

Can Microfinance Reach the Poorest: Evidence from a Community-Managed Microfinance Intervention

**Jonas Helth Lønborg and
Ole Dahl Rasmussen**

**Can Microfinance Reach the Poorest:
Evidence from a Community-Managed Microfinance Intervention**

Study Paper No. 55

Published by:

© The Rockwool Foundation Research Unit

Address:

The Rockwool Foundation Research Unit

Soelvgade 10, 2.tv.

DK-1307 Copenhagen K

Telephone	+45 33 34 48 00
Fax	+45 33 34 48 99
E-mail	forskningsenheden@rff.dk
Home page	www.rff.dk

ISBN 978-87-90199-90-6

ISSN 0908-3979

August 2013

Can Microfinance Reach the Poorest: Evidence from a Community-Managed Microfinance Intervention

JONAS HELTH LØNBORG[†] AND OLE DAHL RASMUSSEN[‡]

***Abstract:** Reaching the poorest is an important objective in many development interventions, and microfinance is no exception. We review performance indicators for effectiveness of targeting described in the literature and suggest a new metric in order to account for extent and severity of poverty as well as the income distribution among the poor. When applying this to a panel dataset from a community-managed microfinance intervention in Northern Malawi, we find regressive outreach: Participants are less poor than the general population in the area. In addition, we provide suggestions as to when and why the poor exit the project.*

Keywords: targeting, outreach, microfinance, Africa, Malawi, savings groups

JEL codes: O16, O13

Acknowledgements: We thank IKI for collecting the data, participants in the seminar at the Rockwool Foundation Research Unit, Nikolaj Malchow-Møller and Thomas Barnebeck Andersen for comments on earlier drafts, and DanChurchAid, the Rockwool Foundation Research Unit, and the Danish Agency for Science, Technology and Innovation for financial support. Any viewpoints and errors are our own.

[†]University of Southern Denmark, johe@sam.sdu.dk, Department of Business and Economics, University of Southern Denmark, Campusvej 55, DK-5230 Odense M, Denmark

[‡] DanChurchAid and University of Southern Denmark, odr@sam.sdu.dk, Department of Business and Economics, University of Southern Denmark, Campusvej 55, DK-5230 Odense M, Denmark

INTRODUCTION

Together with economic growth, poverty reduction is perhaps the most agreed-upon goal for development aid. Targeting interventions and resources toward the poorest is often viewed as a necessary means to reaching this goal. At the same time, however, an increasing number of development interventions require strong involvement from and capacity of participants. It is likely that this limits the outreach of the interventions. To guide the implementation and design of interventions we need more knowledge about participation in interventions which require a high degree of involvement by participants.

Microfinance, which includes the provision of loans, savings, and insurance, is a case in point. Targeting is widely used in microfinance as a means to deepen outreach by, for example, the Grameen Bank and BRAC (Bandiera *et al.* 2011), and there has been increasing focus on avoiding "mission drift", whereby programs include richer people (Christen 2001, Cull *et al.* 2007, Hermes *et al.* 2011) At the same time, however, a common opinion is that microfinance does not work for the poorest, a finding which is confirmed by early studies (Hulme 2000, Navajas *et al.* 2000). Microcredit, in particular, requires a high degree of involvement as well as prior skills from the loan takers. Borrowers need to be able to use a loan to create income, keep track of repayment schedules, and possess basic financial literacy.

In this paper we investigate to what extent it is possible to reach the poorest with a program that requires a high degree of involvement by participants. We analyze participation in the context of one of the most poverty-focused microfinance methodologies, namely community-managed microfinance, in a rural area in one of the world's poorest countries, Malawi.

To do this, we use a panel dataset from a household survey from 885 households in northern Malawi. The first round of data was collected just before the introduction of a large-scale community-managed microfinance project. Two years later, the households were revisited to gather information about their participation in the program.

We contribute to the literature in three ways: In addition to reporting the standard metrics of targeting effectiveness, or outreach, we develop our own metric, inspired by the squared poverty gap commonly used in poverty measurement. Unlike typical outreach metrics, the new measure is sensitive to changes in the depth of poverty as well as the income distribution among the poor. Furthermore, we analyze participation in a sequential framework, assessing the necessary steps preceding participation, including project awareness and interest in participation. We picture these steps as a leaking pipeline and examine where the pipeline is leaking. This framework has been used in developed countries but only sporadically in developing countries. Finally, we illustrate the use of the new metric and the sequential approach by examining outreach in community-managed microfinance, or more specifically savings groups, a type of microfinance intervention that is highly standardized and widely used (Allen and Panetta 2010).

We find that outreach is regressive: participants are less poor than the overall population in the area. This result is even stronger when we use our own metric based on the squared poverty gap and appears in three out of four ways of measuring consumption. The exception is when we measure consumption directly using recall questions on 17 items, in which case the results point in the opposite direction but are statistically insignificant. This might be explained by measurement error on this particular variable. Asked about the reasons for not joining, non-participants report that the problem is lack of cash to fulfill the compulsory savings requirements.

The analysis of the pipeline of participation shows that the awareness campaign initially attracts both the poor and the non-poor, but that the poor are first-movers in the sense that they are more likely to join, given they received the information about the upcoming intervention.. Only later, richer households join and do so in larger numbers. In other words: the awareness campaign seems to attract a different group of people than those who end up joining.

The rest of the article is structured as follows. The following section describes the intervention implemented in the area, the so-called Villages Savings and Loans Associations (VSLAs). Section three provides an overview of the existing targeting literature. Section four reviews targeting measures used in the literature and develops a new metric based on the squared poverty gap. Section five explains our sequential approach. In sections six and seven we present the data and our empirical strategy before turning to the results in section eight. The final section concludes and provides policy recommendations based on the results.

THE INTERVENTION

The microfinance intervention in our study is a community-managed microfinance program called village savings and loan associations (VSLAs). VSLAs are a form of accumulating savings and credit association (following the definitions used by e.g. Bouman 1995), where villagers meet every week and contribute an amount to a common pool of funds. The procedures of setting up and running these groups are thoroughly documented in a set of manuals (Allen and Staehle 2007). Key characteristics are that no external funds are provided, so all loans are made using participants' savings. There are lower and upper limits to the amount that it is possible to save at each meeting. Credit is provided to members at an interest rate set by the group, typically five to ten percent per month with a three month repayment period. An association also includes a welfare fund financed by very small weekly payments by each member. The welfare fund can be invoked on certain occasions, for example the death of a family member, crop failure, or weddings. The formation of the groups is essential for the present study. In our case, the formation of VSLAs was facilitated by a local organization called SOLDEV. The implementing organization approaches the village leaders to get their approval of the project. The village leaders are asked to gather all villagers who might be interested in joining such a group at a designated time

for an awareness meeting. These awareness meetings are held in the villages to inform people about the initiative. People are asked to form groups with other villagers they trust. Subsequently training sessions are conducted by the implementing organization. During the first three months, a field officer participates in every group meeting and trains the groups in various aspects of the methodology: Electing a management committee, administering savings, giving out loans, etc. After the first three months, the group is still supervised by the field officer, although at a lower frequency. After twelve months, the groups “mature” and are no longer supervised by the implementing organization.

The purpose of the intervention is to increase participant’s access to credit, to offer savings with higher return and lower risk than alternatives and to offer a basic insurance product. This can benefit the participants if they are facing credit constraints (Stiglitz and Weiss 1981), it can help them cope with shocks (Dercon 2002) or it can help them save by overcoming their own or their spouse’s self-control problems (O’Donoghue and Rabin 1999, Ashraf *et al.* 2006). Evaluations using randomized control trials are under way and will show if this indeed happens.

As for targeting, the purpose of the intervention was to reach the poorest, but there was no explicit targeting goal. For this reason we focus the analysis on targeting effectiveness, or outreach: Who did the intervention reach? In the next section we situate this question in the broader literature on targeting.

TARGETING AND OUTREACH: A LITERATURE REVIEW

Targeting is when an intervention aims to include only a specific subgroup of the population. Most of the literature on targeting rests on the assumption that interventions must reach the poorest in order to benefit the poorest (Amin *et al.* 2003, Coady *et al.* 2004). The opposing view, i.e. that increased poverty reduction does not follow from better targeting, has also been voiced. In this

section we summarize these arguments before we turn to the approaches used when analyzing targeting.

One reason why targeting can lead to less poverty reduction is that it is costly to the implementer and can include hidden costs to participants in terms of conditions for participation or stigma (Ravallion 2009). Moreover, targeting methods might be inefficient. (Niehaus *et al.* 2013) document how targeting on a large number of indicators in a proxy means test might improve statistical accuracy but at the same time decrease enforceability if implementers are corruptible.

