Financial inclusion of vulnerable households through Savings and Borrowing Groups: Theory and experimental evidence from Uganda*

Alfredo Burlando† *and Andrea Canidio*‡

PRELIMINARY AND INCOMPLETE.
This version: December 31, 2014.

Abstract

Savings and borrowing groups (SBGs) are allowing millions of unbanked people in developing countries to save and borrow without resorting to formal banking systems. However, it is still unclear whether SBGs are also useful in fostering the financial inclusion of the poorest members of local communities. In this paper, we develop a model of SBG. We show that the supply of savings does not necessarily match the demand for loans, and we describe the determinants of saving and borrowing behavior as a function of the ability to save and borrow of other groups' members. We argue that including a person with a low ability to save imposes a negative externality on the group when the funds available to the group are scarce, and a positive externality when funds are in excess of the demand for loans. We test the model using data from an evaluation of a financial inclusion program in rural Uganda, where we exogenously changed the proportion of vulnerable participants with low ability to save in some newly created groups. We show that having members with a greater propensity to save increases borrowing from other members of the group. This is consistent with the prediction that groups are resource constrained, and suggest that including vulnerable members generates negative externalities. We discuss the policy implications.

JEL classification: O12, O16

Keywords: Savings and borrowing groups, Savings groups, Financial inclusion, Local finance, VSLA, Self-help groups.

*We are grateful to SCORE's chief of party Massimo Zucca, as well as to Patrick Walugembe, Noel Nakibuuka, John Paul Nyeko, Michael Muwairwa, Ramadhan Kirunda, and all the staff at FHI360, AVSI, CARE, and TPO for their support in the design of the intervention, implementation of the intervention, and data collection. We received outstanding research assistance from Derek Wolfson, Biraj Bisht, and Attila Gaspar, for which we are extremely thankful. We benefited from comments and suggestions received from the participants to the 2014 GDN Workshop in Prague and the 2014 Northwest Development Workshop. Finally, we acknowledge the financial support of USAID (through project SCORE), CERGE-EI Foundation under a program of the Global Development Network (Regional Research Competition), Central European University (Research Support Scheme), and University of Oregon. All opinions expressed are those of the authors.

†Department of Economics, University of Oregon, Eugene, OR 97403-1285, USA, burlando@uoregon.edu.

‡Department of Economics, Central European University, Nádor u. 9, 1051 Budapest, Hungary; canidioa@ceu.hu.

1 Introduction

Lack of access to financial services has been identified as a key roadblock for the socioeconomic development of poor rural households in the developing world. For instance, savings and microfinance has been shown to foster the creation of micro-enterprises (Duflo et al., 2013, Dupas and Jonathan, 2013), and savings and borrowing groups contribute to consumption smoothing (Beaman et al., 2014, Ksoll and Forskningsenhed, 2013). The critical question of whether access to financial services would be beneficial to the very poorest and most vulnerable in a community remains. This is of critical importance in a context where financial services are quickly spreading to the better-off members of a community, and possibly hurting the poorest members through negative spillovers through reduced financial transfers (Lingon et al., 2002), changes in the structure of financial networks (Comola and Prina, 2014), or from over indebtedness (Karim, 2011).

In this paper, we study the ability of savings and borrowing groups (SBGs) to bring financial inclusion to poor vulnerable households who are usually not reached by traditional banking or microfinance interventions.[1] SBGs are spreading extremely fast in sub-Saharan Africa and elsewhere. In 2014, an estimated 10.5 millions households are members of SBGs, a tenfold increase relative to 2008.[2] While SBGs are being used by many rural unbanked individuals, it remains unclear whether the poorest and most vulnerable members of a community benefit from inclusion into these groups, or whether ultimately this institution will benefit only the better-off part of the community. This concern arises because well-functioning SBGs require a membership that is able to save, to borrow, to utilize the borrowed funds, and to repay the group. In order to gauge whether these groups are effective vehicles of financial inclusion, it is thus important to understand whether the poorest carry out these functions in a different way than the non-poor, and whether these differences affect the performance of the group.

In the first part of the paper, we build a theoretical model representing the functioning of a SBG. We use the model to study how the propensity to save and borrow of the group's members shape the individual saving and borrowing decisions, as well as the collective performance of the group. The kind of group we are interested in modeling is commonly found in sub-Saharan Africa. SBGs are typically composed of twenty to thirty members who save and borrow from the group. Individual savings are accumulated on a weekly basis in a common pool stored in a safe box, and lent out to requesting members over a year-long operating cycle. Loans are paid back to the group with interest, which is set at the beginning of the cycle. At the end of the operating period, all funds in the safe are shared among the group's member in proportion to

[1] Different papers use different names to denote SBGs, partly because there are many slightly different ways in which SBGs are run. See section 2.1 for more details.

[2] Currently, 1.2 million people belong to SBGs in Uganda, where we conduct our intervention. These statistics are taken from the Savings-Led Working Group (SLWG) of SEEP. See www.seepnetwork.org/filebin/docs/SG_Member_Numbers_Worldwide.pdf

the amount saved during the period of operation.

In our model, we assume that the interest rate on loans is established by the group at the beginning of the cycle and is taken as given. The return on savings is determined *ex-post* (at share out) and is equal to the total interest-rate payments collected by the group divided by the total savings collected. The members of the group save and borrow from the group taking as given the return on savings, the cost of borrowing, and the availability of investment opportunities outside of the group. The return on savings is then determined in equilibrium, but–crucially–it may fail to equate demand and supply for funds. As a consequence, within the group, funds may be in excess or may fall short of the demand for loans.

The main result we derive is that savings and borrowing decisions by one member of the group impose an externality on other members of the group. Crucially, whether this externality is positive or negative depends on whether the demand for loans is rationed, that is, whether every person who wishes a loan of a certain size can get it. When there is no rationing, the supply of funds available exceeds the demand for loans, and every additional dollar saved by one reduces the return on savings for everyone else. This negative externality causes the other members of the group to reduce their savings. An agent who increases her borrowing from the group, on the other hand, causes an increase in the equilibrium return on savings, which in turn increases savings by everyone else. However, when the group does not generate enough funds to satisfy the demand for loans, the sign of the externalities generated by an increase in savings or borrowing are reversed. In this case, additional contributions to the group create a positive externality on potential borrowers, because rationing is eased and more people are able to meet their borrowing needs, while additional demands for funds create a negative externality as they worsen rationing.

Our theory has important implications for the financial inclusion of the very poor, as the model suggests that including vulnerable members in a savings group will affect the performance of the group and its ability to provide financial services to its members. The sign of this effect depends on whether this extra member is a net borrower or net saver, and whether the group is rationed out of funds or has excess funds. In the context of the savings groups we analyze in Uganda, the vulnerable have a lower ability to save but share the same demand for loans, and their inclusion will reduce supply of funds available. If inclusion occurs in groups that would have been otherwise rationed, it is welfare decreasing. If inclusion occurs in groups that are not rationed, it is welfare increasing.

We evaluate the theory using data from an evaluation of a particular program that actively enrolls the most vulnerable members of a community in SBGs. This program identifies members of the community with significantly lower socioeconomic characteristics, and organizes new savings groups that include those identified members with other, self-selected members of the population. To study the implications of this inclusion of vulnerable populations, in our exper-

iment we randomly assigned some identified vulnerable participants to groups that had many other vulnerable members, while others were assigned to groups with few vulnerable members. In other words, we created exogenous variation in the vulnerability profile of peers for a random sample of vulnerable group members.

The empirical evidence points to three main results. First, less vulnerable groups saved and borrowed significantly more than more vulnerable groups. Total savings in the less vulnerable groups was, by the end of the first cycle, 30% larger than the comparison. We find no evidence that the intensity of use of resources differed between the two types of groups, suggesting that the returns on savings was similar across groups with different vulnerability profiles. Second, the vulnerable themselves save and borrow less when placed in more vulnerable groups. Again, the magnitude of the difference in savings and borrowing is significant: the vulnerable enrolled in highly vulnerable groups save 20% less and borrow 45% less on average than those enrolled in less vulnerable groups. Since our vulnerable participants were randomly assigned to groups, this difference in savings and borrowing is driven by the characteristics of their group. Third, savings and borrowing from other group participants does not vary significantly with the treatment. That is, while vulnerable groups operate at a smaller scale, this seems to affect differentially the vulnerable and the non-vulnerable. However, because there was no random assignment of non-vulnerable participants, it is hard to establish the extent to which this result is due to unobserved and systematic differences in the demand to save or borrow for these members.

The empirical results together with the theoretical predictions suggest that the groups we study are rationed of funds, and this rationing is less severe in groups with fewer vulnerable members. As a consequence, our model predicts that vulnerable households are better off when placed in groups with many non-vulnerable households.[3] Importantly, the penalty for being included in a vulnerable group does not seem to extend to the non-vulnerable population. This conclusion is relevant for the creation of SBG in areas where there is the possibility of including both vulnerable and non vulnerable households in the same SBG. It also suggests that SBGs created in areas in which the majority of the population is vulnerable may be more acutely unable to meet the demand for loans. The functioning of these groups can be improved by providing outside funds, as well as by changing the rules of functioning of the group so to encourage early savings (that can be lent out multiple times).

The existing literature on savings and borrowing groups has largely focused on measuring the effect of SBG participation. For example, Beaman, Karlan, and Thuysbaert (2014) randomize at village level the creation of SBGs in Mali. They find that treated villages have higher savings (+30%), borrowing, consumption smoothing, food security, livestock holding compared to control villages. They also report that the wealthiest member of each village tend to select into SBGs. Ksoll and Forskningsenhed (2013) finds similar results employing a similar research

[3] We will be able to directly measure the welfare implications for the vulnerable once an in-depth endline data collection process is concluded.

design. Bundervoet (2012) finds large effects of SBG participation on household welfare by randomizing the timing of the provision of the SBG training. Whereas these works establish that SBG participation is overall beneficial to the households who choose to participate, our study is mainly concerned with the determinants of savings and borrowing behavior within the SBG. By investigating these incentives, we are able to show that the benefit of SBG participation depends crucially in the composition of the groups. With this respect, we are close to Greaney, Kaboski, and Van Leemput (2013), who study theoretically and empirically the mechanism of group formation. Their main finding is that when the SBG training is paid by the group members rather than being provided for free, high-risk low-returns agents are driven out of the group, which improves the functioning of the SBG. Greaney, Kaboski, and Van Leemput (2013) abstract away from the specific rules determining the functioning of a SBG, and simply assume that the equilibrium return on savings is such that supply and demand for funds are equalized (i.e. an SBG is a local credit market with limited liability). Instead, we focus on the rules governing the functioning of an SBG, and on how these rules determine the return on funds, the behavior, and the payoffs of the group's members. At the same time, we abstract away from the determinants of group formation.

