Current Strategy

Our annual impact evaluation efforts consist of obtaining wide geographic representation of One Acre Fund farmers and comparing their harvest yields and agricultural profit to those of similarly situated neighboring farmers who are subject to the same micro-climate conditions. We attempt to minimize differences between these groups by (1) excluding farmers who are vastly different (e.g. have large land sizes, or do not rely on agriculture) (2) select farmers who are recommended as "interested" in joining One Acre Fund the following year and (3) matching One Acre Fund and comparison farmers on relevant characteristics during our statistical analysis.

The value of this measurement strategy is that we can do many measurements, over a wide geographic area, over many crops. In 2015 for example, we obtained harvest information directly from 16,000 fields. This gives us plenty of data, and enables us to make programmatic changes in response to that data. However, we also make targeted use of more rigorous approaches (randomized control trials, or RCTs, and Difference-in-Difference, or Diff-in-Diff) to verify the quality of our everyday measurements.

More Rigorous Approaches

One Acre Fund uses both RCT and Diff-in-Diff to verify our everyday measurements.

Our most recent RCT was in 2014 in Kenya. This study estimated a very similar impact to our regular impact estimate. (The RCT actually found a slightly higher impact). There were, however, significant limitations to the study – although 1,200 farmers were involved, we only had 4 test sites and 2 control sites. Statistical significance was 0.01 if we treat every farmer equally, but after properly adjusting for the small number of randomization units (using wild cluster bootstrap to adjust our standard errors), statistical significance reduced to 0.09. See here for a full discussion. Because RCTs involve enrolling farmers in our program and then dropping them randomly, they can be difficult to execute as a frequent evaluation effort.

We also use difference-in-difference as another rigorous method. This memo covers our latest, and most comprehensive, difference-in-difference studies.

- To employ a difference-in-difference approach, we measure a farmer's harvest the year before joining, and then again after joining the One Acre Fund program – a before and after comparison. If, for example, a farmer's harvest makes a big jump in year 2 after joining our program, then this is probably caused by the One Acre Fund program.

- Of course, it is possible the jump in harvest was caused by differences in the actual year – bad rain in year 1 but good rain in year 2, for example. This is why we also measure a comparison group of farmers – those farmers who remained out of our program in both years. If those farmers' harvests trends up or down, we "cancel" that effect out.
- Once this "year effect" is eliminated, we can be reasonably confident that the difference in harvests is caused by the One Acre Fund program. This is why the difference-in-difference methodology is so powerful, because it allows elimination of "individual year effects." (There are other potential biases, which we address below).

What We Did

In 2015, we did small experiments in three of our four core countries: Kenya, Tanzania and Burundi (Rwanda was not possible due to timing reasons). In these countries, we followed up with several hundred farmers, none of whom were One Acre Fund participants in 2014, but some of whom became One Acre Fund farmers in 2015. This allowed us to assess the impact over time for farmers who entered the program compared to similar farmers who stayed out of the program. Because of the difficulty of (1) following up with the same farmers over time and (2) predicting how many control farmers will become One Acre Fund farmers in year 2, it has been difficult for us to get very large samples for our diff-in-diff.

In Burundi and Tanzania, we targeted this quasi-experiment to new program areas, surveying farmers during the prior season's harvest but after they had joined the program. In Kenya (where program enrollment happens after harvest), we simply followed up with all of our comparison farmers from the year before and compared the change in harvest over 2 seasons among those who stayed out of our program and those who became One Acre Fund farmers during that time period.

In 2016, we did yet another experiment, this time following three groups of farmers: (1) those who remained out of the program both years, (2) those who joined the program for the first time in 2016 but were not in the program in 2015 (this is the traditional difference-in-difference sample), and (3) those who stayed in the program both years. In this way we had two cohorts (1 and 3) with which to compare farmers entering the program.

What We Found – Summary

Below we present the findings of these studies in our three countries, comparing the difference-in-difference estimate and the core monitoring and evaluation (M&E) estimate. Overall, the difference-in-

difference impact estimates broadly confirm the validity and general magnitude of our core M&E estimates. (*The P-value corresponds to the likelihood that a result would occur by chance. A p-value of <.01 means there is less than a 1 percent chance that result would occur by chance*.)