Even if the ultimate focus of an intervention is poverty reduction, targeting may not be required. If the goal is poverty reduction through growth, then any intervention must focus on stimulating the economy as such. This might benefit the poorest through trickle down and does not depend on active participation by the poorest. However, whereas there is little doubt that GDI growth reduces poverty on average, there is a large heterogeneity in the existing evidence. The connection between growth and poverty reduction is uneven, and the link is sometimes weak (Ravallion 2001). Even if economic growth is the end goal, there is still a reason to care about targeting as extreme poverty might in itself have adverse effects on growth (Ravallion 2012).

For these reasons, we believe that there is a strong case for investigating targeting effectiveness.

Measuring targeting effectiveness or outreach

One part of the literature on targeting evaluates different practical methods like proxy means tests, geographical targeting, and community-based targeting (Conning and Kevane 2002, Houssou and Zeller 2011, Alatas *et al.* 2012, Lang *et al.* 2013). Another part of the literature focuses on targeting effectiveness and in particular on whether interventions are successful in reaching the subgroup in question. This division of the literature is illustrated in figure 1. VSLAs are designed to reach the poorest but do not use explicit

targeting methods like the ones just mentioned. When there is no specific targeting goal we simply analyze the intervention's outreach: How successful was the program in reaching the poorest? In the following, we use outreach and targeting effectiveness interchangeably. For an overview of the literature mentioned below, see table 1.

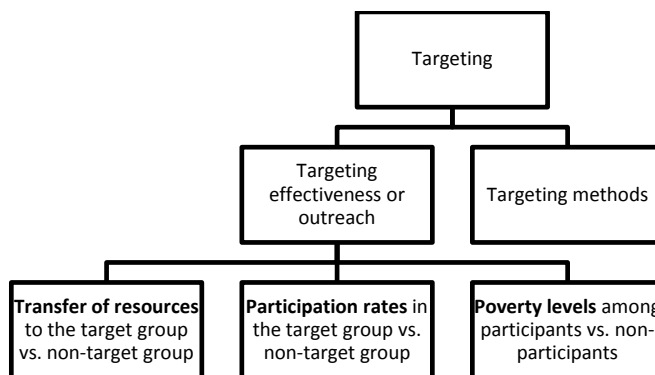


Figure 1. Sub-divisions in the literature on targeting

The literature on targeting effectiveness can be further divided into three categories according to its focus of attention. One type is concerned with the amount of resources transferred to the target group, e.g. the poor, compared to the resources transferred to people outside the target group. A second type looks at participation rates in the two groups. A third part of the literature compares poverty levels among participants and non-participants. This is the approach we use below. Since we draw on lessons from all three approaches, we discuss key contributions from each in the following description before turning to the poverty metrics used in comparing participants to the general population in the area.

Grosh (1994) and Coady *et al.* (2004) are central to the first strand of the literature, which measures targeting by the resources transferred to the poor. Both studies compare multiple interventions across countries. To facilitate comparison they develop a generalized targeting performance indicator, which

we utilize and develop further in the analysis below. Conceptually, they compare the targeting of an intervention to the common reference of neutral targeting where all subgroups of a population receive the same share of the total transfers, irrespective of the income level of the subgroup. The indicator is calculated as the share of funds transferred to the target group, e.g. the poor, divided by the target group's proportion of the overall population. A performance indicator above one indicates progressive targeting – i.e. that the poor are given preferential treatment – whereas an indicator below one indicates regressive targeting. Coady *et al.* (2004) construct a database of 122 targeted anti-poverty programs and find that a quarter of the programs exhibit regressive outreach despite ambitions of the opposite. Grosh (1994) finds that the twenty-three programs in Latin America she has information on exhibit progressive targeting.

Even though there is no transfer of resources in microfinance in general, nor specifically in community-managed microfinance, we can easily adopt the same approach when developing targeting performance indicators. Instead of using the proportion of the amount transferred, we use the average of a number of different poverty metrics for the participants divided by the same metrics of the general population in the area.

The second part of the literature compares participation rates in various ways. This is particularly useful when interventions have clear targeting criteria, e.g. the Zambian maize subsidy given to everyone with an income below K20,500 or the Jamaican food subsidy program for pregnant women, as mentioned by Cornia and Stewart (1993). The authors use these clear targeting criteria to quantify mistargeting by dividing errors into F-mistakes, which is failure to reach the entire target group, and E-mistakes, which is excessive targeting or inclusion of people from outside the target group in an intervention. Investigating primarily food subsidy schemes, Cornia and Stewart (1993) find that a targeting mechanism designed to minimize E-mistakes often increases F-mistakes at the same time. Other studies that

investigate participation rates include Ravallion (2009), Handa *et al.* (2012), and Houssou and Zeller (2011). Coady *et al.* (2004), mentioned above, also use participation rates as the basis of their performance indicator whenever they cannot find information on transfers.

There are two lessons to be learned from this literature. First, participation rates are primarily useful when programs operate with clear targeting criteria. That is not the case for a program like ours, which targets the poor in general. Second, when the target group is the poor, participation rates treat all poor equally, thereby ignoring the severity of poverty. This issue is discussed further in the section on the outreach ratio.

The literature that analyzes targeting by comparing poverty levels among participants and non-participants includes Mohammed *et al.* (1999) and Amin *et al.* (2003). Both studies analyze microfinance in Bangladesh and find that participants are poorer than non-participants on average. Also, Navajas *et al.* (2000) compare poverty levels of participants in microfinance in Bolivia and find that they are just below the national poverty lines but do not belong to the very poor. The advantage in comparing poverty levels is that it allows for flexibility in the definition of poverty apart from the dichotomy poor/non-poor. We use several specific metrics from this literature, as discussed in depth below.

One way of looking at outreach in microfinance in particular is to use loan size as a proxy for the poverty level of clients (Cull *et al.* 2007, Hartarska and Nadolnyak 2007, Hermes *et al.* 2011). This approach is common despite the fact that loan size is likely to differ across sectors, possibly creating a systematic correlation since the provision of large loans may also occur to very poor people. A case in point is agricultural loans, which are likely to be larger than average, while clients might very well be poorer.

A separate issue in the literature is endogeneity since households are often surveyed after an intervention. Comparing non-participants and participants at this stage mixes pre-program differences with any positive or

negative effects of participation on either levels or variances. Mohammed *et al.* (1999) and Navajas *et al.* (2000) are examples of this. Ravallion (2009) also uses post-intervention figures, but argues that pre-intervention income is equal to post-intervention income plus transfers. In the overview by Coady *et al.* (2004) it is not clear whether data is pre- or post-intervention. Our analysis avoids endogeneity since we use poverty measures collected in a survey before roll-out of the program.

THE OUTREACH RATIO

The previous section discussed the literature on targeting in general. In this section we review specific metrics used in assessing outreach, specifically poverty levels of participants. Our starting point is a measure of outreach first introduced by Coady *et al.* (2004), which we call the outreach ratio. The outreach ratio compares the actual targeting in a program with neutral targeting, i.e. a situation where the intervention reaches a representative group of the population. The advantage is that it enables comparison of outreach across different interventions, contexts, and metrics.

The outreach ratio can be based on different measures. If, for example, the basis is poverty headcount, then the outreach ratio is the share of participants falling below the poverty line divided by the share falling below the poverty line in the entire population. If the outreach ratio is above one, targeting is progressive, since the share of poor people participating is greater than in the population as a whole. If it is below one, outreach is regressive.

The key issue is to choose the basis of the outreach ratio. In doing that, we draw on the literature mentioned above, but we add a new type of outreach ratio inspired by the literature on poverty measurement, specifically Sen (1976) and Foster *et al.* (1984). In total, we will include outreach ratios based on four different poverty metrics, one of which we are the first to use in a targeting analysis.

A simple approach is to base the outreach ratio on levels of income or consumption as illustrated in the following equation:

$$OR_c = \frac{\frac{1}{N} \sum_{i=1}^n (y_i)}{\frac{1}{N_p} \sum_{i=1}^{n_p} (y_i)} \quad (1)$$

where OR_c is the outreach ratio based on consumption, N_p is the number of participants, N is the total number of observations, and y_i is the consumption of household i . Amin *et al.* (2003) do not construct the outreach ratio, but answer the same question by comparing incomes among participants and non-participants. In a situation where we want to reach the poor, the limitation of the outreach ratio based on income levels is that the rich contribute to the average to the same extent as the poor. An example of a hypothetical change in income illustrates the problem: A decrease in income among a rich participant, with everything else staying the same, would make targeting more progressive, even though no additional poor people are reached.

To overcome this issue, several studies construct the outreach ratio using poverty headcounts. The corresponding equation is:

$$OR_{ph} = \frac{\frac{1}{N_p} \sum_{i=1}^{n_p} (I_i(y_i < z))}{\frac{1}{N} \sum_{i=1}^n (I_i(y_i < z))} \quad (2)$$

where OR_{ph} is the outreach ratio based on poverty headcount, z is the poverty line and I_i is an indicator which is zero if household i 's income is above the poverty line, and one if it is below. Note that the metric for all participants is in the denominator, and not in the numerator as was the case in equation (1).

