

Spillovers of Community-Based Health Interventions on Consumption Smoothing

Bansi Malde and Marcos Vera-Hernández*

December 7, 2020

*Affiliations: Malde (corresponding author) - University of Kent and IFS (E-mail: B.K.Malde@kent.ac.uk; Postal Address: School of Economics, Keynes College, University of Kent, Canterbury, Kent CT2 7NP, United Kingdom.); Vera Hernández - UCL and IFS (m.vera@ucl.ac.uk). This research was funded by the ESRC/Hewlett Joint Scheme under grant RES-183-25-0008, and the Bureau of International Labor Affairs at the Department of Labor. We are immensely grateful to Emla Fitzsimons for her contributions to the initial phase of the project. We thank Orazio Attanasio, Pierre Dubois, Erica Field, Dilip Mookherjee, Mushfiq Mobarak, Rohini Pande, and other conference and seminar participants at The US Department of Labor, EDePo@IFS, and The 2012 Royal Economic Society Conference (Cambridge) for useful discussions and suggestions. We thank the Mai Mwana team, especially Tambosi Phiri, Andrew Mganga, Nicholas Mbwana, Christopher Kamphinga, Sonia Lewycka, and Mikey Rosato for their advice and assistance in collecting the data. Julia Behrman, Senthuran Bhuvanendra, Lena Lepuschuetz and Carys Roberts provided excellent research assistance. All errors are the responsibility of the authors.

Abstract

Community-based interventions, particularly group-based ones, are considered to be a cost-effective way of delivering interventions in low-income settings. However, design features of these programs may influence behaviors beyond those targeted by the intervention. This paper studies spillover effects of a participatory community health intervention in rural Malawi, implemented through a cluster randomized control trial, on an outcome not directly targeted by the intervention: household consumption smoothing after crop losses. We find that while crop losses reduce consumption growth in the absence of the intervention, households in treated areas are able to compensate for this loss and perfectly insure their consumption. An exploration of mechanisms indicates that the effects are not due to better self-insurance or to labour supply adjustments, suggesting that informal risk sharing must have improved. Health improvements cannot explain the whole effect and instead, suggestive evidence indicates that social interactions, which could have alleviated contracting frictions had a role to play.

Keywords: participatory community interventions, spillovers, consumption smoothing, Sub-Saharan Africa

JEL Classification: E21, G22, O12, O13

1 Introduction

Community-based interventions, particularly group-based interventions, are wide-spread in developing countries (Mansuri and Rao, 2004). They can build and harness social capital to trigger and sustain behaviour change and shift social norms; and deliver interventions at a lower cost than one-to-one delivery. Consequently, they have been used to provide financial services (e.g. microfinance groups), promote collective action (e.g. Community-Led Total Sanitation), deliver early childhood parenting programs, and deliver aid and infrastructure, among other interventions.

Design features of these community-based interventions can influence behaviors beyond those directly targeted. For example, regular group meetings incorporated into many interventions may encourage information sharing and trust building, alleviating contracting frictions constraining other untargeted behaviours. It is important to account for such spillover effects when assessing the welfare impacts of interventions (Angelucci and De Giorgi, 2009). However, empirically disentangling these is challenging since interventions might provide other resources which also influence these outcomes.

In this paper, we study the spillover effects of a participatory community-based health intervention in rural Malawi, implemented within a cluster randomized controlled trial, on household consumption smoothing of idiosyncratic crop losses. The chief goal of the intervention was to improve maternal and child health. Consumption smoothing was not regarded as an issue of importance by the implementation agency; and no direct efforts or resources were provided to influence this. Instead, the mere design of the program could have influenced consumption smoothing outcomes, for reasons that we explain below. Focusing on an outcome not directly targeted by the intervention allows us to shed light on this type of spillover effect. Our contributions are two-fold: First, we study whether participatory group-based community health interventions affect consumption smoothing. Second, drawing on the theoretical literature on informal risk sharing and

self-insurance, we consider ways in which the intervention may have altered consumption smoothing patterns, paying special attention to features of its design.

The intervention encouraged women to form groups which met on a fortnightly or monthly basis to engage in a participatory action cycle to improve reproductive and child health (Lewycka et al., 2010, 2013). The group meetings did not follow a top-bottom approach. A trained facilitator organized the discussion, but it was up to the group members to come up with ideas and take decisions. Initial meetings were discussion-based, while later meetings planned and implemented (along with the wider community) strategies identified by the groups to improve maternal and child health.¹ It provided no monetary resources whatsoever, and as such is extremely cost effective, and increasingly used in developing countries to improve infant mortality and female reproductive health (see Prost et al., 2013 for a systematic review).

Lewycka et al. (2013), who report intervention impacts on the primary and secondary outcomes, found that it reduced maternal mortality by 74%, and peri-natal, neo-natal and infant mortality by 33%, 41% and 28% respectively, in the second and third years of the trial. They also documented that the women's groups improved health seeking practices during pregnancy and at birth (e.g. antenatal care use increased by 50% and 30% fewer births were overseen by traditional birth attendants), and infant care practices - exclusive breastfeeding increased by 74% and complete immunization at 6 months by 266%. If beneficial effects on outcomes beyond those targeted directly by the program are achieved, the case for such already successful programs is strengthened even further.

Though households in the study frequently experience shocks, formal credit and insurance markets are severely lacking. Instead, informal tools such as transfers and informal loans between relatives, friends and neighbors; labour supply adjustments; and liquidating assets such as livestock are the key mechanisms for dealing with the consequences of

¹Strategies included lobbying a range of donors to provide bicycle ambulances to transport pregnant women to health facilities, among others.

risk (Besley, 1995). The intervention could have improved consumption smoothing outcomes through three possible channels: First, by improving health the intervention could have reduced health expenditures, increasing household savings and improving households' abilities to self-insure against crop losses. Improved adult health also increases labour productivity, better supporting the use of labour supply as a consumption smoothing device. Second, the groups could have directly supported households experiencing crop losses by providing gifts, loans and transfers from the proceeds of group activities. Finally, by facilitating the formation of groups that met regularly and aimed to involve a broad swathe of the community, the intervention could have increased social interactions and alleviated contracting frictions thereby improving consumption smoothing.²

We begin our analysis by considering how consumption smoothing – measured as the response of changes in log household (non-durable) consumption to idiosyncratic crop losses, net of village level aggregate shocks – varies in villages receiving the intervention relative to control villages, which did not receive the intervention. To identify causal impacts, we exploit the random allocation of the intervention to groups of villages (or clusters).

