops.

to the opposite correlation. Network information about *another* crop makes that crop relatively more attractive through a reduction in the uncertainty about its yield. The farmer has to trade off the two crops and hence will allocate a lower share to crop b if crop c becomes more attractive. This only holds if crop b does not fully outperform crop c .¹⁴

I can consider crop c to represent the crop portfolio of all other risky crops grown by those of the farmer's informants who do not grow crop b . Then I can draw a final empirical implication of the model:

Model implication 4: An increase in the number of informants *not* growing crop b will decrease the adoption of crop b .

3.5 Imperfect input market

Model implication 1 through 4 are derived under the assumption that the network provides knowledge about the mean yield of the risky crop(s). However, the same implications could be derived from a model with perfect information, but where the inputs are instead very costly and where the social network can lower the cost of inputs.

To see how, I define the profit from growing crop b on the total farm area as $\pi_b = y_b - \kappa_b(N_b)$, where y_b represents the value of the yield, and $\kappa_b(N_b)$ is the cost of the seed input which depends negatively on the number of network members growing crop b . The yield is still risky which implies that the profit follows a normal distribution with mean $\mu_b - \kappa_b$ and variance σ_b^2 . I can derive the optimal share allocated to crop b the same way as in section 3.1, so equation 3.1 is now modified by the input costs:

$$\omega^* = \frac{\mu_b - \kappa_b - y_a}{\lambda \sigma_b^2}$$

and the optimal share becomes

¹⁴In the case where $(\mu_b - \mu_c) / \left(\lambda \left(\sigma_b^2 + \frac{1}{qN_b+k} \right) \right) \geq 1$ then $\omega_b^* \geq 1$. Hence, crop b outperforms crop c , ($\omega_b = 1$), and ω_b is unaffected by changes in N_c .

$$\omega = \begin{cases} 0 & \text{if } \mu_b - \kappa_b - y_a < \lambda\sigma_b^2\omega_{min} \\ \omega^* & \text{if } \lambda\sigma_b^2\omega_{min} \leq \mu_b - \kappa_b - y_a < \lambda\sigma_b^2 \\ 1 & \text{if } \lambda\sigma_b^2 \leq \mu_b - \kappa_b - y_a \end{cases}$$

Consider the case where the input market is distorted by high transaction costs due to poor infrastructure, such that the cost of inputs when purchased through formal channels is so high that it is not optimal to adopt the new crop, $\mu_b - \kappa_b(0) - y_a < \lambda\sigma_b^2\omega_{min}$. It is clear to see that allowing the network members to provide inputs at a reduced or zero cost creates a positive network effect on the propensity to adopt. This corresponds to model implication 1 above. If I furthermore assume that one network member provides a sufficient amount of inputs, then follows implication 2. Actually, there would only be an effect on adoption from an extensive margin change in the network. Alternatively, I could assume decreasing marginal returns to seed inputs for a given level of land and labor inputs which would also yield implication 2.¹⁵ Model implication 3 would require the assumption that RIPAT network members would lower the input costs more than non-RIPAT members or be more likely to provide inputs. Given that they are obliged by the project to pass on seedlings, this is not an unreasonable assumption. Implication 4 could also easily be derived assuming that the network members growing other crops similarly lowered the adoption costs of these other crops making them more profitable than crop b .

Hence, empirical support for the four model implications would be evidence for network effects, but not for the mechanisms through which the network affects the adoption decision. It is a possibility that knowledge and input sharing simultaneously play a role.

¹⁵In the current specification I implicitly assume a linear relationship between seed inputs and yields, but I could instead assume a Cobb-Douglas production function which would exhibit decreasing marginal returns.

4 Data and summary statistics

4.1 Data collection

The empirical analysis is based on cross-sectional data collected in January 2011 as a part of an impact evaluation funded by the Rockwool Foundation and administered by Helene Bie Lilleør and the author (see Larsen & Lilleør, 2014). Household data were collected from RIPAT households, non-RIPAT households in RIPAT villages and households in comparison villages. This paper employs data from the choice-based sample of non-RIPAT households from the eight RIPAT villages. Households growing bananas were oversampled to ensure enough adopting households which are the households of interest for this study. Within the biomedical literature choice-based sampling is known as case-control studies and is widely used for studying infrequent events (Prentice & Pyke, 1979). The sample of non-RIPAT households consists of a random sample of households in the RIPAT villages and additional households who had received banana seedlings from RIPAT households according to RECODA records. The random sample facilitates the calculation of the population share of adopters among non-RIPAT households. The calculation is described in Appendix C. For a detailed description of the sampling scheme, see Appendix A of Larsen (2012).

The main respondent was either the person who took the decision to grow bananas or the person who takes most farming decisions, depending on whether the household had adopted bananas or not. This person was interviewed about his or her personal characteristics and network, and about the members of the household, their farm, crops, livestock, and assets. In addition, the adult female in the household was interviewed about household facilities and food security, and we collected child anthropometrics.

The sample of non-RIPAT households for the analysis of the adoption decision is constructed as follows: 597 non-RIPAT households in the eight RIPAT villages were interviewed in total. Out of these, 62 households are disregarded due to missing data on network or other explanatory variables. The data are not systematically missing

from either adopting or non-adopting households. In addition, 26 households are left out of the analysis because they either claim to have planted their first improved bananas before 2006 or because they moved to the village later than 2006 when RIPAT was implemented.¹⁶ This leaves a final sample of 509 households among which 193 are growing improved bananas as listed in the final row of Table 1.

4.2 Measuring networks

When assessing the role of the network for adoption of technology, it is important how the network is measured. Maertens & Barrett (2013) argue that the network is almost surely misrepresented if the researcher does not have explicit network data, but instead relies on proxy measures such as other farmers in the village (e.g. Foster & Rosenzweig, 1995; Munshi, 2004; Moser & Barrett, 2006), or geographical neighbors of the farmer (e.g. Krishnan & Patnam, 2014). Direct network measures typically capture egocentric networks either by prompting the farmer about links to other farmers in the study (Conley & Udry, 2010; Carter et al., 2014; Magnan et al., 2015) or by open-ended questions about whom they discuss farming issues with (Bandiera & Rasul, 2006). Open-ended questions might only elicit the farmer's "strong" network links because the "weak" links may be forgotten when the farmer is not prompted (Maertens & Barrett, 2013). However, the network size is not captured when the researcher only asks about links to other farmers in the sample, and I show in this paper that the size is important to account for. Also, women may have systematically smaller network measures than men if only the network links within the village are elicited and it is a patri-local society where the women moves to the men's villages when they marry. This is the case in the context I study.

Hence, I employ three open-ended questions to capture the egocentric networks of non-RIPAT respondents. The phrasing below was used for non-adopting farmers while adopting farmers received the questions in past tense and the recall frame in square brackets was used.

¹⁶Including households who planted before 2006 as either adopting or non-adopting households does not alter the results. Neither does the inclusion of immigrants.

Network size: Think about your relatives and friends and other people that you know.

[Before you decided to start growing improved bananas,] how many people do you discuss farming issues with?

Banana network: Among these, how many of them are growing improved bananas [before you decided to grow improved bananas]?

RIPAT banana network: If any of these are RIPAT farmers, could you please give me their names?

The timing of the recall was chosen to avert the potential upward bias due to endogenous network formation: If the network was measured after the adoption decision I could capture links between banana cultivating farmers that were established because they both grow bananas. This would induce an upward bias in the correlation between the banana network and the adoption decision.

The questions are sequential such that the mentioned farmers will be a subset of the response to the preceding question. This implies that only RIPAT farmers who were growing bananas were listed. The listed RIPAT farmers can be linked to detailed household and farmer characteristics, because the data collection also covered all RIPAT farmers. I exploit this information in section 6.1 to study whether the effect of the RIPAT network depend on the socioeconomic characteristics of the RIPAT farmers.

The empirical network measures easily relate to the network in the theoretical model presented in section 3. The banana network net of the RIPAT banana network will capture the number of non-RIPAT banana growers in the network. From model implication 1, I expect that the RIPAT and non-RIPAT network are positively correlated with the propensity to adopt. Model implication 3 suggests that the RIPAT network has a stronger impact than the non-RIPAT network. The network size for a given banana network will capture the number of informants growing other crops than bananas and model implication 4 predicts a negative correlation between the network size and the propensity to adopt conditional on the banana network. Furthermore, controlling for the network size ensures that the impact of the banana network is not confounded by

the network degree.

Is there any correlation between the banana network measured and the propensity to adopt in the raw data? Figure 2 shows the sample share of adopting households depending on the number of people in the farmer's banana network (2a) and the farmer's RIPAT network (2b).¹⁷ The sample share of adopting farmers is clearly larger for the subsets of farmers who discuss farming issues with banana growers which corresponds with model implication 1. The figure suggests that the greatest difference is on the extensive margin, i.e. whether you discuss farming issues with at least one RIPAT farmer or other banana grower. This is in line with model implication 2.

4.3 Summary statistics

Table 1 summarizes farmer and household characteristics for the full sample of non-RIPAT households and for adopting and non-adopting households separately. Due to the choice-based sampling of non-RIPAT households the adoption share in the sample is 38 percent which is almost twice the population share of 20 percent of adopting households among non-RIPAT households.

The explanatory variables of interest are the network variables. The farmers in the sample discuss farming issues with 2.8 people on average, where 0.5 are RIPAT banana growers, 0.3 are non-RIPAT banana growers and the remaining two people do not grow bananas. Adopting farmers are more likely to discuss farming issues with banana growers than non-adopting farmers, and though adopters also have a larger total network size than non-adopters, they discuss farming issues with 0.8 fewer people who are not growing bananas.

Furthermore, I control for a range of farmer and household level characteristics. At the farmer level, I include gender, age, religion and literacy. The reference category for the religion dummies is that the farmer is Protestant. 'Other religion' is a combined group of both traditional religion practitioners, Seventh Day Adventists and other groups that do not fall into the three main religion groups. Adopting and non-

¹⁷Recall that the adoption share in the sample does not correspond to the population share due to choice-based sampling.

adopting farmers are quite similar, though adopting farmers are slightly more likely to be Catholic and less likely to be Muslim than non-adopting farmers.

At the household level I consider different components of the household structure, namely the highest education level obtained, available household labor, whether the household head is a widow(er), the wealth of the household, and the farm size. This range of variables address constraints to adoption with respect to inputs to agricultural production: capital, labor, human capital, and land. The highest level of education achieved within the household will be the level of formal knowledge that the farmer can tap into. Since the highest level is 'completed primary education' (7 years) in 58 percent of the households I use this as the reference point and include indicators for having less or more education than completed primary. Adopting households are more likely than non-adopting households to have a household member with more than primary education. To capture household composition I control for whether the household head is a widow(er) and for the available household labor which is measured as the number of household members who can do hard manual labor to full extent. Adopting households have significantly more household labor and are less likely to be widowed. The level of wealth of the household is measured by a Tanzanian poverty score developed by Schreiner (2011) and I also include the number of acres of land the household employ in 2006. I use a recall measure for the farm size since it may be endogenous to the adoption decision, say if a farmer finds that banana cultivation is lucrative and rent in more land.¹⁸ I do not have a recall measure for the poverty score but in the impact evaluation of RIPAT (Larsen & Lilleør, 2014) we do not find the poverty score to be significantly affected by project participation. Adopting households are significantly more wealthy than non-adopting households, but they do not have more land. In addition I measure remoteness by the distance from the GPS location of the household to the nearest road,¹⁹ and this measure is not significantly

¹⁸As 97.5 percent of the sample owns at least some of their land and 83.9 percent owns all of the crop land they cultivate, inadequate incentives with respect to farm tenure arrangements should not be a constraint. Hence, I do not distinguish between whether the household owns or rent in the land that they cultivate.

¹⁹Data on roads are downloaded from OpenStreetMap (<http://download.geofabrik.de/africa/tanzania.html>) and kilometer distance from household GPS points is calculated in ArcGIS.

correlated with adoption.