This is to ensure that the ratio is above one when targeting is progressive. This exact ratio is used by Coady *et al.* (2004) and Handa *et al.* (2012). Mohammed *et al.* (1999) compare participation rates among the poor and non-poor, which provide similar insights into the outreach. The consumption levels of the non-poor households do not affect the outreach ratio based on poverty headcount. But there is another caveat, again illustrated by a hypothetical example: If we reduce the consumption level of a poor participant, the targeting metric remains the same, even though we now reach deeper than before.

This leads to a third metric, used by Park *et al.* (2002) in an analysis of Chinese counties, based on the so-called poverty gap measure, called the targeting income gap. The targeting income cap is the absolute distance from a county's average income to the poverty line summed over all mistargeted counties, i.e. counties that are in the program, but should not have been, and counties that are not in the program, but should have been. This measure is not subject to any of the critiques we discussed above: If a county with average income below the poverty line experiences a reduction in income and everything else stays the same, then the poverty gap measure increases. Applying this to individual level data, thus calculating the outreach ratio based on the poverty gap (OR_{pg}), is straight forward. Using the same notation as in equation (2), the equation is the following:

$$OR_{pg} = \frac{\frac{1}{N_p} \sum_{i=1}^{n_p} (I_i(y_i < z)(z - y_i))}{\frac{1}{N} \sum_{i=1}^n (I_i(y_i < z)(z - y_i))} \quad (3)$$

One way of comparing the outreach ratios based on the poverty headcount (equation 2) and the poverty gap (equation 3) is how they weight people below the poverty line. Poverty headcount assigns a weight of one to everyone below the poverty line and a weight of zero to households above the poverty line.

The poverty gap uses the distance to the poverty line as weights. One critique of this weighting scheme is that it ignores the depth of poverty in the sense that an increase in income counts the same no matter how poor the household is, as long as it is under the poverty line (Foster 1984). The hypothetical example illustrating this problem is as follows: A transfer from a poor participant to a richer participant, still under the poverty line, and where everything else stays the same, would leave the measure unchanged.

Using the squared poverty gap as a basis for the outreach ratio overcomes this critique. Park *et al.* (2002) suggest this but do not implement it. To the best of our knowledge, this is in fact the first use of the squared poverty gap as a basis for a targeting metric. Given the popularity of the measure in poverty analysis Foster *et al.* (2010), this is peculiar. The outreach ratio becomes:

$$OR_{spg} = \frac{\frac{1}{N_p} \sum_{i=1}^{n_p} (I_i(y_i < z)(z - y_i)^2)}{\frac{1}{N} \sum_{i=1}^n (I_i(y_i < z)(z - y_i)^2)} \quad (4)$$

Again, the notation follows equation (2). Since the different outreach ratios give different weights to the poor, we implement all four in our analysis of targeting effectiveness.

In assessing the metrics, we have implicitly judged them using critiques often raised in the literature on poverty measurement. As is commonly done in this literature, we can now sum this up as three principles that a targeting metric must meet, assuming that the program targets the poor:

- The threshold principle: A change in the income of a person above the poverty line should affect the targeting metric less than a change in the income of a person below the poverty line.

- The poverty principle: A reduction in income among participants below the poverty line must make the targeting metric more progressive.
- The distribution principle: A transfer of income from a participant below the poverty line to any participant who is richer must make the targeting metric more progressive.

The four different outreach ratios discussed above are particularly appropriate for assessing targeting based on a continuous variable, such as total consumption. But we also compute outreach based on two simpler consumption metrics, specifically meals per day and the length of the hungry period. To use the outreach ratio requires defining a poverty line and measuring the distance to this poverty line for each household. We define poverty lines in the section on empirical strategy. Finally, we also include variables on education and health to analyze multidimensional poverty. For these we calculate only outreach ratios based on levels similar to OR_c in equation (1).

After computing the outreach ratios, we present some self-reported reasons for why some people do not join. The survey leaves us with very little data on this issue, but we will nevertheless provide some descriptive results since this issue is important for the usefulness of the results.

Table 1. Overview of the literature on targeting used in the text.

	Country	Interventions analyzed	Method	Dimension	Findings	Endogeneity issues
Microfinance Amin <i>et al.</i> (2003)	Bangladesh	Grameen Bank, ASA, BRAC	Compares consumption and poverty levels for participants and non-participants.	Poverty levels	Participants are poorer than non-participants.	No. Surveys 1992, use membership 1995.
Mohammed <i>et al.</i> (1999)	Bangladesh	BRAC	Compares participation rates across three wealth groups	Participation rates and poverty levels	The poor participate more than the non-poor	Yes. Survey from 1994 only.
Navajas <i>et al.</i> (2000)	Bolivia	Several microfinance institutions	Analyze poverty level compared to national averages	Poverty levels	Participants are just below the poverty line, not poorer.	No
Other Coady <i>et al.</i> (2004)	Multiple countries	111 targeted anti-poverty programs	Divides the share of transfers to a segment of a population with the segment's share of the total population	Transfer of resources	The median intervention is progressive, but a quarter is regressive.	Not certain.
Ravallion (2009)	China	Di Bao transfer program	Compare different targeting metrics' ability to predict differences in poverty reduction due to transfers	Participation rates	Typical targeting metrics do not predict poverty reduction impact	Yes. Assumes pre-program income to be current income minus transfer.
Handa <i>et al.</i> (2012)	Malawi, Kenya,	Social transfer programs	Divides the proportion of participants in the lowest quintile with the proportion of the population in the lowest quintile (20%).	Participation rates	The poor participate more	Not relevant. Compare participants to national averages only, not non-participants in the same area
Park <i>et al.</i> (2002)	China	Various interventions	Unit of analysis is the county and their eligibility for transfers. Measure is the proportion of mis-targeted counties and a measure of the extent of mistargeting	Participation rates	High degree of mistargeting. 22% of counties	No
Houssou and Zeller (2011)	Malawi	Agricultural input subsidies in Malawi	Analyze the performance of alternative targeting metrics	Participation rates	An indicator-based metric is superior.	Not relevant. Compares different metrics only

A LEAKING PIPELINE

As is clear from the previous sections, many studies have looked at who is reached by microfinance, including the poverty status of these. There have been several examples where microfinance has failed to reach the poor. A natural question is: Why are the poor not included? What mechanisms lead to non-participation of the poorest in microfinance and what can be done to prevent this from happening? To investigate this, we borrow the metaphor of a leaking pipeline, which has been used in the literature on gender disparities in education and academia (Barinaga 1992, White 2004). The leaking pipeline illustrates the fact that women exit the academic career path on several steps along the way to professorship. This type of sequential approach has been used to a lesser extent in studying social programs in developed countries (Heckman and Smith 2004) and only recently in developing countries as well (Coady *et al.*).

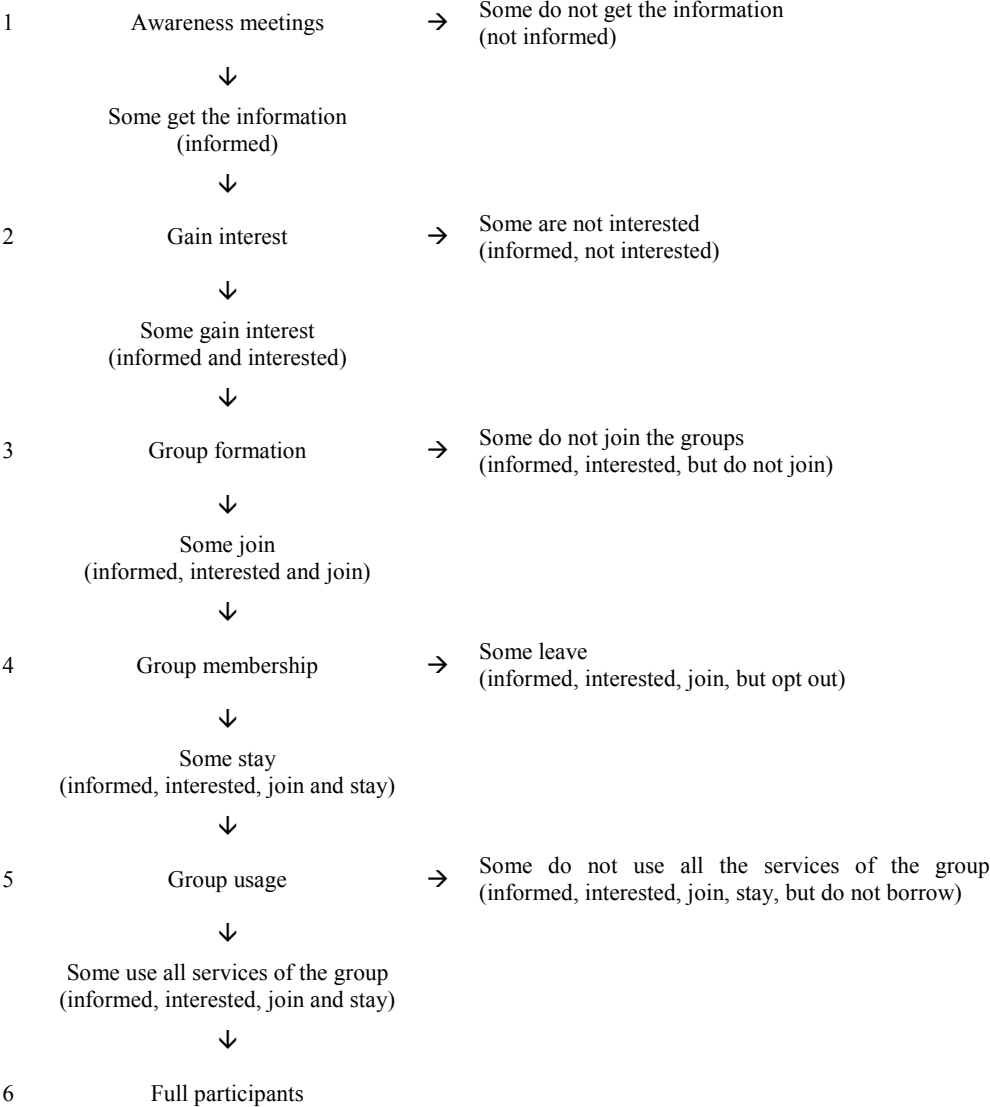
We identify five steps where households can exit the pipeline of participation (see table 2 for an overview). First, the poorest might not receive the information about the upcoming awareness meetings and hence do not know about the groups being initiated. Information about the project is likely to spread via informal channels in addition to the actual meetings, but it is possible that the poorest are excluded from these information networks as well. Second, of all those who get the information, only some turn out to be interested. Some people might not need the services of the groups, or they might already at this stage think that they will be unable to find a group which will accept them. Third, the group formation in itself might leave some out, even if they are interested, as it happens voluntarily—villagers are asked to form groups with other people from the village whom they trust. It could be that assortative matching takes place such that riskier borrowers join groups together, following the traditional theory of information asymmetry (Ghatak and Guinnane). If poverty status is correlated with risk aversion, this might