Household crop losses are an important source of risk in these communities: around 28% of households in our sample experience a crop loss over a 2-year period. Moreover, these crop losses are economically significant: those experiencing a crop loss lose, on average, 1.2 months of household non-durable consumption.

Our empirical strategy relies on two key assumptions. First, crop losses are exogenous and uncorrelated with the treatment. Second, the Stable Unit Treatment Value Assumption (SUTVA) holds, and the treatment does not spillover to control clusters, through for instance crowding out cross-village transfers. We provide evidence and argue that both of these conditions are satisfied. First, we show that the distributions of our measures of

²Contracting frictions include information asymmetries (e.g. Ligon, 1998; Kinnan, 2014) and imperfect enforcement (Kocherlakota, 1996; Ligon et al., 2003; Ali and Miller, 2013).

crop losses are similar between the control and intervention groups. We argue that the crop losses are likely to be exogenous since agricultural production is primarily rain-fed in this context, and thus vulnerable to unpredictable weather, pests and crop diseases. Second, we argue that the experimental design – which included large buffer areas between study clusters – makes it very likely that SUTVA holds across clusters. In addition, the crop losses should represent idiosyncratic (household-level) rather than aggregate (village-level) adverse events: we document substantial variation in our measure of crop losses within village, thereby satisfying this requirement.

Our results indicate that the intervention leads to significant improvements in consumption smoothing: while experiencing a crop loss leads to a reduction in household consumption growth by around 9% for households in control areas, those in intervention areas are able to compensate for this loss, and subsequently perfectly smooth their consumption.

Next, we consider the channels through which the intervention could have improved consumption smoothing. Our data is somewhat limited, so that this analysis is exploratory. We start by considering intervention impacts on how households smooth consumption. First, we study whether households in the intervention areas were more likely to draw down on savings and assets following a crop loss. In fact, we find the opposite effect: households in intervention areas were *less* likely to decumulate their savings, particularly holdings of large animals and cash savings, than households in control areas. This is consistent with the intervention allowing households to protect their savings and livestock following a crop loss. This has important implications for long run poverty alleviation efforts, since forced sales of productive assets not only yield losses due to lower sales prices, but also impact future agricultural production (Rosenzweig and Wolpin, 1993). Next, we analyze whether the intervention increased the use of labour supply as a consumption smoothing mechanism; and find no evidence this is the case for adults or children. Thus, the improvements in consumption smoothing must be a result of better informal risk

sharing.

The second part of our exploration investigates *how* the intervention improved informal risk sharing, focusing on three channels (i) improved health; (ii) direct aid from the groups themselves and (iii) increased social interactions.

We find no impacts of the intervention on general adult and child health or child physical growth. We also show that the significant reductions in maternal mortality documented by Lewycka et al. (2013) are not large enough to explain the magnitude of improvement in consumption smoothing found. There is also no evidence that the groups provided aid to help households cope with adverse crop losses. By contrast, we find evidence that the intervention improved within-village social interactions. Our data is limited, in that we only observe whether an individual chatted with friends and family about various topics, rather than information on informal gifts, loans and transfers between households. We can thus assess whether or not within-village social interactions improved in general, but not precisely how social interactions improved consumption smoothing. Thus this exploration provides speculative evidence.

Nonetheless, we document an increase in the probability that households chat with their friends and family on a one-to-one basis *outside the intervention group*, providing suggestive evidence that social interactions outside the groups increased. Further suggestive evidence indicates that the chats responded to crop losses: households experiencing crop losses in intervention areas were more likely to start asking friends and family about sources of credit, than in control areas. Other work has shown that increased social interactions can improve informal risk sharing, e.g. Feigenberg et al. (2013) show that increasing the frequency of meetings among microfinance clients increases risk sharing within the microfinance group, as measured by a higher propensity to share lottery tickets, which reduces one's own likelihood of winning, with fellow group members.

Several aspects of the intervention likely contributed to the large observed impacts: first, the intervention had already been running for over three years at the time of data col-

lection. Second, the group meetings did not follow a top-bottom approach but - on the contrary - a participatory one in which the role of the facilitator was to simply encourage discussion within a 4-cycle structure which also involved non-group villagers. This participatory approach is more likely to have resulted in increased communication (and ultimately information sharing) among group members. Moreover, the third and fourth parts of the participatory cycle – where groups implemented chosen strategies – also provided opportunities to participants to learn about the traits (e.g. trustworthiness, ability, etc) of their fellow group members. Finally, the intervention required significant interactions among group members, and the wider community outside of the group meetings, particularly during the third and fourth parts of the cycle.

Our findings are important for the design of development programs. In particular, they speak to the debate on whether interventions should be delivered to beneficiaries on a one-to-one basis or within a group. One-to-one delivery of interventions such as health education allow for the provision of tailored services, and hence may be more effective; while group-based interventions might be less costly but may entail a dilution in services. Our findings provide one additional benefit of group-based interventions – improved consumption smoothing as a result of increased social interactions – which not only influences household welfare, but may also affect households’ ability to make use of the services provided by the intervention.

Our paper contributes to a number of strands of the literature. First, it adds to a body of work investigating the impacts and workings of community-driven (CDD) interventions, an immensely popular way of delivering development aid (Mansuri and Rao, 2004). Such interventions aim to improve coordination between community members, and broaden participation by under-represented groups within a community. Existing studies, though, find mixed evidence on their efficacy in improving the ultimate outcomes of interest (typically public goods provision, distribution of funds, or health outcomes) and in improving social outcomes (cohesion, cooperation, collective action). On the former, Casey et al.

(2012), Prost et al. (2013), Pickering et al. (2016) and Cameron et al. (2015) find positive impacts on this type of intervention in increasing local public good provision in Sierra Leone, improving maternal and infant health, and increasing uptake of sanitation in Mali and Indonesia respectively. However, Banerjee et al. (2010) and Abramovsky et al. (2019) find that such interventions were ineffective in improving education outcomes in India, or sanitation take-up in Nigeria. On the latter – social outcomes – an evidence synthesis by White et al. (2018) of CDD programs targeting local public goods concludes that such programs have little effect on social capital. Our paper contributes to this literature by providing evidence of impacts on a novel outcome - consumption smoothing.