I further include variables that capture agricultural practices, entrepreneurship and growing conditions that may correlate with both network and adoption.²⁰ Entrepreneurial households who are open for change could be more likely to participation in an NGO project (other than RIPAT) and to grow more different crops (net of improved bananas), and these two variables are indeed positively correlated with adoption. Whether or not the household has grown traditional bananas indicates if the household has some prior knowledge about or special preferences for banana cultivation and it appears to be an significant determinant of adoption: Adopting households are 18 percentage points more likely to have grown traditional bananas than non-adopters. This is important to control for as farmers who have grown traditional bananas may be more likely to discuss farming issues with each other. The local growing conditions are captured by the number of banana growers within a radius of 0.5 kilometers from the household where the distance is measured as the distance between GPS points taken at the household's compound. As we have not collected census data the measure is not complete, but it is a good proxy for the growing conditions that the household faces.²¹ Indeed, there is geographical clustering in the adoption of banana cultivation with adopting households having more neighbors who also grow bananas than non-adopters. The inclusion of this variable may cause the network estimates to be downward biased if farmers mainly discuss farming issues with their neighbors. However, if I exclude it, the network estimates may be confounded by a correlation in growing conditions within networks, hence I prefer the conservative estimates.

The eight villages in the data have a total of 24 subvillages with two to five subvillages in each village. Between six and 55 households are included in the sample from each subvillage. One subvillage has no adopting farmers and hence, these six observations are excluded when subvillage fixed effects are controlled for.

²⁰I could further include measures to capture access to information such as household ownership of a mobile phone or a radio, but these variables are uncorrelated with adoption and inclusion of them does not alter the results. I leave them out to reduce dimensionality.

²¹Once the number of banana growers within a 0.5km radius is controlled for, use of irrigation channel, historical rainfall at the household level and distance to nearest waterway becomes insignificant. Hence these measures are not included. Including them does not alter the results.

5 Econometric analysis

I estimate a logistic model of the probability of adopting banana cultivation:

$$Pr \{adopt_i = 1\} = \Lambda [\beta_1 R_i + \beta_2 nR_i + \beta_3 N_i + \delta X_i + \gamma G_i + \alpha_s] \quad (5.1)$$

where R_i is the RIPAT banana network of farmer i , nR_i is the non-RIPAT banana network and N_i is the network size. According to model implication 1 from section 3 β_1 and β_2 are positive, while model implication 3 predicts that $\beta_1 > \beta_2$. Model implication 4 suggests that β_3 is negative. In the main specification, the two banana network variables are specified as the number of (non-)RIPAT banana farmers in the network. To investigate the effect of the network on the extensive margin, I use a more flexible specification of the network. Instead of R_i and nR_i I include three indicator variables: Discuss farming issues with at least one; at least two; or at least three banana farmers. To allow for differential effects between RIPAT and non-RIPAT banana growers I also include corresponding indicator variables for at least one, two and three non-RIPAT banana growers in the network. This specification will allow me to test the model implication 2 that the extensive margin network effect is larger than the intensive margin effect.

In addition to the network variables, I control for farmer and household characteristics, X_i , and growing conditions, G_i , as these may both correlate with the network and the adoption decision. These are described in section 4.3. I can further control for local factors that makes adoption behavior correlate within the subvillages by subvillage fixed effects, α_s .²² All standard errors are clustered at the subvillage level.

The logistic model is based on the assumption that the individual unobserved characteristics are logistically distributed. This is a convenient model for choice-based samples because it provides consistent estimates of the parameters—apart from the constant term—as opposed to the linear probability model and the probit model (McFadden,

²²The model including fixed effects are estimated using conditional maximum likelihood where I only use within-subvillage variation in the adoption behavior to estimate the parameters (Chamberlain, 1982).

1973; Prentice & Pyke, 1979). The bias of the constant term can be corrected if the proportion of adopting households in the population is known. This enables the calculation of marginal effects when subvillage fixed effects are not included. Appendix B derives the consistency of the logit estimator for a choice-based sample and subvillage fixed effects. It further describes how the correction of the constant term is calculated.

Marginal effects are calculated for the specification that includes farmer and household characteristics but not subvillage fixed effects. For count variables such as the network variables, the marginal effect is calculated as the change in the propensity to adopt for a discrete change around the mean value of the count variable.²³ The remaining explanatory variables are evaluated at the sample mean and the constant is corrected for choice-based sampling. In the second specification I calculate For indicator variables I provide the marginal effect of a discrete change and for continuous variables I provide the usual marginal effect, still correcting the constant term and using the sample mean of the remaining variables.

The model in section 3 motivates a causal interpretation of the β estimates as network effects. However, the identification of such network effects requires careful scrutiny of all potentially confounding effects and considerations of reverse causality. The estimates of the network effects may be confounded by contextual or correlated effects using the terminology of Manski (1993), and they may be driven by self-selection into RIPAT. I will discuss the issues of identification in section 6 and address them with three falsification tests. I will also discuss whether the network effects are driven by dissemination of information or provision of inputs. But before I address the issues of identification I will present the regression results based on the specification in equation 5.1.

²³For the number of RIPAT and non-RIPAT banana growers in the network this corresponds to a change from zero to one while for the network size it is a change from two to three. For the second specification, I consider discrete changes in the indicator variables. E.g. the marginal effect of knowing at least two RIPAT banana growers is calculated by changing this variable from zero to one while *knowing at least one RIPAT banana grower* is set equal to one and *knowing at least three* is equal to zero.

5.1 Empirical results

Table 2 presents logistic coefficients and marginal effects for the propensity to adopt improved banana cultivation. Column (1) presents the simple logistic regression of the propensity to adopt on the three network variables. Discussing farming issues with a banana grower – whether RIPAT or non-RIPAT – is significantly positively correlated with the decision to adopt which is in line with model implication 1. Knowing an extra RIPAT farmer appears to be five to six times as important as knowing an extra non-RIPAT farmer who grows bananas, given network size, which correspond to model implication 3. The network size is negatively correlated with adoption when the number of banana growers is controlled for as predicted by model implication 4. These strong correlations persist when I include farmer and household characteristics in the regression. Furthermore, they remain unaffected when I account for subvillage fixed effects.²⁴ Since these parameters are only identified by variation within subvillages, factors that cause adoption rates to be correlated within subvillages such as soil quality, distance to markets and village institutions are not confounding the network effects.

The marginal effect show that the RIPAT network is really economically significant: Knowing an extra RIPAT banana grower increases the propensity to adopt by 24 percentage points. The non-RIPAT network appears to provide information of much lower quality as an extra farmer only increases the propensity to adopt by 5 percentage points. On the other hand discussing farming issues with a person not growing bananas reduces the propensity to adopt by 5 percentage points. These results illustrate well a situation of information deficit: Farmers are easily convinced to try a new crop by well-informed farmers, they are less affected by farmers who provide second hand knowledge and if they are in general more informed through a larger network, they are more difficult to persuade. But they might as well be explained by the provision of seedlings through the network which I return to in section 5.2.

Turning to farmer characteristics, female farmers are 14 percentage points more

²⁴The number of observations is reduced by six households because one subvillage does not have any adopting farmers in the sample.

prone to adopt banana cultivation than male farmers. This is well in line with anthropological field work in the area which concludes that women generally have the authority over bananas as compared to beans which is the domain of men (Mogensen & Pedersen, 2013). There appears to be an inverse U-shaped relationship between the propensity to adopt and the age of the farmer, though the two terms are not jointly significant at the ten percent level.²⁵ Though the coefficients are not significant, the pattern is well interpretable. Until the age of 41 there is an increasing relationship between the farmer's age and adoption while the relationship is negative for older farmers. This can be explained by the different phases in a household where a young farmer has to spend time on child rearing, while when the children become older the household can draw on teenage labor force. For older farmers, the children may have left home leaving fewer hands in the family farming activity.

Religion appears to play an important role showing that Catholics are 19 percentage points more likely to adopt than Protestants who constitute the reference group.²⁶ The other religion dummies are not significantly different from zero. The literacy of the farmer does not correlate with the adoption decision.

The highest education attained in the household does also not correlate with the adoption decision.²⁷ The little importance of education suggests that the new technology is so simple that lack of formal education is not a barrier to adoption. On the other hand, household labor appears to have some impact on adoption though it is only significant at the ten percent level. It is measured as the number of household members who are able to do hard manual labor to a full extent. As the establishment of a banana plantation requires a lot of hard manual labor it is intuitive that the available household labor is positively correlated with adoption. Whether the household head is a widow(er) seems to be negatively correlated with the adoption decision as expected. The estimated coefficient is rather large, but quite inaccurate and hence not statistically

²⁵Wald tests of joint significance: $\chi^2_{(2)} = .49, p = 0.11$.

²⁶Catholics are equally represented among RIPAT and non-RIPAT farmer, so the large coefficient can not be explained by Catholics being reluctant to join RIPAT groups. The role of religion in networks would be an interesting topic for future studies.

²⁷Neither does the education of the farmer if that was included instead.

significant. Naturally, a widow(er) household has less available household labor, and the strong negative correlation between household labor and the widow dummy explains the large standard errors. The wealth of the household as measured by a poverty score does not appear to be a determinant of adoption and neither does the number of acres the household has access to. Hence, little wealth or limited access to land does not seem to be important barriers to adoption, supporting the trialability of banana cultivation and further suggesting that network effects are not driven by access to credit. The distance to the road does not correlate significantly with adoption so more remote farmers are not more or less likely to adopt banana cultivation. Bandiera & Rasul (2006) finds participation in other NGO projects to be an important determinant of adoption of sunflower cultivation, and in Larsen & Lilleør (2014) we find project participation in the past to be correlated with participation in RIPAT. However, I do not find project participation to correlate with adoption of banana cultivation among the non-RIPAT participants.

The last three variables cover agricultural practices and conditions. I include the number of crops the household grew in 2010, net of traditional and improved bananas, to control for the combination of entrepreneurship and preference for risk diversification that would induce the farmer to plant many different crops. The number of crops grown in 2010 is indeed positively correlated with the adoption of banana cultivation. Previous or current cultivation of traditional bananas indicates that the household has some prior knowledge about banana cultivation reducing the information gap. It could also capture that the household has adequate growing conditions or special preferences for banana cultivation. Households who grows or have grown traditional bananas are 17 percentage points more likely to adopt improved banana cultivation. Finally, I control for growing conditions such as soil quality and rainfall by including the number of banana growers within a radius of 0.5 kilometers and the parameter is positive and significantly different from zero though less so when subvillage fixed effects are included as they capture some of the same variation in the data.

Among the list of characteristics, the number of RIPAT banana growers in the net-

work of the farmer prevails as one of the most important determinants of adoption both economically and statistically. It represents the highest marginal effect on the propensity to adopt and the t-statistic of the parameter estimate of 6.00 is by far the largest t-statistic of the included controls. As the sample share of adopting households is 24 percent for those farmers who do not know any RIPAT farmers, discussing farming issues with just one RIPAT farmer doubles the propensity to adopt.

I use standard errors clustered at the subvillage level to assess the significance level of the estimates. However, there are 24 subvillages in the data and this rather low number of clusters raises the question of the asymptotic distribution of the test statistics. Following Cameron et al. (2008) I address this question by estimating a linear probability model using ordinary least squares and calculate wild bootstrap-t p-values for the network variables. The coefficients to the RIPAT network and the network size both have p-values below 0.01, however the non-RIPAT network variable is not statistically significant when all covariates are included. Results are shown in panel A of Table A.1 in Appendix.

Even though I address the oversampling of adopting households by using the logit estimator, there might be an additional concern. The logit estimator provides consistent estimates if the additional adopting households are drawn randomly among all adopting households in the population. I identify additional adopting households through RECODA records and may thereby exactly sample households who are connected to RIPAT farmers potentially leading to an upward bias in the network estimates. However, I obtain the same parameters when I only include households drawn randomly, see Table A.2 in Appendix. If anything, the network estimates based on the random sample are larger, so I do not overestimate the effects using the non-random sample.

5.1.1 Extensive and intensive margin effects

Figure 2 suggested that the relationship between the propensity adopt and the banana network is not linear, but rather that the extensive margin change in the number of

banana informants is what matters for adoption. This is also supported by model implication 2. The regression results presented in Table 3 allows for a flexible relationship between the propensity to adopt and the banana network using six indicator variables: Discuss farming issues with at least one; at least two; or at least three banana growers or non-RIPAT banana growers in the network.

The correlation between the banana network and adoption is clearly driven by the extensive margin: Knowing at least one banana grower increases the propensity to adopt by 39 percentage points.²⁸ In the sample of farmers with no banana growers in their network, 22 percent grow bananas on their own farm and hence, this corresponds to a 177 percent increase in adoption! When controlling for having at least one banana grower in the network, the second and third banana grower does not correlate significantly with the adoption decision. The effect of knowing at least one non-RIPAT banana grower is not significantly smaller than knowing a RIPAT banana grower, however it is rather imprecisely estimated.²⁹

If I believe that the network effects are indeed driven by dissemination of knowledge, this result tells me something about the nature of the information constraint that the farmers are facing. It suggests that the informational barrier is relatively easily surpassed as only *one* source of information is needed to relax the constraint. This requires that the agricultural technique for improved banana cultivation is relatively simple to learn, and that just one observation can convince the farmer that improved banana cultivation is very profitable compared to many of the traditional crops grown in the area.