More generally, our paper is related to the literature on financial inclusion. Several existing papers demonstrate that mocroentrepreneurs benefit from access to finance (see, for example, Banerjee, Duflo, Glennerster, and Kinnan, 2013) and access to a safe way to store their money (see, for example, Dupas and Robinson, 2013). For other subpopulations, the evidence remains mixed, with some evidence of overindebtness from microfinance in Bangladesh (Karim, 2011) but not in Mexico (Angelucci, Karlan, and Zinman, 2015). Finally, our paper also contributes to the literature on peer effects in financial markets, as it quantifies, in the particular context of savings groups, the importance of the peer group in generating financial outcomes (Feigenberg et al, 2014).

2 Background information

2.1 Functioning of savings and borrowing groups

We define SBGs as a formal association with limited membership (typically around 20-30 participants) which provides a source of interest-bearing savings and self-generated credit to its members. There are many types of SBGs in operation, each following somewhat different rules. The most common type is the Village Savings and Loans Association (VSLA), which was first introduced by CARE International in Niger in 1991. Other NGOs have introduced their version of savings and borrowing groups, such as Oxfam's *Saving for Change* groups and Catholic Relief Services' *Savings and Internal Lending Communities*. These models are similar to each other with respect of the basic functioning of the groups, but differ in how training is organized, how

records are kept, the level of involvement of NGOs in the workings of the groups (see Ashe, 2009, Vanmeenen, 2010). The data used in this paper originated from VSLAs, but our theoretical and empirical contribution extends to the most common types of SBGs.

Savings and borrowing groups generally operate in the following way. First, an association organizer recruits potential members within a community.[4] Participants attend weekly training meetings for a period of one month, during which the association is explained. Following the training period, the group agrees on the bylaws of the association, which include the maximum weekly savings allowed, the internal interest rate charged on loans and the length of the saving cycle. Savings contributions are measured in *shares*, and the participant can deposit in the account up to a multiple of the share value (in the groups we study, the multiple is five). Savings are kept in a metal safe box, which is open only when the group is in session and is secured with multiple locks whose keys are given to different members of the group.

Borrowing from the safe box takes place some time (typically three months) after the beginning of the savings period. Each participant is limited to borrowing up to three times the amount they saved until that point. Loans are extended only if the group agrees on the purpose of the loan, and subject to a resource constraint (the amount available in the safe box). Loans must be generally repaid within three months; the interest is calculated on the outstanding loan amount after one and two months from the date the loan was given out; no minimum weekly contributions are generally required. Once the loan is paid back, the borrower is eligible to receive another loan.

Loan disbursements end some time before the end of the cycle, and all repayments must occur by the end of the cycle. The last meeting is devoted to the *share out*: the contents of the safe box are emptied and divided among the members of the group in a way that is proportional to the amount each person saved. A new cycle is eventually started. Between the end of the old cycle and the beginning of the new cycle the group composition may change and the rules governing the groups may be modified.

2.2 Project SCORE

Our research project is in partnership with three NGOs based in Uganda (AVSI, CARE, TPO) that provide services to 125,000 Vulnerable Children (VC) and their household in 35 districts across Uganda. Project SCORE (Sustainable COmprehensive REsponses for vulnerable children and their families) was launched in the fall of 2011 with a USD 9 million USAID grant. Project SCORE is a set of interventions implemented over a period of 5 years having the following goals: to identify vulnerable children and their families across communities in Uganda; to improve their socio-economic status, food security, and nutrition status; and to increase the availability and

[4] SBGs are often initiated by NGOs: there are fixed costs that a group needs to pay (such as books required to register the transactions and the safe where the funds are kept), and the group requires intensive training. Once they are set up, SBGs can be sustained by members only, without the need of continuous NGO support.

access to protective, legal and other critical services. A household is included into the program if identified as vulnerable, which is done using a short questionnaire about the socio-economic status and well-being of the household. Households identified as being vulnerable receive a number of interventions, including classes on advanced farming techniques, cooking, nutrition, business training and business development. The program offers no transfers to beneficiary households, either in money or in kind.

The most important intervention carried out under SCORE involves the enrollment of beneficiaries into SCORE-created SBGs, which follow CARE's VLSA model. Such groups are formed by first enrolling a core of SCORE recipients, and then enrolling other interested community residents. To ensure that vulnerable participants remain dominant in the group, SCORE requires that at least half the membership belongs to beneficiary households. As a consequence, compared with SBGs supported by other organizations, SCORE groups are generally much more inclusive of vulnerable and marginal households, but have less freedom to self-select their members.

Importantly, the three NGOs involved in program SCORE do not provide services directly. Rather, they outsource all interventions to a number of smaller community-based local organizations (which we refer to as Implementing Partners or IPs). IPs typically operate in small areas (few villages) and have a varying degree of capability in carrying out the interventions.

3 The Economics of a Savings and Borrowing Group: a Theoretical Model

We start by developing a theoretical model of SBG functioning. We assume that there exist a mechanism that prevents both voluntary or involuntary defaults, such as sufficiently strong peer pressure and/or lack of uncertainty in the production process. Instead, we focus solely on how the rules governing a SBG shape the incentive to save and borrow of each member of the group. We show that, despite the absence of moral hazard or adverse selection, the rules governing a SBG can generate an inefficient supply of funds, which could be in excess or fall short of the demand for loans. After that, our theory describes how each member's welfare changes when the composition of the group changes, and draw implications for the inclusion of poor members.

Consider a group composed of n individuals. The timing of the game is the following:

- In period 0, the group meets and agrees on the interest rate that will be charged on loans r, and on the maximum savings per period $\overline{s}$.
- In periods 1 and 2, each member i:
  - first receives $w_{i,t}$, which is a per-period wage (i.e. non-investment income generated outside of the group).
  - then saves $s_{i,t}$,
  - then borrows $b_{i,t}$,

  - then invest $y_{i,t}$ in an outside project.
  - then earns $f_i(y_{i,t})$ from the funds invested outside of the group,
  - finally repays $(1+r)b_{i,t}$ to the group, saves $a_{i,t} \geq 0$ outside of the group, and consumes the rest.

- In period 3, the total money generated by the group is redistributed to the group members in proportion to the amount saved.

Both $w_{i,t}$ and $f_i(y_{i,t})$ can be stochastic. Independently on the rules agreed upon in period 1, no member is allowed to borrow more than 3 times the total amount saved with the group up until that period, and therefore

$$b_{i,t} \leq 3 \sum_{s=1}^{t} s_{i,s} \quad \text{for } t \in \{1, 2\} \tag{1}$$

which we call *leverage constraint*. In addition, the agent can save with the group up $\overline{s}$:

$$s_{i,t} \leq \min\{w_{i,t} + a_{i,t-1}, \overline{s}\} \tag{2}$$

where we assume $a_{i,0} = 0$, so that the resources available for saving are $w_{i,1}$ in period 1 and $w_{i,1} + a_{i,1}$ in period 2.[5] In addition, the timing described above implies that

$$y_{i,t} \leq b_{i,t} + w_{i,t} - s_{i,2} + a_{1,t-1} \quad \text{for } t \in \{1, 2\} \tag{3}$$

In other words, the resources available for investment are equal to the own funds (either earned during that period $w_{i,t}$ or carried from the previous period $a_{1,t-1}$) minus the savings with the group, plus borrowing with the group. Finally, consumption at the end of each period is:

$$c_{i,t} = f_i(y_{i,t}) - y_{i,t} - rb_{i,t} + w_{i,t} + a_{i,t-1} - a_{i,t} - s_{i,t} \tag{4}$$

which is the agent's *budget constraint*. We assume that consumption must be non-negative for every realization of $w_{i,t}$ and $f_i(y_{i,t})$.

In this section we analyze the functioning of a SBG from period 1 onward, i.e. for given interest rate on loans and maximum amount saved. On the empirical part, we will show that groups of different composition adopted similar rules, mostly determined by the particular Implementing Partner who formed the SBG. Hence, in our data, the rules adopted by the group reflect mostly the default model proposed by the Implementing Partner rather than the group composition.

We conclude the description of the model by introducing our main assumption:

[5] Note that we implicitly assumed that the resources saved outside of the group do not generate any return.

Assumption 1. *The return on savings at the end of the cycle and the funds available to the group in each period are taken as given by the group members but are determined in equilibrium.*

In other words, the group members fail to anticipate that by increasing the amount saved (or borrowed) they will affect the return on savings for the entire group. Similarly, the group members fail to anticipate that by borrowing and repaying in period 1 they will end up increasing the funds available in period 2 and their own ability to borrow. As a consequence, we can treat the return on savings and the funds available to the group in each period as equilibrium quantities.[6]

Individual problem Under these assumptions, member i decides how much to save and borrow within the group by solving the following 2 problems. In period 2:

$$V(a_1, s_1) = \max_{b_{i,2}, s_{i,2}, y_{i,2}, a_{i,2}} \{E\left[u(c_{i,2})\right] + \beta \cdot m\}$$

$$\text{s.t.} \begin{cases} m = (s_{i,1} + s_{i,2})(1+R) + a_{i,2} \\ b_{i,2} \leq s_{i,2} + \tilde{C}_{i,2} & \text{aggregate resources constraint} \\ \text{equations 1 to 4} \end{cases}$$

And in period 1:

$$\max_{b_{i,1}, s_{i,1}, y_{i,1}, a_{i,1}} \{E\left[u(c_{i,1})\right] + \beta E\left[V(a_1, s_1)\right]\}$$

$$\text{s.t.} \begin{cases} b_{i,1} \leq s_{i,1} + \tilde{C}_{i,1} & \text{aggregate resources constraint} \\ \text{equations 1 to 4} \end{cases}$$

where $\beta \leq 1$ is a discount factor, $u(.)$ is a strictly increasing and concave function, and m is the amount of money available to the household at the end of the cycle. We assume that m enters linearly into the agent's utility function, and hence represents the continuation value corresponding to a given amount of money received at share out.[7]

The term R is the implicit return on savings, while $\tilde{C}_{i,t}$ is the cash available to member i of the group at the beginning of each period. Call $B_t = \sum_i b_{i,t}$ and $S_t = \sum_i s_{i,t}$ the aggregate borrowing and savings in each period. We have

$$\tilde{C}_{i,1} = \sum_{j \neq i} s_{j,t} - \sum_{j \neq i} b_{j,t}$$

[6] Given that the group is large, the incentives to influence the return on savings and the resources available to the group by setting a specific s_i or b_i are likely to be negligible.