Second, it also contributes to the literature on social networks and risk sharing or closely related outcomes. Groups with strong social connections provide mutual assistance (e.g. Fafchamps and Ferrara, 2012 for self-help groups, Angelucci et al., 2018 for extended family), and higher social capital can sustain more complex financial transactions (Guiso et al., 2004). Recent literature also shows that stronger social connections share more risk by overcoming contracting frictions: Chandrasekhar et al. (2015) find, within lab-in-the-field experimental games, that socially connected ties are able to overcome enforcement constraints and cooperate in the absence of outside enforcement); while Breza and Chandrasekhar (2015) document that informing socially close, and more central villagers about an individual's savings increases the amount saved. Finally, several theoretical papers show that the network architecture influence informal risk sharing in the presence of frictions such as limited commitment and asymmetric information (Ambrus et al. (2014), Ambrus and Elliott (2018), Ambrus et al. (2019), Bourles et al. (2020)). Our study provides evidence of an intervention that improves consumption smoothing by altering social interactions.

Finally, it contributes to a body of work on the impacts of groups on economic behavior. The microfinance literature has contributed to either justify theoretically or establish empirically a correlation between group formation and positive loan outcomes. Closest to

our paper is the contribution by Feigenberg et al. (2013) who randomly increase individuals' social interactions (in their case micro-finance clients) by experimentally varying the frequency of micro-finance meetings. Our paper complements theirs by looking at different outcomes (consumption smoothing rather than propensity to share lottery tickets or default probability), in a different setting (rural rather than urban), and studying a different type of group meeting - open community-wide meetings rather than micro-finance group meetings (that unlike the women's groups studied here might crowd out informal risk sharing arrangements). In another related study in post-conflict Uganda, Blattman et al. (2016) document that cash grants and training provided to women encouraged to form groups engaged more in informal finance as well as labour sharing, and enjoyed higher earnings, but did not increase their consumption in response.

The rest of the paper is organized as follows. Section 2 provides background on the context, and introduces the intervention and the experimental design. We outline our empirical strategy in section 3. Section 4 describes the data, including our measures of crop losses; Section 5 displays the main results, and investigates the underlying mechanisms. Section 6 concludes.

2 Background and Intervention

2.1 Context

Our study takes place in Mchinji, a rural district in the central part of Malawi. According to the 2008 Malawi Census, socio-economic conditions in Mchinji are similar to or worse than the average for Malawi (in parentheses in what follows), with literacy rates of just over 60% (64%), poor quality flooring materials used by 85% (78%) of households, piped water access for 10% (20%) of households, and electricity access for just 2% (7%) of households. Subsistence agriculture is the most important economic activity in Mchinji,

and provides the main income source for a large majority of households. Key crops include maize, tobacco and ground nuts. Agricultural production is mainly rain-fed, and the use of modern inputs and implements such as fertilizer is very rare; leaving household incomes vulnerable to fluctuations from unpredictable weather as well as pests and crop diseases.

Access to formal financial and insurance products is limited: the 2008 Finscope Malawi Survey documented that only 3% of Malawians held an insurance product, and fewer than 20% had a formal bank account. Social connections such as family, friends and other community members are very important in helping households cope with risk. For example, Trinitapoli et al. (2014) document that older siblings play a role in protecting educational investments of younger siblings, while Munthali (2002) and Peters et al. (2008) show that extended family members foster and care for children orphaned by HIV/AIDS. A number of studies have shown that contracting frictions are present in the context of rural Malawi. Difficulties in verifying an individual's identity, and private information have been shown to introduce frictions in credit and labour markets (Gine et al., 2012; Guiteras and Jack, 2014), and hamper the effective targeting of subsidies (Jack, 2013). Such imperfections not only constrain formal credit and insurance markets, but also influence the terms of informal arrangements, limiting perfect insurance.

2.2 Intervention

We consider a participatory women's group intervention aimed at improving reproductive health, set up in 2005 by the MaiMwana ("Mother and Child") Project. The intervention was implemented as a cluster randomized control trial in order to assess its effectiveness in improving maternal and neonatal health outcomes.

It aimed to impart information and mobilize the local community (village) to act on issues relating to pregnancy, childbirth and newborn health, with the overarching goal of

improving maternal and child health outcomes (Lewycka et al., 2010; Lewycka, 2011). To do so, trained local women (facilitators), visited study villages and encouraged women to form groups to follow a participatory action cycle. Each covered a population of around 3,000 individuals, and was charged with forming and guiding between 6 and 12 groups.³ Groups were guided through a four-stage community mobilization action cycle of fortnightly (in principle; every 4-6 weeks in practice) meetings under the guidance of the facilitator. In the first stage (comprising 8 meetings) they identified and prioritized problems relating to pregnancy, childbirth and newborn health.⁴ In the second stage (4 meetings), the groups shared the results of their discussions with the wider community and devised strategies to overcome identified problems. To take an example, the problem identified might be a lack of transport to health centers in cases of emergency and the strategy to overcome this was to lobby various donors to provide bicycle ambulances.⁵ In the third stage (4 meetings), the groups involved the wider community in implementing the strategies, before evaluating their activities and making future plans in the final stage (4 meetings).

Besides these formal meetings, sub-groups met regularly outside the formal group meetings particularly during the implementation stage of the participatory action cycle (Rosato, 2012). Many women's group established committees to guide the implementation of the strategies, disciplinary committees to resolve conflicts within the groups, and social welfare committees to coordinate the generation and distribution of funds among group members. Moreover, committee leaders would meet on a regular basis outside the committee, and formal group meetings. Thus, the intervention involved significant

³The facilitators, who were aged 20-49 years, literate and mothers, received initial training over 11 days, followed by refresher training every 4 months. They were paid a salary, and provided with a bicycle, T-shirt, umbrella and field bag. They were also given a manual to implement the participatory action cycle, as well as picture cards to help guide discussions. Finally, they received support and feedback from a supervisor.

⁴Rosato et al. (2006) and Rosato et al. (2009) summarize the problems identified by the groups relating to maternal health (anaemia, obstructed labour, malaria, retained placenta, and hemorrhaging, amongst others) and neonatal and infant health (diarrhoea, pre-term births, tetanus, asphyxia, infection and malaria).

⁵Common strategies implemented by groups included vegetable garden cultivation, health education sessions, insecticide treated net distribution, health program radio listening clubs and cleaning the surroundings around one's house.

interaction outside the formal group meetings.

The primary targets of this intervention were women aged 15-49, and particularly pregnant women. However, older women who had completed their fertility, who were considered to be very influential in shaping reproductive health outcomes in the community, were also encouraged to attend to share their experiences. From the third part of the cycle, men joined the groups and attended meetings.