However, the importance of the extensive margin of the network is also very well in line with the idea of input provision through the network. It only requires contact to one banana grower to get hold of the first improved banana seedlings which makes adoption of improved banana cultivation feasible. This would generate the network

²⁸Using the random sample only, the marginal effect is as large as 49 percentage points, see Table A.2 in Appendix.

²⁹When I account for the small number of clusters using wild bootstrap-t p-values from OLS regressions, the extensive margin of the banana network is still significantly different from zero at the one percent level (see panel B of Table A.1 in Appendix).

effects found even if the farmers do not face an information constraint about the agricultural technique or profitability of improved banana cultivation. Both the input and information channel could very well be in play at the same time, and I am not able to fully disentangle the two. In the following section I present evidence that provision of inputs takes place in the networks.

5.2 Provision of inputs

There could be several interpretations consistent with the network effects found. The theoretical model presented in section 3 demonstrates that the network effects found are consistent with a story of social learning. However, as pointed out in section 3.5, the same model implications could be derived from a model where the network members provide free access to inputs, but no information. Hence, the fact that the data are in accordance with the model from section 3 does not allow me to conclude on whether the network effects are driven by learning about the expected yield or reduced costs of adoption.

Data on where adopting farmers got hold of their first improved seedlings can shed some light on the role of input provision. Adopting farmers were asked who gave or sold them the first banana seedlings and if they mentioned a RIPAT farmer I can cross check if the farmer is also in the network. The data are presented in Table 4. The first row covers the full sample of adopting farmers and shows from whom they received their first banana seedling. The majority (37 percent) received it from a RIPAT farmer in their network, while 29 percent received it from a RIPAT farmer who is not in their network.³⁰ Hence, two thirds received the seedling from a RIPAT farmer. A quarter received it from a non-RIPAT farmer and the remaining nine percent either received it directly from the implementing organization, a RIPAT group, or another NGO. No one bought the seedlings through formal channels suggesting that limited access to improved banana seedlings could be a binding constraint for adoption.

Network connections could also affect the price of the banana seedlings, and I

³⁰The person may belong to the higher order network.

can use information on whether the farmer paid anything in cash or in kind for the first seedling(s) to shed light on the issue. The majority (78 percent) received the first seedlings for free and this is statistically independent from whether they received the seedling from a RIPAT or non-RIPAT farmer. On the other hand, 44 percent of those who received their seedlings from the other sources mentioned above had to pay either in cash or in kind for the seedlings. This suggests that being connected to other banana growers reduces the cost of inputs.

The network measure I apply relates to the dissemination of information (“discussing farming issues”). Nevertheless, this network measure could very well be correlated with the provider of the first seedlings. Not surprisingly, discussing farming issues with one or more RIPAT banana growers increases the probability of receiving a seedling from a RIPAT farmer. This can be seen from the second row of Table 4 where I only consider adopting farmers who know at least one RIPAT banana grower (107 observations). Among these, 65 percent receive the first seedlings from a RIPAT network member suggesting that the network plays an important role for the provision of inputs. To test this more formally, I consider whether the farmer received the first banana seedlings from a RIPAT farmer and regress this indicator on the network variables and other covariates using various specifications. The results are presented in Table 5. Panel A shows the main specification where the network variables are entered linearly. When I account for subvillage fixed effects in column (4) none of the network variables are significantly correlated with the provision of seedlings from a RIPAT farmer. However, a very different pattern emerges when I allow for different effects on the extensive and intensive margin in Panel B of Table 5: Now discussing farming issues with at least one RIPAT farmer increases the probability of receiving the first seedlings from a RIPAT farmer with 19 percentage points. Surprisingly, knowing a second RIPAT farmer decreases the probability again by 11 percentage points. Discussing farming issues with at least one non-RIPAT banana grower decreases the likelihood of receiving seedlings from a RIPAT farmer, simply because it increases the probability of receiving seedlings from a non-RIPAT farmer. This can also be seen from

the third row of Table 4.

Hence, the provision of seedlings through networks could be a plausible explanation for why the network effects are found to be so large. In particular, the fact that the network effects are dominant on the extensive margin could be driven by seedling provision. You only need to know *one* banana grower to get your first improved seedlings.

The difference in the network effect for RIPAT and non-RIPAT banana growers could also be explained by seedling provision. Among the adopting farmers who knows both at least one RIPAT and non-RIPAT banana grower, 65 percent receives the first seedling from a RIPAT farmer (either in network or not) while only 29 percent receives the seedling from a non-RIPAT farmer, as presented in the fourth row of Table 4. This suggests that RIPAT farmers are more likely to share seedlings in their network. There could be several explanations for this. It takes approximately one year from the establishment of a banana plant before the farmer can harvest seedlings that can be passed on to other farmers in the network. Hence, if non-RIPAT banana growers have planted recently, they may not be able to share seedlings in their network. Also, RIPAT farmers may be more prone to share seedlings than non-RIPAT farmers. In fact, 61 percent of the RIPAT farmers who grow bananas and have passed on seedlings to other farmers mention “obligation in the project” as one of the reasons for passing on improved banana seedlings to other farmers. This points to the importance of the solidarity chain principle for the diffusion of improved banana seedlings.

This fact raises the question of whether the adopting farmers simply plant a few banana plants because they received the seedlings as gifts which leads to the high impact of RIPAT network on adoption or whether they really learn the new technology and adopt it because they perceive improved banana cultivation to be advantageous. I investigate this issue by considering whether the farmer has a plantation with more than ten banana plants instead of at least one banana plant. I find the same network effects on the propensity to establish a banana plantation as on propensity to grow at least one plant (see Table A.3 in Appendix). Hence, the network effects on adoption are not artificially high because the solidarity chain principle could induce RIPAT farmers to pass

on banana seedlings to other farmers who were not interested in banana cultivation.

Naturally, data on seedling provision can only explain variation within adopting farmers as they do not exist for non-adopting farmers. To assess the constraints faced by the non-adopting farmers, they were asked why they had not planted improved bananas, and Table 6 presents the categories of answers to this open-ended question. Water shortage is the dominant self-reported reason for not planting improved bananas, while land constraint is the second most important reason mentioned by the farmers. Lack of knowledge about production techniques is mentioned as frequently as no access to seedlings suggesting that these two constraints are both in play and are equally important.

These data suggest that provision of improved banana seedlings takes place within the networks. However, I cannot rule out that the farmers also disseminate information about improved banana cultivation through the networks. The fact that I find very strong network effects could suggest that both channels are at work.

6 Identification of network effects

Identification of social interaction effects is inherently difficult because most networks are endogenously shaped through individual choices. This may cause behavior to correlate within the network for other reasons than social interaction. In section 5.1 I found that the network of the farmer is a very strong predictor of the farmer's adoption decision. Regardless of whether this network effect is driven by information or input provision I consider it to be a social interaction effect as the adoption behavior of the network member is a prerequisite for both the information and input channel. In this section I will go through all the different causes of correlated adoption behavior in the networks which cannot be assigned to social interaction and address them one by one.

In order to identify social interaction, Manski (1993) employs the useful vocabulary of endogenous social interaction effects, contextual effects and correlated effects. The *endogenous effects* describe how the behavior of the individual is affected by the behav-

ior in the peer group. These are the network effects I want to identify. The *contextual effects* (or exogenous social effects) cover how the behavior of the individual is affected by the exogenous characteristics of the group such as education or wealth.³¹ I investigate whether the network effects found are driven by characteristics of the network rather than their adoption behavior in section 6.1. The *correlated effects* are covariation in behavior within a group due to similar unobserved individual characteristics or because group members face a similar environment. Endogenous network formation naturally leads to a correlation in individual unobserved characteristics if similar people prefer to share information. In addition, behavior may be spatially correlated due to growing conditions or institutions. I address these confounding factors in section 6.2.

Furthermore, Manski (1993) discuss the *reflection problem* that arises when the researcher wants to determine how the average behavior in a group affects the individual behavior in the group. The simultaneity within the group makes it difficult to identify who affects whom. In the network I study there is a natural ordering of events which circumvents this kind of simultaneity bias. I consider how non-RIPAT farmers are affected by discussing farming issues with RIPAT farmers who have adopted improved bananas. The ordering is created by the fact that RIPAT farmers were the first to be introduced to improved banana cultivation. Data on time of adoption allows me to check that non-RIPAT farmers did indeed plant bananas later than the RIPAT farmers in their networks. In 96 percent of the links, the RIPAT farmer planted before the non-RIPAT farmer and for 98 percent of the adopting non-RIPAT farmers at least one of the RIPAT farmers in their network planted before them. Hence, I do not consider simultaneity bias to be a great concern.

Another type of reverse causality could occur if the network was measured after adoption as banana growing farmers may endogenously form networks after adoption. However, recall that I capture the network *prior* to the adoption decision. Though I cannot rule out the existence of a recall bias, this suggests that the estimated network

³¹Manski (1993) uses the word “exogenous”, however, it should be noted that if networks are endogenously formed then the characteristics of network members may be endogenous too.

effects presented in section 5.1 are not confounded by the creation of links between banana cultivating farmers after they have chosen to adopt.

Even though networks are measured prior to the adoption decision, the self-selection into RIPAT could potentially cause adoption behavior to correlate without the RIPAT farmer affecting the non-RIPAT farmer to adopt. I discuss this issue in section 6.3.

It should be noted that when studying the impact of the network on adoption behavior, I cannot distinguish imitative behavior from learning as pointed out by Foster & Rosenzweig (1995). It would require farm productivity data to distinguish between imitation and learning which are very costly to collect and often subject to a large degree of measurement error. Nevertheless, if I assume that farmers are rational, adoption must indicate that they perceive improved banana cultivation to be relatively advantageous either with respect to profits, household food security or social factors such as prestige.

6.1 Contextual effects

Could the network effects found in section 5.1 be driven by the characteristics of network members? For instance, if banana growers are on average wealthier than other farmers, then knowing several banana growers implies knowing several wealthy people who may provide informal credit or insurance for you. In that case, a positive correlation between the number of banana growers in network and own adoption is not an evidence for learning or input provision but confounded by access to informal credit.

Ideally, I would like to control for the average characteristics of all the information network members to ensure that the correlation between adoption and the adoption behavior in the network is not driven by exogenous characteristics of network members. However, I only have detailed information about the RIPAT farmers in the network and not other network members.³² To the extent that exogenous characteristics are highly correlated within the network, controlling for farmer characteristics, X_i , that

³²To my knowledge, the only paper analyzing adoption of agricultural technologies and networks with detailed information on all network members is that of Van den Broeck & Dercon (2011).

are expected to affect adoption partly resolves the issue. But it is not sufficient in the case of heterogeneous networks. Though I do not have data on all network members, I can exploit the detailed data I have on RIPAT farmers. I split the RIPAT farmers in the network based on five central socioeconomic characteristics: wealth, land, education, gender, and age, and I examine whether the network effects differ dependent on the characteristics of the network members. If the network effects were driven by the characteristics of the farmers in the network rather than their adoption behavior (contextual effects), I would expect to find different network effects for e.g. rich and poor network members. No differential effects would support the hypothesis that the network effects found are not driven by contextual effects.

I measure wealth by a poverty score with a range of 0-100 (Schreiner, 2011) and split the sample of RIPAT farmers in networks at the mean poverty score, 47.4. I split the RIPAT farms into small and large by the average number of acres of 4.4. With respect to education, RIPAT farmers are divided into three groups: less than seven years of education (26.9 percent), seven years of education (68.1 percent) and more than seven years of education (5.0 percent of the sample). The gender split is self-explanatory, while the sample is split into young and old farmers at the mean age of 46.9 years.

Table 7 shows the estimation results together with tests of equal network effects across different characteristics of network members. Column (1) does not provide support for the hypothesis that the network effect is driven by access to credit through network members as the estimated effect of knowing rich RIPAT banana growers is in fact lower than knowing poor RIPAT banana growers. However, this difference is very far from being significant. The same pattern shows when I split the RIPAT network on the size of the farm in column (2) which could also be considered as a measure of wealth. RIPAT farmers with a small farm actually appear to have a stronger network impact, but again the difference is not significant.