affect the outreach, since only a limited number of groups are formed in each village. Alternatively, poorer households could be seen as less advantageous for a group and therefore will have difficulties finding one.

Once the groups have been formed, some attrition will eventually happen, which is the fourth step. It might be the case that the poorest do not find the groups useful or that they are pressurized into leaving the group by other members. Finally, as a fifth step, even if the poorest stay in the groups, it is possible that they do not use all the services of the group. Possibly, the poorest do not have enough non-credit resources to make use of loans, or their income is too volatile to risk an involuntary default making them focus on smoothing income instead (Morduch 1995).

Following these stages where exiting is possible, we arrive at six mutually exclusive groups: One for each stage of exiting and a sixth group containing full participants. We compare the households that exit with those who stay in the pipeline with respect to poverty levels using the metrics proposed in the previous sections.

Table 2. The pipeline



THE DATA

The data were collected in one sub-district of Karonga in northern Malawi during two six-week periods in August 2009 and August 2011 as part of the randomized controlled trial documented in the first paper of this thesis. A total of 3,700 households with 20,800 people live in the villages, which cover an area of approximately 400 square kilometers.

The total sample consists of 890 households from twenty-three villages. The entire survey covered forty-six villages, but VSLAs were only established in a randomly selected half of the villages so the current analysis only uses data from the villages in which the VSLA groups were established.

Interviewed households were sampled randomly from household lists provided by local authorities. Stratified sampling was done using two criteria: Village and initial household interest in participating in the VSLAs. In total, data were collected from forty-six villages with interested and non-interested households in each, i.e. a total of ninety-two strata, but since we use only data from half of the villages, we have 46 strata altogether. A higher propensity of sampling was chosen for households in smaller villages relative to larger villages as well as for households who initially expressed interest relative to households not showing interest. Village stratification was performed because of the village-level randomization in the impact assessment. Oversampling of interested households had the aim of oversampling final participants in treatment villages and potential participants in control villages, also for the purpose of assessing impact. Results reported below are weighted according to the inverse probability of sampling following standard practice in survey research and as recommended by several authors, for example Deaton (1997) and Solon *et al.* (2013).

The randomization of villages into treatment and control groups was done within seven blocks to increase baseline balance. The variables used in

the analysis below are from the baseline survey in 2009 except for the variable indicating participation in a VSLA, which was determined in the follow-up survey in 2011. The attrition rate from 2009 to 2011 was less than three percent, which is low compared to similar surveys (Glewwe and Jacoby 2000). In total, forty-eight percent of households participated in VSLAs in 2011.

The questions allowing us to later identify the pipeline were not asked to all respondents but only to a random half using a questionnaire which was longer than the standard questionnaire. Below, this questionnaire is referred to as the long questionnaire, as opposed to the standard questionnaire, which contained only a subset of the questions in the long questionnaire. The purpose of the long questionnaire was to gather information particularly on future participants, and thus households who indicated interest in participating were oversampled to a larger extent than in the short questionnaires. This leads to a different weighting scheme when we analyze the long questionnaire data only. On the other hand, the standard questionnaire increased the power of detecting an effect for the impact assessment in terms of participants and non-participants alike. For these reasons our sample size decreases when we discuss the pipeline compared to simply looking at whether the poor households participate in the VSLA groups.

We use four consumption measures in what follows: Total consumption calculated using recall questions regarding the consumption of seventeen food items in the past week, total consumption predicted using USAID's Poverty Assessment Tool for Malawi (PAT), meals per day, and finally the length of the hungry period measured as the number of months where the household consumed less than three meals a day.

For the first measure, seventeen food items were identified from the Malawi Second Integrated Household Survey from 2004/5 (IHS2). The seventeen items are the ones consumed the most in rural Karonga, the district of the survey, and made up eighty-nine percent of total food consumption and fifty-five percent of total consumption. The total consumption figure below is

the total value of these seventeen items divided by fifty-five percent. The method used in summarizing the 17 items, in particular the calculation of the prices, is explained in detail in Appendix A in the first paper of this thesis. This appendix also describes some of the limitations of this measure, for example that we were forced to calculate our own conversions for each combination of items and units, e.g. a tin of maize, since the conversions provided in the IHS2 were unreliable. This can lead to measurement error. A common finding is that poorer households spend a larger share of their income on food compared to richer households, and thus one concern when using this method is whether it overestimates total consumption for the poor while underestimating it for the rich. In our case, a regression of the share of consumption spent on food on total consumption using the 2004 integrated household survey data from the area, results in a negative, but insignificant estimate ($t=-1.24$).

The second consumption measure is USAID's Poverty Assessment Tool (PAT), which is 20 questions selected on the basis of their ability to predict total consumption in data from the Malawi Second Integrated Household Survey (IRIS Center 2012). We included these questions in our survey and use them and the parameters provided by USAID PAT to predict total consumption for each household. All USD figures are in 2005 dollars using the exchange rate of 91 MKW/USD, adjusted for inflation and using the poverty-adjusted purchase power parity exchange rate described in Deaton and Dupriez (2011).

Apart from the two consumption measures described above, we include two simpler consumption measures: "Meals per day" and "length of the hungry period". "Meals per day" is simply the number of meals the household consumed the day before the survey, and the hungry period is the number of months within the last year where the household consumed less than three meals per day using recall. This occurs most often in the period just before the green harvest in April. However, the length of the hunger period differs

greatly among households. While these measures are cruder than the two consumption measures mentioned earlier, they are easier to measure and thus may contain less measurement error.

Finally, we include four indicators of education and health, which are components commonly included in multidimensional poverty (Alkire and Foster 2011). For education, we look at years of schooling for the household head as well as the share of children aged sixteen to twenty-five who are in school. We choose this age group since this particular area of Malawi is known for a general high level of primary education, and we therefore do not expect much variation for other ages. For health, we use a subjective health measure indicating whether each individual's health is very good, good, average, bad, or very bad. We include both the household average and an indicator for households having one or more members in "bad" or "very bad" health. raising livestock, or fishing.

Table 3 shows summary statistics. The top rows of the table show measures of consumption and food security. Fifty-one percent of the population lives below the 1.25 USD poverty line and one third eat less than three meals per day in August, which is just after the harvest that occurs between May and July. The average hunger period is four months. The middle rows display household characteristics. An average household had almost six members and 17% were headed by women. The last part of the table shows the livelihood patterns with 80% being involved in farming, twenty-one percent in fishing, and sixty-two percent in some income generating activity other than agriculture, raising livestock, or fishing.

