Targeted use of Randomized Controlled Trials (RCTs)

What is an RCT: A randomized control trial is the gold standard for determining the effectiveness of a program. People sign up for a program, and then are randomly assigned either to receive a program (treated or test), or to not receive the treatment (control). The results are then compared. This is the best-known way to address "selection bias," where those who sign up for the program might be systematically different from those they are measured against. One Acre Fund believes in a world where programs are measured for effectiveness using high-quality methodologies like RCT when possible.

Why it is difficult: RCTs can be expensive, time-consuming, and operationally complex. Once we have marketed our program to farmers and they have signed up to participate, it can be difficult to randomly exclude some that we would have been able to serve. Because RCT is expensive and difficult to execute, the results are limited to a small area for a small number of crops.

One Acre Fund internal M&E strategy: One Acre Fund does not regularly employ RCTs, but rather uses a reasonably high-quality quasi-experimental design to test our impact in a broad variety of contexts, for roughly 5,000 test and control farmers per year. This routine measurement activity allows us to measure across all countries, across regions within countries, and to compare ten different crops. This provides rich, easy-to-use operational data to refine our program. Crucially, we pay a lot of attention to control group selection, and constantly experiment with methods such as finding farmers who are "enrolled but not yet treated" or "likely to enroll," and propensity score matching.

One Acre Fund's increasing use of RCTs: We will increasingly use RCTs to confirm that our regular measurement methods produce correct indications of impact. Below, we summarize one RCT completed in 2009 with independent researchers. We subsequently completed another 2014 RCT which we conducted internally (with the analysis portion provided by an outside firm). Both studies were conducted in Western Kenya and examine impacts on maize production. Maize is relatively simple to measure (as opposed to many other crops), and frequently accounts for the majority of our impact.

We do not believe that either RCT provides definitive evidence of positive harvest impact – they both have important limitations. There is high variability in farmer results, particularly in the 2009 study, resulting in 11 percent of farmers having negative return to farm inputs (although this is lower amongst One Acre Fund-trained farmers). This information was concerning and led to important program changes. As indicated in the above table, the 2009 study also conflicted with our internal 2009 M&E, which led to a significant overhaul in our M&E methodology.

Despite the limitations of any individual study, however, One Acre Fund believes that these two studies largely validate the current internal M&E that we conduct with thousands of farmers every year. Together, we believe that these RCTs plus internal measurement currently add up to a reasonable body of evidence that the One Acre Fund program has positive impact on harvest profits. Here are more detailed write-ups of both studies, including lessons learned by One Acre Fund. Both studies have important drawbacks, but they are also are rich in detail that is helpful for program learning.

In particular, the 2009 study summarized below drove a lot of programmatic learning at One Acre Fund. On the measurement side, we changed a lot of our M&E practice, made assumptions more conservative, and overhauled our M&E agent hiring and development. In 2010, we started an agriculture innovation

team that is dedicated to steadily increasing our agriculture yields. We diversified our programming so that today a majority of our estimated impact comes from non-maize crops and energy products. We increased customer screening to deliberately avoid customers that might have a high propensity to have a negative return to farm inputs. We began a crop insurance program that is now Africa's largest smallholder crop insurance program. We also conducted several studies on non-customers and our own customers, seeking to better understand their comparative wealth levels.

2009 Kenya RCT for Maize

Overview

Independent researchers conducted an RCT for One Acre Fund in 2009, our third year of operation, with the major goals of determining how effective our internal M&E was, and generating program learning early in our history. Our internal M&E results at the time were shown to be substantially different from the RCT results, which led to several important changes to our internal M&E methodology—eventually causing us to restate our impact figures when we had developed confidence in a changed method. The RCT showed some farmers who had a negative return on investment from farm inputs and also showed some evidence that we were selecting wealthier farmers – both of which were cause for concern and led to program change.

Heterogeneity of results: One of the variables of interest was the level of return on investment (ROI) experienced by the farmer in response to labor and farm input expenditure. This study showed high variance in ROI from using farm inputs, with some farmers experiencing negative returns to inputs. This was concerning and led to important programmatic changes. Some of these results were driven by an important drawback to the study design and execution, discussed in detail below (see "lessons learned").

Low profit compared to 2009 internal One Acre Fund M&E: One Acre Fund conducted this RCT in our third year of operation, when we were a young organization. A major purpose of the study was to determine the correctness of One Acre Fund's internal M&E practice. The RCT results showed percent and dollar profit impact from the One Acre Fund program that were lower than our 2009 internal M&E. In the RCT, farmer profit levels were only 40 percent higher on "Trained/Input" plots (reflecting the full One Acre Fund program) vs. "Not Trained/No Input" plots (reflecting the control). This corresponded at the time to total incremental profit of only $30 USD per farmer planting 0.5 acres with One Acre Fund. This contributed to the decision to change our internal M&E process, and to eventually restate our impact numbers once we had built confidence in the results.

