Myanmar's LIFT project and earnings: a quantitative analysis

Hans R.A. Koster & Peter Das

ZOA
from relief to recovery

WWW.ZOA.NL and WWW.URBANECONOMICS.NL

Version 1.0, February 2023

Contents

1. Introduction *(Executive summary)*

The Livelihoods and Food Security Trust Fund (LIFT) implemented projects and activities that support pro-poor policy development and took place between 2016 and 2020 in Myanmar. It covered three sub-townships of Thandaunggyi Township in Northern Kayin State and was in 2019 extended with a few more villages. The project targeted over 5,000 small-holder farmer households. This target area comprises a mixture of displaced persons, incumbent households and returnees with a focus on smallholder farmers as direct beneficiaries (LIFT Uplands Programme Project 2017).

The project helped the rural poor in both areas to "step-up" and aimed to improve their position in the value chain, to improve market access and sustainable access to credit and other financial and material inputs. One of the main aims of the project was to improve earnings of the beneficiaries of the programme as to free them for the vicious cycle of increasing debt. Next to improving earnings, the project has focused on reducing malnutrition, as a lack of availability and access to adequate drinking water and nutritious food was considered to be a serious issue.

In order to improve living standards, several activities referred to as *outputs* have been rolled out. These include supervision and training of mothers to improve knowledge on nutrition (*output 1*), home gardening training and inputs (*output 2*), water and sanitation trainings (*output 3*), trainings on agricultural methods (*output 4*), the provision of agricultural inputs (*output 5*), improvements to irrigation infrastructure (*output 6*), the joint construction of motorcycle paths to improve accessibility (*output 7*), as well as loan-and-savings trainings (*output 8*). The main idea is that a consistent set of interlinked activities, trainings, and support, should improve the living conditions of the beneficiaries.

Although the premise of the project is clear, only exploratory analyses have shown that the LIFT project was effective in improving earnings and the nutrition of young children. However, much is to be learned here as the evidence is circumstantial, mostly anecdotal, or based on qualitative methods. This report aims to fill this gap by quantitatively evaluating the impact of the various activities undertaken under de umbrella of the LIFT project. We focus on supposedly the single most important outcome – earnings (or income). In several waves of surveys undertaken to monitor the project, households were asked to report their annual income. Using multivariate regression techniques, we investigate to what extent households that have participated in more LIFT-related activities witness higher earnings after project participation.

We emphasise that the evaluation is not based on an experimental setting as advocated by, among others, Banerjee & Munshi (2004), Banerjee & Duflo (2011) and Banerjee et al. (2018). An experimental setting with a treatment group and control group is considered to be ideal from the viewpoint of identifying causal effects, because the selection effect into participation is fully addressed. For example, particularly richer and more able households may participate in the several activities undertaken by the LIFT project. In such a situation one does not measure a causal effect of the programme but instead the sorting of richer/able households into the project. Unfortunately, randomised experiments are usually very costly and hard to implement in practice.

Still, we think that we come as close as reasonably possible to an experimental setting by adopting a so-called *difference-in-differences* set up (see Angrist & Pischke 2008). This implies that we compare *changes* in household earnings and compare them to the *intensity* of participation. Hence, we expect that households that participate more intensely in the various LIFT activities will witness a larger increase in earnings.

A downside of standard quantitative methods, including randomised experiments and difference-in-differences techniques, is that only an *average* treatment effect is identified. However, we would hope to see the largest incremental changes in the initially poor households, while households performing already reasonably well at the start of the programme are supposed to benefit less from the various trainings and activities offered. To test this hypothesis, we use an innovative new method referred to as *Unconditional Quantile Regressions* (see Firpo et al. 2009). This technique enables us to estimate treatment effects at various points in the earnings *distribution*. For example, we may investigate whether the effects are larger in the lower end of the earnings distribution (*i.e.* for the low-income households).

Ideally, we would also have investigated in detail the effects of the different outputs, instead of only analysing the aggregate impacts of the LIFT activities. However, it appears that our sample, consisting of about 500 households, is too small to obtain enough statistical power to identify those effects. Indeed, LIFT Uplands Programme Project (2017) understandably argued for a relatively low number of households in the surveys to save costs. Future projects should probably increase the sample size if more detail on the exact workings of the investigated project is warranted.

Our results show that LIFT activities have increased annual earnings. The results indicate that, on average, having participated in one of the programme's activities generated a 3-4% higher income. This estimate is robust to various methodologies, including a cross-sectional approach with household control variables and a difference-in-differences approach. Then, turning to our heterogeneous estimates, we show that particularly low-income households have benefited from LIFT activities. The effect for poor households who are in the lowest 10% of the earnings-distribution is about10%, while we do not find statistically significant positive earnings effects for the 50% richest households in the sample. Hence, the LIFT programme seems to have contributed not only to increases in earnings, but also to reductions in earnings inequality.

This report proceeds as follows. In Section 2 we outline the LIFT project and the intended activities/outputs. Section 3 discusses the preparation and cleaning of the data used for the analysis. We also provide some initial descriptive statistics for the studied sample. Section 4 outlines the methodology, which is followed by the results in Section 5. Section 6 concludes.

2. Project background and project aims

From 2016-2019 ZOA with its partners implemented the programme titled '*Improved economic and nutritional outcome of poor rural people in Myanmar*'. It was a project funded by the Livelihoods and Food Security Fund (LIFT) and covered three sub-townships of Thandaunggyi Township in Northern Kayin State (see Figure 2.1). The project targeted over 5,000 smallholder farmer households (HHs) with commercial potential in Thandaunggyi Township. This area emerged from conflict and consists of a mixture of internally displaced persons (IDPs), IDP returnees and incumbent households.

The overall purpose of the project was formulated as "improved economic status and nutritional outcomes for poor rural people in Myanmar with increased income and stable access to food for vulnerable households". This was in line with the purpose that was articulated by LIFT's Uplands Programme.

