

## 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 and another 2014 RCT 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 in 2014, Kenya maize harvest accounted for almost half of our Kenya impact.

These two RCTs estimate a percent harvest impact that is reasonably consistent with our 2014 internal M&E findings for Western Kenya. However, there are important study limitations and low statistical significance in the most recent study. Below is One Acre Fund’s interpretation of some of the core study results of interest.

	<b>Estimated % improvement in maize profit from One Acre Fund program</b>	<b>Translated \$ (USD) impact on farm profit, from maize alone</b>
2014 RCT	31% <sup>(1)</sup>	\$91 <sup>(3)</sup>
Internal 2014 M&E	21%	\$87
2009 RCT	40% <sup>(2)</sup>	\$30 <sup>(4)</sup>
Internal 2009 M&E	100%	\$120

(1) Significant at the 0.09 level using wild-cluster bootstrap to adjust for low number of randomization units (we believe this to be the correct regression specification). Naive regression specification without adjustment is significant at the 0.01 level.

(2) Cost assumptions use OAF-obtained farm input prices. One Acre Fund calculated results to be significant at the 0.001 level.

(3) Baseline profit \$ was slightly lower in the RCT study area, compared to our internal 2014 M&E, which is why dollar impact is lower than internal 2014 measurement despite higher percent profit improvement.

(4) One Acre Fund (distinct from the independent researchers) believes that 2009 exhibited simultaneous drought and lower maize prices compared to 2014, which caused baseline profits to be low, resulting in lower dollar impact. There was also a study limitation that may have affected \$ impact. See below for more details.

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. Our 2014 internal M&E is now more consistent with both RCTs.

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 rich in detail that is helpful for program learning.

## **2014 Kenya Mini-RCT for Maize**

### Overview

The major purpose of this study was to evaluate the accuracy of our internal M&E when compared to a more rigorous RCT design. Overall, the estimate of the harvest impact from the One Acre Fund program is reasonably consistent compared to what we normally see in our regular M&E in Kenya. The study showed 32 percent increase in revenue corresponding to 31 percent profit improvement, or \$91 USD profit improvement for maize in Kenya in 2014. This compares to our internal estimate of 21 percent increase in profit corresponding to \$87 USD per farmer for our Kenya program in 2014. (Baseline profit was lower in the RCT study area, resulting in lower dollar impact per percentage point improvement in profit<sup>1</sup>).

We also learned several important lessons for future RCTs. These lessons include the need to increase the number of randomization units, the importance of extremely intense harvest follow-up to avoid sample attrition, and the need to collect more thorough intercropping data outside of just the core crop of interest (maize). These lessons will be discussed in more detail below.

### Study Design and Objectives

In late 2013, 1,195 farmers in six “cells” in Busia District, Kenya signed up for the 2014 One Acre Fund program. After enrollment was concluded, One Acre Fund randomly selected two cells to not receive the One Acre Fund program. Instead, these farmers were given a free mobile phone, and a health program that included free enrollment in national health insurance. Out of 1,195 total enrollees, 772 were in the treatment group, and 423 were in the control group.

Farmers not receiving the program were deeply disappointed, and our program reputation suffered in this area. However, we felt that *total* program exclusion was important to get a truly representative result. To help mitigate disappointment, we felt it ethically important to offer the health program of similar value to control farmers. Control farmers were also offered the opportunity to join the One Acre Fund program in the next year (2015).

---

<sup>1</sup> See “Results” section below for more context on the precision of these results.

At harvest, field agents were sent to farmers' fields. Two different 8x10 meter boxes were placed in each farmer's field, the harvests weighed, and revenue calculated. Costs were determined by surveying input purchase and reported labor.

After data collection was complete, IDinsight conducted the analysis to generate the estimated impact on maize production (see [here](#) for a full report and [a summary of study limitations](#)) and provide an impartial view on study limitations. [Idinsight](#) is an NGO devoted to helping international development managers generate and use rigorous evidence to improve at-scale impact. One Acre Fund was responsible for study design, data collection, and converting yield numbers to profit without external support.