Table 3. Descriptive statistics

	Mean	SD
Consumption (USD/capita/day, log)	0.20	0.48
Household living below the 1.25 USD poverty line	0.51	0.44
Meals per day	2.61	0.49
Household consumed less than three meals yesterday	0.36	0.42
Hungry period in months	4.24	3.58
At least one hungry month during the last year	0.74	0.38
Number of household members at time of interview	5.71	2.05
Household head is a women	0.17	0.33
Years of education of household head	7.06	2.82
Household health score (1=good)	1.62	0.50
Anyone in the household has bad health	0.11	0.27
Share of children age 16-25 currently in school	0.10	0.17
Household member of VSLA group in 2011	0.44	0.43
Household does petty trade or small business	0.55	0.43
Household does fishing	0.21	0.36
Household does any farming (subsistence, cash crop, or livestock)	0.80	0.35
Any income-generating activities (excluding agriculture and livestock)	0.63	0.42
Agriculture is the most important income source	0.54	0.43
Land ownership in acres	2.48	1.80
Number of rooms	2.73	1.07
Number of observations	2.73	

Note: The number of observations is lower than in the main analysis due to missing observations on the age of household head. All statistics are on the same sample and computed using sampling weights.

EMPIRICAL STRATEGY

As mentioned in the section on targeting measures, our primary measure of targeting effectiveness is the outreach ratio which compares the poverty status of participants – irrespective of whether they use the loan feature or not – to the poverty status of the population as a whole. In this way, it compares the actual targeting with neutral targeting, i.e. the situation where households participate irrespective of their poverty status.

In effect, the outreach ratios are all ratios of averages taken over four corresponding household metrics in the participant group and among all households, respectively. The four household level poverty metrics are:

Consumption (used in OR_c):	$pm_i^c = y_i$
Poverty headcount (used in OR_{ph}):	$pm_i^{ph} = I_i(y_i < z)$
Poverty gap (used in OR_{pg}):	$pm_i^{pg} = I_i(y_i < z)(z - y_i)$
Squared poverty gap (used in OR_{spg}):	$pm_i^{spg} = I_i(y_i < z)(z - y_i)^2$

where, as above, y_i is consumption, I_i is an indicator which equals one if the household is poor and zero otherwise, and z is the poverty line. For the measures based on 2005 USD we follow Chen and Ravallion (2010) in choosing 1.25 USD as the poverty line. For the simple measures there are levels which can be considered a natural choice in identifying the poor. Regarding meals per day we classify households consuming two meals or less as poor. For hungry months, we label households as poor if they indicate one or more hungry months. The distance to the poverty line is then simply the number of hungry months.

When it comes to the multi-dimensional poverty measures, there is no agreement on how to sum up across the different measures. One attempt was made by Alkire and Foster (2011) and used in the Human Development Report 2010 (UNDP 2010), but this approach was later criticized particularly because the authors add up different measures with arbitrary weights (Ravallion 2011). Because of this, we do not sum up the multidimensional measures but simply display them one by one.

As already discussed, the issue is whether or not the outreach ratio is different from one. An outreach ratio smaller than one corresponds to regressive targeting, whereas a ratio larger than one is progressive targeting. To assess if the ratio is indeed statistically significantly different from unity,

we compare the average of the participants with the average of non-participants by fitting the following very simple regression model:

$$pm_i^k = \alpha + \beta D_i + \varepsilon_i \quad (5)$$

where pm_i^k is the individual poverty metric z as listed above for household i , D_i is a dummy which is equal to one when anyone from the household participates in a VSLA in 2011, and ε_i is the error term. We also estimate the total mean, $\mu = \frac{1}{N} \sum_{i=1}^N (w_i * pm_i^k)$ where w_i is the sampling weight of household i . The average level for participants is $\alpha + \beta$. We assess whether α/μ is different from one by testing $H_0: \beta=0$, which is effectively a weighted t-test for equality of means between the two groups. We always use the inverse probability of sampling as weights as explained in the first paper of this thesis.

When investigating the pipeline, we use a simple version of the approach used by Heckman and Smith (2004), who analyze participation in the job training program JTPA as a pipeline. As mentioned above, we analyze leakages in five sections of the pipeline, which can be formulated as conditional probabilities: (1) the probability of gaining awareness of the project, (2) the probability of interest given awareness, (3) the probability of joining a group given interest and awareness, (4) the probability of staying in a group, i.e. not dropping out, given joining a group, awareness and interest and (5) the probability of utilizing the full range of services provided by the group, i.e. both loans and savings given (1) to (5). More formally, we have:

$$\begin{aligned} P_1 &= \Pr(\text{aw}=1) \\ P_2 &= \Pr(\text{int}=1 \mid \text{aw}=1) \\ P_3 &= \Pr(\text{join}=1 \mid \text{int}=1, \text{aw}=1) \\ P_4 &= \Pr(\text{stay}=1 \mid \text{join}=1, \text{int}=1, \text{aw}=1) \\ P_5 &= \Pr(\text{use}=1 \mid \text{stay}=1, \text{join}=1, \text{int}=1, \text{aw}=1) \end{aligned} \quad (6)$$

To compute these probabilities we use data from several sources. Steps one and two draw on the baseline survey. Steps three, four, and five use data from the follow-up survey two years later. Apart from estimating the actual probabilities, we also estimate average poverty metrics within each of the five nested groups. We test if there is a difference between the group that stays in the pipeline and the one that exits. We do not calculate the outreach ratio since we do not compare the groups to the entire population. For example, to test for difference in poverty among those who join and those who do not join a group is defined as the following conditional expectation:

$$E(y_i | \text{join}=1, \text{int}=1, \text{aw}=1) - E(y_i | \text{join}=0, \text{int}=1, \text{aw}=1) \quad (7)$$

We test the H_0 that this expectation is equal to zero using the same regression model as (5) above, where D_i in this case is a dummy for joining a group, and the only households included are the ones who were aware and interested. A natural challenge in doing this is that the sample size, and thus the power, decrease throughout the pipeline. The chance of not rejecting H_0 , even when it is in fact false, increases.

RESULTS

There are clear signs of regressive outreach across almost all the different poverty metrics and the four different outreach ratios described in our analysis above as displayed in table 4. The only notable exception is when the outreach ratio is based on directly measured consumption calculated from 17 food items, in which case we find that the outreach ratio is not different from one.

When it comes to the outreach ratio based on the USAID PAT, it is lower than one, and the difference is significant at a ten percent level for ratios based on levels, poverty gaps, and squared poverty gaps. Using the PAT consumption level results in a ratio of 0.8, since log of consumption is 0.16

among participants and 0.20 in the whole population. This difference of 0.04 corresponds to a 4% difference in consumption levels or approximately 5 cents, using the formula provided in Kennedy (1981).¹ The outreach ratios based on poverty gap and squared poverty gap are 0.85 and 0.90, respectively, indicating a 10-15% difference between participants and the whole population in these poverty metrics. Since results are significant only at a ten percent level, we interpret this merely as indications of differences.

The results on the simpler measures are stronger, however. Within these the outreach ratios are consistently below one, meaning that outreach is regressive, and participants are systematically better off compared to the population as a whole. For the length of the hungry month and meals consumed yesterday results are significant at 1% or 5% levels, except for the poverty headcount with regard to the hungry period. This is one limitation of setting the “poverty line” at zero, since it leads to almost all households being poor, and there is thus little difference in the poverty levels of participants and non-participants. However, the outreach ratios that measure the distance to the created poverty line overcome this problem and find significant differences.

Interestingly, the outreach ratio worsens as we apply greater weights to the poorest. The outreach ratio based on the average number of meals per day is 0.97. Based on the share of households eating less than three meals it is 0.86. Looking at the poverty headcount, poverty gap and squared poverty gap it is reduced to 0.86, 0.79, and 0.68, respectively, all significant at a 1% level. An illustration of this is provided in figure 2, which shows that almost none of the households consuming one meal per day are members. The hungry period follows the same pattern, where participants face a period which is half a month shorter than the one the overall population experiences. Figure 3 displays the entire distribution and shows a clear difference among the poorest, for example that twice as many non-members as members have a constant hungry period, i.e. consume less than three daily meals in all twelve of the last twelve months. Among the households experiencing a relatively

short hunger period – between one and four months – the members are over represented. This also explains why the outreach ratio based on the poverty headcount is not different from one: Since the poverty line is set at one hungry month, households with short and long hungry periods all count the same.

Regarding our measures of multidimensional poverty, we cannot reject neutral targeting with respect to health, but we find regressive outreach when it comes to children’s education levels, where both indicators are statistically significant, albeit only at a ten percent level.


In sum, these results indicate that not even the most pro-poor microfinance can reach the poorest. As a final note, though, we compare the outreach of VSLA to other providers of microfinance. Several established microfinance institutions operate in the area, for example FINCA, Pride Africa, Malawi Rural Finance, and Opportunity International, and 0.5% of the households in our sample have an account in these institutions. If we extend the pool to friends and relatives, which was the most common source of finance before the intervention, the percentage with loans is five percent. We interpret these results in the conclusion.