Turning to the split on farmers' education in column (3), it appears that there is a smaller effect from knowing RIPAT banana growers with less than seven years of education, but there is no significant difference on the three estimated network effects

for different education categories. The coefficient for knowing RIPAT farmers with a high education is imprecisely estimated as the group is fairly small.³³ To ensure that the acceptance of the null hypothesis is not driven by large standard errors induced by the high education category, I combine the high and medium education category, and again I accept the null of no difference in network effects across education of network members.³⁴

As can be seen in column (4), the impact of knowing a male RIPAT farmer seems to be larger than knowing a female RIPAT farmer. Nevertheless I again reject that the network effects are differential across gender. Neither do the estimates and test results in column (5) provide evidence for a difference in the network effects across age.

Hence, I conclude that the network effects appear to be rather homogenous across these five socioeconomic characteristics which indicates that the network effects are not driven by the characteristics of the farmers in the network, i.e. contextual effects. At least, the network effects do not seem to be driven by access to informal credit or e.g. by knowing older RIPAT farmers who are maybe more respected and influential in the village.

6.2 Correlated effects

I distinguish between correlated effects due to environment (growing conditions and location institutions) or individual unobserved characteristics.

Farmers within a network may behave similarly because they face the same environment. Agricultural activities may be correlated for neighboring farmers due to similar growing conditions rather than social network effects. If the subvillage leadership is supporting and promoting banana cultivation in a particular subvillage then a correlation in adoption behavior within networks in the subvillage would not necessarily indicate the existence of social network effects. In the empirical specification in section 5.1 I have addressed these issues in several ways.

³³Only 13 farmers in the sample know a RIPAT farmer with high education.

³⁴Estimation results not reported. Wald test of equal effects for knowing RIPAT farmers with low and medium/high education: $\chi^2_{(1)} = 1.00$, $p = 0.318$.

I capture the growing conditions of a farmer by the number of banana growers in my sample within a radius of 0.5 kilometers from the farmer's dwelling as measured by GPS,³⁵ and the results in Table 2 show that this measure is an important determinant of adoption. As all RIPAT farmers are interviewed and in some villages all identified adopting farmers are interviewed, this measure almost corresponds to the actual number of banana growers within a radius of half a kilometer. However, in the villages where adoption is very wide spread so that the sample does not include all adopting households in the village, it understates the number of adopters within the radius. This is somewhat problematic since it will not capture the full effect of growing conditions in these villages. To mitigate this problem, I could additionally control for the historical rainfall within one square kilometer of the household,³⁶ the distance from the household to the nearest waterway³⁷ and whether the household uses an irrigation channel. However, I do not find any of these measures to be important for adoption once the number of neighboring adopters is controlled for, neither does inclusion of them affect the estimated network effects. Hence, I consider the number of adopters within a small radius to be a good measure for the growing conditions of the farmer. To ensure that institutional effects are not driving the results I show that the network estimates are invariant to the inclusion of subvillage fixed effects. The fixed effects also capture general equilibrium effects such as the effect of wide spread adoption in the local market price of bananas.³⁸

Another important correlated effect stems from the likely correlation of unobserved individual characteristics within networks which are formed by individual choices. Entrepreneurial farmers may first of all have larger networks and hence, be more likely

³⁵The distance is calculated using the 'geodist' package in Stata. The GPS measure is taken at the household dwelling and not at the farmer's plot(s), but this should not add too much noise as the majority of households have plots that are contiguous to their dwelling.

³⁶For historical rainfall I use interpolated data on yearly precipitation measured in mm from the period 1950-2000 available from <http://www.worldclim.org/> and link it to the households using GPS coordinates.

³⁷Data on waterways is downloaded from OpenStreetMap available from <http://download.geofabrik.de/osm/africa/> and the kilometer distance to household GPS points is calculated using ArcGIS.

³⁸However, it should be noted that the majority of farmers face periods of food insecurity and mainly grow bananas for home consumption.

to know adopting farmers. Thus, I control for network size in all regressions. In addition, an entrepreneurial farmer may choose to discuss farming issues with other farmers who are themselves entrepreneurial. Hence, a correlation between their adoption behavior may simply reflect that they are of the same type rather than being an indication of social interaction effects. If eligibility into RIPAT had been randomized, I could have used the random variation in the network of the non-eligible farmers to circumvent this problem (see e.g. Kremer & Miguel, 2007). But because participation in RIPAT was voluntary I must address the potential correlation of unobservables within the network. I do that by performing the following placebo study.

6.2.1 Placebo study

To examine if the strong correlation I find between adoption behavior and adopters in network is driven by a correlation in unobservables I consider adoption of three other crops: vegetables, sunflowers and sugarcane which are all profitable cash crops.

Cultivation of vegetables (e.g. onions, tomatoes) is very profitable but also requires access to water and intensive seasonal labor input. Sunflowers can be grown under rather dry conditions and the sunflower oil can be extracted from the seeds with a simple hand press. Sugar cane is a perennial grass that can be grown under varying conditions but access to irrigation water increases yields.³⁹ If the profitability of banana cultivation dominates the profitability of vegetables, sunflowers and sugarcane for all farmers then the placebo test has no bite. However, I would argue that this is not the case. Farmers who have access to plenty of water would profit more from vegetables than bananas whereas it might be more beneficial to grow sunflowers for farmers who have very limited access to water. I have chosen these three crops because their profitability relative to banana cultivation varies across farmers conditional on their available inputs. Banana cultivation is not very likely to dominate the cultivation of all of these crops.

If the correlation between the number of banana growers in network and adoption

³⁹Information on cultivation of vegetables, sunflowers and sugarcane is based on conversations with Jens Vesterager, Programme Manager, Rockwool Foundation.

of banana cultivation is driven by a correlation of entrepreneurship in the network I would expect the number of banana growers in network to explain variation in the adoption of vegetables, sunflowers and sugarcane. If knowing more RIPAT farmers is simply a proxy for being more open to new ideas then it should be correlated with the adoption of other crops too. However, if the RIPAT banana growers in the network also grew vegetables a positive correlation would occur in the case of social interaction within vegetables production. Hence, I control for the number of RIPAT banana growers in network who *also* grow the placebo crop. Since only 13 percent of the farmers grow sunflowers and eight percent grow sugarcane there are several subvillages with no variation in the adoption of the placebo crop which results in a fewer number of observations.

Table 8 presents the estimates from logistic regressions of adoption of vegetables, sunflowers and sugarcane, respectively, on the network variables and farmer and household characteristics. None of the network estimates are significantly different from zero and they appear small in magnitude. These results show that the number of banana growers in the network cannot explain adoption of any of the three placebo crops.

The lower number of observations in particular in column (2) and (3) raises the question of whether the regressions have enough power to explain the variation in adoption of sunflower and sugarcane. If the effects of the network variables are not significantly different from zero simply due to large standard errors caused by smaller sample sizes then the placebo test has no bite. Running the regressions without fixed effects can serve a double purpose: First, it allows me to keep the full sample in the regression to obtain smaller standard errors. Second, it allows me to calculate marginal effects and thereby assess the magnitude of the network effects found on the adoption of placebo crops. However, exploiting variation across subvillages may bias the network estimates if bananas and placebo crops require the same growing conditions which are correlated within subvillages, and if farmers generally discuss farming issues with their neighbors. Table A.4 in Appendix shows the logit estimates and the marginal effects for the network variables when subvillage fixed effects are excluded.

In this specification, the adoption of sunflower is significantly correlated with the RIPAT network and the network size, however the marginal effects are miniscule compared to those presented in Table 2. I conclude that the large network effects found cannot be explained by a correlation in entrepreneurship within the networks of farmers who adopt banana cultivation.

Entrepreneurship may not be the only unobserved factor which is correlated within networks. It seems plausible that farmers who prefer a certain crop are more prone to discuss farming issues with each other. Preferences for banana cultivation are most likely captured by the indicator for traditional banana cultivation, but what about other crops? If improved banana cultivation dominates the cultivation of beans, say, then the adoption behavior among bean growers could be correlated, and this would cause a spurious correlation between adoption of banana cultivation and adoption in the network. I can investigate this for the farmers who adopt in the second half of 2009 or later as I have data on the crops grown in 2009. I include indicator variables for the four most popular crops (improved maize, traditional maize, beans and vegetables) to see if the cultivation of any of these crops in 2009 is driving the subsequent adoption and hence the network effects found. The estimates are shown in Table A.5 in Appendix. The effects of the RIPAT and non-RIPAT banana network on adoption are unaffected when I control for the cultivation of improved and traditional maize, beans and vegetables, respectively.⁴⁰ When subvillage fixed effects are included, the coefficients to these four cultivation indicators are jointly insignificant.

This further supports the existence of network effects on adoption behavior since it does not provide support for a potential correlation in unobserved farmer characteristics within networks as the main driver of the network effects found.

6.3 Self-selection into RIPAT

The self-selection into RIPAT creates an additional concern. A farmer who knew many farmers who signed up for RIPAT could have chosen not to sign up simply because

⁴⁰Column (1) shows that the network effects for this subsample are slightly smaller compared to the full sample estimates from Table 2, before controlling for the cultivation of the other crops.

she knew that she would learn about the new technologies anyway. Since participation in RIPAT required weekly participation in meetings and joint cultivation of the demonstration plot and hence, many work hours, it is a reasonable concern that some farmers who were initially interested in improved banana cultivation could have chosen not to sign up for RIPAT, because several of their network members had done so. This corresponds to the idea of strategic delay derived from the target input model where a farmer would choose to postpone adoption if she knows sufficiently many adopters allowing her to learn from their experimentation without incurring the cost of experimenting herself (Bandiera & Rasul, 2006). Similarly, a farmer could avoid the opportunity cost of labor related to RIPAT participation if one or more network members had chosen to participate from whom she could get improved banana seedlings and instructions.

If I assume that these farmers would adopt banana cultivation relatively early since they were interested in banana cultivation already at the start of the project, I can split the sample of adopters into early and late adopters to see if there are differential effects. If the network effects only persist among early adopters, they may simply be generated by self-selection mechanisms into RIPAT.

I split the sample of adopting farmers on early and late adopters, where 'early adopters' planted their first improved bananas in 2006 or 2007. Within these two years RIPAT farmers would have time to plant bananas on their own farm and the plants would grow sufficiently to produce seedlings that can be shared in the network. I assume that the type of farmers who self-select out of RIPAT because they have network members in a RIPAT group would adopt as soon as possible and hence they would fall in the category of early adopters. If they do believe initially that banana cultivation is more profitable than other crops they grow, the optimal strategy would be to adopt as soon as possible.

Table 9 show the logit estimates with and without farmer and household characteristics and subvillage fixed effects for the two subsamples. There is indeed a larger parameter estimate for the RIPAT network for early adopters but the parameter es-

timate for the late adopters is very close to the full sample estimate and still highly significant. So even though some of the network effects found in this paper may be explained by self-selection out of RIPAT groups, it appears that social learning also takes place through the perception of the relative profitability of banana cultivation.

7 Conclusion

This paper studies how networks can relax constraints to adoption of a new agricultural technology. The existing literature on networks and adoption of technologies focus on the provision of information through networks (Conley & Udry, 2010; Bandiera & Rasul, 2006; Munshi, 2004; Foster & Rosenzweig, 1995). I contribute to this literature by showing that networks can affect the adoption of a new crop not only through information provision, but also by providing necessary inputs for adoption.

I set up a theoretical model to illustrate how a farmer's network can impact the adoption decision through the two different channels. The model has equivalent implications for both information and input provision through networks which calls for caution when interpreting empirical estimates of network effects.

Empirically, I study the adoption of improved banana cultivation which has been introduced by a project called RIPAT in Tanzania. I study how the adoption among non-RIPAT farmers is affected by their self-reported links to RIPAT farmers and other farmers and find that knowing at least one banana growing farmer increases the propensity to adopt by 39 percentage points. I carefully investigate whether I can consider this estimate to be a causal network effect and I find no evidence for contextual or correlated effects confounding the network estimate. The estimated network effect is most likely a compound effect of information and input provision through the network. Though I cannot fully determine their separate channels, I document that the input provision channel is playing an important role. A solidarity chain principle imbedded in RIPAT obliges participants to pass on thrice as many improved banana seedlings as they have received, and I find that 65 percent of the adopting farmers who knows a

RIPAT farmer have received their first improved banana seedlings from him or her.

I furthermore add to the literature on networks and adoption by extending the typical measure of the egocentric network to also include network members who are not growing the crop of interest. The estimates show that they have a negative impact on the adoption decision. The model provides a theoretical explanation for this finding: Network members growing other crops provide information or inputs that makes other crops more attractive, reducing the relative profitability of the crop of interest. For a given amount of land, the farmer is then less likely to adopt the crop of interest. When the total network size is not controlled for in an adoption regression, the effect of the network members who grow the crop of interest may be confounded.