Country	Sample	Diff-in-diff estimate	Core M&E estimate
Tanzania 2015	*104*	**490****	**478*****
Burundi 2015	*148*	**56*****	**49*****
Kenya 2015	*281*	445	476***
Kenya – Western	*205*	**641*****	**559*****
Kenya – Nyanza	*76*	Not sig	331***
Kenya 2016			
Joiners vs non-joiners	*278*	**515*****	**359*****
Joiners vs. veterans	*327*	**464*****	**359*****
*** p-value> .01; ** p-value >.05; * p-value >.01			

Tanzania 2015

Program context: Tanzania is in its third year of operation. The program provides fertilizer, seed, and credit, focusing almost exclusively on maize. Farmers in Tanzania have slightly larger land sizes than our other core programs.

Sample: A sample of 104 farmers were followed from 2014 new Kilolo district sites. Farmers that were newly enrolled in One Acre Fund (treatment group) and farmers that had never joined (control group) were visited and their maize yields measured. These same One Acre Fund and control farmers were re-visited in 2015 and maize yield measured. The total sample size and distribution of these farmers can be seen in the tables below.

Results: Those who joined One Acre Fund had slightly lower per-acre maize harvests in 2014 than those who stayed out of the program. 2015 was likely a worse agricultural year, as maize yields for the group as a whole went down. However, farmers who joined the program saw a slight increase in their maize harvests on their One Acre Fund land relative to 2014, while those who stayed out of the program saw a significant decline on average. The difference-in-difference estimation was 555 kg/acre and was highly statistically significant.

	2014 maize yield (kg/acre)	2015 maize yield (kg/acre)	Sample	Diff	p-value
Did not join One Acre Fund	1,756.9	1,245.9	*68*	-511.04	
Joined One Acre Fund	1,602.0	1,646.9	*36*	44.90	
		Diff-in-diff		**555.94**	**0.000**

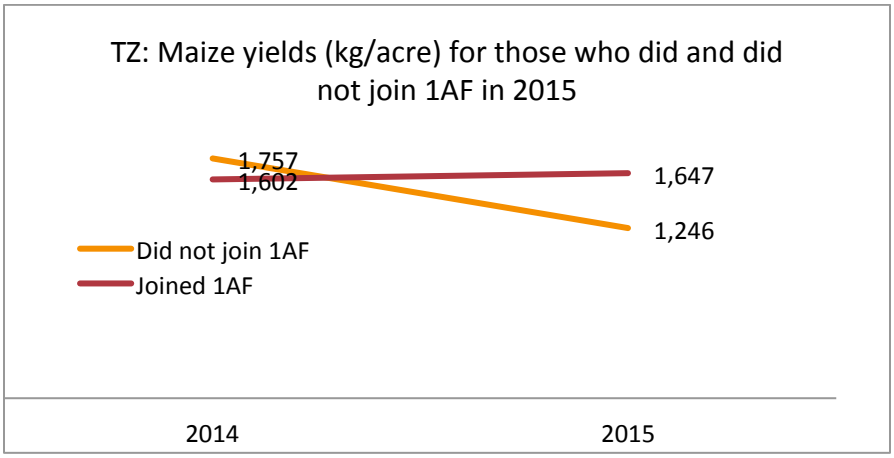

In order to more rigorously analyze these findings, we ran an OLS regression accounting for land size differences (a proxy for wealth), as well as location fixed effects from the four sites we collected data from. These results show a difference of 490 kg/acre at a p-value of .047. In 2015, our core M&E found a difference of 478 kg/acre.

OLS Regression: Dependent variable = change in harvest 2014-2015

	Coefficient	Clustered Std. Err.	t	p-value
Joined OAF	**490.61**	**149.84**	**3.27**	**.047**
Total maize acres in 2014	4.31	49.50	.09	.939
Site specific effects	Included 4 clusters			
Constant	-350.30	100.88	-3.47	.040
N	104			

Burundi 2015

Program Context: Burundi is in its sixth year of operation. The program offers fertilizer on credit and trains on proper planting techniques, covering a variety of crops. Burundian farmers are among the poorest farmers in our program and have relatively small land sizes. They grow a variety of crops but mainly apply fertilizer to potatoes, beans, and maize.