Changed M&E method is now reasonably consistent, although there are differences in the dollar profit results: This study did contribute to changes to our M&E method, and our regular M&E is now somewhat consistent with most of the results. In 2014, our internal M&E in Kenya showed that our farmers obtained a 21 percent increase in maize net revenue (what we call "profit" internally) when compared to the point estimate of 40 percent in the 2009 RCT (using One Acre Fund-obtained farm input prices). Our internal dollar estimate of $87 USD increased net revenue in 2014 is higher than the estimate of the 2009 researchers, who translate harvest increases to only $30 USD (using One Acre Fund-obtained farm input prices and normalized to 0.5 acres average planted acreage). One Acre Fund believes this dollar discrepancy is potentially bridged by updating maize selling prices to 2014 levels; by the fact that One Acre Fund farmers plant 0.7 acres, not 0.5 acres; by unusually low baseline yield in 2009 (perhaps due to drought); and a study limitation that caused unusually high labor and farm input expenditure of farmers on their "No Inputs" field (see study limitations).

Although maize is our largest source of impact in Kenya, our program has now significantly diversified to other crops and products such as solar lamps, making maize responsible for less than half of estimated impact in Kenya.

This study contributed to changes to our M&E practices, the diversification of our programs into other crops and energy products, the creation of an agriculture innovation team to improve impacts, increased customer screening practices, non-customer studies, and our crop insurance program. See "What We Learned" below for more detail.

Study Design and Objectives

In late 2008, farmers in ten "cells" in Chwele District, Kenya signed up for the One Acre Fund program. Each "cell" was served by a One Acre Fund field officer. There were 1,155 enrollees in total. Every farmer identified two side-by-side, quarter-acre plots of land within their farm. On one plot, they were given a full One Acre Fund loan for farm inputs. On the other plot, they were instructed to buy farm inputs as they "normally" would do. Farmers were randomized by "group," a unit of roughly nine farmers each, and randomly split into "Trained" and "Not Trained" groups.

This formed four types of plots: (A) Not Trained/ No Inputs, (B) Not Trained/ Inputs, (C) Trained/ No Inputs, (D) Trained/ Inputs. In theory, (A) vs. (D) should represent a measurement of the impact of the One Acre Fund program, and this was One Acre Fund's primary comparison of interest. Although the experimental design may seem odd (and indeed had problems as discussed below in the study limitations section), the design was jointly selected with the researchers because we would not have to exclude anybody from receiving at least some treatment.

Harvests for all plots were then weighed by a field team hired by the researchers, supervised by two successive research associates, and then a One Acre Fund program associate. 826 farmers planted in the program, and 757 were included in the dataset. A five-by-five meter box was constructed in a random portion of each quarter-acre plot, and the physical harvest weighed. Costs were determined using a survey of both farm inputs and labor expended. This determined farm profit. There were several dependent variables of interest, but One Acre Fund was most interested in farmers' return on investment (ROI) for their expenditure on the One Acre Fund program. Our primary variable of interest was percent profit improvement = (Plot D revenue – Plot D cost) / (Plot A cost – Plot A cost) - 1.

Results and What We Learned

Low dollar Impact compared to 2009 internal measurement. The key finding of interest to One Acre Fund was that Plot D profit of 3,616 Kenya shillings was only 40 percent higher than Plot A profit of 2,583 Kenya shillings. When normalized to 0.5 acres, this amounted to only about $30 USD of incremental profit for a One Acre Fund farmer in 2009. This was significantly lower than our regular M&E at the time.

One Acre Fund believes that some of the low dollar impact is accounted for by unique 2009 conditions, including drought, low maize prices, and farmers changing practice on "control" plots (see study

limitations below). However, we did learn many important lessons from the RCT results, as outlined below.

- *Change M&E practices and re-state impact figures:* The 2009 study caused us to make several important changes to our internal M&E, including more thorough control group selection, better M&E staff selection and training, and increased rigor to the actual day-to-day measurements (see this memo for more details). This also contributed to our restating our impact figures more conservatively in 2013. The delay was caused by uncertainty around how unique conditions in 2009 may have influenced the results, and because agriculture measurement can only occur once per year. It took time to build certainty in the new figures. Today our internal M&E is roughly consistent, as discussed above.
- *Agriculture innovation team*: In mid-2010, One Acre Fund created an agriculture innovation team which today is primarily supported by the Bill and Melinda Gates Foundation. In 2014, this team ran dozens of trials together with thousands of farmers, testing everything from fertilizer dose, to new crops, to differential plant spacing. The products of this team have led to steadily improving agriculture recommendations.
- *Diversified programming*: Today, an increasing proportion of our impact comes from non-agriculture products. For example, we are now one of the largest sellers of solar lamps in Africa (see below for two RCTS that we completed around our solar impact). Although maize is our single-largest source of estimated impact, it now accounts for less than half of estimated impact in Kenya and even less than that in other countries.