To achieve this high level outcome and in consideration of the needs of the target groups, the project focused on *(i)* farm advisory services and Producer Groups, *(ii)* nutrition, and *(iii)* Social protection and access to collective / public services. Sustainable natural resource management (NRM) and Gender were so-called cross-cutting issues, meaning that they were integrated into all activities of the programme.

The direct aim of the project was to help the rural poor in the targeted areas to 'step-up' and improve their position in the value chain (VC), to get access to markets and to credit and other inputs. Besides improving their income and getting them out of the circle of debt, the project included nutrition and WASH activities as a lack of availability and access to adequate drinking water and nutritious food was identified.

Improved economic and nutritional outcome of poor rural people in Myanmar was implemented in the northern township of Kayin State, Thandaunggyi Township (see map below), and targets a

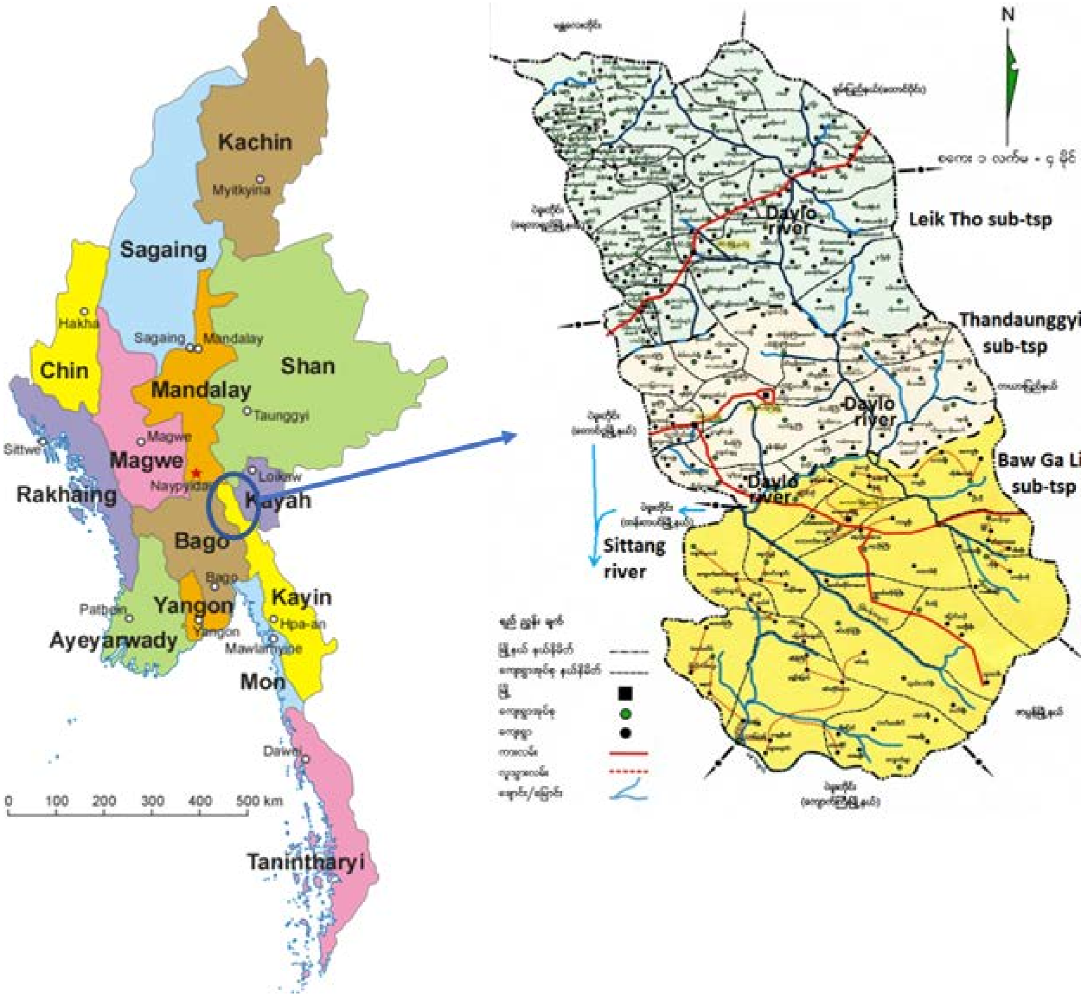

FIGURE 2.1 – MAP OF THE TARGETED AREA

total of 100 villages. Kayin State is located in the southeast of Myanmar and is bordered by the Mandalay Region and Shan State to the north, Kayah State to the northeast, Mon State and Bago Region to the West, and Thailand to the East.

Thandaunggyi Township consists of three sub-townships: Leiktho sub-township, Thandaunggyi sub-township and Bawgali sub-township. There are two types of villages: Core-villages and Value Chain-only villages (VC). Core-villages are the villages targeted with all project components. There are 40 core-villages. VC-only villages are only targeted with the activities described under output 4 and 5. There are 60 VC-only villages.

3. Data

3.1 Data preparation

We use several surveys undertaken to evaluate the LIFT project. The first is the baseline survey with 373 unique households from LIFT Uplands Programme Project (2017), which was undertaken before the LIFT project started in the final quarter of 2016. The second survey captures another 250 households in early 2019. This second baseline survey was undertaken because another set of villages was added to the 40 core villages in the sample (see LIFT Uplands Programme Project 2019). Then we have two 'endline' surveys after the project finished capturing 393 households for the initial villages and 250 for the villages added in 2019.

For each of these surveys, we keep the variables that are consistently measured across the different waves in our data. The key *outcome* variable of interest is annual earnings, which is measured in Myanmarese Kyat (*i.e.* Ks.10,000 is about $4.80). One may be worried that earnings are measured with error, because earnings are self reported. Fortunately, as earnings is our dependent variable, with random measurement error, the estimated effects are not impacted (see Koster & Van Ommeren 2020). We think it is reasonable to assume random measurement error because people are unlikely to make systematic errors in reporting their earnings.