Sample: In 2014B, Burundi M&E measured bean harvest weights among enrolling clients and controls in new program area. The enrolling clients in 2014B had not received One Acre Fund inputs that season, but were enrolling for the following season in 2015A. In 2015B, Burundi M&E measured bean harvests of the *same farmers* who either (1) had been enrolled with One Acre Fund in 2015A and 2015B for our treatment group or (2) had remained controls over the course of the same time period for our control group.

Results: Note that the following harvest weights in the table below are weighted by the same land sizes per field type (with fertilizer and without fertilizer) used in 2014B and 2015B. Those who joined the program in 2015 started off with slightly higher per-acre beans yields in 2014B (500 kg/acre vs. 466 kg/acre). However, over this time period, those who joined the program increased their bean yields by 56.8 kg/acre, whereas those who stayed out of the program increased by only 0.4 kg/acre. This is a difference-in-difference estimation of 56 kg/acre. This compares with an annual program impact for beans in 2015B of 48.8 kg/acre.

	2014B beans yield (kg/acre)	2015B beans (kg/acre)	Sample	Diff	p-value
Did not join One Acre Fund	466	466.4	*54*	0.4	
Joined One Acre Fund	500	556.8	*94*	56.8	
		Diff-in-diff		**56.4**	**.00**

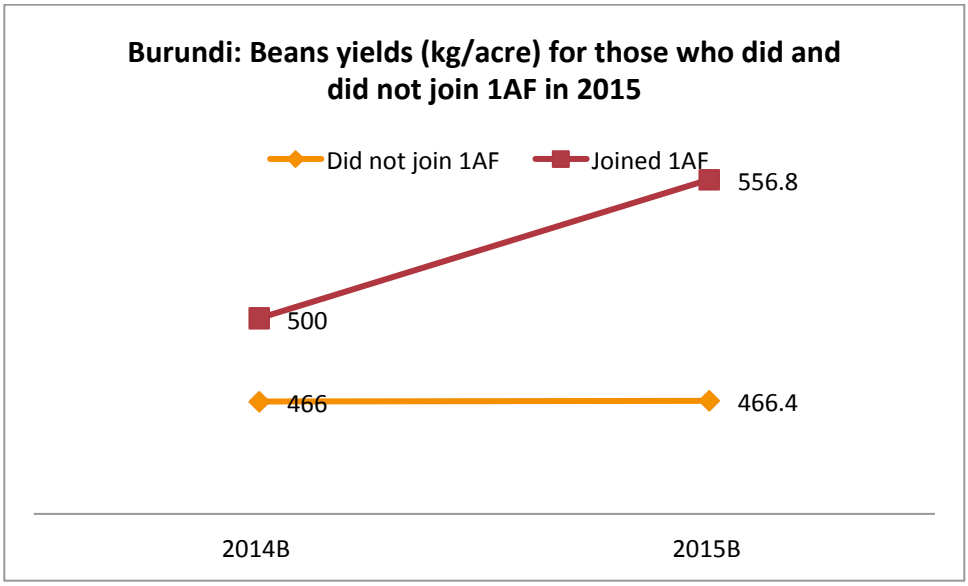

Published April 2017 | www.oneacrefund.org

Note: We were not able to run more rigorous regression equations on the Burundi data. This is because we have a number of field types in this analysis (client exemplary, client fertilized, client non-fertilized, control fertilized, control non-fertilized), and some of our observations switched categories over the time period. For example farmer x would have an exemplary field surveyed in 2014 and a fertilized field surveyed in 2015. We were able to weight each observation by the amount of land dedicated to each field type in each time period to do get a typical per-farmer average in the above analysis. However, we were not able to run any properly weighted regression analyses using the difference in yields over time, given the fact that field categories changed over time.

Kenya 2015

Program Context: Kenya is in its 10th year of operation. The program offers fertilizer on credit and trains on proper planting techniques, focused mostly on maize and beans. Control farmers in Kenya do use significant quantities of fertilizer and hybrid seed on their maize plots. However, compared to OAF farmers, they do under-dose their fertilizer per acre.