Some farmers had a negative return to farm inputs. The study allowed for a direct side-by-side evaluation of the return on investment from using farm inputs. 11 percent of farmers in the study had *lower* farm profit on their "Inputs" side when compared to their "No Inputs" side. Although this proportion was lower for trained farmers, this is still cause for concern. Buying farm inputs is a major risk for households, and we must not make them worse off financially. This result caused several program changes, as described below.

- *Increased customer screening and prepayment.* One Acre Fund has now instituted two important program changes that limit our program to farmers that we believe will more consistently benefit from our program. First, we began visiting most farmers' fields at enrollment, ensuring that they have a suitable piece of land on which to plant. Second, we now require farmers to prepay ten percent of their loan to join, demonstrating their ability to participate in a market-like program.
- *Increased attention to "tails" in our measurement.* We began paying greater attention in our monitoring and evaluation studies to farmers in the "low tail." When we evaluate a product, we look both at the point estimate of impact and also the dispersion of results. Risk for the farmer must be minimized.
- *Trial work to reduce fertilizer usage recommendations.* One theoretical problem is that fertilizer could have diminishing returns to profit, as farmers use more. We have now conducted many trials in several countries on fertilizer levels, subsequently reducing fertilizer recommendations for some crops, such as climbing bean in Rwanda.
- *Crop insurance.* One Acre Fund began working with crop insurers starting in the 2010 season, and by 2013 made crop insurance a mandatory component of our program. Today, we are the largest retailer of crop insurance to smallholder farmers in Africa. In 2014, we insured the harvests of more than 200,000 farm families.
- *Initial farmers appeared less poor than average.* According to the RCT results, One Acre Fund farmers averaged $900 USD in assets comparable to shopkeepers in a similar area. Also, One

Acre Fund farmers average 8.7 years of schooling, nearly two years higher than the area average. We respectfully disagree with the researchers that higher assets necessarily mean our farmers are less poor; we feel large assets are explainable because farming is simply an asset-intensive business. The greater education finding caused us to increase our scrutiny of the poverty factor.

- *Non-customer studies.* One Acre Fund now regularly conducts non-customer studies that enable us to better understand who is selecting into our program and who is not. Customers are surveyed on demographic characteristics as well as asked why they did not join our program. This has enabled us to better target our program.
- *Longer repayment window.* One finding of these studies was that farmers felt a sense of "sticker shock" at how much farm inputs cost (usually in the range of $50-$80 USD). One Acre Fund has worked hard to extend our repayment window in Kenya by nearly 20 percent, enabling farmers more time to pay their loans.

Study Limitations

Single-year effects. There were several factors unique to the 2009 study year. First, 2009 was affected by drought. The Kenya Ministry of Agriculture estimates the maize harvest in 2009 at 27 percent lower than the preceding four-year average. Second, the economics of maize farming in Kenya have changed since 2009. That year, the price of maize (the source of revenue) was 2,000 Kenya shillings per 90-kilogram (kg) bag, or $28 USD at late-2009 exchange rates. As of late 2014, maize prices were approximately 3,000 Kenya shillings per 90-kg bag, or roughly $33 USD at late 2014 exchange rates. At the same time, farm input prices are actually slightly lower in 2014 compared to 2009. Revenue is higher, and costs are lower.

Mathematically, these factors would not affect the point-estimate on percent profit improvement (they impact test and control equally) but they would probably increase the dollar impact level to possibly be more consistent with what we find in our program today.

Farmers changed their practices on the No Inputs plots. Compared to the "rest of their field," farmers spent a lot more money and time on their "No Inputs" plot, which is a major study limitation. It is not clear why this is, but the high visibility of the plots played a role. Follow-up surveys also indicate that farmers believed it was important to their One Acre Fund field officer to perform well on both Inputs and No Inputs plots. The failure to communicate with field staff about the importance of impartiality was an important lesson incorporated into future trials. This contributed to a major overhaul in staff selection and training in our M&E department.

Because farmers changed their practices on the No Inputs plots, the amount of time and money spent on those plots was artificially boosted. Farmers spent 61 percent more money and 105 percent more time on the supposed "No Inputs" plot when compared to a similar quarter-acre on the rest of their farm that was not included in the study. We believe that this likely would cause the study to underreport the magnitude of the dollar effect of the One Acre Fund program, because there was an artificially low difference between Inputs (test) and No Inputs (control). This also could contribute to the high heterogeneity of farmer results. However, this is speculative.

Further reading

The above summary has been completed by One Acre Fund in an attempt to succinctly summarize the study results while also including our most important program learning. Data collection and analysis was operated by independent researchers unless otherwise indicated.