### Results and What We Learned

*Reasonable consistency with internal M&E.* The RCT results were overall consistent with our internal M&E, actually showing higher percent profit improvement than normal. The point estimate for harvest revenue increase for One Acre Fund farmers compared to control was 31 percent improvement, corresponding to a profit improvement of \$91 USD. This compares to our internal estimate of 21 percent profit improvement, corresponding to \$87 USD in maize profit improvement per farmer in our full Kenya program. We also saw similar improvements on our sorghum and millet profitability, although we focused on maize for simplicity of analysis and write-up of results.

*Low precision.* Although the raw data show high precision with significance  $<0.01$ , it would not be correct to accept this result considering the low number of randomization units (only 2 control cells and 4 test cells). IDinsight attempted to correct for this factor using both the clustered sandwich estimator, showing significance levels of 5 percent, and then wild cluster bootstrap, showing a significance level of 9 percent. Therefore, there was a small but not insignificant probability that the program impact reported was due to the random selection of test and control cells. This is an important lesson.

*Randomization units.* In future RCTs, One Acre Fund will consider ways to increase the number of randomization units – for example, randomizing by farmer group. We would need to significantly increase randomization units to make a difference, which presents a challenge. If we randomized by farmer group, for example, there would be neighboring test and control farmers, creating many opportunities for program farmers to impact the control group. An alternate strategy is to simply accept lower significance levels and simply repeat these studies regularly to build confidence in results.

### Study Limitations

*Sample attrition.* Due to the operational challenge of placing 1,000+ “harvest boxes” in farmer fields all at once, One Acre Fund randomly selected 287 treatment farmers and 286 control farmers out of 1,141 farmers surveyed at baseline to collect maize harvest data. Of these farmers, production data was successfully collected from 210 (73 percent) treatment farmers and 203 (71 percent) control farmers. In general, sample attrition was caused the sheer number of harvest measurement boxes that needed to be created during the extremely narrow window of time for this activity. (One Acre Fund shifted a dozen extra enumerators into the area but that was not enough). Some checks lead us to believe that there was not any significant bias created by this. However, there is still a valuable lesson: Put extraordinary effort into harvest surveys. Especially for studies where we have put a lot of effort into setting up the experiment, it is worth putting extra-heavy focus on harvest enumeration.

*Intercropping.* More control farmers when compared to treated farmers intercropped beans together with their maize fields. This is a source of bias, making the estimated impact of the One Acre Fund program slightly too high. We believe this factor is not very large and in the range of, at most, a 5 percent bias – intercropped beans generally produce very low harvests in Kenya. Starting from 2015,