[7] All results derived are robust to more general specifications of the utility function (for example, utility non-separable in period-1 and period-2 consumption). We choose to focus on the above formalization because it generates choices that are dynamically consistent. Hence, we abstract away from potential disagreements between period-1 and period-2 self.

and

$$\tilde{C}_{i,2} = \sum_{j \neq i} s_{j,t} - \sum_{j \neq i} b_{j,t} + (1+r)B_1 + S_1 - B_1 = \sum_{j \neq i} s_{j,t} - \sum_{j \neq i} b_{j,t} + S_1 + rB_1$$

i.e. the funds available to the group at the beginning of period 2 are given by the return on the money that was lent out in period 1, and the excess funds from period 1. The implicit return on savings is given by the total amount of cash available to the group at the end of the cycle, divided by the total amount saved within the group:

$$R = \frac{(1+r)B_2 + (S_1 + rB_1 + S_2 - B_2)}{S_1 + S_2} - 1 = \frac{r \cdot (B_1 + B_2)}{(S_1 + S_2)}$$

By assumption 1, both R and $\tilde{C}_{i,t}$ are taken as given by the group's members. Call $s_{i,t}(R, \tilde{C}_{i,t})$ the optimal savings and $b_{i,t}(R, \tilde{C}_{i,t})$ optimal borrowings of agent i in period t.

Lemma 1. *$s_{i,t}(R, \tilde{C}_{i,t})$ and $b_{i,t}(R, \tilde{C}_{i,t})$ are upper hemicontinuous. In addition, $s_{i,t}(R, \tilde{C}_{i,t})$ is weakly increasing in R. If the aggregate resource constraint is binding $b_{i,t}(R, \tilde{C}_{i,t})$ is weakly increasing in $\tilde{C}_{i,t}$. If the aggregate resource constraint is not binding, $s_{i,t}(R, \tilde{C}_{i,t})$ and $b_{i,t}(R, \tilde{C}_{i,t})$ are independent on $\tilde{C}_{i,t}$.*

Proof. Because in the utility maximization problem the objective function as well as all constraints are continuous, by the theorem of the maximum $s_{i,t}(R, \tilde{C}_{i,t})$ and $b_{i,t}(R, \tilde{C}_{i,t})$ are upper hemicontinuous. In addition, note that $s_{i,t}$ and R are complements in the objective function, and that $s_{i,t} \cdot R$ is fully deterministic in equilibrium. Therefore by Topkins's theorem $s_{i,t}(R, \tilde{C}_{i,t})$ is weakly increasing in R. Finally, $b_{i,t}(R, \tilde{C}_{i,t})$ is weakly increasing in $\tilde{C}_{i,t}$ because increasing $\tilde{C}_{i,t}$ relaxes the aggregate resources constraint and allows for higher borrowing.

□

Note that, whereas the supply for funds is increasing in R, the demand for loans may be somewhere increasing and somewhere decreasing in R. It may be increasing because an agent needs to save if she wishes to borrow, and therefore, conditional on being a net borrower, increasing the return on savings decreases the cost of borrowing. However, the demand for loans may be somewhere decreasing if, as R increases, an agent switches from being a net borrower to being a net saver.

In addition, if the resource constraint is binding, then both the demand and the supply for loans are a function of the resources available to the group. We say that a member is *rationed out* if her demand for loans is strictly increasing in $\tilde{C}_{i,t}$. Intuitively, the fact that the aggregate resource constraint is binding does not imply that all members are rationed. It may be the case that some members can fully meet their demand for loans, while the burden of rationing falls disproportionately on some other members. Hence, as additional resources are introduced into

the group, those who are rationed out increase their borrowing, while those who are not rationed out are unaffected.

Lemma 2. *If* $b_{i,t}(R, \tilde{C}_{i,t}) > 0$ *and either*

- $r < R$, *or*
- β *sufficiently low,*

then constraint 1 holds with equality and therefore

$$b_{i,t}(R, \tilde{C}_{i,t}) = 3 \sum_{s=1}^{t} s_{i,s}(R, \tilde{C}_{i,s})$$

.

Proof. Suppose that the demand and supply for funds of an agent is such that:

$$b_{i,t} < 3 \sum_{s=1}^{t} s_{i,t}.$$

Clearly, the agent could reduce $s_{i,t}$ and $b_{i,t}$ by the same amount, satisfy all constraints (including the aggregate resource constraint), and maintain the same outside investment $y_{i,t}$. Doing so would reduce the interest payment on loans, but also reduce the end-of-cycle payment from the group.

Note, however, that resources can be costlessly transfered from period 1 or 2 to period 3 by setting an appropriate $a_{i,t}$. Therefore, as long as $r > R$ the agent prefers to reduce both $s_{i,t}$ and $b_{i,t}$ up until constraint 1 is binding or $b_{i,t}$ is zero (and the project is fully self financed).

Finally, note that when $\beta = 0$, no matter the values of r or R, money in period 3 has no value, so the agent want to minimize the amount saved with the group. By continuity, for β low, the agent saves with the group the minimum amount required to borrow. □

The above lemma shows that, when $r > R$, a borrower will save the minimum amount required to meet the leverage constraint. When $R > r$ instead, whether constraint 1 is binding depends on the return on the external project, on the desire to smooth consumption, and on the discount factor. If the agent discounts heavily the future, period-3 payout has little value in period 1 or 2, and the agent only saves enough to borrow. Hence, constraint 1 is binding.

For the sake of the argument assume that there is no discounting ($\beta = 1$). Assume also that the aggregate resource constraint is not binding. The key observation is that, in this case, the maximum investment scale is reached when constraint 1 is binding. Therefore if, for example, the return on the outside investment is linear, then the agent will always invest the maximum

amount and 1 is binding.[8] On the other hand, if the only investment opportunity outside of the group has a fixed scale, then the agent may want to save in excess to what is necessary in order to borrow. Doing so, the agent is able to simultaneously start the project and capture some of the return on savings paid by the group.

Suppose instead that the aggregate resource constraint is binding. It follows that any change in the amount saved causes an equal change in the amount borrowed. As a consequence, the scale of the investment is independent on the amount saved with the group. The decision to save is determined by solving for a trade off between higher returns tomorrow, and desire to smooth consumption over time. Remember that $a_{i,t}$ should be non negative. Hence, it is always possible to transfer resources from the present to the future, but the opposite is not true. When $R > r$ total cash is higher if the agent borrows and saves at the maximum. However, saving the maximum may cause consumption in a given period to drop relative to consumption in the future, which may not be optimal.

Equilibrium Despite being taken as given by the group's members, R and $\tilde{C}_{i,t}$ are determined in equilibrium. In particular, the equilibrium $R = R^\star$ solves:

$$R^\star \left(S_1(R^\star) + S_2(R^\star)\right) = r \left(B_1(R^\star) + B_2(R^\star)\right) \tag{5}$$

where $S_1(R)$, $S_2(R)$, $B_1(R)$ and $B_2(R)$ are the aggregate demand and supply of funds, which depend on R because the individual demand and supply depend on R.

Note that, whereas the individual demand and supply for funds depend both on R and on $\tilde{C}_{i,t}$, the aggregate demand and supply for funds only depend on R. The reason is that, in the individual maximization problem, $\tilde{C}_{i,t}$ matters only if the aggregate resources constraint is binding. Furthermore, the aggregate resource constraint is either binding for everybody or not binding for anybody. Therefore, when looking at the aggregate, we can simply distinguish between R for which the aggregate resource constraints is binding and R for which the aggregate resource constraint is not binding. Whenever the aggregate resource constraint is not binding, aggregate borrowing depends on aggregate savings only through the equilibrium $R^\star$. Instead, in periods in which the aggregate resource constraint is binding, aggregate borrowing depends on aggregate savings directly. In particular, if the period-1 aggregate resource constraint is binding, then $B_1(R) = S_1(R)$. If the period-2 aggregate resource constraint is binding, then $B_2(R) = rB_1(R) + S_1(R) + S_2(R)$.

Distinguishing between periods in which the aggregate resource constraint is binding or not will be relevant when performing our comparative statics analyses. For example, assume that a member of the group drops out and is replaced by another person with a higher propensity

[8] Because, in our groups, the scale of outside investments is small, a linear return on investment can be seen as a local approximation of a concave return on investment.

to save.[9] If the resource constraint is binding, we can solve for the new equilibrium simply by shifting upward $S_1(R)$ and $S_2(R)$. If instead the aggregate resource constraint is binding, then aggregate borrowing will also respond to an increase in overall funds available to the group.

Proposition 1. *An equilibrium $R^\star$ always exists. If multiple $R^\star$ exist, the largest $R^\star$ Pareto dominates all the other $R^\star$. At the Pareto optimal $R^\star$, if both $S_1(R^\star)$ and $S_2(R^\star)$ are unique, then the LHS of equation 5 crosses the RHS of equation 5 from below.*

Proof. As a preliminary step, note that the aggregate demand for savings and aggregate demand for loans inherit the properties of the individual demand for savings and loans derived in lemma 1. Also, the largest $R^\star$ Pareto dominates all other $R^\star$. The reason is that $S_1(R)$ and $S_2(R)$ are increasing in R, which implies that at the largest $R^\star$, net borrowers can borrow the same amount as at the lower $R^\star$. At the same time at the largest $R^\star$ everybody enjoys a higher return on their savings compared to lower $R^\star$.[10]

Note also that both $S_1(R)$ and $S_2(R)$ are bounded above by $\sum \min_i\{w_{i,t}, \overline{s}\}$. It follows that $B_1(R)$ and $B_2(R)$ are also bounded above. Hence, for R sufficiently large:

$$R(S_1(R) + S_2(R)) > r(B_1(R) + B_2(R))$$

Therefore, if an equilibrium exists, then at the largest equilibrium R we have:

$$\frac{\partial\left[R(S_1(R) + S_2(R))\right]}{\partial R}|_{R=R^\star} > \frac{\partial\left[r(B_1(R) + B_2(R))\right]}{\partial R}|_{R=R^\star}$$

To prove the existence of the equilibrium, consider an R arbitrarily close to zero. If R is sufficiently small, then $R(S_1(R) + S_2(R))$ is approximately zero. However, both aggregate savings function are strictly positive, implying that $B_1(R) + B_2(R)$ are positive (if the return on the outside investment exceeds r for some agents) or zero. Hence, for R small

$$R(S_1(R) + S_2(R)) \leq r(B_1(R) + B_2(R))$$

We previously have established that, for large R:

$$R(S_1(R) + S_2(R)) > r(B_1(R) + B_2(R))$$

Together with the fact that all functions are upper hemicontinuous, these results imply that an

[9] If the rules of the group are chosen by majority voting, then changing the composition of the group does not affect the rules adopted by the group as long as the median member of the group does not change. Hence, we can analyze changes in the demand and supply of funds due to a change in the group's composition keeping the rules adopted by the group constant.