Table 4. Main results

Basis for calculating the outreach ratio	Outreach ratio	Averages	
		Participants	All households
<i>Consumption poverty: Directly measured</i>			
Consumption (USD/capita/day, log)	1.10	0.18	0.20
Poverty headcount	1.03	0.51	0.50
Poverty gap	1.02	0.17	0.17
Squared poverty gap	0.98	0.08	0.08
<i>Consumption poverty: Indirectly measured</i>			
Consumption using PAT (USD/capita/day, log)	0.80	0.20	0.16
Poverty headcount using PAT	0.94	0.52	0.55
Poverty gap using PAT	0.90*	0.15	0.16
Squared poverty gap using PAT	0.84*	0.05	0.06
<i>Consumption poverty: Simple measures</i>			
Meals per day	0.97**	2.69	2.61
Poverty headcount (less than three meals per day)	0.86**	0.31	0.35
Poverty gap (meals)	0.79***	0.10	0.13
Squared poverty gap (meals)	0.70***	0.04	0.05
Hungry period in months	0.89**	3.76	4.20
Poverty headcount (at least one hungry month last year)	0.99	0.72	0.73
Poverty gap (hungry period)	0.89**	3.76	4.20
Squared poverty gap (hungry period)	0.82***	28.44	34.69
<i>Multidimensional poverty</i>			
Years of education of household head	0.96	7.34	7.07
Share of children age 16-25 currently in school	0.81*	0.12	0.09
Anyone in the household has bad health	0.95	0.11	0.11
Household health score (1=good)	0.97	1.58	1.62
Observations	883		


Note: This table shows that there is regressive targeting in VSLA. Stars on the outreach ratio are a result of a t-test of whether there is a difference among participants and non-participants, and thus whether or not the outreach ratio is different from one. The test is performed as a regression of the poverty parameter on a dummy for non-participation where $H_0: \beta=0$. * Significant at 10%, ** Significant at 5%, *** Significant at 1%. All estimations use sampling weights and robust standard errors. Clustering standard errors at the village level yield similar results.

Figure 2. Food consumption (meals/day)



Note: Proportions are estimated using sampling weights and sum to 100% across non-members and members.

Figure 3. Hungry period in months



Note: Proportions are estimated using sampling weights and sum to 100% across non-members and members.

What Prevents the Poor From Joining?

Our results from the previous section show that regressive outreach takes place in the intervention. But what are the reasons for this? Since there is self-selection into the groups, the issue is whether the poor are left out because the other members do not want them, or because they themselves do not want to join a group. In the latter case, why do they not want to join? Due to the importance of the issue, we give some preliminary answers to these questions below before investigating the participation pipeline described earlier. However, the following section draws on data from only seventy-three households.

The data we use come from an additional survey administered halfway between the baseline and the follow-up, i.e. after one year, which is the only time we asked questions regarding this issue. Two selection criteria make the sample size small: We ask only respondents who are not already members of a VSLA (564/801), and we limit the questions to the non-members who, at the time before start-up, actually had knowledge that groups were starting up (73/564). Eight of these seventy-three initially joined groups but dropped out, and further thirteen were not interested in joining. The remaining fifty-two were interested, but did not join. The reasons stated by all seventy-three are summarized in table 5.

The primary reason is lack of demand with 34.7% answering “Didn’t have money to save”, followed by lack of supply: “Not enough groups formed” (15.1%). The latter could both reflect a lack of supply from the NGO or the inability of the household to find a group willing to let the household join. This indicates that the intervention design itself might prevent some people from participating. Unfortunately, the sample size is too small to investigate whether the poorest fall in either of these categories.

Practitioners and academics have repeatedly refuted the claim that someone can be too poor to save (Bouman 1995, Armendariz de Aghion and Morduch 2005, Rutherford 2009). Given the responses above, this is not as

straightforward as these authors claim. If by savings we understand cash savings, then some people might actually have too little cash to save, and VSLA is not suitable for them. Another explanation could be that the dynamics of group formation leads to groups with too high minimum savings requirements at least in some villages. Following this strain of thought, VSLA in and of itself might be suitable even for the poorest, but the way the groups are formed does not ensure that groups can accommodate small savings.

Table 5. Reasons for not joining a VSLA group

No need for the services in the group	6.6%
Other members in the group did not want me to join	5.0%
I didn't know time/location of meetings	3.5%
Not enough groups formed	15.1%
My husband did not approve	9.3%
Didn't have money to save	34.7%
Temporarily away at the time	13.1%
Other	12.6%
Total	100.0%
N	73

Note: The table show that lack of money to save is an important constraint preventing people from joining VSLAs. Respondents are people who knew VSLA was starting before it started and therefore 73. Control village with VSLAs are included. Proportions estimated using sampling weights.

Investigating the Pipeline

We now turn to the pipeline of participation. As mentioned in the section on data, the analysis includes only half of the respondents as opposed to the main analysis, and as we move down the pipeline the number decreases.

Table 6 shows how membership develops through the pipeline. The first column shows the conditional probabilities mentioned in equations (6.1) to (6.5).

Table 6. Pipeline results

Section of the pipeline	Conditional probability		
	$P_x = \Pr(j=1 \mid \text{previous steps})$	Share of total in the pipeline	Share of total leaving the pipeline
1) Gain awareness	60%	60%	40%
2) Express interest	88%	53%	7%
3) Join group	59%	31%	22%
4) Stay in group	95%	30%	2%
5) Use both savings and loans	64%	19%	11%

Note: The table shows that only 60% of households got information about the intervention prior to start-up and no more than 59% of those who got information and were interested, actually joined. All percentages are estimated using sampling weights.

Two interesting facts emerge from the table: The awareness campaign reached sixty percent of the population. If the ambition is to reach everyone, and considering that the area is fairly small, then this is low. Moreover, even though almost all of those who hear about the project are interested, only fifty-nine percent of these choose to join a group. For some reason, forty-one percent of those who hear about the project *and* express interest end up not joining two years down the line. In order to understand how participation happens, this is important. The flip side of this is that a large share of those who initially said they were not interested ended up joining. Awareness campaigns are needed, but in their current design they are not necessarily good at communicating costs and benefits of participating in ways that match the intervention.

After looking at the general pipeline, we now analyze the poverty profile of the pipeline. We start out by looking at actual participation on this smaller sample available for the analysis of the pipeline. Since the questions regarding the pipeline were administered to only half of the households, there may be differences between this randomly chosen half of the sample and the full sample. To investigate this, we perform the general participation analysis on this subset of households (the first column in table 7). The results from the general analysis disappear, and in one case—that of directly measured

consumption—we find a significant result where the main analysis found no significant difference. Apparently, although the sample was randomly chosen, there are differences between the whole sample and the subset of the sample that was given the long questionnaires, although there are no cases where the results have opposite signs and are statistically significant at the same time.

Going into details in the different elements of the pipeline we estimate the regression model in equation (5) for each of the steps in the pipeline. This tells us if those in the pipeline differ from those exiting the pipeline. The results are shown in table 7, columns 2 to 6. In most of the steps there are no systematic differences between those who exit and those who stay in the pipeline. There is, however, one important difference: The people who join a group, conditional on getting information about the project and indicating awareness, are significantly poorer than those who do not join. The households who join have lower average consumption, higher poverty headcount, and higher squared poverty gap, indicating that even when looking at the very poorest and taking the distribution of poverty into account, there is a difference.

Since we find no difference on the awareness campaign, we conclude that information about the project reaches a wide, and poverty-wise representative, selection of households, but that the actual activities initially attract the less well-off and appeal less to the richer. Also, the richer join later since we find the opposite difference in the analysis of final participation above—a result of what has been termed “injections” into the pipeline (Soe and Elaine 2008). One reason for this, which is sometimes voiced by practitioners, is that in areas with aid dependency, the well-connected are attracted by the expectation that they can receive transfers of money or goods. Since VSLA does not involve transfer of resources to households, the richer households, who first expressed interest, might be discouraged from spending the time group membership takes.

For the same reason, these results do not support a common assertion by practitioners, i.e. that the poorest are at first reluctant and then join later, possibly because they are risk averse and afraid to place the little money they have outside of the home. If anything, the results point in the opposite direction: The poorest join first.