Diffusion of knowledge in networks has received a lot of attention in the literature and recent work by Banerjee et al. (2014) explores the policy question of whom to target to increase the diffusion of knowledge. This is relevant for societies where information flows are hampered by limited access to information technologies.

However, lack of access to information is not the only barrier to adoption of agricultural technologies. In societies with poor infrastructure input markets suffer from high transportation costs which can be a barrier to adoption even if the gross return is high (Suri, 2011; Shiferaw et al., 2008). Hence, diffusion of agricultural inputs in networks is a highly relevant topic to study in that context. To my knowledge, Emerick (2013) provides the only empirical study of network trading of agricultural inputs. He finds that input provision through networks is inefficient compared to door-to-door sales of inputs as door-to-door sales lead to a larger degree of adoption in the context of a new rice variety in India. However, the low road density in Sub-Saharan Africa may impede such a market based input distribution due to high transportation costs. The first best solution may be to increase the road density, but according to Spencer (1996) the corresponding costs are so high that it is a more viable strategy to develop agricultural technologies that rely on local input provision.⁴¹

This is one of the motivations for the solidarity chain principle implemented in

⁴¹The efficiency of road construction may have improved since the 1990s but poor infrastructure remains an important challenge in Sub-Saharan Africa (Carruthers et al., 2009).

RIPAT: When farmers are obliged to pass on improved banana seedlings it creates a local supply of inputs that are necessary for adoption of the new technology. This study documents that the project has successfully fostered diffusion of improved banana cultivation, and that input provision through networks has played an important role. However, the design of the project does not allow me to assess the contribution of the solidarity chain principle to the diffusion of technology. How to best design agricultural projects to foster diffusion of new technologies remains an interesting topic for future research. Not only diffusion of knowledge, but also access to inputs must be addressed, and their separate contributions to the diffusion of new agricultural technologies should be assessed. More research is needed in order to understand when and how network based approaches to the distribution of new inputs can improve on the existing distribution systems.

8 Bibliography

- Bandiera, O. & Rasul, I. (2006). Social networks and technology adoption in northern mozambique. *The Economic Journal*, 116, 869–902.
- Banerjee, A., Chandrasekhar, A. G., Duflo, E., & Jackson, M. O. (2014). Gossip: Identifying central individuals in a social network. NBER Working Paper No. 20422, National Bureau of Economics Research.
- Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3), 797–817.
- Bardhan, P. & Udry, C. (1999). *Development Microeconomics*. Oxford University Press.
- Besley, T. & Case, A. (1993). Modeling technology adoption in developing countries. *The American Economic Review*, 83(2), pp. 396–402.
- Bramoullé, Y., Djebbari, H., & Fortin, B. (2009). Identification of peer effects through social networks. *Journal of Econometrics*, 150(1), 41 – 55.

- Brock, W. A. & Durlauf, S. N. (2007). Identification of binary choice models with social interactions. *Journal of Econometrics*, 140(1), 52 – 75.
- Cai, J., De Janvry, A., & Sadoulet, E. (2015). Social networks and the decision to insure. *American Economic Journal: Applied Economics*, 7(2), 81–108.
- Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3), 414–427.
- Carruthers, R., Krishnamani, R. R., & Murray, S. (2009). *Improving Connectivity: Investing in Transport Infrastructure in Sub-Saharan Africa*. Africa infrastructure country diagnostic, World Bank.
- Carter, M. R., Laajaj, R., & Yang, D. (2014). Subsidies and the persistence of technology adoption: Field experimental evidence from mozambique. NBER Working Paper No. 20465, National Bureau of Economic Research.
- Chamberlain, G. (1982). *Analysis of Covariance With Qualitative Data*. Working Paper 325, National Bureau of Economic Research.
- Christiaensen, L., Demery, L., & Kuhl, J. (2011). The (evolving) role of agriculture in poverty reduction—an empirical perspective. *Journal of Development Economics*, 96(2), 239 – 254.
- Conley, T. G. & Udry, C. R. (2010). Learning about a new technology: Pineapple in ghana. *The American Economic Review*, 100(1), 35–69.
- Davis, K., Nkonya, E., Kato, E., Mekonnen, D., Odendo, M., Miiro, R., & Nkuba, J. (2012). Impact of farmer field schools on agricultural productivity and poverty in east africa. *World Development*, 40(2), 402–413.
- de Janvry, A. & Sadoulet, E. (2010). Agricultural growth and poverty reduction: Additional evidence. *The World Bank Research Observer*, 25(1), 1–20.
- Diao, X., Hazell, P., & Thurlow, J. (2010). The role of agriculture in african development. *World Development*, 38(10), 1375 – 1383.

- Emerick, K. (2013). *The efficiency of trading in social networks: Experimental measures from India*. Job market paper, Department of Agricultural and Resource Economics, University of California Berkeley.
- Evenson, R. E. & Westphal, L. E. (1995). Chapter 37 technological change and technology strategy. In J. Behrman & T. Srinivasan (Eds.), *Handbook of Development Economics*, volume 3, Part A (pp. 2209 – 2299). Elsevier.
- Feder, G., Just, R. E., & Zilberman, D. (1985). Adoption of agricultural innovations in developing countries: A survey. *Economic Development and Cultural Change*, 33(2), 255–298.
- Foster, A. D. & Rosenzweig, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, 103(6), 1176–1209.
- Foster, A. D. & Rosenzweig, M. R. (2010). Microeconomics of technology adoption. *Annual Review of Economics*, 2(1), 395–424.
- Gollin, D., Lagakos, D., & Waugh, M. E. (2014). Agricultural productivity differences across countries. *The American Economic Review*, 104(5), 165–170.
- Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological change. *Econometrica*, 25(4), 501–522.
- Kremer, M. & Miguel, E. (2007). The illusion of sustainability. *The Quarterly Journal of Economics*, 122(3), 1007–1065.
- 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.
- Larsen, A. (2012). Adoption of banana cultivation and information networks: An empirical study of northern rural Tanzania. Master's thesis, Department of Economics, University of Copenhagen.

- Larsen, A. F. & Lilleør, H. B. (2014). Beyond the field: The impact of farmer field schools on food security and poverty alleviation. *World Development*, 64, 843–859.
- Ligon, E. A. & Sadoulet, E. (2011). *Estimating the effects of aggregate agricultural growth on the distribution of expenditures*. CUDARE Working Paper 1115, Department of Agricultural & Resource Economics, UC Berkeley.
- Lybbert, T. J. & Sumner, D. A. (2012). Agricultural technologies for climate change in developing countries: Policy options for innovation and technology diffusion. *Food Policy*, 37(1), 114 – 123.
- Maertens, A. & Barrett, C. B. (2013). Measuring social networks' effects on agricultural technology adoption. *American Journal of Agricultural Economics*, 95(2), 353–359.
- Magnan, N., Spielman, D. J., Gulati, K., & Lybbert, T. J. (2015). *Information Networks among Women and Men and the Demand for an Agricultural Technology in India*. IFPRI discussion paper 01411, International Food Policy Research Institute (IFPRI).
- Manski, C. F. (1993). Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3), 531–542.
- Manski, C. F. (2000). Economic analysis of social interactions. *The Journal of Economic Perspectives*, 14(3), 115–136.
- McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers of Econometrics* (pp. 105–142). Academic Press.
- Mogensen, H. O. & Pedersen, E. K. (2013). Household dynamics and gender politics: Female farmers in ripat 1. In H. B. Lilleør & U. Lund-Sørensen (Eds.), *Farmers' Choice: Evaluating an approach to agricultural technology adoption*. Practical Action Publishing.
- Moser, C. M. & Barrett, C. B. (2006). The complex dynamics of smallholder technology adoption: the case of sri in madagascar. *Agricultural Economics*, 35(3), 373–388.
- Munshi, K. (2004). Social learning in a heterogeneous population: technology diffusion in the indian green revolution. *Journal of Development Economics*, 73(1), 185 – 213.

- Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica*, 32(1/2), 122–136.
- Prentice, R. L. & Pyke, R. (1979). Logistic disease incidence models and case-control studies. *Biometrika*, 66(3), 403–411.
- Saha, A., Love, H. A., & Schwart, R. (1994). Adoption of emerging technologies under output uncertainty. *American Journal of Agricultural Economics*, 76(4), 836–846.
- Sargent, T. J. (1979). *Macroeconomic Theory*. Academic Press New York.
- Schreiner, M. (2011). *A Simple Poverty Scorecard for Tanzania*. Technical report, Microfinance Risk Management. Available online at <http://progressoutofpoverty.org>, accessed 6 May, 2015.
- Shiferaw, B., Kebede, T., Kassie, M., & Fisher, M. (2015). Market imperfections, access to information and technology adoption in uganda: challenges of overcoming multiple constraints. *Agricultural Economics*, 46, 1–13.
- Shiferaw, B. A., Kebede, T. A., & You, L. (2008). Technology adoption under seed access constraints and the economic impacts of improved pigeonpea varieties in tanzania. *Agricultural Economics*, 39(3), 309–323.
- Spencer, D. S. (1996). Infrastructure and technology constraints to agricultural development in the humid and subhumid tropics of africa. *African Development Review*, 8(2), 68–93.
- Sunding, D. & Zilberman, D. (2001). Chapter 4 the agricultural innovation process: Research and technology adoption in a changing agricultural sector. In B. L. Gardner & G. C. Rausser (Eds.), *Agricultural Production*, volume 1, Part A of *Handbook of Agricultural Economics* (pp. 207 – 261). Elsevier.
- Suri, T. (2011). Selection and comparative advantage in technology adoption. *Econometrica*, 79(1), 159–209.

Van den Broeck, K. & Dercon, S. (2011). Information flows and social externalities in a tanzanian banana growing village. *The Journal of Development Studies*, 47(2), 231–252.

Waddington, H., White, H., & Anderson, J. (2014). Farmer field schools: From agricultural extension to adult education. *Systematic Review Summary*, 1.