[10] In case the resource constraint is binding, the borrowing profile within the group is not fully specified. However, if the aggregate resource constraint at the largest $R^\star$ is binding, there exists a borrowing profile at the largest $R^\star$ that Pareto dominates any other borrowing profile at any other $R^\star$.

equilibrium exists. □

In what follows, when multiple equilibria exist, we focus on the Pareto preferred equilibrium.

Increase in aggregate savings Suppose that the aggregate savings in both periods increases. This could be the case if, as discussed earlier, a member of the group who only saves drops out of the group and is replaced by another agent who also only saves but has a larger propensity to save at every R. Clearly, if the resource constraint is never binding, then, for given R, the increase in aggregate savings has no effect on aggregate borrowing. Hence, the behavioral responses of the group's members is driven by the fact that, by proposition 1, when $S_1(R) + S_2(R)$ shifts upward $R^\star$ decreases.

Corollary 1. *If the aggregate resource constraint is never binding, then an increase in aggregate savings leads to lower $R^\star$ and lower individual savings. Everybody in the group is worse off.*

If instead the aggregate resources constraints is binding in both periods, adding resources to the group has also a direct effect on the borrowing levels that are possible within the group.

Lemma 3. *Suppose that the aggregate resource constraint is binding in both periods. Furthermore, suppose that $S_1(R)$ and $S_2(R)$ increase by the same factor. Each member's borrowing (weakly) increases and everybody in the group is (weakly) better off.*

Proof. If both aggregate resources constraints are binding, then at $R = R^\star$:

$$R(S_1(R) + S_2(R)) = r((2 + r)S_1(R) + S_2(R))$$

$$S_2(R)(R - r) = S_1(R)(r(2 + r) - R)$$

Which implies that increasing both $S_1(R)$ and $S_2(R)$ by the same factor does not change $R^\star$. At the same time, higher $S_1(R)$ and $S_2(R)$ relaxes the aggregate resource constraint. Net savers and borrowers who are not rationed out are indifferent, while borrowers who are rationed out increase their borrowing and are better off. □

Corollary 2 and lemma 3 illustrate one of the main results of the model: that exogenously increasing the funds available to the group (for example, by replacing one of the members of the group) will impose an externality on the other members of the group. The key determinant of the sign of this externality is whether the group is resource constrained. Quite intuitively, when the resources within the group are scarce, adding more resources is beneficial to the other members of the group. More interestingly, when the group is not resources constrained, adding resources to the group hurts everybody else in the group because it decreases the return on savings.

When the resource constraint is binding in only one period, the overall welfare effect of adding resources to the group is ambiguous. All members are made worse off by the addition of extra funds. However, net borrowers who are rationed out benefit from the availability of extra funds.

Increase in aggregate borrowing We can similarly analyze what happen when the group composition changes in a way that increases aggregate borrowing leaving unchanged aggregate savings. This would be the case if, for example, a net saver is replaced with a net borrower who saves the same amount, but uses these savings to actually borrow funds from the group.

If the aggregate resource constraint is never binding, by proposition 1, $R^\star$ increases following an increase in borrowing, leading to the following corollary.

Corollary 2. *If the aggregate resource constraint is never binding, then an increase in aggregate borrowing leads to an increase in $R^\star$, higher individual savings and borrowing. Everybody in the group is better off.*

If instead the aggregate resources constraints is binding, then the impact of an increase in aggregate borrowings depends on how the funds are rationed among borrowers. For example, if the new demand for funds goes completely unmet, then the existing member of the group are indifferent to the increase in the demand for funds. If instead the addition of a borrower decreases the amount of funds available to the other borrowers, then the existing borrowers are made worse off by the increase in the demand for funds.

Remark 1. *If the resource constraint is binding in both periods, an increase in the demand for loans has no effect on $R^\star$. As a consequence, everybody in the group is weakly worse off.*

Similarly, if the resource constraint is binding in one period but not the other, the welfare effect of increasing the demand for funds is ambiguous. On the one hand, $R^\star$ increases and everybody benefits. On the other hand, net borrowers may be hurt by the fact that rationing is now worse.

Overall, increasing aggregate borrowing and increase savings have opposite effects on the group. When the aggregate resource constraint is binding, increasing savings makes the group better off while increasing borrowing makes the group (weakly) worse off. When the aggregate resource constraint is not binding, increasing savings makes the group worse off, while increasing borrowing makes the group better off.

Supply of funds over time There is an additional dimension that is relevant in determining the efficiency of the group: the timing of saving. Suppose that cumulative aggregate savings $S_1(R) + S_2(R)$ are constant, but the group can substitute one or more members, so that the timing of these borrowing and savings changes. It is quite immediate to see that, as long as the aggregate resource constraint is never binding, this reallocation of borrowing and savings has no impact on the return on savings.

Instead, suppose that the aggregate resource constraints is binding in period 1, and savings are reallocated from the second to the first period. If the period-1 demand for loans is rationed, then this reallocation increases the loans given out in period 1. In addition, all these loans will

be repaid at the end of period 1. So, for every dollar that is reallocated from period 2 to period 1, $1 + r$ dollars become available in *period 2*. Hence, if the resource constrain is binding also in period 2, this reallocation eases rationing in period 2 as well.

Remark 2. *Suppose the resource constraint is binding in period 1. Suppose that $S_1(R)$ increases and $S_2(R)$ decreases by the same amount. It follows that $R^\star$ increases, and all agents increase their level of borrowing and savings. All agents are better off. If instead the resource constraint in period 1 is not binding, reallocating funds from one period to the other has no impact on $R^\star$ and no impact on the group members' welfare.*

Hence, contrary to changing the level of savings or borrowing, changing the timing of savings or borrowing has a unambiguous welfare effect.

Empirical implications In our experiment, which we fully describe below, we exogenously varied the number of vulnerable participants in a SBG. The vulnerable differ from the non-vulnerable in a number of important dimensions, which will be described in greater detail in section 7.1. However, for the purposes of our theory, the key observation is that vulnerable households save less than non-vulnerable households. Among other relevant dimensions, including demand for loans and likelihood of default on loans, the vulnerable and the non-vulnerable are indistinguishable. In our model, more vulnerable participants should therefore reduce the supply of funds available to the group. If the aggregate resource constraint is not binding, the reduction in savings should lead to an improvement in realized returns on savings, and all members who are not saving the maximum amount should save more.

If, on the other hand, the resource constraint is binding, the reduction in savings exacerbates the scarcity of funds. The response of savings and borrowing depends crucially on the rationing mechanism. Since the realized return on savings is unchanged, those whose borrowing decisions are not rationed will see no difference in their savings and borrowing behavior. On the other hand, those who are rationed will further reduce their own borrowing. Because they can borrow less for a given amount saved, they have an additional disincentive to save. That is, they exacerbate the already present scarcity of funds by cutting back on additional savings.

4 The Intervention

We now empirically test some of the predictions of our model using data collected as part of a randomized experiment in rural Uganda. Our experiment takes place in the context of an expansion of project SCORE in 90 mostly rural villages not previously served by the project. This expansion was widely geographically dispersed, involving 28 districts in the Western, Central, Eastern and Northern regions, starting in January 2013. In each study village, SCORE representatives worked with local administrators to identify households that were categorized

as "vulnerable" by the local government. Identified households were further screened through a questionnaire to ensure that they met the vulnerability criteria set by the project, and their willingness to join a savings group was also assessed. Among the selected households, 14 in each village were chosen to be part of the study. We refer to these participants as the *pre-selected* participants. All pre-selected participants enrolled were given access to other SCORE services, none of which involved transfer of material or financial resources.

Using the information contained in the screening questionnaire, the research team randomly assigned 30 villages to a *sparse treatment*, and the remaining 60 villages to a *dense treatment*. In dense treatment villages, one savings group was formed, comprising of all 14 pre-selected participants. In the sparse treatment village, pre-selected households were evenly divided into two separate savings groups. Since the size of the group was initially capped at 27, the experiment generated variation in the quantity of community members allowed into the groups: only 13 community members were allowed in dense groups, whereas 20 were allowed in sparse groups. See figure 1 for a simple schematic of the protocol. To avoid treatment spillovers, villages located in the same larger administrative unit (the parish) were assigned to the same treatment.

Within the constraints set by the randomization, group formation followed the standard procedure established by SCORE. Field officers engaged the local community to generate interest in the savings groups through presentations in local markets, churches, and community-based organizations. Interested members of the community were then given the opportunity to join the savings group. Apart from the quota and the geographical separation of dense and sparse groups, we imposed no further restrictions on the formation of study VSLAs in either of the two treatment arms. In the end, we expected to have 60 VSLAs in which at least 50% of the members belong to a vulnerable household, and 60 VSLAs in which around 25% of the members belong to a vulnerable household.

The protocol for the intervention was followed closely by our implementing partners: villages assigned to the low-threshold treatment indeed had two sparse VSLAs, and villages assigned to the high-threshold treatment had one dense VSLA. We can see in figure 2 that the makeup of the groups follow the intended assignment.