Around 31% of women aged 15-49 in our data report having attended at least one women's group meeting. Conditional on attending at least one meeting, women report having attended on average 12 women's group meetings over a 4-year period, indicating active engagement in the intervention among this group.⁶ The groups were successful in reaching less affluent women and households: attendees were poorer than non-attendees (as indicated by an asset index and the quality of housing materials), older on average, and were less likely to have completed primary schooling. Attendees were also more likely to have been married at least once; while those from households with more than one adult woman were less likely to attend.

Potential Intervention Impacts on Consumption Smoothing The intervention could have influenced consumption smoothing through three key channels: First, by improving health, the intervention might have reduced health expenditures and/or the need to decumulate savings and assets in response to adverse health events. This, in turn, could have improved the household's financial position, allowing it to better withstand crop losses through self insurance. Improved health could also improve informal risk sharing by reducing the likelihood of health shocks, thereby increasing the capacity of risk sharing partners to provide support following adverse crop losses. Second, numerous groups implemented strategies that generated food, or income, which could have been shared with households experiencing crop losses.

⁶Unfortunately, we do not have information on participation in committee and other sub-group meetings. Thus, the average attendance reported here is likely to be a lower bound of attendance to all meetings.

Third, the groups provided women with a forum within which to interact with other women regularly. Such repeated interaction can alleviate contracting frictions, as indicated by research on repeated games (Spear and Srivastava, 1987 among others). Information sharing among group members on topics other than health would alleviate hidden income frictions, and notify members of job opportunities and sources of credit. Participation in intervention activities – which encouraged active participation in tasks they would not otherwise engage with in the course of other social interactions – could have allowed participants to learn more about the ‘types’ of others, and have made it easier to monitor effort, and reduce transaction costs associated with making informal transfers. Finally, the repeated interactions might have also made it easier to punish renegers, thereby alleviating frictions emerging from imperfect enforcement. Thus, by reducing contracting frictions, the intervention might have allowed women (and their households) to either form new social connections that could support informal risk sharing; and/or strengthen existing arrangements to allow for more risk sharing.⁷ Further, information sharing on non-health topics such as job opportunities would have alleviated barriers faced by households in adjusting their labour supply in response to adverse crop losses.

2.3 Experimental Design

The evaluation is based on a cluster randomized controlled trial designed in the following way (see Lewycka et al., 2010, Lewycka, 2011). Mchinji District was divided into 48 clusters by combining enumeration areas of the 1998 Malawi Population and Housing Census in a systematic way so that each cluster contained approximately 8,000 individuals.⁸ Within each cluster, the 3,000 individuals (equating to 14 villages on average) living

⁷However, the increased interactions with group members could crowd out other or weaken relationships with connections who were not group members, potentially crowding out existing risk sharing; and making the effects on overall consumption smoothing ambiguous.

⁸The District Administrative Centre was excluded because it is relatively more urbanized and hence less comparable to the rest of the District.

closest to the geographical centre of the cluster were chosen to be included in the study.⁹ This leaves a buffer area between adjacent clusters created purposefully to limit contamination. As we explain in Section 4.4, the presence of these buffer areas makes it more likely that the Stable Unit Treatment Value Assumption (SUTVA) holds across clusters. 12 clusters were randomly selected to receive the intervention. A further 12 serve as controls.¹⁰

The intervention began in 2005, and was still active when the data used in this study were collected (2008-09, and 2009-10). Though groups began in 2005/06, the participatory action cycle was not completed within the originally intended 2 years. This was primarily due to delays arising from groups trying to involve men in their activities at the start of the second stage of the cycle. Most groups were in either the third or fourth stage of the cycle at the time of the first follow-up survey, with many in the fourth stage of the cycle at the time of the second follow-up survey.

3 Empirical Model

To identify whether the intervention improves consumption smoothing, we exploit its cluster randomized allocation and assess how household consumption smoothing varies by whether or not the household lived (prior to the random assignment) in a cluster that received the intervention. Our empirical specification is derived from a standard model of inter-temporal consumption smoothing in which households have a utility function of the constant relative risk aversion (CRRA) form (as is conventional - see for instance

⁹The geographic centre was chosen to be the most central village in the cluster as shown on a cartographic map from the National Statistical Office, Malawi, and whose existence was corroborated with the District Commissioner's records. See Lewycka (2011), pp. 122 for more details.

¹⁰A further 12 were assigned an Infant Feeding intervention, and the final 12 were assigned both the Infant Feeding Intervention and the Women's Group one. We do not exploit these here; see instead Lewycka et al. (2013) and Fitzsimons et al. (2016) who study the impacts of the infant feeding intervention on child health and household behaviors. MaiMwana Project also improved health facilities across the District, which benefitted both intervention and control clusters equally.

Cochrane, 1991, and Townsend, 1994), and face idiosyncratic and aggregate risk.¹¹

The specification is as follows:

$$\Delta \log(c_{hvt}) = \alpha \Delta \text{crop}_{hvt} + \gamma \Delta \text{crop}_{hvt} * D_v + \theta \Delta X_{hvt} + \mu_{vt} + \Delta \varepsilon_{hvt} \quad (1)$$

where $\Delta \log(c_{hvt})$ is the consumption growth rate for household h living in village v , between periods $t - 1$ and t , Δcrop_{hvt} refers to changes between periods $t - 1$ and t in household-level crop loss indicators (incidence and severity, separately), D_v is an indicator variable that equals 1 if, in 2004 (before the start of the intervention), household h lived in a cluster that was subsequently randomly selected to receive the women's group intervention, X_{hvt} is a vector of household level time-varying characteristics such as household size and demographics,¹² and μ_{vt} is a set of village-specific time dummies which will absorb village-specific aggregate shocks and any village-level effects of the treatment on $\Delta \log(c_{hvt})$ through channels other than consumption smoothing of crop losses.¹³

Under perfect consumption smoothing (or full insurance), the household consumption growth rate should move one-to-one with the aggregate consumption growth rate (Cochrane, 1991, Townsend, 1994), and be uncorrelated with household-level crop losses (Δcrop_{hvt}).¹⁴ In our case, we take the village to be the unit within which risk sharing takes place, and so changes in aggregate consumption are captured by the village-time dummies. We consider village-level risk sharing since the participatory groups were generally formed

¹¹The specification is derived from the first-order condition of a problem where a social planner chooses the consumption levels for members of a risk sharing group by maximising a weighted sum of household expected utility subject to an aggregate budget constraint. The the social planner weights can take any form; taking first differences allows us to difference them out.