Sample: In the 2014 long rains season, One Acre Fund measured harvests for control farmers for Kenya's annual impact assessment activity. Hundreds of farmers from that group were followed through to 2015, and we collected data regarding whether or not they joined One Acre Fund in 2015 and once again measured their yield outcomes.

We found that the average yields for those who did not join One Acre Fund went down over the two years by an average of 252 kg/acre, which is a trend similar to our annual M&E harvest estimates. (Our annual harvest survey found a decrease in yields among control farmers between 2014 (*n=573*) and 2015 (*n=1206*) of nearly 200 kg/acre.) Those who did join our program over this time saw an increase in harvests of nearly 200 kg/acre, for a diff-in-diff estimate of 450 kg/acre.

Results:

Cohort	2014 Maize yield (kg/acre)	2015 Maize (kg/acre)	Sample	Diff	p-value
Did not join One Acre Fund	1264.7	1100.9	*113*	-251.8	
Joined One Acre Fund	1292.8	1501.5	*168*	198.7	
			Diff-in-diff	450 kg/acre	0.00

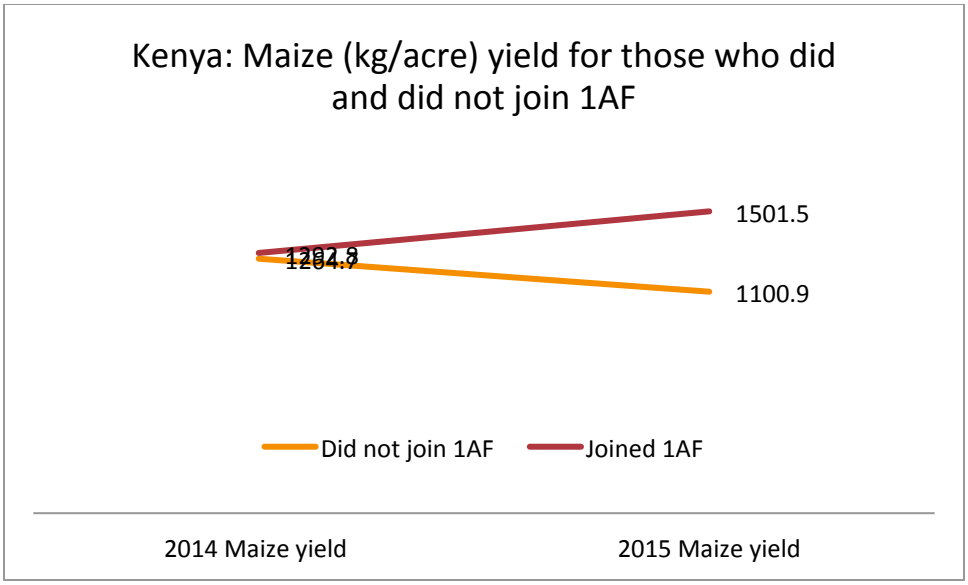

The output above does not take into account the province in which the farmer is located, so we checked that these results generally stay the same when accounting for provincial location, and clustering the standard errors. When we do this, we no longer see a statistically significant result (p=.20).

OLS Regression: Dependent variable = change in harvest over time

	Coefficient	Robust Std. Err.	T	p-value
Joined One Acre Fund	445.16	147.20	3.02	.20
Province effects	Included			
Constant	-259.87	42.61	-6.1	.103
N	281			

Because Nyanza and Western provinces have very different agro-ecological conditions, we further investigated to see if we could see a robust impact looking at the results disaggregated by province. While this further limited our sample size, we were able to detect a highly statistically significant effect in Western province, even when controlling for district-level effects. However, Nyanza did not show any program impact. This might be due to (1) a very small sample size of only 76 farmers or (2) a truly poorer program impact in that region.

We see below that in Western, even controlling for district-level effects, we have a highly significant 641 kg/acre program effect.

OLS Regression in Western: Dependent variable = change in harvest 2014-2015

	Coefficient	clustered Std. Err.	t	p-value
Joined One Acre Fund	**641.83**	**91.04**	**7.05**	**.000**
District effects	Included			
Constant	-490.21	81.28	-6.03	.000
N	205			

However, in Nyanza, we no longer have a statistically significant effect. This is likely due to the very small sample size of 76 farmers, and the fact that we generally find a smaller program effect in Nyanza province, which would make it even harder to find an effect without a large sample.