Table 4 provides some measures of attrition. 113 (56 dense and 57 sparse) of the expected 120 VSLAs were formed. We find that 75% of the preselected households that were assigned to a study VSLA are indeed enrolled in a study VSLA. While attrition is high, there is no evidence that it was differentially worse in one of the two treatment arms. We also find that many of the preselected leavers decided to pull out from the VSLA and, to compensate, the field officers in charge of forming the groups replaced those who pulled out with other vulnerable individuals.[11]

Our intervention can be thought of generating several differences between dense and sparse groups. First, sparse groups have a larger number of self-selected members and, provided that

[11] The average characteristics of those who pull out are generally similar to those who are drawn in as replacement. Results are available upon request.

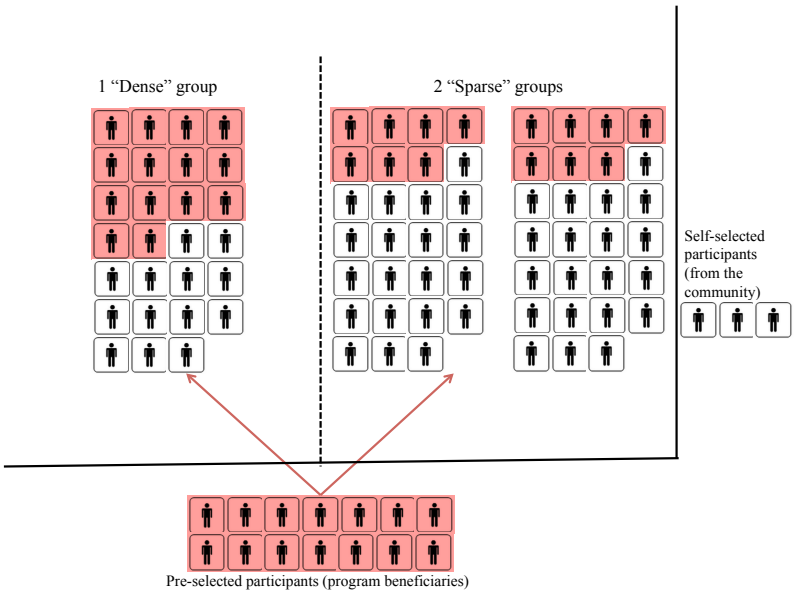

Fig. 1: Randomization strategy–Group formation stage

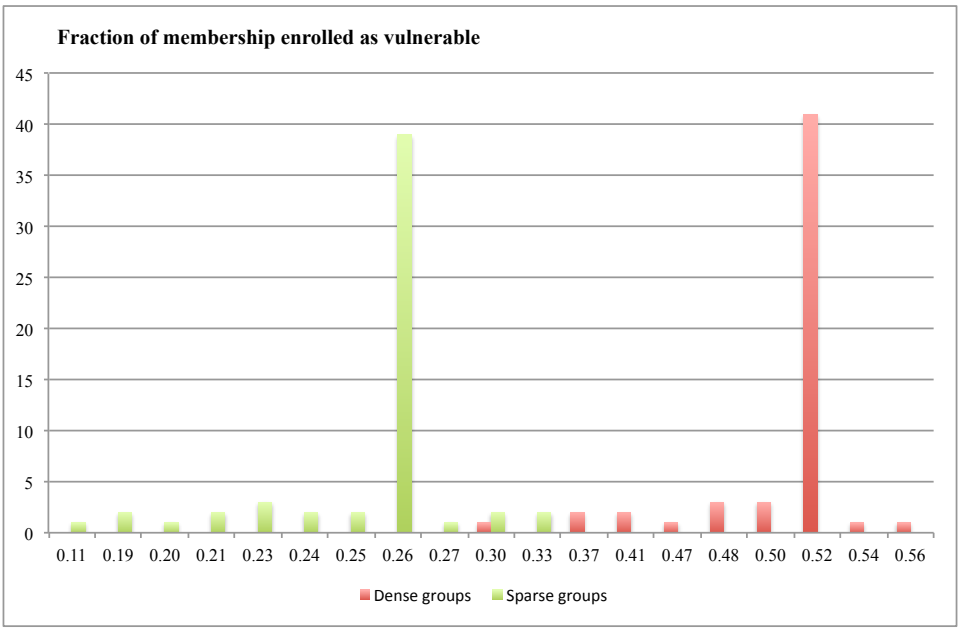

Fig. 2

these are less vulnerable than the preselected participants, the overall vulnerability in these groups should be lower. As a consequence, the peer group of the preselected participant who is randomized into a sparse group should be better off than that of the equivalent person randomized into a dense group. Second, it is conceivable that sparse groups have more social cohesion than dense groups, because there is more self-selection into dense groups, and this could drive some of the results we see. While we abstract from this issue for now, we do have data to verify this and plan to study this issue in future versions of the paper. Third, our intervention may have generated differential general equilibrium effects in dense and sparse villages, because the number of SBG in dense villages should be smaller than in sparse villages. In the presence of within-village spillovers from savings groups, this could differentially influence the decision to save or to borrow of pre-selected participants. For instance, if VSLA participants are conduits of loanable funds to the rest of the community, the larger number of participants involved in savings groups in sparse villages could potentially reduce the average amount of funds borrowed in those groups. However, it is important to remember that SCORE created many additional non-study SBG in the study areas. In addition, other non-SCORE SBGs are operating in most villages. Using data from a census of SBGs carried out after our study groups were established, we show that the total number of SBGs is not statistically different between dense and sparse areas. In addition, dense and sparse areas have the same amount of access to financial services (i.e., no local access to MFI and other banks, and statistically insignificant differences in terms of local credit unions or SACCOs). See table 5 for details.

5 Data

Our data consists of household surveys from a random sample of group participants, and financial records (amounts saved, borrowed, and defaulted) of all participants during the first cycle.

Household surveys Household data was collected in two periods. First, at the onset of the study, pre-selected participants were administered a standard questionnaire developed by project SCORE to identify and screen in families with vulnerable children. This vulnerability tool contains questions regarding the family socioeconomic status (such as the condition of the main house) and the well being of the children belonging to the household (including disability and HIV status, as well as history of physical or mental abuse). The information contained in this vulnerability tool represents the baseline data in our study, and this is the data that was used to randomize villages into the two treatment arms.

The rest of the study participants were identified and interviewed after the groups were formed. Two months after most groups were created, a team of enumerators carried out in-depth interviews with a random sample of participants. The questionnaire included the vulnerability assessment tool used for screening, as well as additional questions aimed at determining respon-

dents socioeconomic status.

By combining the data from the preselected participants with those collected from the self-selected participants we are able to construct a profile of characteristics for groups in the two treatment arms. Note that this profile is constructed from some pre-intervention data, and some post-intervention data. To the extent that the variables measured are time invariant, this should make little difference. Nonetheless, one should be careful in interpreting differences across treatment arms as being purely baseline differences.

VSLA administrative records The main outcome variables of interest come from administrative records collected by SCORE field officers during their regular financial audit of all groups. Auditing records contain information on how much each group member saved, borrowed, and repaid up until the audit date. In addition, administrative records include information regarding the rules set by the group, such as the interest rate charged on loans, minimum and maximum savings allowed per meeting, and length of the borrowing cycle.

The present study combines information from two audit rounds. Round 1 occurred approximately half way through the cycle (September through November 2013), while round 2 took place mostly from February to April 2014, when almost all groups had stopped disbursing loans and most completed or were close to completing their cycle.

Summary statistics Table 3 reports the summary statistics for 20 household characteristics obtained from the screening questionnaire among all pre-determined participants. The table shows mean values for dense and sparse VSLAs separately, and conducts a t-test of the difference in the means (last column). Predetermined participants (who are enrolled in SCORE program) clearly represent very poor and very disadvantaged members of the community. For instance, approximately one fifth of participants report having a child who is involved in child labor, 17% have children with disabilities, 70% report missing meals, and the average reported monthly household income is USH 41,000 (approximately $16). The table also indicates that average characteristics do not differ significantly between the two treatment arms, indicating that the groups are balanced in the way it was intended.

Panel B considers an expanded set of financial variables from the household surveys carried out after the groups were formed. Since the survey was conducted on a random selection of savings group members, the sample size utilized is much smaller. The comparison suggests some marginally significant differences in savings groups participation and in land ownership: In particular, those randomly assigned to dense villages were less likely to participate in other savings groups, and had larger land sizes. Inclusion of these controls in the regressions does not significantly alter the estimates of the paper.

6 Empirical strategy

Our empirical objectives are two fold. First, we would like to establish whether our intervention changed the composition of the savings groups, and if so, in which way. Based on the empirical observation that fully self-selected groups generally attract better-off members, we expect that the self-selected population have lower vulnerability scores than our preselected participants. If indeed we find differences in the composition of the groups across treatment arms, we turn to testing some implications of the theoretical model by studying how vulnerable households' saving and borrowing behavior differ across treatment arms. We measure saving and borrowing behavior as the total amount saved and total amount borrowed by person i in SBG group g from the start of the cycle until audit period t, and denote them by y_{ig}^t. The empirical specification is

$$y_{ig}^t = \alpha Dense_g + X_g^t \beta_1 + X_i \beta_2 + \epsilon_{ig}^t. \tag{6}$$

In the equation, *Dense* is an indicator variable that identifies those who joined groups designated as more vulnerable; X are group level and individual level controls which include the interest rate, the share price, the cycle start date and the number of meetings since the start date until the audit period t (where, as mentioned above, we have two audit periods). α is the coefficient that describes the difference in outcomes between those enrolled in dense groups relative to those enrolled in sparse groups. This coefficient then gives the overall effect of the intervention, but owing to the fact that some members are self-selected, it does not represent a causal relationship.

Our main specification estimates (6) on the subsample of preselected individuals. Since the preselected participants were randomly assigned to dense or sparse groups, they should not have observed or unobserved differences in their willingness to save or to borrow. Controlling for the interest rate and the share value, then, any difference between dense and sparse (captured by the coefficient α) measures the individual members' response to the composition of the group. In addition, given our randomization, we expect individual-level characteristics X_i not to be significant once we control for the type of group and its rules.

We finally estimate (6) on the subsample of self-selected participants. This group includes individuals who may or may not be vulnerable, but which we expect to be on average less vulnerable. More importantly, we do not think that the observed and unobserved characteristics of this group to be independent of the assignment to dense groups. While the estimated α thus cannot be considered causal, it will be informative to show how the estimate changes with the inclusion of controls X_i.