Figures

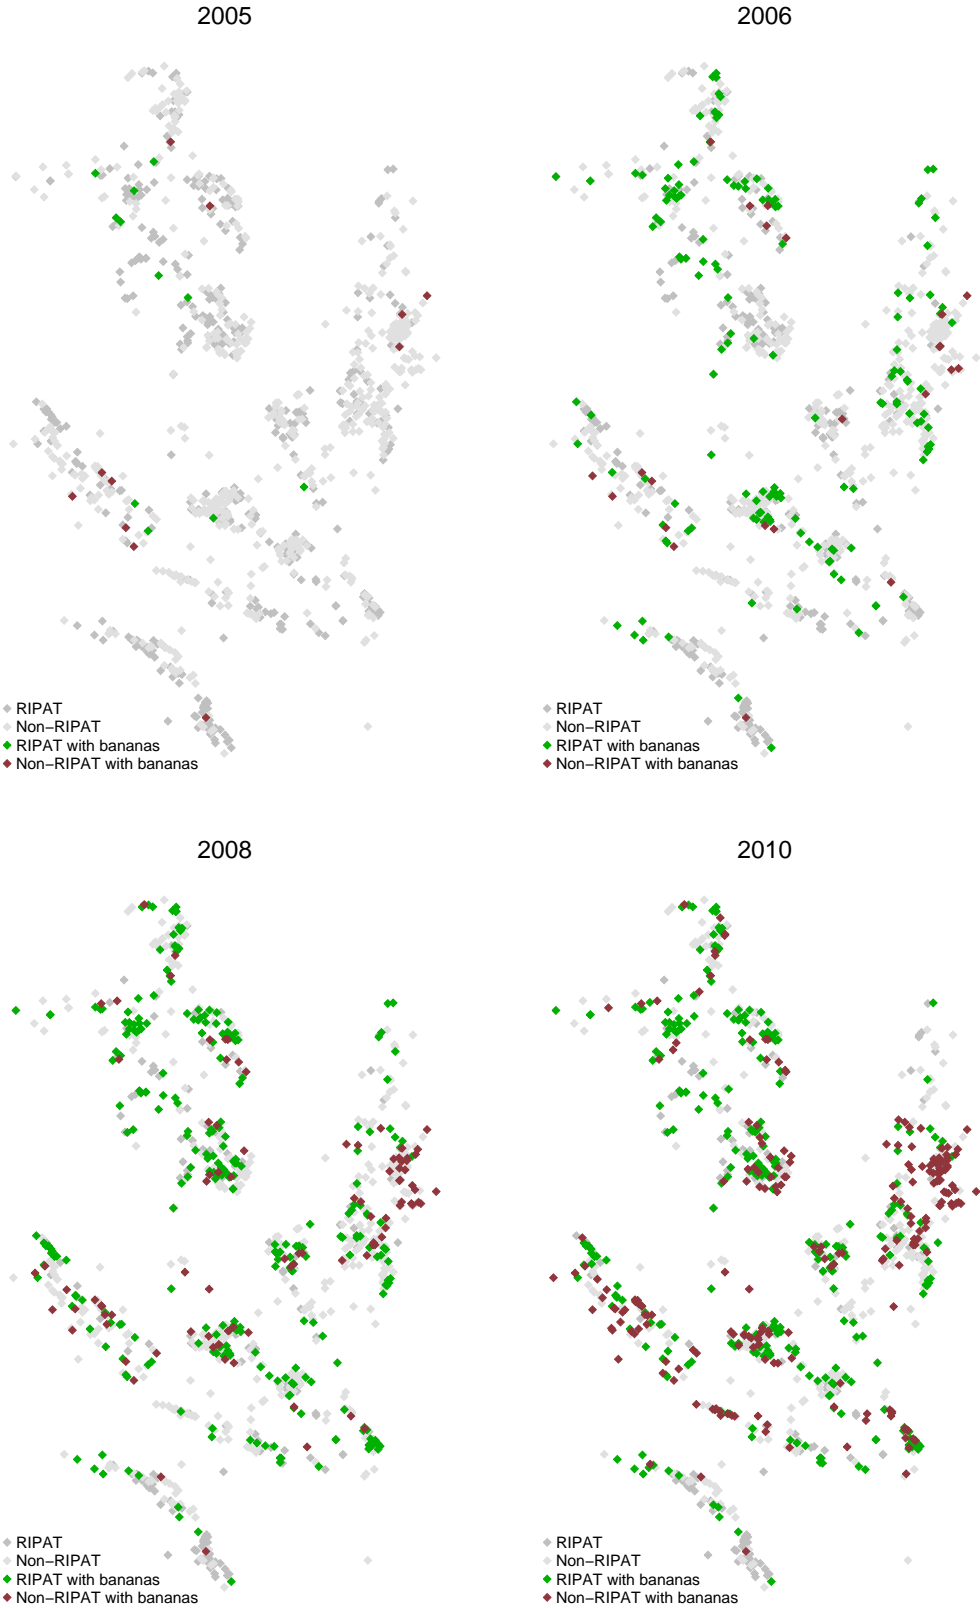
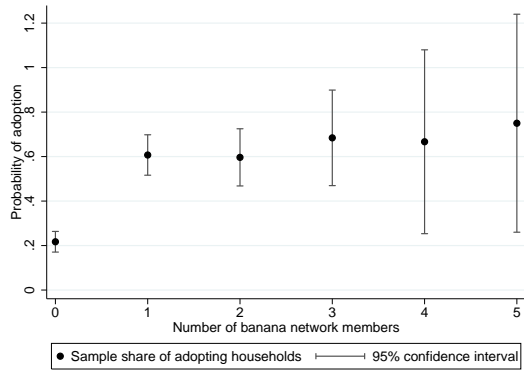
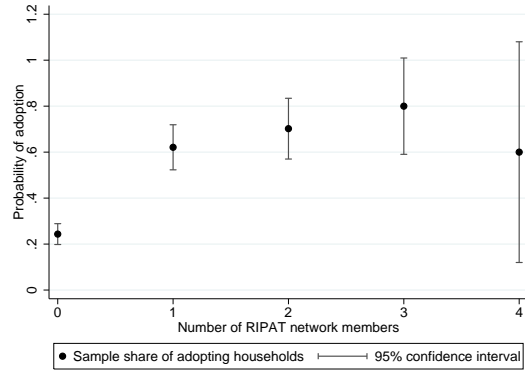


Figure 1: Diffusion of improved banana cultivation



(a) Banana network



(b) RIPAT network

Figure 2: Network measures and sample share of adopting households

Tables

Table 1: Summary statistics

	All		Adopting		Non-adopting		P-value
Grew improved bananas in 2010	0.38	(0.49)	1.00		0.00		
NETWORK VARIABLES							
RIPAT banana growers in network	0.52	(0.91)	0.95	(1.10)	0.26	(0.64)	0.000
Non-RIPAT banana growers in network	0.32	(1.87)	0.52	(2.73)	0.20	(1.04)	0.056
Network size	2.83	(4.12)	2.95	(3.79)	2.76	(4.32)	0.613
FARMER CHARACTERISTICS							
Farmer is female	0.22	(0.41)	0.23	(0.42)	0.22	(0.41)	0.736
Age of farmer	44.77	(15.54)	44.06	(13.31)	45.20	(16.76)	0.420
Farmer is Catholic	0.08	(0.28)	0.11	(0.32)	0.07	(0.25)	0.062
Farmer is Muslim	0.05	(0.22)	0.03	(0.16)	0.07	(0.25)	0.033
Farmer has other religion	0.25	(0.43)	0.25	(0.44)	0.25	(0.43)	0.859
Farmer can read	0.20	(0.40)	0.18	(0.38)	0.22	(0.41)	0.287
HOUSEHOLD CHARACTERISTICS							
Highest education, less than primary	0.07	(0.25)	0.05	(0.22)	0.08	(0.27)	0.238
Highest education, more than primary	0.35	(0.48)	0.44	(0.50)	0.30	(0.46)	0.002
Household labor	2.54	(1.51)	2.84	(1.63)	2.35	(1.40)	0.000
Household head is widow(er)	0.09	(0.29)	0.06	(0.23)	0.11	(0.31)	0.040
Wealth (poverty score)	44.30	(15.14)	46.32	(14.77)	43.07	(15.26)	0.019
Acres of land 2006	4.21	(6.04)	4.12	(4.36)	4.27	(6.87)	0.780
Distance to nearest road, km	1.45	(1.22)	1.49	(1.26)	1.42	(1.20)	0.526
Participate in other project	0.24	(0.42)	0.28	(0.45)	0.21	(0.41)	0.068
Number of crops grown, 2010	3.96	(1.93)	4.44	(1.95)	3.67	(1.86)	0.000
HH grows/has grown traditional bananas	0.36	(0.48)	0.47	(0.50)	0.29	(0.45)	0.000
No. banana growers within radius of 0.5km	10.83	(8.37)	13.62	(8.61)	9.12	(7.76)	0.000
Observations	509		193		316		

Notes: Sample means and standard deviations in parantheses for all farmers in the sample and for adopting and non-adopting farmers, respectively. The last column presents p-values from double-sided t-tests of equal means for adopting and non-adopting farmers.

Table 2: Adoption of improved banana cultivation

	Logit estimates				Marg.eff.
	(1)	(2)	(3)	(4)	(5)
NETWORK VARIABLES					
RIPAT banana growers in network	1.126*** (0.14)	1.105*** (0.14)	0.999*** (0.17)	0.989*** (0.17)	0.241***
Non-RIPAT banana growers in network	0.202** (0.10)	0.210** (0.10)	0.201** (0.10)	0.203** (0.08)	0.049**
Network size	-0.124* (0.08)	-0.127 (0.08)	-0.163** (0.08)	-0.207*** (0.06)	-0.040**
FARMER CHARACTERISTICS					
Farmer is female		0.399 (0.29)	0.592** (0.27)	0.581* (0.35)	0.141**
Age of farmer		0.058** (0.03)	0.053 (0.04)	0.043 (0.05)	0.007
Age of farmer, sq./100		-0.062** (0.03)	-0.064* (0.04)	-0.053 (0.05)	
Farmer is Catholic		0.835** (0.41)	0.817** (0.39)	0.880* (0.45)	0.185**
Farmer is Muslim		-0.788 (0.50)	-0.211 (0.53)	-0.060 (0.61)	-0.052
Farmer has other religion		0.006 (0.22)	-0.055 (0.28)	-0.044 (0.31)	-0.014
Farmer can read		-0.092 (0.29)	0.147 (0.33)	0.309 (0.39)	0.036
HOUSEHOLD CHARACTERISTICS					
Highest education, less than primary			0.218 (0.35)	-0.109 (0.66)	0.053
Highest education, more than primary			0.382 (0.26)	0.411 (0.29)	0.093
Household labor			0.132 (0.08)	0.188* (0.10)	0.032
Household head is widow(er)			-0.739 (0.49)	-0.585 (0.51)	-0.183
Wealth (poverty score)			0.011 (0.01)	0.011 (0.01)	0.002
Log of acres, 2006			-0.119 (0.19)	-0.127 (0.17)	-0.017
Log distance to road			0.016 (0.11)	-0.103 (0.14)	0.002
Participate in NGO project			-0.057 (0.24)	0.113 (0.29)	-0.014
Number of crops grown, 2010			0.176** (0.07)	0.204*** (0.07)	0.044**
HH grows/has grown traditional bananas			0.689** (0.28)	0.625** (0.26)	0.165**
No. banana growers within radius of 0.5km			0.055*** (0.02)	0.033* (0.02)	0.008***
Subvillage fixed effects	No	No	No	Yes	No
Observations	509	509	509	503	509

Notes: The dependent variable is an indicator equal to one if the farmer grows improved bananas in 2010. Logit coefficient estimates are presented in column (1)-(4), constant not reported. Standard errors in parentheses are clustered at the subvillage level. Marginal effects presented in column (5) are calculated as described in the text based on estimates from column (3). * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table 3: Extensive and intensive margin network

	Logit estimates				Marg.eff.
	(1)	(2)	(3)	(4)	(5)
1+ banana grower in network	1.908*** (0.26)	1.890*** (0.26)	1.694*** (0.29)	1.495*** (0.32)	0.392***
2+ banana growers in network	0.248 (0.41)	0.231 (0.41)	0.496 (0.40)	0.659 (0.43)	0.078
3+ banana growers in network	0.579 (0.47)	0.520 (0.49)	-0.182 (0.53)	-0.212 (0.64)	-0.027
1+ non-RIPAT banana grower in network	-0.384 (0.30)	-0.328 (0.31)	-0.061 (0.36)	-0.499 (0.43)	-0.011
2+ non-RIPAT banana growers in network	-0.934 (0.73)	-0.930 (0.72)	-1.440* (0.82)	-1.227 (0.88)	-0.294*
3+ non-RIPAT banana growers in network	1.189 (1.24)	1.353 (1.20)	2.027 (1.35)	1.847 (1.23)	0.387
Network size	-0.078 (0.06)	-0.081 (0.06)	-0.113* (0.07)	-0.123*** (0.05)	-0.028*
Farmer characteristics	No	Yes	Yes	Yes	Yes
Household characteristics	No	No	Yes	Yes	Yes
Subvillage fixed effects	No	No	No	Yes	No
Observations	509	509	509	503	509

Notes: The dependent variable is an indicator equal to one if the farmer grows improved bananas in 2010. Logit coefficient estimates are presented in column (1)-(4), standard errors in parentheses are clustered at the subvillage level. Column (5) presents marginal effects calculated as described in the text. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table 4: Where did the adopting farmers get the first banana seedlings?

Sample	Received seedling from				Total	Obs.
	RIPAT in network	RIPAT not in network	Non- RIPAT	Other		
All adopting farmers	36.9%	29.4%	25.3%	8.4%	100%	193
Know 1+ RIPAT banana grower	65.4%	14.0%	15.9%	4.7%	100%	107
Know 1+ non-RIPAT banana grower	25.7%	25.7%	37.1%	11.5%	100%	35
Know both	52.9%	11.8%	29.4%	5.9%	100%	17

Notes: The numbers in the first five columns are percentages while the last column gives the number of observations for the given row. Percentages in each row sum to 100. The category "Other" covers the implementing organization, a RIPAT group, or another NGO. The last row includes adopting farmers who discuss farming issues with at least one RIPAT banana grower *and* at least one non-RIPAT banana grower.