OLS Regression in Nyanza: Dependent variable = change in harvest over time

	Coefficient	clustered Std. Err.	t	p-value
Joined One Acre Fund	-262.49	493.85	*-.53*	.623
District effects	Included			
Constant	-1012.98	24.69	-41.02	.000
N	76			

Checks on Results

Because our control groups for the difference-in-difference analyses were quite small, we wanted to check that their trends over time were not anomalous and were in fact in line with the trends over the same time period from other groups. We therefore looked at other year-over-year data sources to estimate the directional "year effect."

The tables below (one for each country) compare the 2014 vs. 2015 measurements for the difference-in-difference sample and our average 2014 and 2015 samples for non-One Acre Fund farmers in each country, which include hundreds of observations each. We should note that this is not a perfect check as (1) these do not follow the same farmers over time, so the sample may have changed in composition year over year and (2) these farmers were drawn from a much larger geographic area and may have possibly experienced different yield patterns overall. Still, their pattern over time shows roughly similar trends to the difference-in-difference control farmer trends.

In Kenya between 2014 and 2015:

- Difference-in-difference controls shows a **19 percent** reduction in yields
- M&E controls saw a **13 percent** reduction in yields.

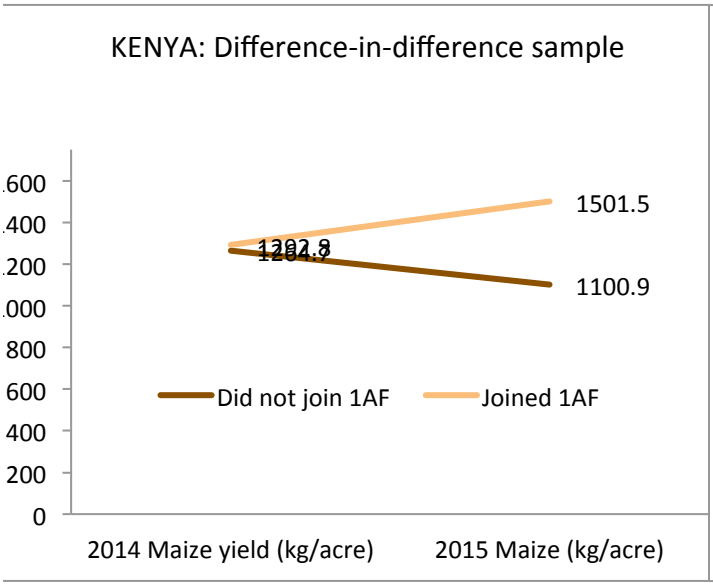

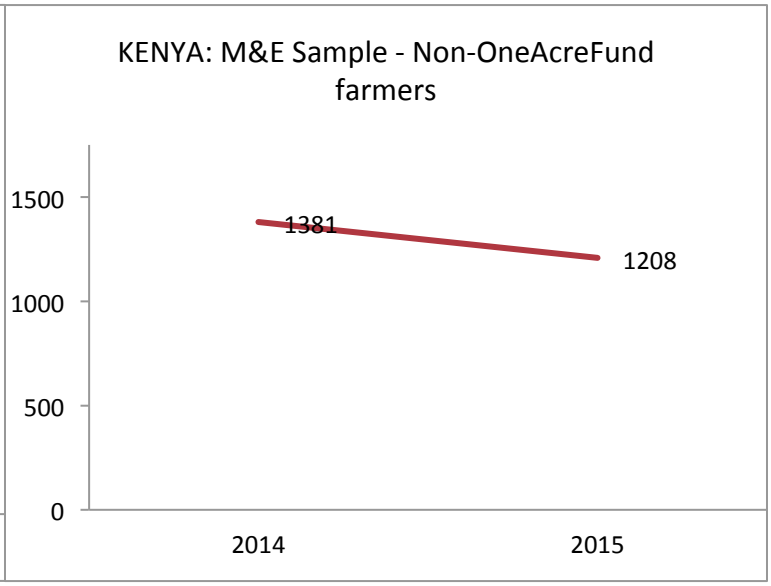

In Burundi[1] between 2014 and 2015:

- The difference-in-difference controls stayed **the same**
- The M&E controls increased by **4 percent**

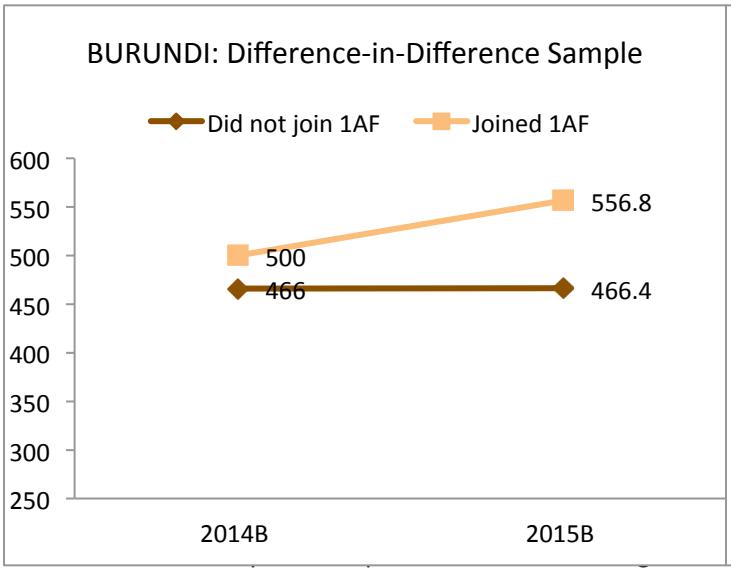

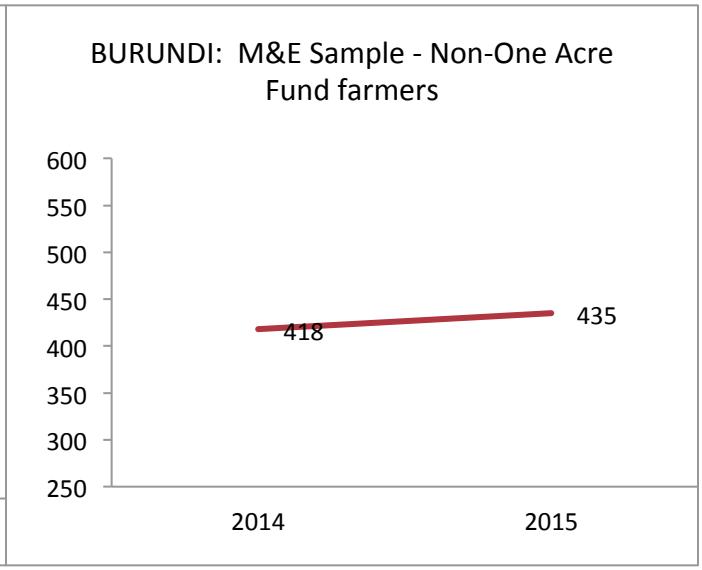

Kenya 2016 (Diff-in-Diff-in-Diff)

To further check the results from 2014-2015, in 2015-2016 we did a "diff-in-diff-in-diff" following three groups of farmers: (1) farmers who stayed out of the program in both time periods, (2) farmers who stayed in the program over both time periods and (3) farmers who were not program farmers in 2015 but became program farmers in 2016. This allowed us 2 comparison groups (the first and second) against which to compare the trend for farmers who enter the program.

Sample: In the 2015 Long Rains season, One Acre Fund measured harvests for control farmers for Kenya's annual Impact Assessment activity. Hundreds of farmers from that group were followed through to 2016, and we collected data regarding whether or not they joined One Acre Fund in 2016 and once again measured their yield outcomes. We also collected data from farmers who remained in the program over this period. In total, we have harvest measurements for 71 farmers that joined One Acre Fund in 2016, 256 farmers who stayed in One Acre Fund, and 207 farmers who remained controls. Using these data, we can compare farmers who joined One Acre Fund with both those who remained controls and those who remained in the program.