To construct our main *treatment* variable, we count the type of activities a household participates in in each study period, which is half a year. For example, a household may have participated in mother groups en in have received trainings on agricultural methods in the first half of 2017. Then, the treatment variable is equal to 2. Alternatively, we consider to count the frequency of activities. For example, a household may have participated 5 times in mother groups, and has followed 3 activities related to business groups. Then, this alternative treatment variable equals 8. We refer to the form variable as *count of activities participated*, while to the latter variable as *frequency of activities participated*.

Further, we have information on the village where the household lives, which enables us to later control for trends in the productivity of certain villages. From the surveys, we also obtain information on the ethnic group where the head of the household belongs to, as well as the reported religion, whether the respondent is married, owns land, and is a farmer. Using a household identifier we can trace households over time to see how earnings have developed. Because we think it is unlikely that ethnicity or religion changes, we take the ethnicity and religion from the baseline survey for each household.

3.2 Descriptive statistics

In this subsection we further illustrate the characteristics of the data. In Table 3.1 we show what we call descriptive statistics for the dependent variable and the control variables. We show represent the mean, standard deviation (*i.e.* the spread), minimum and maximum values per variable. This is useful to see if there are any outlier values. The average annual earnings in the first baseline survey from 2016 are Ks. 909 thousand, which is about \$ 430, which is only \$ 1.18 per day. This clearly indicates that households in this survey and participating in the support programme are very poor. The spread, however, is substantial. The household with the highest earnings is about \$ 4 thousand per year. The average earnings considerably increased over the years. In the final survey, average earnings have increased by almost 50% to Ks. 1.3 million (approximately \$ 620)

Looking at the control variables, most participants are married (about 95%). Almost always the head of the household is male (also about 95%). Further, almost all households own land that they use for agricultural activities. Further, most households are Christian, with Baptists and Catholics being the largest groups. However, please note that the share of Baptist households is considerably larger in the first baseline survey, as compared to the second baseline survey. The most dominant ethnic group is Keba Karen (respectively 36% and 77% in the first and second baseline survey).

In Table 3.2 we show the descriptives of the participation in various activities (*i.e.* outputs). For completeness, we show also the first baseline. Obviously, as the first baseline survey took place before any activities where launched, all variables equal zero. As the second baseline survey took place early 2019, households already participated in various activities. As not all activities were rolled out in the villages that were part of the second baseline, the outputs 1, 2 and 3 are equal to zero.

We think Panel C in Table 3.2 is the most interesting. On average, households at the end of the sample period participated in almost 8 *different* activities, which are counted each half a year. Despite the second baseline group not participating, the first output, providing knowledge on nutrition to mothers, has the highest participation across households. There are households that have participated a lot in various activities, as the maximum is 25. Other activities that have high participation rates are output 4 (training on agricultural methods) and output 5 (provision of agricultural inputs). We also report the frequency of participating in activities. One household has participated in a stunning number of 109 activities. We see that the frequency of activities is dominated by output 1 (knowledge on nutrition) and output 2 (home gardening and inputs).

Next, we will investigate in Figure 3.1 whether the main variables of interest are normally

TABLE 3.1 – DESCRIPTIVE STATISTICS: DEPENDENT VARIABLE AND CONTROLS

PANEL A: Baseline 1	(1) mean	(2) sd	(3) min	(4) max
Annual earnings *(Ks.)*	909,070	989,968	16,000	8,500,000
Head of household is male	0.974	0.159	0	1
Married	0.971	0.167	0	1
Household owns land	0.963	0.190	0	1
Occupation – Agriculture	0.966	0.183	0	1
Religion of head of household – Baptist	0.644	0.480	0	1
Religion of head of household – Catholic	0.218	0.414	0	1
Religion of head of household – Anglican	0.109	0.312	0	1
Religion of head of household – Nature worship	0.0201	0.141	0	1
Ethnicity of head of household – Bwe Karen	0.276	0.448	0	1
Ethnicity of head of household – Paku Karen	0.227	0.420	0	1
Ethnicity of head of household – Maw Nay Bwa Karen	0.0201	0.141	0	1
Ethnicity of head of household – Keba Karen	0.359	0.480	0	1
Ethnicity of head of household – Kayan Keko Karen	0.0431	0.203	0	1
Ethnicity of head of household – Other	0.0460	0.210	0	1
PANEL B: Baseline 2	(1) mean	(2) sd	(3) min	(4) max
Annual earnings *(Ks.)*	969,385	578,891	202,500	2,375,000
Head of household is male	0.935	0.246	0	1
Married	0.948	0.223	0	1
Household owns land	0.984	0.126	0	1
Occupation – Agriculture	0.980	0.142	0	1
Religion of head of household – Baptist	0.298	0.458	0	1
Religion of head of household – Catholic	0.500	0.501	0	1
Religion of head of household – Anglican	0.0202	0.141	0	1
Religion of head of household – Nature worship	0.173	0.379	0	1
Ethnicity of head of household – Bwe Karen	0.00403	0.0635	0	1
Ethnicity of head of household – Paku Karen	0.0242	0.154	0	1
Ethnicity of head of household – Maw Nay Bwa Karen	0.177	0.383	0	1
Ethnicity of head of household – Keba Karen	0.766	0.424	0	1
Ethnicity of head of household – Kayan Keko Karen	0.00403	0.0635	0	1
Ethnicity of head of household – Other	0.0242	0.154	0	1
PANEL C: Endline	(1) mean	(2) sd	(3) min	(4) max
Annual earnings *(Ks.)*	1,341,215	1,076,640	100,000	8,840,000
Head of household is male	0.948	0.223	0	1
Married	0.960	0.195	0	1
Household owns land	0.992	0.0887	0	1
Occupation – Agriculture	0.775	0.418	0	1
Religion of head of household – Baptist	0.509	0.500	0	1
Religion of head of household – Catholic	0.334	0.472	0	1
Religion of head of household – Anglican	0.0712	0.257	0	1
Religion of head of household – Nature worship	0.0791	0.270	0	1
Ethnicity of head of household – Bwe Karen	0.161	0.368	0	1
Ethnicity of head of household – Paku Karen	0.157	0.364	0	1
Ethnicity of head of household – Maw Nay Bwa Karen	0.0807	0.273	0	1
Ethnicity of head of household – Keba Karen	0.525	0.500	0	1
Ethnicity of head of household – Kayan Keko Karen	0.0364	0.187	0	1
Ethnicity of head of household – Other	0.0237	0.152	0	1

Notes: The number of observations in de first baseline is 348, it is 248 in the second baseline and 623 in the endline survey.