¹²The inclusion of changes in household demographics allows us to account for changes in $\log(c_{hvt})$ due to changes in household composition between survey rounds. The results are robust to the exclusion of these controls.

¹³Since we have only two time periods, the term μ_{vt} is essentially a village fixed-effect.

¹⁴Shocks could also affect long term outcomes, such as education and health, see for instance Maccini and Yang (2009). More generally, poor households who are close to subsistence could be smoothing consumption at the expense of investments, which would affect their long-term wellbeing (Chetty and Looney, 2006). In analysis not presented here, we find no evidence that the crop losses negatively affected the physical growth of children aged less than 5 years in our sample.

within a village. Any impacts on consumption smoothing, including on social interactions would be concentrated within the village. This specification means that we capture the net effects of the groups on *village-level* risk sharing and consumption smoothing.¹⁵

If households are able to perfectly smooth consumption through existing mechanisms, and particularly in the absence of the participatory intervention, then we would expect the coefficient α to equal zero. If, however, households are unable to smooth consumption following a crop loss, α would be negative. γ identifies the additional consumption smoothing available to households in intervention areas as a result of the intervention. The sum of the coefficients $\alpha + \gamma$ reveals whether households in intervention areas are perfectly able to smooth their consumption following an idiosyncratic crop loss. Our test for whether the intervention improves risk sharing therefore entails testing whether γ is positive and statistically significantly different from 0.

Note that this specification resembles that for estimating heterogeneous treatment effects, where the objective would be to test whether the treatment effect on $\Delta \log(c_{hvt})$ varies by whether or not the household experienced a crop loss. Our objective is different, in that we test whether consumption smoothing – measured following the existing literature – is different in treatment and control villages. It is straightforward to show that equation 1 can be derived from specifications that separately estimate the extent of village-based consumption smoothing in control and treatment clusters, while restricting the coefficient θ to be the same.¹⁶ Our interpretation of the coefficients α and γ thus follows the consumption smoothing literature.

In terms of the dependent variable, we use two measures of household consumption: non-durable household consumption and food consumption. This choice is motivated by the fact that different types of consumption might be more or less sensitive to the effects

¹⁵It is possible that households that are part of a group may reduce risk sharing with those outside the group leading to some households being left worse off by the intervention. In the absence of exogenous variation in within-group participation decisions, we are unable to study such heterogeneous effects.

¹⁶Relaxing this assumption to allow θ to vary by treatment does not change our findings.

of shocks. Food consumption is a necessity, accounting for a large share of households' budgets in this context. Thus, households might choose to cut back on other types of consumption in order to protect food consumption following a crop loss.

Our measures of crop loss, $\Delta crop_{hvt}$, include incidence of the loss, and the size of the loss, as measured by the monetary income lost in Malawi Kwacha. The size of the loss is continuous and provides information on the intensity of the shock. When using the size of the loss directly as the measure of shock intensity, we need to make adjustments for the fact that the same income loss in levels will be a larger share of consumption for a poor household than for a wealthier household. Hence, the wealthier household might find it easier to smooth consumption, which would mechanically dampen the correlation between changes in log consumption and changes in crop loss income (in levels).¹⁷ To avoid this, we normalize the size of the loss by the household's predicted consumption, which is calculated as a function of pre-intervention education.

This normalization variable deserves further discussion: ideally we would like to normalize the loss by the household's pre-intervention permanent income, which would not be altered by transitory shocks, or by the intervention itself. However, we do not observe this in our data. We instead proxy for permanent income using the value of a household's consumption that is explained by the pre-intervention education – a variable that is considered to be an important determinant of permanent income – of the survey respondent.^{18,19} Moreover, pre-intervention education is, by design, uncorrelated with the

¹⁷We would not encounter this issue if the intensity of the loss was also measured in logs. However, it is not possible to use logs of the loss since a significant proportion of households in our data experience no crop loss, i.e. a loss of 0 MK. In a robustness check presented in Section 5.3, we use the inverse hyperbolic sine transformation, which is defined for 0, as an alternative way of dealing with this issue. The results are qualitatively similar to those found with the normalization procedure used here.

¹⁸At baseline, around 31.9% of survey respondents had no education, 62.4% had some or completed primary schooling. The remaining 5.7% had some secondary schooling or higher. Table 4 in the Appendix shows the correlations between household non-durable consumption and education.

¹⁹We also experimented with the value of a household's consumption that can be proxied by pre-intervention assets. However, this proxy could generate spurious correlation since households could choose to liquidate assets in response to shocks. A further concern with this normalization strategy is that it might generate some spurious correlation since consumption appears as both the dependent variable and a function of it as an independent variable. To alleviate these concerns, in Section 5.3, we present a number of

intervention and cannot be altered by transitory shocks. Thus, this proxy for permanent income will be exogenous to both these variables.

A critical identification assumption underlying our empirical strategy is that crop losses are exogenous, and uncorrelated with the treatment. Moreover, crop losses should represent idiosyncratic (household-level) rather than aggregate (village-level) adverse events. We provide evidence indicating that both of these requirements are likely to be satisfied in our setting in Section 4.

Important consideration needs to be given to inference. Since the treatment was allocated at the cluster (group of villages) level, we cluster standard errors at the cluster level. In doing so, since the cluster is a higher level of aggregation than the village, we also account for the fact that risk sharing makes within-village observations interdependent so that SUTVA does not hold within village. The cluster-correlated Huber-White estimator for clustered standard errors have been shown to provide standard error estimates which are too small (and thereby over-reject hypothesis tests of significance) when the number of clusters is less than 30 (Donald and Lang, 2007, Wooldridge, 2004, Bertrand et al., 2004, Cameron et al., 2008). This is important here since the number of clusters is 24. To ensure correct inference, we implement a wild cluster bootstrap-t procedure recommended by Cameron et al. (2008) to compute the correct p-values for hypothesis tests of significance. Their Monte Carlo simulations suggest that this method performs relatively well with as few as 5 clusters.²⁰ Corresponding studentized bootstrap confidence intervals can be constructed by inverting the test (Cameron and Miller, 2015). We present the main results in Figures, which display the coefficients α (the effect of the crop loss in the control group), $\alpha + \gamma$ (the effect of the crop loss in the treatment group) and γ (the treatment effect) and the associated studentized bootstrap confidence intervals.²¹

robustness checks where we use alternative adjustments which do not include any function of consumption. We obtain qualitatively similar results with these alternative strategies, suggesting the absence of such a spurious correlation.

²⁰Monte Carlo simulations reported in Fitzsimons et al. (2016) also indicate that the wild bootstrap method performs well in this setting.