Table 5: Receive first banana seedling from RIPAT farmer

	Logit estimates				Marg.eff.
	(1)	(2)	(3)	(4)	(5)
PANEL A: MAIN SPECIFICATION					
RIPAT banana growers in network	0.260** (0.12)	0.265** (0.12)	0.253** (0.13)	0.101 (0.20)	0.036**
Non-RIPAT banana growers in network	-0.028 (0.10)	-0.013 (0.10)	-0.074 (0.11)	-0.067 (0.14)	-0.010
Network size	0.069 (0.08)	0.059 (0.08)	0.113 (0.10)	0.097 (0.10)	0.015
PANEL B: EXTENSIVE AND INTENSIVE MARGIN					
1+ RIPAT banana grower in network	1.651*** (0.35)	1.703*** (0.39)	1.668*** (0.40)	1.797*** (0.55)	0.192***
2+ RIPAT banana growers in network	-1.114* (0.66)	-1.131* (0.64)	-1.162* (0.63)	-1.287** (0.65)	-0.109*
3+ RIPAT banana growers in network	0.548 (0.77)	0.584 (0.72)	0.764 (0.63)	0.341 (0.81)	0.082
1+ non-RIPAT banana grower in network	-0.858** (0.37)	-0.735* (0.41)	-0.957** (0.42)	-1.312** (0.62)	-0.147**
2+ non-RIPAT banana growers in network	0.302 (1.28)	0.285 (1.20)	0.154 (1.41)	0.134 (1.27)	0.029
3+ non-RIPAT banana growers in network	0.506 (2.57)	0.683 (2.52)	0.798 (2.57)	1.456 (1.91)	0.117
Network size	0.063 (0.07)	0.047 (0.07)	0.092 (0.08)	0.040 (0.09)	0.011
Farmer characteristics	No	Yes	Yes	Yes	Yes
Household characteristics	No	No	Yes	Yes	Yes
Subvillage fixed effects	No	No	No	Yes	No
Observations	193	193	193	172	193

Notes: The dependent variable is an indicator equal to one if the farmer received the first banana seedlings from a RIPAT farmer. Only the sample of adopting farmers is used to produce the estimates. Panel A presents network estimates based on the main specification as presented in Table 2. Panel B presents the extensive and intensive margin estimates as presented in Table 3. Logit coefficient estimates are presented in column (1)-(4), standard errors in parentheses are clustered at subvillage level. Marginal effects presented in column (5) are calculated as described in the text based on estimates from column (3). * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table 6: Reasons for not adopting improved bananas

Why have you not planted any improved bananas?	percent
Water shortage	39.2
I don't have enough land	17.4
I do not know the production techniques	13.3
Never got seedlings	13.0
The work is too hard	10.1
Weather is not suitable	5.7
I have never grown bananas	4.4
My soil is inadequate for banana cultivation	4.1
I prefer growing traditional bananas	3.2
I think it will not be remunerative	0.6
Other	0.6
Observations	316

Notes: Multiple answers were possible so the percentages do not sum to 100.

Table 7: Adoption of improved banana cultivation, split on characteristics of RIPAT network

	Wealth (1)	Land (2)	Education (3)	Gender (4)	Age (5)
Poor RIPAT banana growers	1.144*** (0.25)				
Rich RIPAT banana growers	0.842*** (0.23)				
RIPAT banana growers w. small farm		1.255*** (0.24)			
RIPAT banana growers w. large farm		0.719*** (0.24)			
RIPAT banana growers with low edu.			0.474 (0.36)		
RIPAT banana growers with medium edu.			1.184*** (0.22)		
RIPAT banana growers with high edu.			0.926 (0.75)		
Male RIPAT banana growers				1.066*** (0.20)	
Female RIPAT banana growers				0.764** (0.31)	
Young RIPAT banana growers					1.000*** (0.26)
Old RIPAT banana growers					1.009*** (0.26)
Non-RIPAT banana growers in network	0.205** (0.09)	0.215** (0.09)	0.208** (0.09)	0.206** (0.09)	0.214** (0.09)
Network size	-0.204*** (0.06)	-0.195*** (0.06)	-0.210*** (0.06)	-0.198*** (0.06)	-0.204*** (0.06)
Observations	503	503	503	503	503
χ^2 test of equality ^a	0.84	2.56	3.00	0.71	0.00
P-value	0.360	0.109	0.223	0.398	0.981

Notes: The dependent variable is an indicator equal to one if the farmer grows improved bananas in 2010. Conditional logit coefficient estimates controlling for farmer and household characteristics and accounting for subvillage fixed effects. Standard errors in parentheses are clustered at the subvillage level. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

^a The following is tested: Column (1): Poor = rich (df = 1), Column (2): small = large (df = 1), Column (3): Low edu. = medium edu. = high edu. (df = 2), Column (4): Male = female (df = 1), Column (5): Young = old (df = 1).

Table 8: Placebo study results, adoption of vegetables, sunflowers and sugarcane

	(1) Vegetables	(2) Sunflower	(3) Sugarcane
RIPAT banana growers in network	0.261 (0.26)	0.222 (0.20)	-0.030 (0.25)
Non-RIPAT banana growers in network	0.139 (0.19)	-0.003 (0.15)	-0.064 (0.19)
Network size	0.047 (0.03)	0.024 (0.08)	-0.031 (0.10)
RIPAT growing vegetables	-0.169 (0.34)		
RIPAT growing sunflowers		0.295 (0.58)	
RIPAT growing sugarcanes			0.763 (0.97)
Observations	487	337	305
Mean of dependent variable	0.489	0.134	0.081
Std.dev. of dependent variable	(0.500)	(0.341)	(0.272)

Notes: The dependent variables are a indicators equal to one if the farmer grows vegetables (column 1), sunflowers (column 2) and sugarcane (column 3) in 2010. Conditional logit coefficient estimates accounting for subvillage fixed effects. Standard errors in parentheses are clustered at the subvillage level. Farmer and household characteristics are included in all specifications, but in column (2) religion dummies are excluded as Catholic dummy predicts non-adoption perfectly leading to drop of 43 observations. Results are robust to inclusion of religion dummies. The number of crops grown in 2010 is subtracted traditional and improved bananas and the placebo crop. Due to lack of variation in adoption within some subvillages, the number of observations is lower than 509.

Table 9: Adoption of banana cultivation, early and late adopter subsamples

	Early adopters			Late adopters		
	(1)	(2)	(3)	(4)	(5)	(6)
RIPAT network	1.497*** (0.24)	1.393*** (0.22)	1.394*** (0.32)	1.049*** (0.13)	0.964*** (0.17)	0.903*** (0.18)
Non-RIPAT network	0.386** (0.19)	0.401** (0.18)	0.318 (0.30)	0.195* (0.10)	0.186* (0.10)	0.191** (0.09)
Network size	-0.399** (0.16)	-0.504*** (0.16)	-0.536*** (0.17)	-0.108 (0.07)	-0.136* (0.07)	-0.180*** (0.06)
Farmer characteristics	No	Yes	Yes	No	Yes	Yes
Household characteristics	No	Yes	Yes	No	Yes	Yes
Subvillage fixed effects	No	No	Yes	No	No	Yes
Observations	358	358	256	464	464	458

Notes: The dependent variable is an indicator equal to one if the farmer grows improved bananas in 2010. Logit coefficient estimates, standard errors in parentheses clustered at the subvillage level. Column (1)-(3) is based on data from non-adopting households and households who adopted in 2006 or 2007. Column (4)-(6) is based on data non-adopting households and households who adopted in 2008 or later. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Appendix

A Derivation of mean-variance expected utility

To see how the expected utility $E[U(y)] = E\left[-e^{-\lambda((1-\omega)y_a + \omega y_b)}\right]$ can be rewritten to depend on the expected mean and variance of y_b , I first write the expected utility as

$$E[U(y)] = \frac{-e^{-\lambda(1-\omega)y_a}}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} -e^{-\lambda\omega y_b} e^{-\frac{(y_b-\mu)^2}{2\sigma^2}} dy_b \quad (\text{A.1})$$

I rewrite the exponent within the integral into two terms where one does not depend on y_b :

$$\lambda\omega y_b + \frac{(y_b - \mu)^2}{2\sigma^2} = \frac{(y_b - \mu + \omega\lambda\sigma^2)^2}{2\sigma^2} + \lambda\left(\mu\omega - \frac{\omega^2\lambda\sigma^2}{2}\right)$$

Inserting this exponent in equation A.1 and rearranging gives

$$E[U(y)] = \frac{-e^{-\lambda((1-\omega)y_a + \mu\omega - \frac{1}{2}\omega^2\lambda\sigma^2)}}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{(y_b - \mu + \omega\lambda\sigma^2)^2}{2\sigma^2}} dy_b$$

Now, I exploit that for all $\tilde{\mu}$ (including $\tilde{\mu} = \mu - \omega\lambda\sigma^2$)

$$\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{(y_b - \tilde{\mu})^2}{2\sigma^2}} dy_b = 1$$

because the left hand side is just the total area under the density function when the mean is $\tilde{\mu}$ and the standard deviation is σ . This implies that the expression for the expected utility simplifies to

$$E[U(y)] = -e^{-\lambda((1-\omega)y_a + \mu\omega - \frac{1}{2}\omega^2\lambda\sigma^2)}$$

B Choice-based sampling in a logit model

This section shows that the logit model provides consistent estimates of the parameters in the case of choice-based sampling.⁴²

Assume that the probability of adoption in the population, $\tilde{P}(a_i = 1)$, is logistically distributed and depends on a range of farmer and household characteristics, Z_i and subvillage fixed effects, α_s :

$$\tilde{P}(a_i = 1|Z_i, \alpha_s) = \Lambda(\theta Z_i + \alpha_s) = \frac{\exp(\theta Z_i + \alpha_s)}{1 + \exp(\theta Z_i + \alpha_s)} \quad (\text{B.1})$$

For simplicity assume that all the covariates are discrete, such that I can consider probabilities instead of distributions. The result generalizes to continuous covariates.

The probability of adoption in the sample, $P(a_i = 1)$, conditional on covariates and subvillage fixed effects can be rewritten using Bayes rule:

$$P(a_i = 1|Z_i, \alpha_s) = \frac{P(a_i = 1, Z_i, \alpha_s)}{P(Z_i, \alpha_s)}$$

I now use that the sample of adopting farmers is a random sample such that the probability of the covariates given that the farmer is adopting is the same in the sample and in the population, $P(Z_i, \alpha_s|a_i = 1) = \tilde{P}(Z_i, \alpha_s|a_i = 1)$, and correspondingly for non-adopting farmers. In addition, I use the law of iterated expectations in the denominator:

$$P(a_i = 1|Z_i, \alpha_s) = \frac{\tilde{P}(Z_i, \alpha_s|a_i = 1) \cdot P(a_i = 1)}{\sum \tilde{P}(Z_i, \alpha_s|a_i = 0) \cdot P(a_i = 0) + \sum \tilde{P}(Z_i, \alpha_s|a_i = 1) \cdot P(a_i = 1)} \quad (\text{B.2})$$

Applying Bayes rule and using equation B.1, I can rewrite:

$$\tilde{P}(Z_i, \alpha_s|a_i = 1) = \frac{\tilde{P}(Z_i, \alpha_s, a_i = 1)}{\tilde{P}(a_i = 1)} = \frac{\tilde{P}(Z_i, \alpha_s) \cdot \Lambda(\theta Z_i + \alpha_s)}{\tilde{P}(a_i = 1)} \quad (\text{B.3})$$

⁴²I would like to thank Professor Bo Honoré, Princeton University, for indispensable help with the following derivation.

Correspondingly,

$$\tilde{P}(Z_i, \alpha_s | a_i = 0) = \frac{\tilde{P}(Z_i, \alpha_s) \cdot (1 - \Lambda(\theta Z_i + \alpha_s))}{\tilde{P}(a_i = 0)} \quad (\text{B.4})$$

I now insert equation B.3 and B.4 in equation B.2:

$$P(a_i = 1 | Z_i, \alpha_s) = \frac{\frac{P(a_i=1)}{\tilde{P}(a_i=1)} \cdot \tilde{P}(Z_i, \alpha_s) \cdot \Lambda(\theta Z_i + \alpha_s)}{\frac{P(a_i=0)}{\tilde{P}(a_i=0)} \tilde{P}(Z_i, \alpha_s) \cdot (1 - \Lambda(\theta Z_i + \alpha_s)) + \frac{P(a_i=1)}{\tilde{P}(a_i=1)} \cdot \tilde{P}(Z_i, \alpha_s) \cdot \Lambda(\theta Z_i + \alpha_s)}$$

I divide numerator and denominator with $\tilde{P}(Z_i, \alpha_s)$ and insert the definition of the logistic distribution:

$$P(a_i = 1 | Z_i, \alpha_s) = \frac{\frac{P(a_i=1)}{\tilde{P}(a_i=1)} \cdot \frac{\exp(\theta Z_i + \alpha_s)}{1 + \exp(\theta Z_i + \alpha_s)}}{\frac{P(a_i=0)}{\tilde{P}(a_i=0)} \cdot \frac{1}{1 + \exp(\theta Z_i + \alpha_s)} + \frac{P(a_i=1)}{\tilde{P}(a_i=1)} \cdot \frac{\exp(\theta Z_i + \alpha_s)}{1 + \exp(\theta Z_i + \alpha_s)}}$$

Finally, I divide both numerator and denominator with the first term of the denominator and rearrange:

$$P(a_i = 1 | Z_i, \alpha_s) = \frac{\exp(\theta Z_i + \alpha_s + \ln(c))}{1 + \exp(\theta Z_i + \alpha_s + \ln(c))}, \quad c \equiv \frac{P(a_i=1)/\tilde{P}(a_i=1)}{P(a_i=0)/\tilde{P}(a_i=0)} \quad (\text{B.5})$$

Comparing the probability of adoption in the sample (equation B.5) with the probability of adoption in the population (equation B.1), it is evident that the choice-based sampling only affects the estimation of the subvillage fixed effects (or the constant in the case of no fixed effects) and hence, the estimated parameters of the covariates (θ) are unaffected by the sampling method.