TABLE 3.2 – DESCRIPTIVE STATISTICS: TREATMENT VARIABLES

PANEL A: Baseline 1	(1) mean	(2) sd	(3) min	(4) max
Count of activities participated	0	0	0	0
Output 1 – participated	0	0	0	0
Output 2 – participated	0	0	0	0
Output 3 – participated	0	0	0	0
Output 4 – participated	0	0	0	0
Output 5 – participated	0	0	0	0
Output 6 – participated	0	0	0	0
Output 7 – participated	0	0	0	0
Output 8 – participated	0	0	0	0
Frequency of activities participated	0	0	0	0
Output 1 – frequency	0	0	0	0
Output 2 – frequency	0	0	0	0
Output 3 – frequency	0	0	0	0
Output 4 – frequency	0	0	0	0
Output 5 – frequency	0	0	0	0
Output 6 – frequency	0	0	0	0
Output 7 – frequency	0	0	0	0
Output 8 – frequency	0	0	0	0

PANEL B: Baseline 2	(1) mean	(2) sd	(3) min	(4) max
Count of activities participated	3.242	1.940	0	10
Output 1 – participated	0	0	0	0
Output 2 – participated	0	0	0	0
Output 3 – participated	0	0	0	0
Output 4 – participated	1.298	0.834	0	4
Output 5 – participated	1.004	0.879	0	4
Output 6 – participated	0.214	0.411	0	1
Output 7 – participated	0.290	0.464	0	2
Output 8 – participated	0.435	0.777	0	3
Frequency of activities participated	4.831	4.460	0	23
Output 1 – frequency	0	0	0	0
Output 2 – frequency	0	0	0	0
Output 3 – frequency	0	0	0	0
Output 4 – frequency	1.710	1.848	0	16
Output 5 – frequency	1.020	0.924	0	5
Output 6 – frequency	0.238	0.480	0	2
Output 7 – frequency	0.290	0.464	0	2
Output 8 – frequency	1.573	3.080	0	13

PANEL C: Endline	(1) mean	(2) sd	(3) min	(4) max
Count of activities participated	7.931	4.244	0	25
Output 1 – participated	2.756	1.726	0	7
Output 2 – participated	0.927	1.363	0	3
Output 3 – participated	0.324	0.584	0	3
Output 4 – participated	1.308	0.912	0	4
Output 5 – participated	1.216	1.082	0	8
Output 6 – participated	0.302	0.574	0	3
Output 7 – participated	0.259	0.456	0	2
Output 8 – participated	0.840	0.991	0	3
Frequency of activities participated	23.39	20.79	0	109
Output 1 – frequency	9.468	8.898	0	36
Output 2 – frequency	4.398	7.862	0	32
Output 3 – frequency	0.327	0.604	0	5
Output 4 – frequency	2.797	2.951	0	20
Output 5 – frequency	1.268	1.235	0	10
Output 6 – frequency	0.375	0.776	0	6
Output 7 – frequency	0.259	0.456	0	2
Output 8 – frequency	4.497	7.269	0	26

Notes: The number of observations in de first baseline is 348, it is 248 in the second baseline and 623 in the endline survey.