²¹The corresponding estimation tables are provided in the Online Appendix 8. In all estimation tables,

4 Data

4.1 Sample Description and Balance

The analysis in the paper is based on two rounds of survey data collected in 2008-09 and 2009-10, respectively 3.5 and 4.5 years after the intervention started in May 2005. We combine this with a more limited set of variables collected in a pre-intervention baseline census of all women of childbearing age by Mai Mwana in 2004 (Mai Mwana census hereon). A large majority of the groups were still in operation at the time of our follow-up surveys (Rosato et al., 2012), partly due to late completion of the participatory action meeting cycle (the median group completed the action cycle in October 2008), as well as continuation of meetings beyond the initial participatory action cycle. The sample for the follow-up surveys was drawn from the census: we selected a random sample of 104 women per cluster to survey.

This census contains basic socio-economic and demographic characteristics, displayed in the first column of Table 1, for households in the control clusters. Households contain an average of just under 6 members, living with poor housing conditions, as demonstrated by the roofing and flooring materials. Less than 2% of households have access to piped water and less than 1% have access to electricity. Among assets, most households own a paraffin lamp, 63% own a radio, and 51% own a bicycle. Ownership of other assets such as cars, motorbikes and oxcarts is rare. Almost all households are involved in agriculture. The second column in the left panel of Table 1 shows the difference in means between treated and control clusters, with the p-value of this difference displayed in the third column. No statistically significant differences are detected for these variables, indicating that the randomization worked well.

In the 2008-09 (“first follow-up”, hereon) survey, we were able to interview around two

we report the clustered standard error, clustered standard error p-value along with the wild cluster bootstrapped p-value.

thirds of the sample drawn of women of child bearing age (aged 17-43 at the time of survey) and their households. In addition to the time-lapse of 3.5 years between the baseline and the first follow-up, two additional factors contributed to the attrition. First, collecting longitudinal data in Mchinji is particularly challenging since respondents are known to report ‘ghost’ – or fictitious – household members in an effort to increase future official aid/transfers (that may increase with household size) (Miller and Tsoka, 2012). Hence, it is possible that some women listed in the baseline census did not actually exist, and hence could not be found by the field team in 2008. Second, an unexpected sharp drop in the British Pound against the Malawi Kwacha in 2008 resulted in fewer resources to track women who had moved.

Reassuringly, this attrition rate is very similar across intervention and control clusters, with 67% and 63% respectively being successfully interviewed. The balance on a range of observed baseline characteristics is mostly maintained, as shown in the middle panel of Table 1, with a small statistically significant difference appearing on only 1 variable, indicating that attrition was not significantly different between intervention and control clusters. The 2009-10 (“second follow-up”) survey followed those women (and their households) who had been successfully interviewed in the first survey. 91% of these were reached: 92% and 90% in intervention and control areas respectively.

As is common practice in the consumption smoothing literature, we restrict the sample to households that were resident in the same village over both survey rounds. This is in order to control accurately for aggregate (village-level) adverse events through the term μ_{vt} in (1). This results in a smaller analysis sample: 1636 vs. 1249 households. The balance for this sample, based on observed baseline characteristics, is displayed in the last three columns of Table 1. We observe no statistically significant differences by treatment in these characteristics.

4.2 Survey Data

The follow-up surveys contain detailed socio-economic information such as non-durable consumption, education, labour supply, self-reported health and anthropometric measurements for young children, information on adverse events experienced by the household and social interactions as measured by one-to-one chats on various topics.

The extensive consumption module required respondents to report, at the household level, the quantity consumed and purchased, and the amount spent on the purchase of 25 different food items in the week prior to the interview. It also elicited expenditures on other important household items including clothing, health, education, and housing improvements, among others. The latter items were collected for recall periods of 1 month (for items such as fuel, utilities and transport), and 12 months or since the last survey (9-11 months) (for items such as house improvements, clothing, health, and education expenditures). In the 2009-10 survey, information was also collected on conversion units, allowing us to convert non-standard units of measurements (such as a cup of beans) to standard units (such as kilograms and grams). Total household food consumption is computed by summing expenditures on food and imputed values of non-purchased food.²² The average monthly food consumption for households in our sample is about 9874 MK (~US\$70), while average total monthly non-durable consumption is 11,808MK (~US\$84).

²²To value food that was not purchased, we first use the conversion units collected in 2009-10 to convert foods measured in non-standard units into standard units. Where available, we used household specific unit-values to value non-purchased household consumption. For households who did not purchase any quantity of the food item in the same unit as that in which it consumed the item, we use median unit-values (computed by dividing expenditure on a certain good by the quantity purchased, and taking the median at the cluster and district levels) to value this non-purchased consumption. An alternative method is to use market prices, which were also collected from local markets and trading centers most regularly visited by sampled households. This is not our preferred method, since most households rarely purchase the foods they commonly consume from the markets, and we may over-value their consumption this way. Reassuringly though, valuing consumption by either method yields the same food consumption share of total non-durable consumption of 0.86. Total non-durable consumption is computed by converting all consumption and expenditure values into monthly terms and summing them up.

4.3 Measuring Adverse Events

Crop losses are a particularly important adverse event in our setting, where practically all households are involved in subsistence, rain-fed agriculture. These could arise for a number of reasons including poor weather conditions (which could be idiosyncratic within a village as documented by, among others, Udry, 1995), localized crop diseases, pests, fires, and so on.²³ Information on crop losses was collected from two questions: first from a question asking households whether they had experienced a crop loss in the year (9-11 months) prior to the first (second) follow-up survey.²⁴ If yes, households were asked to report the size (severity) of the crop loss, in terms of the estimated income loss associated with it. The bottom 4 rows of Table 2 display the prevalence and severity of crop losses in our sample, with the latter variable normalized by predicted pre-intervention consumption.

On average, 35% and 22% of households experienced a crop loss in each wave respectively. Among households that experienced a crop loss, the average size of the loss was 116% and 84% of predicted monthly pre-intervention consumption in waves 1 and 2 respectively. The higher incidence and severity of crop loss in the first round can be explained by two reasons. First, the incidence of crop losses in the second round relates to a shorter period (9 to 11 months as opposed to 12 months). Second, data collection in the second (first) round took place during (after) the main growing season, and so not all crop losses would have been realized by the time of the second round interview. In our analysis, we control for the differing household-level recall periods.