Results:

Cohort	2015 Maize yield (kg/acre)	2016 Maize (kg/acre)	Sample	Diff	p-value
Did not join 1AF	1246	983	207	262.3	
Joined 1AF	1090	1287	71	-197.0	
Stay OAF	1583	1334	256	248.5	
Diff-in-Diff Did Not Join Versus Joined 1AF				***459.3***	***p<.01***
Diff-in-Diff Stayed 1AF Versus Joined 1AF				***445.5***	***p<.01***

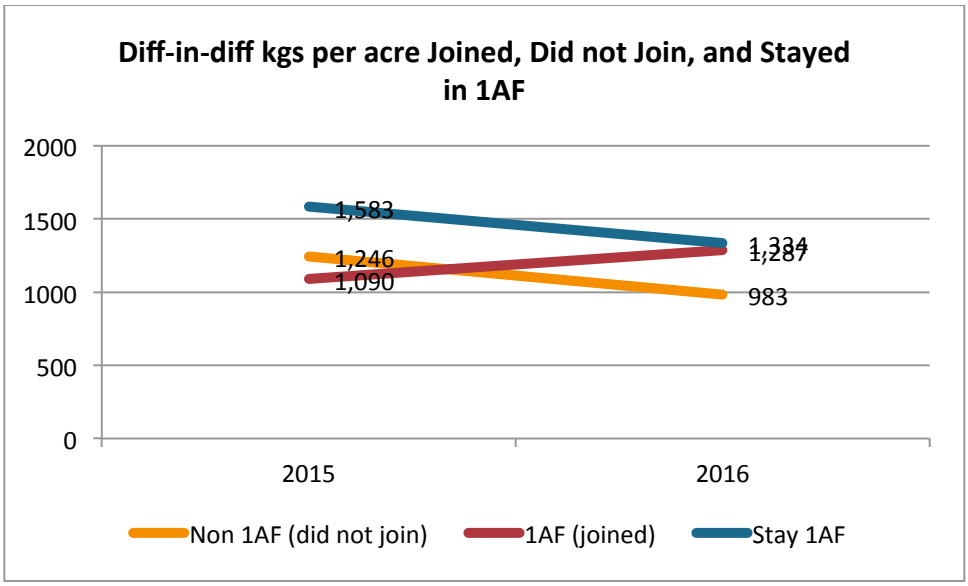

We find that the average yields for those who do not join 1AF went down over the two years by an average of 262 kg/acre. Those who did join our program over this time saw an increase in harvests of 197 kg/acre, for a diff-in-diff estimate of 459 kg/acre, which is statistically significant. Farmers who joined One Acre Fund were thus somewhat behind other farmers initially but increased their yields above farmers who remained controls.

Those who remained in One Acre Fund experience a loss in average maize yields over this period, which is similar in magnitude to that experienced by control farmers (249 kg/acre compared to 262 kg/acre). Comparing the difference experienced by those who joined One Acre Fund versus those who remained in the program, gives us a statistically significant increase in yields of 446 kg/acre for those who joined the program ($p<.01$).

Most farmers' harvests decreased in the study area, largely owing to the drought. However, for those farmers who joined One Acre Fund, their average yields increased over the same time period. In fact, the average yields for those who joined One Acre Fund nearly converges with those who remained in the program.

OLS Regression estimates. We ran additional checks to see whether the differences between control farmers and those who joined One Acre Fund could be due to some other underlying differences. First, we ran an OLS regression of maize yields per acre on whether the farmer joined One Acre Fund in 2016, and control for gender, age, household size, education level, marital status, land ownership, and province. After controlling for all of these factors, we found a larger and statistically significant effect of **515 kg/acre ($p<.01$).** Results are shown in the Stata output below.

diff_maize_pa	Coef.	Robust Std. Err.	t	P>|t|	[95% Conf.	Interval]
new_client_2016	514.7559	142.1888	3.62	0.002	216.0283	813.4836
gender	-253.2315	114.3642	-2.21	0.040	-493.5017	-12.96127
age	6.753564	3.063343	2.20	0.041	.3177189	13.18941
household_number	48.17905	23.29267	2.07	0.053	-.7570373	97.11514
agrictotal_acres_owned	-54.45942	27.17884	-2.00	0.060	-111.56	2.641198
edu_diploma	58.50163	307.2042	0.19	0.851	-586.9104	703.9137
edu_formal	-157.6356	186.6903	-0.84	0.410	-549.8573	234.5861
edu_none	49.29249	89.55361	0.55	0.589	-138.8527	237.4376
polygamous	-185.2343	186.0733	-1.00	0.333	-576.1598	205.6912
married	18.46552	151.5405	0.12	0.904	-299.9093	336.8404
single	-562.9723	213.403	-2.64	0.017	-1011.315	-114.6293
province_nyanza	41.65615	242.3374	0.17	0.865	-467.4758	550.788
_cons	-594.3759	253.5754	-2.34	0.031	-1127.118	-61.63378