A few technical notes on the estimation strategy are required. First, due to noncompliance, we do not consider whether the final distribution in a VSLA was dense or sparse, but rather we focus on the initial determination of a group as being dense or sparse. That is, we rely

on intent to treat estimates. Second, groups were formed by different implementing partners (IPs). IPs are an important input in VSLA outcomes, and our sample includes several IPs with different degrees of expertise in forming and managing groups. Since groups depend crucially on the expertise of the IP to ensure that the accounts are kept correctly and meetings develop according to the rules, we should expect that savings and borrowing rates will differ significantly across IPs. Thus, all regressions are estimated with and without IP fixed effects. Third, as we will show in the next section, not all pre-selected households joined their assigned groups, and some self-selected households are vulnerable. When describing the data, we often compare vulnerable households enrolled in SCORE with the other members of the groups. When making causal statements, we compare per-selected households who joined their assigned groups across different type of VSLA. Finally, it is not only likely that outcomes are autocorrelated within each group, and therefore errors are clustered at the group level.

Our theory suggests that those who are matched with people who borrow more should also save and borrow more. Here we test this prediction by provide preliminary evidence on the impact of the intervention on savings and borrowing four to six months after the group has been formed.

7 Empirical analysis

7.1 Difference between vulnerabile members enrolled in SCORE and all other members.

The definition of vulnerability followed by project SCORE is multi-dimensional and not limited exclusively to economic vulnerability. In table 6 we show that project beneficiaries have, on average, higher measures of vulnerability than those who are not directly enrolled in project SCORE. In other words, the screening process adopted by SCORE seems to capture some coarse dimensions of vulnerability, albeit not all.

For the analysis, we restrict our sample to those participants who were enrolled in a group at the time of the first meeting and to those who enrolled later; we exclude those who initially intended to participate but ended up not enrolling officially. Table 6 regresses thirteen coarse characteristics from the vulnerability screening tool and our household survey on an indicator variable that identifies those members of the group enrolled in SCORE.[12] Across almost all measures, vulnerable households have higher measures of poverty and lower socioeconomic outcomes. They have higher rates of physical or mental disability in the household, are more likely to miss meals, have lower access to latrines, and are more likely to be considered in need by the assessor who carried out the interview;[13] they are less connected financially to other savings groups or

[12] We also ran the same analysis with being a preselected dummy as the independent variable, and the results are very similar.

[13] Assessor scale is a subjective evaluation of the interviewer, with higher numbers corresponding to worse

bank accounts (columns 8-11); and have both lower levels of assets (as measured by an index of assets, column 12), and a lower reported per capita income (column 7)[14] In addition, point estimates suggest that the average amount of land owned by vulnerable households is larger than the average for other participants. While this difference is not significant, it is suggestive that vulnerable households are more likely to be farmers than non vulnerable households. Finally, for reasons that are not entirely clear, beneficiaries are significantly more likely to report to have access to a safe source of water.

Table 7 studies the determinants of borrowing and savings, and shows how the demand and supply of funds varies between program and non-program beneficiaries. Here, we use the administrative records of each account holder. For each person i in accounting period $t \in \{1, 2\}$, we regress her accumulated amount borrowed, accumulated amount saved, whether the person ever borrowed, and whether the person is late with the payment of the loan on whether the person was enrolled as a vulnerable SCORE beneficiary. To isolate individual factors as opposed to the conditions in the group, we control for group characteristics by including VSLA fixed effects. As can be seen in panel A, beneficiaries save less than non-beneficiaries (columns 1 and 2). They also borrow somewhat less, but this is never statistically significant. In addition, they are neither less likely to take up loans, nor more or less likely to default. Thus, the vulnerable seem to have a demand for loans that is similar to non-beneficiaries, and–given that both loan amounts and default rates are almost identical to non-beneficiaries–they are as credit-worthy. In conclusion, the main difference seems to be that the vulnerable save less.

7.2 First stage: Savings group composition

Our estimation strategy rests on the assumption that, because we assigned more vulnerable participants to dense groups, these groups are on average more vulnerable than sparse groups. This assumption must be tested. The reason is that different people may self-select into sparse group than into dense group, potentially making the vulnerability profile of the two types of groups indistinguishable.

Table 8 compares the averages of the thirteen characteristics between the dense and sparse groups. The coefficients in this table should be related with those in table 6 in a very specific way: absent any differential composition of self selected members in dense and sparse groups, all coefficients should be 1/4th the size of the coefficients in table 6. It is indeed the case that many coefficients have this rough proportion, even though only a few are statistically significant. On average, participants of dense groups are more likely to skip meals, have higher subjective measures of vulnerability as reported by assessors, and have a higher amount of land. Impor-

socio-economic situation.

[14] The difference with the comparison group is less than 2,000 shillings per person, or approximately 75 US cents. On average, the reported monthly income for non-beneficiaries is 10,800UGX, so beneficiaries' reported income is 18% lower. It is very likely that all income is underreported.

tantly, some coefficients (like the wealth index coefficient, land, access to safe water, and access to savings groups) have a very different proportion. This is evidence that community members who join a dense group are somewhat different than those who join a sparse group.

7.3 Intervention outcomes

7.3.1 Share price and interest rates

At the group formation stage, groups must decide on the interest rate and on their share price, that implicitly determine a limit on weekly savings. Qualitative interviews suggest that groups make this decision with very limited information on what would be a reasonable rule, with many groups eventually regretting their choice and adjusting the rules in the following cycle. If that were the case, we would expect our intervention to have no systematic effect on the decision over the rules. Table 9 rigorously tests this. Positive coefficients indicate higher interest rates on average in dense groups, but the results are not statistically significant. Alternative specifications, including logit models for setting high interest rates, yield the same insignificant results (not shown). We are left to conclude that, in the short run, our intervention does not impact the interest rates set by the group. Our preferred explanation is that the groups resorted on the default set of rules proposed by the organization training the group (the Implementing Partner or IP). In table 9, the R^2 of the regression increases dramatically when including IP fixed effect, which implies that a large fraction of the variation in interest rate and share price can be attributed to the IPs. The expectation is that groups will adjust their interest rates in following cycles.

7.3.2 Group performance

We are now ready to explore the effects of selection on overall group performance. Table 10 estimates total group savings on the explanatory variables using regression (1), with data aggregated at the group level. Across the four specifications (columns 1-4), dense groups have significantly lower levels of aggregate savings. The point estimate in column 4 indicates 767,453 fewer shillings in dense groups ($280, or over $10 per person), or approximately one third of average savings in the sample. Dense groups have, therefore, a lot fewer resources to lend out.

What are the implications for loan transactions? The additional savings either generate new loans for the group, or remained unused. If all resources are lent out, then on average we should expect the loan to savings ratio to remain constant, and the implied returns to savings would not change. If, on the other hand, the additional savings are not used to generate new loans, then the loan to savings ratio would decrease and so would the implied returns to savings. Observing systematic movements in the loan to savings ratio would therefore shed light on whether the increase in savings erodes the value of those savings. Columns 5-8 show that, if anything, loan to

savings ratios are lower in dense groups than in sparse groups, which suggests that, if anything, sparse groups have higher rates of utilization of their loans. However, the result is not robust to the inclusion of controls. It is reasonable to conclude that, at best, all groups are utilizing funds at the same rate.

7.3.3 Individual savings and borrowing choices

Our experiment allows us to shed direct light on the impact of group composition on the individual incentive to save (and borrow). We report in our next two tables the results of regressing (1) on individual savings and loan positions, taken in mid cycle (wave 1) and towards the end of the cycle (wave 2). Panel A of table 11 shows that participants in the dense VSLAs save significantly less than participants in sparse VSLAs, once the direct effect of the implementing partner and the length of time the VSLA has been in operation are taken into account. As expected, the magnitude of effects largely reflect those from table 10–on average, members of more vulnerable groups save 30% less.

In panel B, we thus consider only preselected participants, who (as we have seen) have similar average characteristics across dense and sparse groups. If group composition had no impact on their savings behavior, or if the limiting factor is simply the amount saved by self-selected members, we should see no difference in the savings of this particular subgroup. Columns 3 and 4 show that early on the pre-selected in these dense groups save approximately 6,200 shillings less than their counterparts in sparse groups, but these results are not statistically significant. However, by the end of the cycle, the difference between the two treatment arms increases to 14,700 shillings and becomes statistically significant. Given that the preselected in sparse groups have saved an average of 74,200 shillings by the end of the cycle, this represents a 19.9% “penalty".

In panel C we focus on self-selected participants. Perhaps surprisingly, their rates of savings are somewhat lower in dense groups, but only in wave 1 we see some level of significance in the estimates. Point estimates are somewhat lower for this group than for the preselected; however, it is important to remember that these groups are not directly comparable across treatment arms because they are differently composed. Nonetheless, at best, savings decisions for the self-selected are marginally affected by selection.

Table 12 regresses total borrowing on the explanatory variables. On average, dense groups generate fewer loans than sparse groups, although this shows up with some significance only in the second wave (panel A), and in our most robust specification the coefficient is only close to significance. However, borrowing decisions for the preselected differ sharply from those of the self-selected, as panel B and C make clear. The preselected in dense groups are rationed out of loans immediately and by the end of the cycle they have borrowed 65,000 fewer shillings ($23), or 44% less, than their counterparts in sparse groups. This is a severe reduction in their borrowing, one that is more than proportional to the one in savings. In sharp contrast, self-

selected participants are not borrowing at lower levels. Thus, the entire burden of operating at lower scale seems to befall on the vulnerable preselected participants.

Several possibilities could explain these results. Our preferred explanation is that lending is rationed in these groups, and members cannot borrow as much as they would like. In groups that are more liquid, the borrowing constraint is relaxed and, on average, members borrow more. In addition, increasing borrowing requires increasing savings, then consequently they will save more. In addition, the fact that we observe an heterogeneous response to the tightness of the credit constraint between preselected and self-selected suggests that the rationing mechanism is such that the vulnerable are penalized. In fact, this is entirely consistent with informal interviews, which suggest that when funds are scarce, priority is given to those who save more.

We can explore and discard other alternatives. One possible explanation is that beneficiaries are able to save, but their demand for loans is lower than non-beneficiaries. If that is the case, having more preselected participants should reduce the implicit rate of return from savings, and everyone should save less. We do not observe this. Another possibility is that potential borrowers, unable to get loans of the necessary scale in dense groups, choose to forgo loans. This could occur for instance in cases where investment opportunities are indivisible goods. If this were the case, then we should observe that fewer people take up loans in dense groups relative to sparse groups. We carry out two checks for this. First we regress an indicator variable for taking up a loan of any size on the usual set of controls. Estimated coefficients are negative and are statistically significant at the 10% level for pre-selected participants, but only if the interest rate and the price of a share are not taken into account. If there is an extensive margin, it certainly does not seem particularly important. We next re-run tables 8 and 9 after reducing the sample to only those who borrowed. The coefficient estimates are very similar to those in table 8 and 9. Again, this confirms that the intensive margin is the relevant margin.