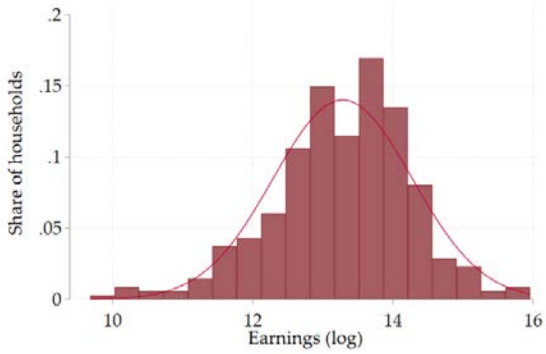

(A) EARNINGS, FIRST BASELINE

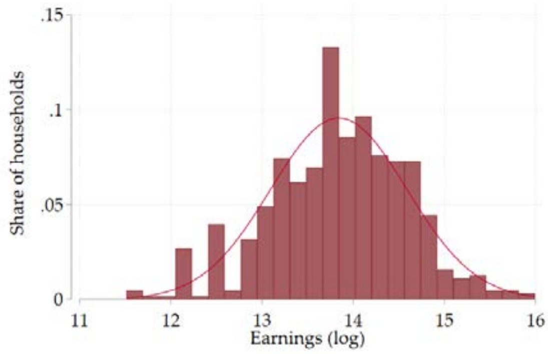

(B) EARNINGS, ENDLINE

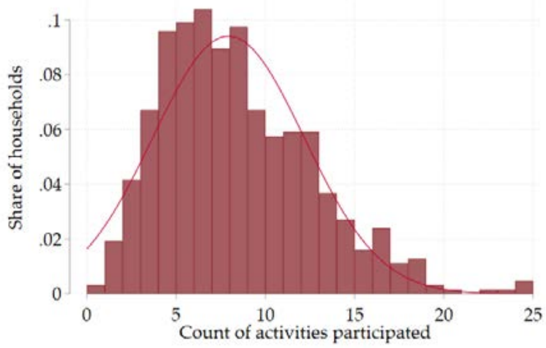

(C) COUNT OF ACTIVITIES, ENDLINE

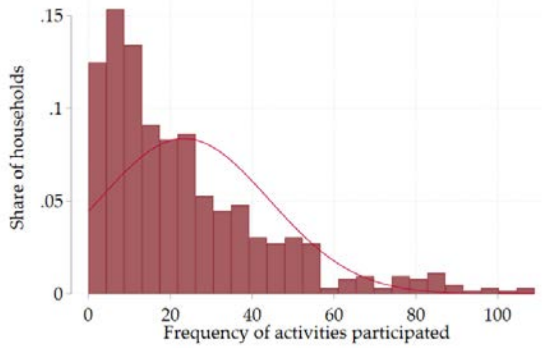

(D) FREQUENCY OF ACTIVITIES, ENDLINE

Notes: The red line indicates a normal distribution.

FIGURE 3.1 – HISTOGRAMS

distributed. In Figure 3.1a we show the distribution of earnings in the first baseline survey. Because earnings are likely a so-called skewed distribution, we take the logarithm of earnings (see Koster & Van Ommeren 2020). It is shown that the distribution is more or less log-normally distributed.

Figure 3.1b shows the distribution of earnings for the endline survey, which is again more or less log-normally distributed. Please note that the distribution shifted to the right compared to the baseline earnings distribution, which is in line with the strong average increase in earnings between the baseline and endline survey of about 50%.

In Figure 3.1c we show the count of activities participated. It is more or less normally distributed, apart from a few outliers beyond 20 activities. Please note that we cannot take the log here, because the count of activities can be zero (which is particularly true for the first baseline survey) and one cannot take the logarithm of zero.

Finally, the frequency of activities (see Figure 3.1d) is strongly skewed. We therefore prefer to focus on the count of activities as the main treatment variable to avoid the issue that outliers have a disproportionate impact on the results. Still, we will provide ancillary analyses where we analyse the impact of the frequency of activities on earnings.

4. Methodology

4.1 A multivariate regression approach

A standard way to investigate the impact of a treatment or independent variable on an outcome or dependent variable is to estimate linear regressions.[1] Usually, the relationship between treatment variable and outcome variables is described by a simple equation of the following form:

$$\log y_{ivt} = \alpha + \beta c_{ivt} + \varepsilon_{ivt}, \tag{4.1}$$

where y_{ivt} are the annual earnings of household i living in village v in time t and c_{ivt} is the count of activities that the households participated in so far. The latter corresponds to either the first baseline survey, the second baseline survey or the endline survey. Further, α and β are regression parameters to be estimated, while ε_{ivt} is the residual, or the part of earnings that we cannot explain by c_{ivt}.

Please note that β here is the key parameter of interest and depicts what happens to earnings, in percentage terms, if the count of activities increases by 1. Say, for example, that $\beta = 0.01$, then for each activity the household participated in, the annual earnings increase by 1%.

However, the above model is likely a too simplistic description of reality as not only the count of activities has an impact on earnings, but also the ethnicity, religion, and marital status of the household, as well as the village where the household lives determines income. If for example households with a certain ethnicity or religion have higher earnings *and* participate more in activities, we may falsely attribute the impact of ethnicity and religion to the treatment variable.

[1] In linear regressions, the relationship between the treatment variable(s) and outcome variable are modelled using linear predictor functions whose unknown model parameters are estimated from the data. It appears that the best linear unbiased estimator of the coefficient of interest – so the impact of the treatment variable on the outcome variable – is obtained by minimising the squared residuals.

To address this issue it is important to *control* for household and location characteristics that may also determine the earnings of the household. An extended equation then looks as follows:

$$\log y_{ivt} = \alpha + \beta c_{ivt} + \gamma x_{ivt} + \delta_v + \varepsilon_{ivt}, \tag{4.2}$$

where x_{ivt} are control variables, such as ethnicity and religion, and γ is a set of parameters capturing the impacts of these variables on earnings. We also include so-called village *fixed effects*, denoted by δ_v. This implies that we include a dummy variable for each village as to control for factors influencing households' earnings at the village level that are the same for everyone. For example, some villages may have better access to fertile grounds, which will lead to higher yields and in turn higher earnings. By including these village fixed effects we control for all those factors.

4.2 Addressing the selection effect

A major concern with the above approach is the presence of a potential selection effect Angrist & Krueger (2001), Angrist & Pischke (2008). This selection effect entails that households that have a higher earnings potential are more likely to participate in (many) activities offered. Hence, in this way, we do not measure the causal effect of the intervention, but instead capture the fact that households that would have experienced higher earnings growth absent of the programme participate more intensely in the activities offered.

A common solution is to apply randomisation, meaning that the treatment (*i.e.*, the activities offered) is randomised across households. Randomisation obviously addresses the selection effect because household could not self-select into the programme Angrist & Pischke (2008).

Unfortunately, we cannot rely on randomisation in the current setting because participation in the activities was voluntary. Instead, we use a version of an approach that is often used in applied economics, which is referred to as the *difference-in-differences* (DID) approach (Bertrand et al. 2004, Angrist & Pischke 2014). We provide an example of this method in Figure 4.1. The idea is the following: one should have observations of an outcome variable (in our case: earnings) before and after treatment. Assume that a certain group of households does not receive treatment. Absent of the treatment, those people witness an increase in earnings of B. Then there are the households that receive treatment. In Figure 4.1 it can be seen that those households initially have higher earnings (*i.e.*, have a higher baseline earnings level). However, this is not an issue because we look at the change in earnings over time, which is equal to A. What is then the causal effect of the treatment? Well, that is $A - B$, because absent of the treatment the treatment group would have experienced an increase in earnings equal to B so the *additional* effect of the treatment effect is equal to $A - B$.