Our empirical strategy relies on the fact that crop losses represent idiosyncratic (household-level) and not aggregate (village-level) adverse events. Crop losses could affect all house-

²³We believe that these losses are likely to be unanticipated. If shocks were anticipated, we would overestimate the amount of consumption smoothing. As we explain in Section 4.4, what matters for identification in our case is that the distribution of crop losses, and whether they were anticipated or not, is similar between the treatment and control clusters.

²⁴In the second follow-up survey households were asked about crop losses since the first follow-up survey, between 9 and 11 months for most households. Our analysis controls for the different durations asked across households.

holds in a village if they are caused by common weather shocks or natural disasters affecting the whole village. To assess the idiosyncratic variation in crop losses, Figure 1 plots a histogram of the village level proportion of households experiencing a crop loss in our sample. The left (right) panel of the figure shows the density of the village level proportion of crop losses in the first (second) round. Both show substantial variation within the village in the incidence of crop losses, though there are some villages where no household experienced a crop loss and others where all households experienced a crop loss.²⁵

Further evidence that there is idiosyncratic variation in the crop loss measures comes from examining the R^2 of a regression of the crop loss measure (either the dummy for any shock, or the amount lost normalized by predicted pre-intervention consumption) on village-time dummies. This would provide an indication of the amount of the variation that can be explained by village-specific shocks, with values close to 1 indicating that the crop losses are mostly aggregate, while values close to 0 would indicate that they are mostly idiosyncratic. Estimating these regressions, we obtain an adjusted R^2 of 0.359 for the dummy for any crop loss, and of 0.0695 for the crop loss normalized by predicted consumption, indicating that there is significant idiosyncratic variation in the crop losses in our data. In any event, our empirical specification nets out the effect of any aggregate (village-wide) events through the village-time dummies, so the crop loss picks up the effects conditional on these.

4.4 Identification Assumptions

Our empirical strategy seeks to identify spillovers of the women's group intervention on consumption smoothing by comparing the responses in changes in consumption of households that experienced crop losses in treated communities with those experiencing

²⁵Reassuringly, this variation does not simply reflect variation in the distribution of occupations across villages, as can be seen in Figure 9 in the Online Appendix, which plots the same distribution as Figure 1 but is restricted to households where the head or spouse report agriculture as their main income generating activity.

crop losses in control communities. Two key identification assumptions underly this empirical strategy: First, the crop losses are exogenous and uncorrelated with the treatment; and second, that the stable unit treatment value assumption (SUTVA) holds across clusters.²⁶ In this sub-section, we carefully examine how likely it is that these assumptions hold in practice.

We argue that the crop losses are likely to be exogenous. Agriculture is mainly rain-fed in the study context, and thus vulnerable to unpredictable weather, pests and crop diseases. To assess whether the crop losses are uncorrelated with the treatment, we first examine the distributions of the changes in crop loss by treatment group, displayed in Figure 2. The top left panel shows the histogram for the change in crop loss using the incidence measure, the top right panel displays the kernel density plots for the income lost measure, while the bottom left panel displays the distributions for the income loss normalized by predicted consumption measure. The Figures show that the distributions of the changes in crop losses are very similar across both treatment arms for all three measures. Indeed, we are unable to reject the null hypothesis that the distributions are the same in a Kolmogorov-Smirnov test of equality of distributions.

We also confirm this in a regression framework, where we regress the changes in crop loss measures on an indicator for being in a treatment village using the following specification:

$$\Delta crop_{hvt} = \lambda_0 + \lambda_1 D_v + \psi_{hvt} \quad (2)$$

where $\Delta crop_{hvt}$ is the measure of change in crop loss incidence or intensity for household h in village v at time t , and D_v is an indicator variable that equals 1 if, in 2004 (before the start of the intervention), household h lived in a cluster that was subsequently randomly selected to receive the women's group intervention. $\lambda < 0$ would indicate that the

²⁶Our empirical strategy also assumes that there are no time varying omitted variables which have a differential effect on consumption smoothing between treatment and control, and which are correlated with the propensity for crop losses.

intervention indeed reduced the incidence and intensity of crop losses.

Table 2 displays the coefficients from an OLS regression of regression 2. All coefficients are very small in magnitude and statistically insignificant from 0, suggesting that the incidence and intensity of crop losses is very similar in the treatment and control groups.²⁷

The second key assumption is that the stable unit treatment value assumption (SUTVA) holds, that is, the treatment should not spill over to the control group. This could happen, for example, if social connections that help treated households cope with the consequences of risk based in control clusters are crowded out by new connections formed as a result of the intervention. The experimental design makes such violations of SUTVA unlikely. In particular, the large buffer areas between adjacent study clusters means that households in adjacent treatment and control clusters are significantly geographically distant from one another. Given poor transport infrastructure and unavailability of mobile money in the study area at the time of the study, they would have faced large transaction costs in interacting with each other (Jack and Suri, 2014). Thus, it is unlikely that such connections, of which we expect there to be few, would have been important risk sharing partners to help households cope with the idiosyncratic crop losses we study.²⁸

5 Results

5.1 Consumption Smoothing

We now turn to the main findings, as estimated using regression 1 shown in Section 3. These are presented in Figure 3. The Figure plots the response of log total non-durable consumption (top two panels) and log total food consumption (bottom two panels) to

²⁷While we have no reason to believe that these losses could have been anticipated, our identification strategy also requires that the crop losses were similarly anticipated (or unanticipated) in the treatment and control clusters. We believe that this is likely to hold in this context.

²⁸Such connections could be very valuable risk sharing partners for aggregate village-wide shocks, and particularly large shocks that village-based connections could not fully absorb.

crop losses in the control and treatment groups, the treatment effect; and the associated 95% studentized bootstrapped confidence intervals. The corresponding estimation tables are presented in the Online Appendix. The Figure indicates that crop losses result in a reduction in household consumption growth in the control group. A crop loss equivalent to a month's predicted pre-intervention consumption results in a reduction in consumption growth of around 6%. However, there is also strong evidence that the intervention improved consumption smoothing in intervention villages: the coefficients on the interaction term, γ , are positive and large so that households in intervention clusters are almost completely protecting their consumption against crop losses. This finding is particularly strong for the 'intensity' measure of crop losses (crop loss as a proportion of predicted consumption), as can be seen from the second panel of the Figure.

From this evidence, it appears that participatory community interventions can help households to completely smooth even significant crop losses. Our results also suggest that the crop loss incidence (most commonly used in the literature) is a rather blunt measure of the adverse event, and data on its severity also provides useful information.