Table 7. Poverty profile of the pipeline

	(1)		(2)		(3)		(4)		(5)		(6)							
	Membership (long q only)		Gain awareness		Express interest		Join group		Stay in group		Use both savings and loans							
	Pipel	Diff	Pipel	Diff	Pipel	Diff	Pipel	Diff	Pipel	Diff	Pipel	Diff						
Consumption poverty: Directly measured																		
Consumption (USD/capita/day, log)	0.09	0.20	2.15**	0.20	0.01	0.17	0.21	-0.07	-0.50	0.10	0.25	2.51**	0.10	0.11	0.72	0.13	-0.09	-0.68
Poverty headcount	0.58	-0.14	-1.31	0.47	0.06	0.79	0.46	0.15	1.24	0.54	-0.20	-2.18**	0.54	-0.06	-0.34	0.48	0.17	1.15
Poverty gap	0.20	-0.06	-1.48	0.17	-0.01	-0.34	0.16	0.06	1.01	0.19	-0.07	-1.48	0.19	-0.03	-0.35	0.17	0.06	0.96
Squared poverty gap	0.09	-0.03	-1.41	0.08	-0.01	-0.73	0.08	0.01	0.40	0.09	-0.03	-0.95	0.09	-0.03	-0.83	0.09	0.01	0.14
Consumption poverty: Indirectly measured																		
Consumption using PAT (USD/capita/day, log)	0.10	0.06	0.76	0.12	0.04	0.64	0.12	0.00	0.01	0.04	0.20	2.38**	0.05	-0.05	-0.45	0.07	-0.07	-0.58
Poverty headcount using PAT	0.59	0.07	0.75	0.60	0.00	0.00	0.59	0.07	0.51	0.63	-0.11	-1.16	0.63	0.09	0.54	0.58	0.14	0.98
Poverty gap using PAT	0.19	-0.04	-1.06	0.17	-0.01	-0.35	0.17	-0.04	-1.05	0.21	-0.09	-2.27**	0.21	-0.01	-0.13	0.20	0.03	0.45
Squared poverty gap using PAT	0.07	-0.02	-1.03	0.06	0.00	-0.25	0.07	-0.03	-1.95*	0.09	-0.05	-2.23**	0.09	-0.02	-0.49	0.09	0.00	-0.02
Consumption poverty: Simple measures																		
Meals per day	2.72	0.04	0.39	2.77	-0.03	-0.45	2.77	-0.01	-0.04	2.71	0.15	1.77*	2.71	-0.07	-0.38	2.72	-0.03	-0.16
Poverty headcount (less than three meals per day)	0.28	-0.03	-0.35	0.24	0.02	0.25	0.24	-0.01	-0.07	0.30	-0.14	-1.59	0.29	0.07	0.36	0.28	0.03	0.14
Poverty gap (meals)	0.09	-0.01	-0.33	0.08	0.01	0.39	0.08	0.00	0.06	0.10	-0.05	-1.59	0.10	0.02	0.36	0.10	0.01	0.14
Squared poverty gap (meals)	0.03	0.00	-0.29	0.03	0.01	0.63	0.03	0.00	0.30	0.03	-0.02	-1.59	0.03	0.01	0.36	0.03	0.00	0.14
Multi-dimensional poverty measures																		
Years of education of household head	7.17	-1.02	-1.36	7.09	-0.55	-0.99	7.09	0.03	0.04	7.13	-0.09	-0.14	7.25	-2.30	-3.08**	7.14	0.30	0.35
Share of children age 16-25 currently in school	0.13	-0.07	-1.52	0.10	0.01	0.18	0.11	-0.06	-1.93*	0.12	-0.04	-0.87	0.13	-0.05	-0.83	0.09	0.11	1.39
Household has had health	0.06	0.04	0.9	0.05	0.06	1.55	0.06	-0.03	-0.76	0.04	0.05	1.68*	0.03	0.08	1.02	0.05	-0.05	-2.32**
Household health score (=good)	1.61	0.07	0.59	1.55	0.26	2.69**	1.54	0.08	0.62	1.58	-0.11	-1.06	1.58	0.06	0.34	1.61	-0.08	-0.50
Observations	414			414			251			221			128			117		

Note: This table demonstrates that there is difference in poverty levels among those who join VSLA and those who do not, when the analysis is made on the group who got information and expressed interest (column 4). Compared to the main analysis, only half of the observations are included since the information used was in half of the questionnaires, which were longer than the remaining. Column (1) shows that overall results on this part of the sample diverge from the main analysis with respect to directly measured consumption. "Pipel" is the value for the households still in the pipeline. "Diff" is the value for the households not in the pipeline minus the value for the households still in the pipeline. "†" is the t-value of the difference. * denotes statistical significance at the 10% level, ** denotes statistical significance at the 5% level, and *** denotes statistical significance at the 1% level. All estimates are calculated using sample weights.

CONCLUSION

Developing country governments, donors, and NGOs all want to reduce poverty through interventions that reach the poorest. Within microfinance, one way of achieving this goal has been to develop and implement community-managed methods particularly suited for people living on less than the ubiquitous “dollar a day.” Out of the 207m clients in microfinance worldwide, at least two million are members of some 87,000 savings groups similar to the ones we analyze in this paper (Maes and Reed 2012 and savingsgroups.com).

Our goal in this paper is to assess whether these methods are indeed successful at reaching the poorest. The answer is provided by analyzing panel data from one typical community-managed intervention in northern Malawi. Overall, we find regressive outreach: The participants are not as poor as the average population in the area.

Before looking at the data, we review the literature and identify the outreach ratio as a common and useful way of analyzing outreach and targeting. The outreach ratio is simply the share of poor people reached divided by the share of poor people in the population. As an extension of this, we develop our own outreach ratio based not on poverty headcount but on the squared poverty gap—a common metric in poverty measurement, which takes the extent and severity of poverty as well as the income distribution among the poor into account. No previous targeting metric has to our knowledge taken all these three issues into account. Our suggested metric does. Finally, we build a framework that allows us to investigate the participation decision sequentially—an approach we label the participation pipeline. The analysis thus serves not only as an analysis of targeting effectiveness in a particular intervention, but also as an application of these methodological developments.

Turning to the data, we find that fifty-five percent of the general population fall below the 1.25 USD PPP poverty line using USAID’s Poverty Assessment Tool, whereas fifty-two percent of the participants are poor using

this definition. The outreach ratio is below one in almost all cases, which means that outreach is regressive. Using our newly developed metric, which gives higher weights to the poorest, we find the same results: participants are not as poor as the average household in the area. The only exception is consumption measured directly using recall questions on seventeen food items, where we find no significant differences, and we provide some suggestions as to why that might be the case. Importantly, the results show the usefulness of our newly developed metric, since regressive outreach is even more widespread here compared to conventional metrics.

To study why the poor end up participating less than the non-poor, we look at reasons for not joining groups. Asking people why they do not join groups reveals that the savings requirements might be one reason. This is in opposition to the often stated claim that nobody is too poor to save. When investigating the participation pipeline, we find that the awareness campaign reaches less than two-thirds of the population in the area and that only about forty percent of those initially indicating interest end up joining. Somehow the awareness campaign is out of touch with the intervention.

As for the poverty profile through the pipeline, the awareness campaign reaches poor and non-poor alike, but out of those who get the information, the poorest join first. Only later do the non-poor join. For practitioners the last fact is likely to be new. A common assertion is that the poorest join later, whereas this points to the opposite.

What are the practical consequences of these findings? On the one hand, the results suggest that microfinance cannot, in fact, reach the poorest. The glass is half empty. On the other hand, however, a large fraction of the poorest does participate, and VSLA certainly reaches a much poorer group than conventional microfinance, which serves 0.5% of the population in the area. In other words: The glass is half full. The verdict on the program depends on the standard by which it is judged.

Whether or not we should adjust the VSLA model to reach poorer groups depends on local trickle down, i.e. whether or not non-participants benefit from VSLAs in their village even though they do not participate. If non-participants can benefit, then the current model could be sufficient. There is very limited knowledge on this, so it is a topic for future research. Implementation manuals could also be focused on creating benefits in the local community beyond participants, for example through community engagement by groups. This could, however, affect the performance of the groups.

If there is no or little scope for local trickle down, then real targeting methods are needed in order for VSLAs to benefit the poor. One bold way to improve targeting effectiveness is to enforce means testing or targeting through indicators. A drawback is that it would interfere with the self-selection mechanisms that many believe is necessary for this type of intervention to work. A softer approach would be to change the way the model works to better fit the poorest. Implementers could ensure that very low savings amounts are possible in all groups. Or they could supplement VSLA intervention with other types of activities aimed at including the poorest, for example, the mimicking the graduation programs used by e.g. BRAC to enable the poorest to join regular microfinance (Halder and Mosley 2004, Bandiera *et al.* 2011, Hashemi and de Montesquiou 2011).

In discussions of microfinance there is a need to be more specific whenever we talk about reaching the poorest, as well as when we claim that microfinance is not for the poorest. If we want to reach a large fraction of the poor in rural Malawi, microfinance institutions are not a good idea. Community-managed microfinance is. But if we want to reach more of the poor than of the non-poor, then VSLA in its present form is not the right intervention, at least not in northern Malawi.