C Correction of the constant term for calculation of marginal effects

In order to calculate marginal effects I need to correct the constant term with the factor $\ln(c)$.

$$\begin{aligned}
\ln(c) &= \ln \left(\frac{P(a_i=1)/\tilde{P}(a_i=1)}{P(a_i=0)/\tilde{P}(a_i=0)} \right) \\
&= \ln (P(a_i = 1)) + \ln (\tilde{P}(a_i = 0)) - \ln (\tilde{P}(a_i = 1)) - \ln (P(a_i = 0))
\end{aligned}$$

The population and sample probability differ due to non-random sampling. I will simply estimate the sample probabilities by sample proportions:

- $P(a_i = 1)$ is equal to the share of adopting farmers in the final sample of nonRIPAT farmers
- $P(a_i = 0)$ is equal to the share of non-adopting farmers in the final sample of nonRIPAT farmers

The population probabilities are more complicated to calculate due to a complex sampling design. It is briefly described in the following, for more details see Appendix A in Larsen (2012).

We drew a random sample of size R in each village among which the village leader identified the adopting households. Let L_i be the indicator for being identified as adopting household by the village leader. This is however not a perfect measure of adoption, $P(L_i = 1) \neq P(a_i = 1)$. Within the random sample we only interviewed a stratified subsample, which was stratified to achieve the same share of adopting and non-adopting farmers based on village leader identification. The population probability of being an adopting household can be calculated as:

$$\tilde{P}(a_i = 1) = P(a_i = 1|L_i = 1) \cdot P(L_i = 1) + P(a_i = 1|L_i = 0) \cdot P(L_i = 0) \quad (\text{C.1})$$

All elements can be estimated by sample shares:

- $\hat{P}(a_i = 1|L_i = 1)$: the share of adopting households in the final random sample among those identified as adopting by the village leader

- $\hat{P}(a_i = 1|L_i = 0)$: the share of adopting households in the final random sample among those identified as non-adopting by the village leader
- $\hat{P}(L_i = 1)$: the share of adopting households as identified by the village leader in the random sample
- $\hat{P}(L_i = 0)$: the share of non-adopting households as identified by the village leader in the random sample

These sample shares are then inserted in equation C.1 to calculate the population probability of being an adopting household. The probability of being a non-adopting household can be calculated in a corresponding way.

D Appendix tables

Table A.1: OLS results with wild bootstrap-t p-values in square brackets

	(1)	(2)	(3)
PANEL A: Linear network variables			
RIPAT banana growers in network	0.192 (0.030) [0.000]	0.192 (0.031) [0.000]	0.149 (0.028) [0.000]
Non-RIPAT banana growers in network	0.041 (0.015) [0.038]	0.042 (0.015) [0.030]	0.026 (0.015) [0.302]
Network size	-0.024 (0.008) [0.002]	-0.024 (0.008) [0.002]	-0.023 (0.007) [0.002]
PANEL B: Extensive and intensive margin			
1+ banana grower in network	0.333 (0.070) [0.000]	0.331 (0.070) [0.000]	0.260 (0.070) [0.002]
2+ banana growers in network	0.044 (0.099)	0.043 (0.097)	0.061 (0.098)
3+ banana growers in network	0.081 (0.190)	0.087 (0.187)	-0.032 (0.175)
1+ non-RIPAT banana grower network	-0.113 (0.099) [0.286]	-0.106 (0.099) [0.330]	-0.075 (0.093) [0.454]
2+ non-RIPAT banana growers network	0.037 (0.131)	0.030 (0.127)	-0.066 (0.121)
3+ non-RIPAT banana growers network	-0.094 (0.212)	-0.088 (0.217)	0.033 (0.206)
Network size	-0.022 (0.008) [0.002]	-0.022 (0.008) [0.002]	-0.022 (0.008) [0.002]
Farmer characteristics	No	Yes	Yes
Household characteristics	No	No	Yes
Observations	509	509	509

Notes: The dependent variable is an indicator equal to one if the farmer grows improved bananas in 2010. OLS estimates, standard errors in parentheses clustered at the subvillage level, wild bootstrap-t p-values are presented in square brackets calculated as suggested by Cameron et al. (2008). Observations are weighted with inverse sampling probability weights.

Table A.2: Network estimates using random sample only

	Logit estimates				Marg.eff.
	(1)	(2)	(3)	(4)	(5)
PANEL A: MAIN SPECIFICATION					
RIPAT banana growers in network	1.481*** (0.22)	1.546*** (0.23)	1.298*** (0.27)	1.291*** (0.32)	0.279***
Non-RIPAT banana growers in network	0.356** (0.15)	0.347** (0.14)	0.273* (0.16)	0.186 (0.22)	0.056*
Network size	-0.358*** (0.14)	-0.365*** (0.13)	-0.392*** (0.14)	-0.397*** (0.11)	-0.078***
PANEL B: EXTENSIVE AND INTENSIVE MARGIN					
1+ banana grower in network	2.645*** (0.47)	2.697*** (0.47)	2.322*** (0.55)	1.980*** (0.50)	0.490***
2+ banana growers in network	0.273 (0.53)	0.416 (0.59)	0.561 (0.63)	0.763 (0.67)	0.120
3+ banana growers in network	0.655 (1.01)	0.582 (0.97)	-0.680 (1.29)	-0.460 (1.21)	-0.149
1+ non-RIPAT banana grower in network	-0.824 (0.53)	-0.883* (0.53)	-0.547 (0.61)	-0.939 (0.66)	-0.132
2+ non-RIPAT banana growers in network	0.015 (0.90)	-0.007 (0.90)	-0.146 (1.23)	0.064 (1.21)	-0.034
Network size	-0.399*** (0.14)	-0.404*** (0.14)	-0.405*** (0.14)	-0.407*** (0.11)	-0.077***
Farmer characteristics	No	Yes	Yes	Yes	Yes
Household characteristics	No	No	Yes	Yes	Yes
Subvillage fixed effects	No	No	No	Yes	No
Observations	356	356	356	309	356

Notes: The dependent variable is an indicator equal to one if the farmer grows improved bananas in 2010. Only the random sample of farmers is used to produce the estimates. Panel A presents network estimates based on the main specification as presented in Table 2. Panel B presents the extensive and intensive margin estimates as presented in Table 3, however as only 4 farmers know 3 or more non-RIPAT banana growers in the random sample, this variable is excluded. Logit coefficient estimates are presented in column (1)-(4), standard errors in parentheses are clustered at subvillage level. Marginal effects presented in column (5) are calculated as described in the text based on estimates from column (3). * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table A.3: Having a banana plantation with ten or more plants

	Logit estimates				Marg.eff.
	(1)	(2)	(3)	(4)	(5)
RIPAT banana growers in network	1.093*** (0.15)	1.061*** (0.16)	0.985*** (0.19)	0.881*** (0.17)	0.235***
Non-RIPAT banana growers in network	0.272** (0.11)	0.280** (0.11)	0.310*** (0.10)	0.308*** (0.10)	0.076***
Network size	-0.163 (0.11)	-0.176* (0.10)	-0.206** (0.09)	-0.255*** (0.07)	-0.051**
Farmer characteristics	No	Yes	Yes	Yes	Yes
Household characteristics	No	No	Yes	Yes	Yes
Subvillage fixed effects	No	No	No	Yes	No
Observations	509	509	509	496	509

Notes: Dependent variable is an indicator variable equal to one if the farmer has a banana plantation with ten banana plants or more. Logit coefficient estimates in column (1)-(4), standard errors in parentheses, clustered at subvillage level. Marginal effects presented in column (5) are calculated as described in the text based on estimates from column (3). * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table A.4: Placebo results without subvillage fixed effects and with marginal effects

	Vegetables		Sunflower		Sugar cane	
	(1) logit	(2) marg.eff.	(3) logit	(4) marg.eff.	(5) logit	(6) marg.eff.
RIPAT banana growers in network	-0.046 (0.25)	-0.010	0.263** (0.10)	0.034**	0.115 (0.14)	0.008
Non-RIPAT banana growers in network	0.131 (0.15)	0.028	-0.036 (0.09)	-0.005	-0.007 (0.08)	-0.000
Network size	0.027 (0.03)	0.006	0.049** (0.02)	0.006**	-0.060 (0.09)	-0.004
RIPAT growing vegetables	0.401 (0.33)	0.083				
RIPAT growing sunflowers			0.758** (0.38)	0.123**		
RIPAT growing sugarcanes					1.628*** (0.59)	0.216***
Observations	509		509		509	
Mean of dependent variable	0.489		0.134		0.081	
Std.dev. of dependent variable	0.500		0.341		0.272	

Notes: The dependent variables are a indicators equal to one if the farmer grows vegetables (column 1-2), sunflowers (column 3-4) and sugarcane (column 5-6) in 2010. Conditional logit coefficient estimates are presented in column (1), (3) and (5), standard errors in parentheses are clustered at the subvillage level. The remaining columns present marginal effects calculated at the sample mean. Farmer and household characteristics are included in all specifications, but in column (3)-(4) religion dummies are excluded as Catholic dummy predicts non-adoption perfectly leading to drop of 43 observations. The number of crops grown in 2010 is subtracted traditional and improved bananas and the placebo crop. * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.

Table A.5: Adoption of banana cultivation in second half of 2009 or later, controlling for other crops grown in 2009

	Logit estimates					Marg.eff. (6)
	(1)	(2)	(3)	(4)	(5)	
RIPAT banana growers in network	0.783*** (0.23)	0.863*** (0.20)	0.873*** (0.19)	0.855*** (0.22)	0.807*** (0.23)	0.183***
Non-RIPAT banana growers in network	0.123 (0.11)	0.134 (0.11)	0.130 (0.11)	0.118 (0.12)	0.128 (0.11)	0.025
Network size	-0.120* (0.07)	-0.076 (0.09)	-0.074 (0.08)	-0.088 (0.09)	-0.129* (0.07)	-0.018
Grows improved maize in 2009		0.233 (0.29)	0.296 (0.30)	0.125 (0.36)	0.109 (0.42)	0.026
Grows traditional maize in 2009		-0.060 (0.33)	-0.094 (0.34)	-0.174 (0.35)	0.104 (0.41)	-0.037
Grows beans in 2009		0.853** (0.41)	0.828** (0.40)	0.929** (0.40)	0.768 (0.50)	0.168**
Grows vegetables in 2009		0.233 (0.28)	0.283 (0.27)	-0.086 (0.35)	-0.243 (0.40)	-0.018
Farmer char.	Yes	No	Yes	Yes	Yes	Yes
Household char.	Yes	No	No	Yes	Yes	Yes
Subvillage fixed effects	Yes	No	No	No	Yes	No
Observations	318	392	392	392	318	392
P-value (testing other crops = 0)		0.045	0.015	0.068	0.587	

Notes: The dependent variable is an indicator equal to one if the farmer grows improved bananas in 2010. Farmers who planted bananas before the second half of 2009 are excluded from the sample. Column (1) provides logit estimates using the same specification as column (4) in Table 2 on this subsample. Column (2)-(5) provides logit estimates while column (6) presents marginal effects based on column (4) estimates. Standard errors in parentheses are clustered at the subvillage level. The bottom row presents p-values from χ^2 tests of joint insignificance of the coefficients to the four other crops (improved and traditional maize, beans and vegetables). * denotes significance at 10 pct., ** at 5 pct., and *** at 1 pct. level.