7.4 Discussion

Our research finds evidence that the composition of savings groups matters in the decision of saving and borrowing; in particular, VSLAs with poorer members save less and borrow less. This is not simply due to the large number of poor people in poor groups: when we control for selection, we find that poor people save and borrow less when they are in poor groups. That is, selection creates a negative incentive to save and borrow.

It should be emphasized at this point that there are alternative mechanisms that could be driving the results. For instance, in dense groups the scope for self-selection is quite limited, which would possibly mean that there are fewer pre-existing social connections within those groups. The lower savings rates/borrowing rates might thus be a reflection of the lack of embedded trust at this early stage. As part of our baseline activities, we collected information on existing social ties within the group, so we will be able to shed light on this mechanism. It should be noted

that, if the driving mechanism is indeed the (lack of) within-group social trust, then over the long run we should see the patterns of savings and borrowing reverse. This is because program beneficiaries are involved in other joint activities under SCORE, and if this leads to more trust within program beneficiaries (Feigenberg, Field, Pande, Rigol, and Sarkar (2014)), the buildup is proportionally larger in dense group.

A second possibility is that members of the savings groups operate as mere depositors for other community members' savings. Given that villages assigned to sparse groups have more community members directly involved in a savings group, fewer community members channel their savings through other VSLA participants. Of course, this channel would arise only if dense villages had fewer savings groups than sparse villages; in reality, we showed that this is not the case.

Some alternative channels are tougher to separately identify. For instance, cross-group learning spillovers could increase the savings and borrowing rates in sparse groups. Even absent such spillovers, it is possible that vulnerable group members are more confident about the safety of their own savings in groups with many non-vulnerable. We will attempt to find ways to address these issues later on.

References

Angelucci, M., D. Karlan, and J. Zinman (2015). Microcredit impacts: Evidence from a randomized microcredit program placement exper iment by compartamos banco. *American Economic Journal: Applied Economics Forthcoming*.

Ashe, J. (2009). Saving for change: Low cost, mass-scale, self-replicating and profitable microfinance for the rural poor. Technical report, Oxfam america.

Banerjee, A. V., E. Duflo, R. Glennerster, and C. Kinnan (2013). The miracle of microfinance? evidence from a randomized evaluation.

Beaman, L., D. Karlan, and B. Thuysbaert (2014). Saving for a (not so) rainy day: A randomized evaluation of savings groups in mali. Technical report, National Bureau of Economic Research.

Bundervoet, T. (2012). Small wonders? a randomized controlled trial of village savings and loans associations in burundi. Technical report, International Rescue Committee.

Comola, M. and S. Prina (2014). Do interventions change the network? a panel peer-effect model accounting for endogenous network changes. *IZA Working Paper 8641*.

Duflo, E., C. Kinnan, A. Banerjee, and R. Glennerster (2013). The miracle of microfinance? evidence from a randomized evaluation. *NBER working paper* (18950).

Dupas, P. and R. Jonathan (2013). Savings constraints and microenterprise development: Evidence from a field experiment in kenya. *American Economic Journal: Applied Economics 5*(1), 163–192.

Dupas, P. and J. Robinson (2013). Savings constraints and microenterprise development: Evidence from a field experiment in kenya. *American Economic Journal: Applied Economics 5*(1), 163–92.

Feigenberg, B., E. Field, R. Pande, N. Rigol, and S. Sarkar (2014). Do group dynamics influence social capital and female empowerment? experimental evidence from microfinance. *Journal of Policy Analysis and Management 33*(4), 932–949.

Greaney, B., J. P. Kaboski, and E. Van Leemput (2013). Can self-help groups really be" self-help"? Technical report, National Bureau of Economic Research.

Karim, L. (2011). *Microfinance and its discontents: Women in debt in Bangladesh*. University of Minnesota Press.

Ksoll, C. and R. F. Forskningsenhed (2013). *Impact of Village Savings and Loans Associations: Evidence from a Cluster Randomized Trial*. Rockwool Foundation Research Unit.

Lingon, E., J. P. Thomas, and T. Worrall (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *The Review of Economic Studies 69*, 209–244.

Vanmeenen, G. (2010). Savings and internal lending communities (silc), voices from africa: The benefits of integrating silc into development programming. Technical report, Catholic Relief Services.

Balance of treatment arms: comparison of preselected in dense and sparse villages

	Dense villages		Sparse villages		Difference	
	Mean	Std. Dev.	Mean	Std. Dev.	Difference	T-Statistics
A: Baseline (screened)						
Child labor	0.247	0.432	0.189	0.392	0.059	1.172
Drug abuse at home	0.105	0.307	0.058	0.233	0.048	1.291
Chronic disease	0.173	0.379	0.141	0.348	0.033	0.874
Disability in household	0.403	0.491	0.368	0.483	0.035	0.617
Sometimes missing meals	0.705	0.456	0.649	0.478	0.056	0.811
Quality diet	2.256	0.872	2.159	0.795	0.097	0.774
Number of meals per day	1.141	1.029	1.071	1.019	0.069	0.319
Informal employment	0.541	0.499	0.646	0.479	-0.105	-1.420
Household unemployed	0.158	0.365	0.149	0.357	0.009	0.155
Orphaned child	0.492	0.500	0.515	0.500	-0.023	-0.402
Safe source of water	0.673	0.469	0.632	0.483	0.041	0.393
Access to latrines	0.769	0.422	0.766	0.424	0.004	0.081
Income per capita	7442	9674	6509	7783	932.5	0.843
Assessor scale	1.906	0.612	1.966	0.527	-0.060	-0.971
Total vulnerability score	57.934	10.389	57.736	9.672	0.198	0.159
Child protection sub-score	6.468	5.713	5.874	4.935	0.594	0.819
Food security sub-score	11.684	5.636	11.736	5.447	-0.052	-0.066
Critical services sub-score	23.915	2.375	23.779	1.948	0.135	0.520
Economic situation sub-score	8.928	4.375	9.023	4.437	-0.095	-0.169
Household size	6.340	2.657	6.628	2.736	-0.288	-1.152
B: Expanded survey (subset of preselected post-intervention)						
Has bank account	0.054	0.225	0.078	0.268	-0.024	-0.824
Has Mobile Account	0.085	0.279	0.099	0.299	-0.014	-0.544
Other savings group	0.132	0.339	0.236	0.425	-0.103*	-1.870
Had Loan Bank	0.039	0.195	0.021	0.144	0.018	1.234
index of wealth	-0.037	2.327	-0.439	2.050	0.402	1.107
Size of land owned	3.233	9.616	1.821	3.078	1.412*	1.949

Panel A: variables described here are from the screening tool. The listed variables were used in the randomization.
Panel B: variables from the follow-up survey on a random sample of VSLA participants. The sample listed here include those preselected who were interviewed.
T-statistics clustered at the village level.

Fig. 3

Noncompliance

	Total	Dense VSLAs	Sparse VSLAs	Difference
Panel A: Group Composition				
Number of VSLAs formed	113	57	56	
Number of SCORE beneficiaries	1158	765	422	
Number of non-beneficiaries	1940	771	1256	
Fraction SCORE	**37.4%**	**49.8%**	**25.1%**	
Panel B: Noncompliance				
Pre-selected participants	1234	798	436	
Pre-selected participants not in group	264	164	100	
Attrition rate	*21.4%*	*20.6%*	*22.9%*	*-2.4%*
SCORE beneficiaries that were not pre-selected	204	131	86	
Fraction of beneficiaries not pre-selected	*16.5%*	*16.4%*	*19.7%*	*-3.3%*

Fig. 4

Title: Financial access in study villages after one year

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Parish Population	Number of vslas	Number of Non-SCORE Vslas	Number of SCORE Vslas	Number of SACCOS
dense	-991.8	1.196	1.594	-0.398	0.0658
	(664.1)	(1.781)	(1.785)	(0.249)	(0.114)
Constant	4,681***	4.594***	2.406***	2.188***	0.250***
	(525.9)	(0.419)	(0.399)	(0.187)	(0.0774)
Observations	89	89	89	89	89
R-squared	0.025	0.003	0.005	0.026	0.003

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Fig. 5

Vulnerability profile of savings groups

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Outcome variables	Disability in hhld	Hhd size	Sometimes missing meals	Safe source of water	Access to latrines	Assessor scale	Per capita income	Has Savings Account	Has Mobile	Other savings group	Had bank loan	Index of wealth	Size of land owned
Comparison between SCORE and non-SCORE members of savings groups													
SCORE beneficiary	0.0937***	0.0127	0.195***	0.0777*	-0.117***	0.533***	-1,941**	-0.0326*	-0.0980***	-0.255***	-0.0178*	-0.334**	0.677
	(0.0292)	(0.168)	(0.0448)	(0.0450)	(0.0274)	(0.0608)	(777.1)	(0.0182)	(0.0251)	(0.0634)	(0.0102)	(0.146)	(0.418)
Obs.	2,731	2,746	2,682	2,382	2,757	2,472	2,715	2,619	2,615	2,620	2,620	2,430	2,620
R-squared	0.141	0.042	0.205	0.348	0.122	0.200	0.149	0.113	0.104	0.327	0.060	0.229	0.139

Beneficiary dummy identifies a person enrolled through program SCORE.
Regressions include implementing partner fixed effects. Errors clutered at the parish level in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