In our setting we only have very few households that did not participate in any of the activities. However, we can use the *intensity* of treatment to compare changes in earnings. Hence, if we compare the differences in the trends households that have participated in 2 versus 1 activities, we can trace the effect of attending one extra activity.

Formally, the application of this DID approach is straightforward: we just include so-called

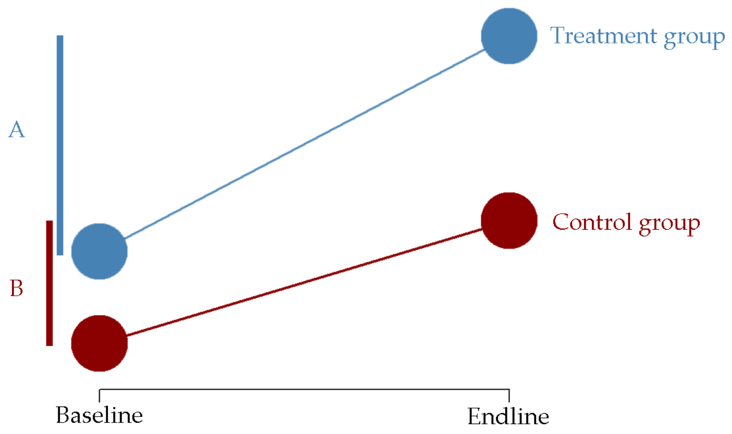

FIGURE 4.1 – DIFFERENCE-IN-DIFFERENCES: EXAMPLE

household *fixed effects*, which are essentially dummy variables for each household:

$$\log y_{ivt} = \alpha + \beta c_{ivt} + \gamma x_{ivt} + \zeta_i + \varepsilon_{ivt}, \tag{4.3}$$

where ζ_i captures the household fixed effects. The inclusion of household fixed effects ensure that we control for baseline differences in earnings across households and only use variation in the trends in earnings of time across households.

One still may be concerned that different villages may be on different trends, implying that, for example, due to changes in climatic conditions some villages may face lower growth in earnings. If these changes in climatic conditions are correlated to the treatment, our estimate β is still biased. To address this issue, in a final specification we include village-by-year fixed effects to absorb all trends in earnings at the village level:

$$\log y_{ivt} = \alpha + \beta c_{ivt} + \gamma x_{ivt} + \delta_{vt} + \zeta_i + \varepsilon_{ivt}, \tag{4.4}$$

where δ_{vt} captures village-by-year fixed effects.

4.3 Allowing for heterogeneity in the treatment effect

The difference-in-difference design is useful in identifying the *average* effect of the treatment. In many applications, however, one is interested in heterogeneity in the treatment effect across households. In our setting we are particularly interested whether households that are initially poor have experienced larger benefits from the programme than initially slightly richer households. A very useful recent innovation to disentangle these effects is what Firpo et al. (2009) refer to as *unconditional quantile regressions*. This method allows us to measure the treatment effect at each quantile of the earnings distribution. Quantiles are cut points dividing the range of the earnings distribution into continuous intervals with equal probabilities. Hence, a lower quantile means that a household has low earnings (*i.e.* is on the left side of Figures 3.1a or 3.1b), while a higher quantile means that a household is relatively richer (*i.e.* is on the right side of Figures 3.1a or 3.1b).[2]

It is generally convenient to apply unconditional quantile regressions because it is shown by Firpo et al. (2009) that this just entails a transformation of the dependent variable, $\log y_{ivt}$, by means of a so-called recentered influence function (RIF). We aim to estimate the following specification:

$$\text{RIF}\big(\log y_{ivt}; q_z, F_{\log y_{ivt}}\big) = \alpha_z + \beta_z c_{ivt} + \gamma_z x_{ivt} + \zeta_{i,z} + \varepsilon_{ivt,z}, \tag{4.5}$$

where $\text{RIF}(\,\cdot\,)$ is the RIF for a given quantile z of earnings and β_z captures the effect of interest for a given quantile z of the earnings distribution.

[2]Let us give an example. Say we have data on 100 households and we range these households with earnings from high to low. Then, the first observations is said to be the first quantile of the earnings distribution, the 50^{th} quantile is the middle observation, which is also called the median, while the 100^{th} quantile is the household with the highest earnings in our data.

5. Results

5.1 Baseline results

We now proceed to report the results. We report the main results in Table 5.1. In column (1) we estimate a very simple specification only with survey wave fixed effects to control for the overall positive trend in earnings over time. We find that participating in another additional activity in half a year implies an increase in earnings of 4.2%.

Column (2) further includes housing characteristics, such as dummies capturing ethnic groups and religion, whether the household is a farmer, owns land and is married. The coefficient regarding the count of activities participated is hardly affected by the inclusion of those controls.

TABLE 5.1 – REGRESSION RESULTS: EFFECTS OF TREATMENT

Dependent variable: the logarithm of earnings

	(1) OLS	(2) OLS	(3) OLS	(4) OLS
Count of activities participated	0.0420***	0.0428***	0.0405***	0.0324**
	(0.0064)	(0.0069)	(0.0140)	(0.0150)
Housing characteristics included		✓	✓	✓
Household fixed effects			✓	✓
Village fixed effects		✓	✓	✓
Village×survey wave fixed effects				✓
Survey wave fixed effects	✓	✓	✓	
Number of observations	1,222	1,219	810	808
R^2	0.1050	0.3101	0.7037	0.7517

Notes: Household characteristics include 6 ethnicity group dummies, 4 religion dummies, and dummy variables indicating whether the occupation is farming, whether the head of the household is male, whether they own land and whether they are married. Standard errors are clustered at the household level and in parentheses; *** $p < 0.01$, ** $p < 0.5$, * $p < 0.10$.