REFERENCES

- Alatas, V., Banerjee, A., Hanna, R., Olken, B.A. and Tobias, J. 2012. Targeting the Poor: Evidence from a Field Experiment in Indonesia. *American Economic Review*, 102(4), pp. 1206-1240.
- Alkire, S. and Foster, J. 2011. Counting and multidimensional poverty measurement. *Journal of Public Economics*, 95(7), pp. 476-487.
- Allen, H. and Panetta, D. 2010. *Savings Groups: What Are They?* Washington DC: The SEEP Network.
- Allen, H. and Staehle, M. 2007. *Programme Guide - Field Operations Manual version 3.1*. Stuttgart: VSL Associates, available from <http://vsia.net> [Accessed October 2012].
- Amin, S., Rai, A.S. and Topa, G. 2003. Does microcredit reach the poor and vulnerable? Evidence from northern Bangladesh. *Journal of Development Economics*, 70(1), pp. 59-82.
- Armendariz de Aghion, B. and Morduch, J. 2005. *The economics of microfinance*. Cambridge, Mass.: MIT Press.
- Ashraf, N., Karlan, D. and Yin, W. 2006. Tying Odysseus to The Mast: Evidence From a Commitment Savings Product in the Philippines. *Quarterly Journal of Economics*, 121(2), pp. 635-672.
- Bandiera, O., Burgess, R., Das, N.C., Gulesci, S., Rasul, I., Shams, R. and Sulaiman, M. 2011. Asset Transfer Programme for the Ultra Poor: A Randomized Control Trial Evaluation *CFPR Working Paper*, No. 22.

- Barinaga, M. 1992. Profile of a field: neuroscience. *Science*, 255(5050), pp. 1366-1367.
- Bouman, F. 1995. Rotating and accumulating savings and credit associations: A development perspective. *World Development*, 23(3), pp. 371-384.
- Chen, S. and Ravallion, M. 2010. The developing world is poorer than we thought, but no less successful in the fight against poverty. *The Quarterly Journal of Economics*, 125(4), pp. 1577-1625.
- Christen, R.P. 2001. *Commercialization and mission drift - the transformation of microfinance in Latin America*. Washington: Occasional Paper No 5, CGAP.
- Coady, D., Grosh, M. and Hoddinott, J. 2004. Targeting outcomes redux. *The World Bank Research Observer*, 19(1), pp. 61-85.
- Coady, D., Martinelli, C. and Parker, S.W. 2013. Information and participation in social programs. *The World Bank Economic Review*, 27(1), pp. 149-170.
- Conning, J. and Kevane, M. 2002. Community-based targeting mechanisms for social safety nets: A critical review. *World Development*, 30(3), pp. 375-394.
- Cornia, G.A. and Stewart, F. 1993. Two errors of targeting. *Journal of International Development*, 5(5), pp. 459-496.
- Cull, R., Morduch, J. and Demirgüç-Kunt, A. 2007. Financial performance and outreach: a global analysis of leading microbanks. *The Economic Journal*, 117(517), pp. F107-F133.

- Deaton, A. 1997. *The analysis of household surveys: a microeconomic approach to development policy*. Baltimore: Johns Hopkins University Press.
- Deaton, A. and Dupriez, O. 2011. Purchasing power parity exchange rates for the global poor. *American Economic Journal: Applied Economics*, 3(2), pp. 137-166.
- Dercon, S. 2002. Risk, Coping Strategies, and Safety Nets. *The World Bank Research Observer*, 17(2), pp. 141-166.
- Foster, J., Greer, J. and Thorbecke, E. 1984. A class of decomposable poverty measures. *Econometrica: Journal of the Econometric Society*, 52(3), pp. 761-766.
- Foster, J., Greer, J. and Thorbecke, E. 2010. The Foster–Greer–Thorbecke (FGT) poverty measures: 25 years later. *Journal of Economic Inequality*, 8(4), pp. 491-524.
- Ghatak, M. and Guinnane, T.W. 1999. The economics of lending with joint liability: theory and practice. *Journal of Development Economics*, 60(1), pp. 195–228.
- Glewwe, P. and Jacoby, H. 2000. Recommendations for Collecting Panel Data. In M. Grosh and P. Glewwe (eds.) *Designing household survey questionnaires for developing countries : lessons from 15 years of the Living Standards Measurement Study*. Washington DC: World Bank, pp. 275-314.
- Grosh, M.E. 1994. *Administering targeted social programs in Latin America: From platitudes to practice*. Washington D.C.: World Bank.

- Halder, S. and Mosley, P. 2004. Working with the ultra-poor: learning from BRAC experiences. *Journal of International Development*, 16(3), pp. 387-406.
- Handa, S., Huang, C., Hypher, N., Teixeira, C., Soares, F.V. and Davis, B. 2012. Targeting effectiveness of social cash transfer programmes in three African countries. *Journal of development effectiveness*, 4(1), pp. 78-108.
- Hartarska, V. and Nadolnyak, D. 2007. Do regulated microfinance institutions achieve better sustainability and outreach? Cross-country evidence. *Applied Economics*, 39(10), pp. 1207-1222.
- Hashemi, S.M. and de Montesquiou, A. 2011. *Reaching the Poorest: Lessons from the Graduation Model*. CGAP Focus Note, no. 69. Washington D.C.: CGAP.
- Heckman, J.J. and Smith, J.A. 2004. The Determinants of Participation in a Social Program: Evidence from a Prototypical Job Training Program. *Journal of Labor Economics*, 22(2), pp. 243-298.
- Hermes, N., Lensink, R. and Meesters, A. 2011. Outreach and Efficiency of Microfinance Institutions. *World Development*, 39(6), pp. 938-948.
- Houssou, N. and Zeller, M. 2011. To target or not to target? The costs, benefits, and impacts of indicator-based targeting. *Food Policy*, 36(5), pp. 627-637.
- Hulme, D. 2000. Is microdebt good for poor people? A note on the dark side of microfinance. *Small Enterprise Development*, 11(1), pp. 26-28.

- IRIS Center 2012. *Poverty Assessment Tool Accuracy Submission USAID/IRIS Tool for Malawi*. Maryland: University of Maryland.
- Kennedy, P.E. 1981. Estimation with correctly interpreted dummy variables in semilogarithmic equations. *American Economic Review*, 71(4).
- Lang, C., Barrett, C.B. and Naschold, F. 2013. Targeting Maps: An Asset-Based Approach to Geographic Targeting. *World Development*, 41(1), pp. 232-244.
- Maes, J.P. and Reed, L.R. 2012. State of the microcredit summit campaign report 2012. *Microcredit Summit Campaign: Washington, DC*.
- Mohammed, R., Norris, A.H., Evans Alayne, M. and Timothy, G. 1999. Demystifying nonparticipation in microcredit: a population-based analysis. *World Development*, 27(2), pp. 419-430.
- Morduch, J. 1995. Income Smoothing and Consumption Smoothing. *The Journal of Economic Perspectives*, 9(3), pp. 103-114.
- Navajas, S., Schreiner, M., Meyer, R., Gonzalez-Vega, C. and Rodriguez-Meza, J. 2000. Microcredit and the Poorest of the Poor: Theory and Evidence from Bolivia. *World Development*, 28(2), pp. 333-346.
- Niehaus, P., Atanassova, A., Bertrand, M. and Mullainathan, S. 2013. Targeting with Agents. *American Economic Journal: Economic Policy*, 5(1), pp. 206-238.
- O'Donoghue, T. and Rabin, M. 1999. Doing It Now or Later. *The American Economic Review* 89(1), pp. 103-124.

- Park, A., Wang, S. and Wu, G. 2002. Regional poverty targeting in China. *Journal of Public Economics*, 86(1), pp. 123-153.
- Ravallion, M. 2001. Growth, inequality and poverty: Looking beyond averages. *World Development*, 29(11), pp. 1803-1815.
- Ravallion, M. 2009. How Relevant Is Targeting to the Success of an Antipoverty Program? *World Bank Research Observer*, 24(2), pp. 205-231.
- Ravallion, M. 2011. On multidimensional indices of poverty. *Journal of Economic Inequality*.
- Ravallion, M. 2012. Why Don't We See Poverty Convergence? *American Economic Review*, 102(1), pp. 504-523.
- Rutherford, S. 2009. *The Poor and Their Money*, 2nd ed. London: Practical Action Publishing.
- Sen, A. 1976. Poverty: an ordinal approach to measurement. *Econometrica: Journal of the Econometric Society*, pp. 219-231.
- Soe, L. and Elaine, K.Y. 2008. What's wrong with the pipeline? Assumptions about gender and culture in IT work. *Women's Studies*, 37(3), pp. 176-201.
- Solon, G., Haider, S.J. and Wooldridge, J. 2013. What Are We Weighting For? *NBER Working Paper*.
- Stiglitz, J.E. and Weiss, A. 1981. Credit Rationing in Markets with Imperfect Information. *American Economic Review*, 71(3), pp. 393-419.

UNDP 2010. Technical note 1. Calculating the Human Development Index. In *Human Development Report 2010 - 20th Anniversary Edition. The Real Wealth of Nations: Pathways to Human Development*. New York: UNDP.

White, K. 2004. The leaking pipeline: Women postgraduate and early career researchers in Australia. *Tertiary Education & Management*, 10(3), pp. 227-241.

ⁱ The formula is $\Delta y = e^{\left(\hat{\beta} - \frac{\text{var}(\hat{\beta})}{2}\right)}$.