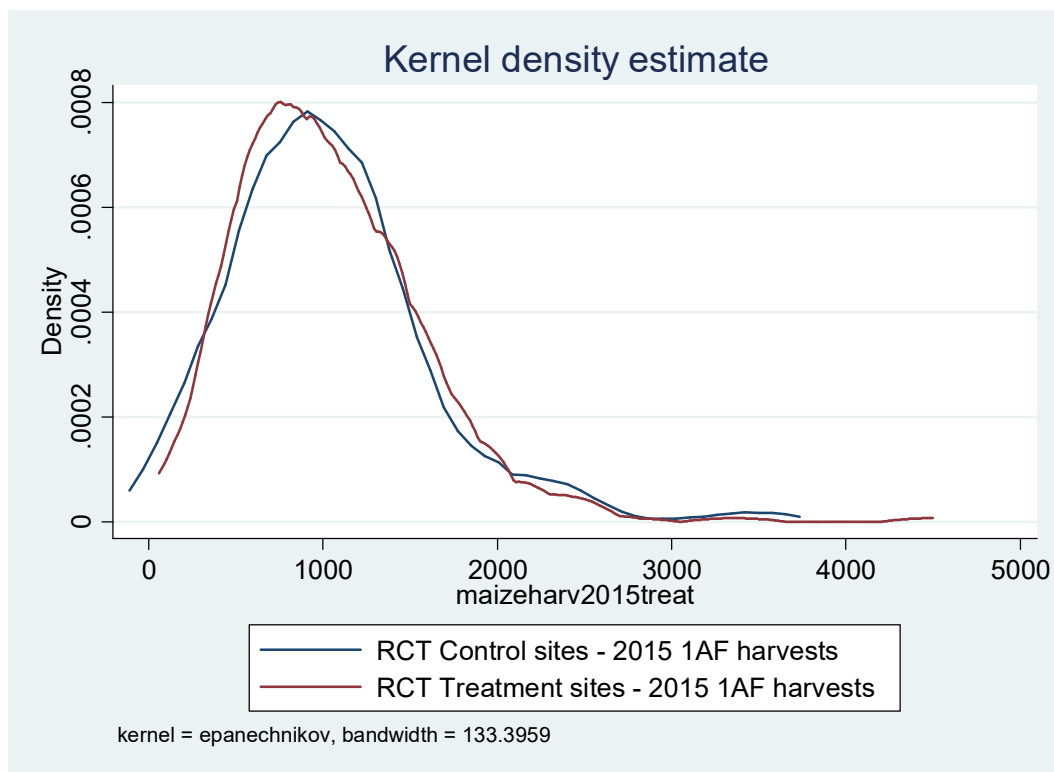
One Acre Fund will begin separately measuring intercropped fields against maize-only and bean-only fields to get a more precise estimate of this factor.

### 2015 Follow-up

One of the known flaws in our 2014 RCT was randomization over only 6 areas (2 control, 4 test), although the absolute number of farmers in the study was high, with 1,200 participating. Because we administer our program at the “sub-location” level, increasing the number of randomization units would have forced us to enroll and then de-enroll many thousands of farmers. So we settled for fewer units. As a result, this reduced the statistical significance of the study. To put it simply, with such few clusters of randomization it was difficult to say with certainty that the harvest impact was due to program participation or because the 2 sites which were randomly de-selected were simply poorer performing to begin with. Although our raw result showed 1% statistical significance of treated vs. untreated farmers, after adjusting with wild cluster bootstrap for the low levels of randomization units, this lowered our statistical significance to 9% - reflecting the small chance that random area selections explained our treatment effect, not our actual program.

We followed up the next year in 2015, to determine if in fact there was anything systematically different about the two control areas versus our four treatment areas. We collected self-reported harvest data for farmers in all six areas, in 2015. (In the four treatment sites, about two-thirds of the study population re-enrolled and in the two control sites 80% of de-selected farmers enrolled in One Acre Fund in 2015).

The data we collected show that the harvests from those who joined our program in 2015 from the four 2014 treatment sites and the two 2014 control sites are remarkably similar, with both averaging around 1030 kg/acre (see table below for the distribution of harvest yields). While this finding cannot completely overcome the limits of an underpowered RCT, it lends support to the idea that the observed effects in 2014 were due to our program and not the results of location-specific effects in the 6 areas under study.



In addition, we took the opportunity to compare the currently enrolled farmers to those who stayed out of the program in 2015 (all farmers had chosen to enroll in the program in 2014, so are a highly comparable sample). Controlling for total value of livestock, cash income, age, gender, marital status and location specific effects, we found a highly statistically significant ( $p=.012$ ) program impact of 32% increase in yields comparing farmers from our original RCT who chose to re-enroll in 2015, with those from the same group who chose not to. This is roughly similar to what we find in our internal program-wide M&E, which found a country-wide 40% improvement in maize yields.

This finding is particularly valid as it is over a study population in which the full sample self-selected into the program in 2014. There might still be unobservable differences between those farmers who did not re-enroll in 2015 or stayed out of the program in 2015 and those who chose to enroll, which we are not able to control for.

Taken together, we believe these two follow-up analyses provide strong confirming evidence that the effect shown in our 2014 Kenya maize RCT is due to program participation rather than some idiosyncrasy related to the random assignment of sites.

See here for a [more complete report](#).