TABLE 5.2 – REGRESSION RESULTS: EFFECTS OF TREATMENT, FREQUENCY

Dependent variable: the logarithm of earnings

	(1) OLS	(2) OLS	(3) OLS	(4) OLS
Frequency of activities participated	0.0074***	0.0067***	0.0050*	0.0035
	(0.0013)	(0.0015)	(0.0028)	(0.0031)
Housing characteristics included		✓	✓	✓
Household fixed effects			✓	✓
Village fixed effects		✓	✓	✓
Village×survey wave fixed effects				✓
Survey wave fixed effects	✓	✓	✓	
Number of observations	1,222	1,219	810	808
R-squared	0.0975	0.3011	0.7005	0.7500

Notes: Household characteristics include 6 ethnicity group dummies, 4 religion dummies, and dummy variables indicating whether the occupation is farming, whether the head of the household is male, whether they own land and whether they are married. Standard errors are clustered at the household level and in parentheses; *** $p < 0.01$, ** $p < 0.5$, * $p < 0.10$.

In column (3) we apply our 'differences-in-differences' approach that controls for the selection of households in participating in the programme. Surprisingly, we find that the selection effect is really important, as the effect is not materially impacted by the inclusion of household fixed effects. Because the so-called 'degrees of freedom' are considerably lower, the standard error (*i.e.* a higher standard error indicates that the effect is less precisely estimated) is somewhat higher. Still, we find that the effect is statistically significant at the 1% level.

Finally, in column (4) in Table 5.1 we display the most comprehensive specification in which we control for all differences in earnings over time between villages. For example, one village may respond more favourably towards changes in seasonal weather conditions so that villagers have more time to join programme activities. We show that these village trends play some role as the coefficient is about 20% lower: participating in another activity increases earnings now by 3.2%. Still, the effect is highly statistically significant (at the 1% level).

One may argue that it is not the count of activities that matter but also the *frequency* or intensity of activities participated in. We have shown in Figure 3.1d that the distribution of frequency of activities is more skewed and yields more outliers. Moreover, it is questionable whether one should treat the frequency of participation in the same way between different outputs (*e.g.* participating in one mother group activity *(output 1)*, is treated in the same way as the provision of motorcycle paths *(output 7)*). In any case, as a sensitivity check, we present the results in Table 5.2.

We follow the same set-up as in the baseline results table. We observe statistically significant effects of attending an additional activity in column (1), where we only include survey-wave fixed effects. Attending an additional activity increases earnings by 0.7%. This effect reduces in size when we add household characteristics and village fixed effects in column (2) and household fixed effects in column (3). In the latter specification, the coefficient is only statistically significant at the 10% level. Note that the quantitative magnitude is somewhat comparable to the results where we use the count of activities, but because the frequency of activities is noisier we find somewhat lower estimates.[1] In column (4), in the specification with all the controls and fixed effects, we do not find a statistically significant effect, although the point estimate is still positive.

Further, we investigate in Table 5.3 the effects of different outputs. The issue is that we have too little power (*i.e.*, our sample is too small) to measure the effects of all outputs at the same time. We take an alternative approach where we count the total of activities participated minus each of the activities in each specification. For example, in column (1) we count the total of activities participated for outputs 2-8, and separately whether the household participated in output 1. In column (2) we count the total of activities participated for output 1, and 3-8, and separately whether the household participated in output 2, etc.

We do not find clear-cut results, which confirms that we lack power to make a decisive answer on what outputs yielded effective increases in earnings. Most of the standard errors are too large, but we find positive point estimates for supervision and training of mothers to improve knowledge on nutrition (*output 1*), water and sanitation trainings (*output 3*), improvements to irrigation infrastructure (*output 6*), the joint construction of motorcycle paths to improve accessibility (*output 7*), as well as loan-and-savings trainings (*output 8*). Especially the latter seems to have yielded large positive effects, which are also statistically significant at the 5% level, although the standard error is too large to make precise statements.

5.2 Allowing for heterogeneity in the treatment effect

In this subsection we aim to allow for heterogeneity in the effect of participating in the various activities offered. In order to do so we estimate unconditional quantile regressions, which implies that we estimate an effect for a given quantile of the earnings distribution (see equation (4.5)). A lower quantile means that households are poorer, while higher quantiles refer to richer households in the sample. In this way, we investigate whether the treatment has been more effective for poorer households, which are arguably the intended beneficiaries of the programme.

We report results for different quantiles in Figure 5.1. We repeat the same specifications as reported in column (3) in Table 5.1 so we control for the selection effect by applying the differences-in-differences approach. We observe a clear downward pattern in the positive effect of the treatment. We find that for the 10% poorest households, the effect of participating in an additional activity increases earnings by about 10%. This effect reduces essentially to zero for the 50% richest households.

Hence, Figure 5.1 seems to suggest that the programme was particularly effective in increasing earnings of the poor, while the richer household did not benefit or at least benefit less from the various activities. In this way, the LIFT programme seems to have contributed not only to increases in earnings, but also to reductions in earnings inequality.

[1] Say that we compare a *standard deviation* increase in the frequency of activities, earnings increase by $0.0050 \times 20.79 = 10.4\%$, while for the count of activities it is $0.0405 \times 4.24 = 17.5\%$

TABLE 5.3 – REGRESSION RESULTS: EFFECTS OF DIFFERENT OUTPUTS

Dependent variable: the logarithm of earnings

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) OLS	(7) OLS	(8) OLS
Count of activities participated – respective output	0.0405** (0.0180)	0.0546*** (0.0208)	0.0399*** (0.0143)	0.0440** (0.0180)	0.0481*** (0.0176)	0.0412*** (0.0159)	0.0399*** (0.0142)	0.0186 (0.0171)
Output 1 – participated	0.0405 (0.0482)							
Output 2 – participated		0.0095 (0.0351)						
Output 3 – participated			0.0549 (0.1060)					
Output 4 – participated				0.0161 (0.0824)				
Output 5 – participated					-0.0146 (0.0827)			
Output 6 – participated						0.0235 (0.1620)		
Output 7 – participated							0.1700 (0.2288)	
Output 8 – participated								0.1600** (0.0624)
Housing characteristics included	✓	✓	✓	✓	✓	✓	✓	✓
Household fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Village fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Survey wave fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Number of observations	810	810	810	810	810	810	810	810
R^2	0.7037	0.7046	0.7038	0.7038	0.7042	0.7038	0.7042	0.7072

Notes: Household characteristics include 6 ethnicity group dummies, 4 religion dummies, and dummy variables indicating whether the occupation is farming, whether the head of the household is male, whether they own land and whether they are married. Standard errors are clustered at the household level and in parentheses; *** $p < 0.01$, ** $p < 0.5$, * $p < 0.10$.