We also ran OLS regressions among the set of farmers who joined One Acre Fund and those who remained in the program. We again found that even controlling for all the factors as above, in a small sample, the coefficient on joining One Acre Fund shows a positive average increase in yields of **464 kg/acre (p<.01)** compared to those farmers who remained in One Acre Fund.

diff_maize_pa	Coef.	Robust Std. Err.	t	P>|t|	[95% Conf.	Interval]
new_client_2016	464.3697	123.3422	3.76	.001	205.2375	723.502
gender	-51.89444	121.3926	-0.43	0.674	-306.9308	203.1419
age	5.375398	4.120532	1.30	0.208	-3.281519	14.03232
household_number	-5.515156	24.68024	-0.22	0.826	-57.36641	46.3361
agrictotal_acres_owned	-7.429119	21.4324	-0.35	0.733	-52.45692	37.59868
edu_diploma	-307.0869	288.2801	-1.07	0.301	-912.741	298.5672
edu_formal	-115.1689	224.8685	-0.51	0.615	-587.6001	357.2623
edu_none	177.9539	98.40091	1.81	0.087	-28.77874	384.6866
polygamous	-431.7034	252.602	-1.71	0.105	-962.4006	98.99372
married	-44.28072	134.9061	-0.33	0.747	-327.7078	239.1464
single	-9.670462	354.3598	-0.03	0.979	-754.1528	734.8119
province_nyanza	357.2395	221.293	1.61	0.124	-107.6799	822.1589
_cons	-533.0663	260.4428	-2.05	0.056	-1080.236	14.10378

Robustness checks. We also tried using propensity score matching (PSM) with difference-in-difference. We did this for robustness, as we are concerned about possible omitted variables in the regression analysis. PSM allows us to compare the effect of program participation among a group of comparable controls farmers with similar characteristics to treated farmers. Using one-to-one nearest neighbor matching, we again found only an increase of 558 kg/acre, which is statistically significant ($p<.01$). However, the sample size is very small with this analysis. Using the four nearest neighbors, we found an increase of 472 kg/acre ($p<.01$). Given the small sample size with PSM, we think the OLS model results are more reliable.

Finally, we also ran a PSM analysis matching newly enrolled farmers to similar farmers who remained in One Acre Fund, and matched along the same characteristics as above. Using one-to-one matching, we found an impact of 311 kg/acre ($p<.10$) and using one-to-four matching we found an impact of 386 kg/acre ($p<.05$). Given the small sample and the discrepancy in the estimates across the models, we cannot be sure of the exact magnitude of the actual effect. Nevertheless, there does appear to be evidence for an increase in maize yields for newly enrolled farmers, showing that newly enrolled farmers are converging with other One Acre Fund farmers maize yields, whereas those who remain out of the program are lagging behind.

Limitations

A common criticism of the difference-in-difference approach is that those participants who had entered the program were already on an upward trajectory, and may have realized some of the increase relative to controls *even without* program participation. This is the known "parallel trends" assumption, which is important to satisfy for a rigorous difference-in-difference analysis. It means that we have to assume that both treatment and control samples were headed on parallel change trajectories even if they started at different starting points.

In reality, historical trends for a study sample are often very difficult to come by, as it requires having followed the sample for several seasons before the program was even available in their area. We do acknowledge that the parallel trends hypothesis remains untested.

Finally, the sample sizes available to us were relatively small and limited in geographic scope. It is possible that our difference-in-difference results might not be representative of our program country-wide.

However, we feel that because we found results quite similar to our regular M&E in so many different country contexts, that this support is, on balance, a strong confirmation of our regular M&E impact assessment. We plan to continue to periodically conduct difference-in-difference analyses to check the validity of our regular M&E impact estimates.