Fig. 6

Savings behavior of SCORE beneficiaries relative to nonbeneficiaries

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Savings		Loan amount		Ever had loan		Has loans past due	
VARIABLES	Wave I	Wave II	Wave I	Wave II	Wave I	Wave II	Wave I	Wave II
Panel A: No household controls								
SCORE	-4,435***	-7,544***	-6,848	-3,557	0.018	-0.006	-0.001	0.005
beneficiary	(1,103)	(2,454)	(4,132)	(10,387)	(0.019)	(0.015)	(0.005)	(0.013)
VSLA f.e.	YES	YES	YES	YES	YES	YES	YES	YES
Observations	3,109	2,811	3,114	2,811	3,109	2,516	3,114	2,811
R-squared	0.671	0.527	0.349	0.246	0.310	0.266	0.186	0.216
Panel B: Full controls, reduced sample								
SCORE	-3,067**	-4,990	-7,876	6,864	0.005	-0.017	-0.006	0.013
beneficiary	(1,296)	(3,513)	(7,496)	(11,091)	(0.028)	(0.018)	(0.007)	(0.020)
Land Owned	-147.3	-85.61	22.82	1,473	-0.002	0.000	-0.000	0.001
	(119.4)	(331.5)	(615.8)	(1,552)	(0.004)	(0.003)	(0.000)	(0.001)
Asset index	1,838***	3,955***	2,220	7,990**	-0.002	-0.002	-0.004***	-0.010***
	(254.5)	(914.6)	(1,427)	(3,889)	(0.004)	(0.004)	(0.001)	(0.003)
Has Account	-21.80	220.3	1,315***	411.3	0.003***	0.000	0.009***	-0.004***
	(52.02)	(159.3)	(386.7)	(625.5)	(0.001)	(0.001)	(0.000)	(0.000)
Has Mobile	76.71	-456.4***	404.7	-305.9	-0.002	0.002*	-0.001	0.005
Account	(150.0)	(49.26)	(733.6)	(477.6)	(0.003)	(0.001)	(0.001)	(0.003)
Other Savings	-60.43	-1,786	3,233	438.7	0.036***	0.029	-0.001	0.007
Groups	(568.5)	(1,371)	(5,004)	(4,212)	(0.011)	(0.018)	(0.005)	(0.008)
Disability in hhld	307.9	1,083	-3,651	7,060	0.029	0.017	0.010	0.007
	(1,461)	(2,599)	(6,161)	(9,156)	(0.024)	(0.018)	(0.008)	(0.012)
Income per capita	0.0696	0.223	0.315	-0.0934	0.000	0.000	-0.000	-0.000
	(0.0447)	(0.136)	(0.321)	(0.393)	(0.000)	(0.000)	(0.000)	(0.000)
Food insecure	1,640	5,182**	6,337	22,754**	0.046*	0.012	-0.005	-0.028
	(1,329)	(2,466)	(5,522)	(9,364)	(0.023)	(0.014)	(0.009)	(0.022)
Assessor scale	-1,084	-2,345	168.8	-2,308	0.010	0.023**	-0.002	-0.010
	(1,180)	(1,846)	(3,079)	(6,040)	(0.015)	(0.010)	(0.005)	(0.010)
VSLA f.e.	YES	YES	YES	YES	YES	YES	YES	YES
Observations	1,811	1,689	1,814	1,705	1,811	1,504	1,814	1,705
R-squared	0.680	0.900	0.327	0.373	0.333	0.348	0.259	0.257

All regressions include start date for VSLA and date of FOVA. Savings and Loan measured in Ugandan shillings. Ever had loan and Has loan past due are dummy variables estimated with a linear probability model
Robust standard errors clustered at the parish level in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

Differences in vulnerability profile between dense and sparse groups

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Outcome variables	Disability in hhld	Hhd size	Sometimes missing meals	Safe source of water	Access to latrines	Assessor scale	Per capita income	Has Savings Account	Has Mobile	Other savings group	Had bank loan	Index of wealth	Size of land owned
Dense VSLA	0.0266	-0.307	0.0697*	0.130**	-0.0286	0.156***	-735.1	-0.00312	-0.00120	-0.0265	-0.00785	0.294	0.592***
	(0.0331)	(0.200)	(0.0399)	(0.0623)	(0.0220)	(0.0489)	(745.0)	(0.0149)	(0.0233)	(0.0541)	(0.0109)	(0.186)	(0.222)
Obs.	2,731	2,744	2,759	2,695	2,394	2,483	2,728	2,634	2,630	2,635	2,635	2,444	2,635
R-squared	0.529	0.133	0.043	0.171	0.355	0.086	0.144	0.109	0.087	0.305	0.059	0.226	0.133

Dense VSLA dummy identifies groups assigned to the dense treatment. Assessor scale is a 1-4 scale with higher numbers being worse conditions. Regressions include implementing partner fixed effects. Errors clutered at the parish level in parentheses.

*** p<0.01, ** p<0.05, * p<0.1

Fig. 8

Impact of selection on interest and price

	(1)	(2)	(3)	(4)
Dep var:	Interest rate (percent)		Share price (UGX)	
Fixed effects	None	IP	None	IP
Dense VSLA	0.668	0.332	83.09	20.71
	(0.545)	(0.591)	(59.28)	(58.85)
Constant	8.351***	9.668***	850.9***	979.3***
	(0.328)	(0.591)	(43.31)	(58.85)
Observations	110	110	110	110
R-squared	0.014	0.661	0.018	0.491

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Fig. 9

Impact of group composition on the loans to savings ratio - Wave II

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep var:	Total aggregate savings (mean: 2,381,480)				Loans to savings ratio (mean: 1.52)			
Dense VSLA	-550,702***	-542,841***	-811,981***	-767,453***	-0.344**	-0.294*	-0.198	-0.220
	(163,367)	(121,946)	(159,485)	(141,157)	(0.169)	(0.164)	(0.183)	(0.174)
Interest rate		-52,535**		-62,213*		-0.0680**		-0.0194
		(25,433)		(36,412)		(0.0336)		(0.0389)
Share price		1,204***		762.1**		-0.000514*		-0.000982***
		(293.2)		(330.8)		(0.000261)		(0.000209)
IP fixed effects	No	No	Yes	Yes	No	No	Yes	Yes
Observations	104	104	104	104	104	104	104	104
R-squared	0.208	0.453	0.604	0.651	0.057	0.138	0.496	0.587

All regressions include start date for VSLA and date of FOVA.
Standard errors clustered at the parish level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Fig. 10

Impact of group composition on amount saved

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Wave of data collection	Wave I				Wave II			
Fixed effects	None		IP		None		IP	
Panel A: All participants								
Dense VSLA	545.7	-1,989	-7,378*	-6,901*	-6,982*	-5,801	-9,685***	-8,529**
	(5,911)	(4,303)	(4,387)	(3,622)	(4,007)	(3,498)	(3,653)	(3,429)
Interest rate		-622.4		-1,175*		-1,639		-3,480*
		(797.9)		(606.2)		(1,593)		(1,986)
Share price		54.20***		34.67***		61.04***		35.17*
		(6.475)		(6.229)		(18.67)		(18.60)
Observations	2,979	2,979	2,979	2,979	2,889	2,889	2,889	2,889
R-squared	0.200	0.406	0.518	0.560	0.092	0.210	0.303	0.326
Panel B: Preselected participants								
Dense VSLA	2,777	-378.6	-6,735	-6,222	-222.6	-817.5	-18,731**	-14,734*
	(5,399)	(4,200)	(4,553)	(3,895)	(8,729)	(7,783)	(8,313)	(8,544)
Interest rate		-1,429*		-1,729***		-1,648		-4,017**
		(746.3)		(577.3)		(1,871)		(1,821)
Share price		47.51***		30.80***		53.61***		28.65*
		(7.307)		(6.777)		(16.46)		(15.99)
Observations	904	904	904	904	834	834	834	834
R-squared	0.136	0.333	0.492	0.535	0.041	0.135	0.243	0.262
Panel C: Self selected participants								
Dense VSLA	531.3	-1,882	-6,613	-6,455*	1,821	1,863	-11,858	-9,782
	(6,344)	(4,536)	(4,494)	(3,848)	(9,783)	(7,394)	(7,148)	(6,895)
Interest rate		-144.5		-867.9		-1,795		-2,625
		(851.1)		(682.3)		(1,702)		(2,238)
Share price		56.86***		36.27***		63.69***		30.92
		(6.502)		(6.705)		(20.69)		(19.56)
Observations	2,075	2,075	2,075	2,075	2,028	2,028	2,028	2,028
R-squared	0.224	0.436	0.539	0.580	0.111	0.239	0.364	0.379

All regressions include start date for VSLA and date of FOVA.
Standard errors clustered at the parish level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Fig. 11

Impact of group composition on total amount borrowed

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Wave	Wave 1				Wave 2			
Fixed effects	None		IP		None		IP	
Panel A: All participants								
Dense VSLA	4,646	923.6	-11,391	-10,725	-24,619**	-19,438*	-17,266*	-15,175
	(12,623)	(12,091)	(7,537)	(7,801)	(9,988)	(9,969)	(9,591)	(9,639)
Interest rate		-751.4		-1,672		-6,809*		-10,466*
		(2,160)		(1,459)		(3,772)		(5,620)
Share price		77.94***		47.77**		61.53		10.09
		(22.63)		(18.75)		(38.61)		(38.56)
Observations	2,982	2,982	2,982	2,982	2,889	2,889	2,889	2889
R-squared	0.118	0.158	0.261	0.269	0.023	0.041	0.120	0.125
Panel B: Preselected participants								
Dense VSLA	2,491	-4,163	-20,738**	-19,713**	-45,385*	-45,465*	-69,141***	-64,824***
	(12,332)	(11,599)	(8,208)	(8,731)	(26,663)	(24,670)	(22,359)	(23,191)
Interest rate		-1,043		-3,479**		-3,658		-5,186
		(2,147)		(1,571)		(4,619)		(3,911)
Share price		82.63*		58.38		86.96		6.801
		(42.30)		(36.94)		(56.95)		(47.28)
Observations	904	904	904	904	834	834	834	834
R-squared	0.072	0.137	0.295	0.314	0.023	0.052	0.224	0.226
Panel C: Self selected participants								
Dense VSLA	8,793	5,849	-4,006	-3,885	-5,958	-3,040	-9,162	-4,769
	(13,770)	(13,347)	(8,886)	(9,231)	(21,471)	(20,903)	(22,523)	(22,850)
Interest rate		-352.8		-560.0		-7,754*		-11,530
		(2,381)		(1,845)		(4,056)		(7,663)
Share price		75.77***		38.60**		51.72		-0.729
		(17.35)		(16.76)		(36.01)		(40.17)
Observations	2,078	2,078	2,078	2,078	2,028	2,028	2,028	2028
R-squared	0.134	0.166	0.259	0.263	0.026	0.042	0.109	0.115

All regressions include start date for VSLA and date of FOVA.
Standard errors clustered at the parish level in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Fig. 12