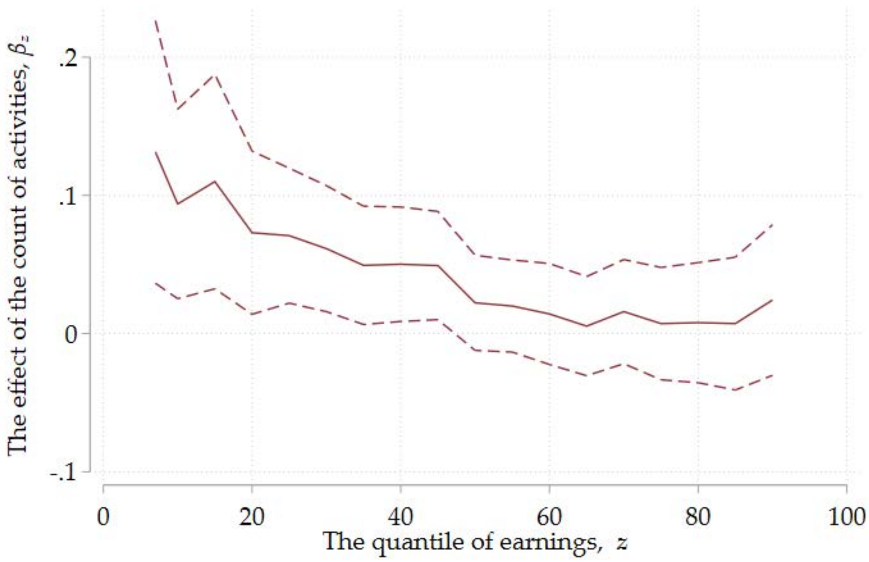

Notes: The solid red lines indicate the effect, β_z, while the dotted lines indicate 95% confidence bands. We include household characteristics, survey wave and village and household fixed effects. Household characteristics include 6 ethnicity group dummies, 4 religion dummies, and dummy variables indicating whether the occupation is farming, whether the head of the household is male, whether they own land and whether they are married. Standard errors are clustered at the household level and in parentheses

FIGURE 5.1 – HETEROGENEOUS EFFECTS

6. Conclusions

The Livelihoods and Food Security Trust Fund (LIFT) was implemented in Myanmar between 2016 and 2020 and targeted over 5,000 small-holder farmer households in three sub-townships of Thandaunggyi Township. The project aimed to improve the position of the rural poor in the value chain, increase market access and access to credit, and reduce malnutrition. Several activities such as supervision and training, home gardening, water and sanitation, agricultural methods, and loan-and-savings trainings were rolled out to improve living standards. Although only exploratory analyses have shown the effectiveness of the project in improving earnings and nutrition, this report aims to fill the gap by evaluating the impact of the various activities through multivariate regression techniques. The focus is on earnings as the single most important outcome, with households reporting their annual income in several waves of surveys.

Our findings reveal that the LIFT activities have resulted in a boost in yearly earnings. On average, participating in the program's activities led to a 3-4% increase in income, which was consistent across various analytical methods including a cross-sectional approach with household control variables and a difference-in-differences approach. More importantly, our analysis shows that the program has particularly benefited low-income households. The effect on poor households, in the bottom 10% of the earnings distribution, was approximately 10%, while no statistically significant positive earnings impact was observed for the top 50% richest households. As a result, it appears that the LIFT program has not only elevated earnings but also reduced earnings inequality.

Literature

Angrist, J. D. & Pischke, J. S. (2014), *Mastering' Metrics: The Path from Cause to Effect*, Princeton University Press, Princeton.

Angrist, J. & Krueger, A. (2001), 'Instrumental Variables and the Search for Identification: From Supply and Demand to Natural Experiments', *The Journal of Economic Perspectives* **15**(4), 69–85.

Angrist, J. & Pischke, J. (2008), *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, Princeton.

Banerjee, A., Barnhardt, S. & Duflo, E. (2018), 'Can Iron-Fortified Salt Control Anemia? Evidence from Two Experiments in Rural Bihar', *Journal of Development Economics* **133**, 127–146.

Banerjee, A. & Munshi, K. (2004), 'How Efficiently is Capital Allocated? Evidence from the Knitted Garment Industry in Tirupur', *Review of Economic Studies* **71**(1), 19–42.

Banerjee, A. V. & Duflo, E. (2011), *Poor economics: Rethinking poverty and the ways to end it*, Random House India.

Bertrand, M., Duflo, E. & Mullainathan, S. (2004), 'How Much Should We Trust Differences-In-Differences Estimates?', *The Quarterly Journal of Economics* **119**(1), 249–275.

Firpo, S., Fortin, N. M. & Lemieux, T. (2009), 'Unconditional Quantile Regressions', *Econometrica* **77**(3), 953–973.

Koster, H. R. A. & Van Ommeren, J. N. (2020), Identification and Estimation in Spatial Economics, Technical report, Amsterdam.

LIFT Uplands Programme Project (2017), Improving the Economic and Nutritional Outcomes of Rural Poor in Northern Kayin State: Baseline Study 2016, Technical report.

LIFT Uplands Programme Project (2019), Improving the Economic and Nutritional Outcomes of Rural Poor in Northern Kayin State: Baseline Study 2019 New Core Villages, Technical report.