The Figure also presents the results for total household food consumption. The results are qualitatively similar to those reported for household non-durable consumption. Households reduce food consumption in order to deal with more severe crop losses, and the intervention appears to aid households in protecting their food consumption, as evidenced by the positive coefficients on the interaction terms. All coefficients on the interaction terms have the expected sign and are statistically significant (at the 10% level) in most specifications.

5.2 Possible Mechanisms

This multi-faceted intervention could have improved consumption smoothing through a number of possible channels, as described in Section 2.2. In this sub-section, we in-

investigate the underlying channels by first considering the impacts of the intervention on the tools used by households to smooth consumption. Thereafter, we explore how any improvements were realised, by considering intervention impacts intermediate outcomes consistent with each channel. Data limitations mean that we are unable to observe all possible intermediate outcomes, and so this exploration of mechanisms provides suggestive and speculative evidence only.

5.2.1 Impacts on Tools through which Consumption Smoothing is Achieved

In this setting, households rely on three key channels to smooth consumption: self-insurance through drawing down savings and assets (Paxson, 1992; Rosenzweig and Wolpin, 1993; Fafchamps et al., 1998); adjusting labour supply (Kochar, 1999; Rose, 2001); and informal risk sharing through informal transfers, loans and gifts from social connections (Townsend, 1994). The intervention could have improved households' ability to smooth consumption through any of these channels.

We start by considering the intervention impacts on these channels. Unfortunately, our data does not include reliable measures of informal risk sharing. However, we can investigate its importance indirectly, by studying the extent to which the intervention facilitated consumption smoothing by either improving households' capacities to self-insure, and/or to adjust labour supply.

Savings and Assets If households rely mainly on self-insurance to smooth consumption, the reduction in maternal mortality due to the intervention could have left them with a larger stock of savings and assets to draw down to protect consumption following adverse crop losses. If this is the case, we should observe *larger* reductions in savings and asset stocks in the treated communities following an adverse crop loss. Since production and consumption decisions for farm households in developing countries are known to be interlinked, the use of productive assets for consumption smoothing has significant

consequences for long-run well-being and efforts to alleviate poverty (Banerjee and Newman, 1993; Aghion and Bolton, 1997).

We thus investigate whether and how flows of savings and asset holdings respond to crop losses in the treatment and control groups by estimating regression 1 with changes in savings and assets (large and small animals) as the dependent variable.²⁹ To alleviate the possibility that findings are driven by false positives detected as a result of using many indicators for savings/assets, we combine the three measures into a single index following Anderson (2008).³⁰

Figure 4 displays the results, starting with impacts on the overall index, before reporting impacts on each measure. The magnitudes of the index are difficult to interpret; however, the sign and statistical significance provides indication of whether the intervention altered the use of savings and assets for consumption smoothing. The left panel shows that following a crop loss, the savings index falls in the control areas (p-value = 0.038), suggesting that households cope with crop losses, at least partially, using self-insurance. However, the coefficient on the interaction term with the treatment indicator is positive and statistically significant (p-value = 0.02), indicating that households in treatment areas do not draw down assets and savings as much to cope with the crop loss.

Breaking this down by cash savings, large animals (e.g. bullocks and cows) and small animals (e.g. chickens and rabbits), we see that the impacts are driven primarily by cash savings and large animals: Households experiencing a crop loss in the treated areas hold one more large animal than those in control areas (p-value = 0.049), suggesting that these households no longer rely on livestock sales to smooth their consumption. These results

²⁹The surveys asked about the numbers of animals owned, but didn't collect any information on their values since survey piloting indicated that respondents struggled to estimate the value of their livestock. We are thus unable to calculate the total monetary value of livestock. Around 17% of households have any cash savings in this sample. By contrast, livestock ownership is more common with around 62% of households owning at least one small animal, and 43% owning at least one large animal.

³⁰The summary index is calculated as a weighted average of the standardized values of the outcome variables (re-defined so that higher values imply a better/more desirable outcome). The weights are calculated to maximize the amount of information captured by the index so that outcomes that are highly correlated with each other receive less weight in the calculation.

have important implications for future household well-being since higher savings are linked with higher agricultural input use and yields in this context (Gine et al., 2012). Thus, the improvements in consumption smoothing did not arise as a result of improved self-insurance.

Labour Supply Households can also smooth consumption by adjusting ex-post labour supply by seeking other jobs and/or increasing hours worked to cover shortfalls from the crop loss. The intervention could have increased the scope for households to use labour supply to smooth consumption by improving health, thereby allowing household members to work more, and by facilitating information sharing about job opportunities through additional social interactions fostered by the regular group meetings.

We assess whether this is the case by testing whether there is a larger increase in labour supply of male and female adults (aged 16-65 years) in response to the crop loss in treatment communities than in control communities. We do so by estimating specifications similar to equation (1), with changes over time in indicators of labour supply for male and female adults as the dependent variable. If the intervention facilitates the use of labour supply as a consumption smoothing device, we would expect the coefficient γ to be positive and statistically significantly different from zero.

We consider impacts on both the extensive and intensive margins, with the following dependent variables: (i) changes over time in whether or not an individual has a job, (ii) changes over time in whether the individual takes a second job; and (iii) changes in hours worked in the main and second jobs respectively.³¹ The results, presented in Figure 5, show no evidence that labour supply is used as a consumption smoothing device in this context for either male or female adults.³²

³¹We also find no evidence that the intervention increased labour supply for adults or children. These results are available on request.

³²In Appendix Table 11, we analyze whether child labour could have been used as a consumption smoothing device. While the intervention sought to improve the health of infants, who would not be old enough to engage in these activities themselves, better infant health could free up the time of older siblings (who potentially would have contributed to child care or household chores) to work following an adverse

The evidence presented in this section shows that the impacts on consumption smoothing are not driven by improved self-insurance through drawing down of savings and assets, or through adjusting labour supply. This suggests that informal risk sharing is likely to have improved. A question then is how did the intervention improve risk sharing, which we investigate next.

5.2.2 How intervention might have improved informal risk sharing

We consider three possible explanations for how the intervention might have improved informal risk sharing. First, it could have improved health, thereby reducing health shocks and increasing the capacity of a household's social connections to help in the event of crop losses. Second, though this was not a core objective of the groups, it is possible that they provided food from the group gardens, or made monetary transfers or loans from income raised through various activities to households experiencing crop losses. Finally, by requiring group members to interact frequently, the intervention design could have reduced barriers to informal risk sharing. We examine the relevance of each of these explanations in turn.

Improved Health Lewycka et al. (2013) document a reduction in maternal mortality of 74% due to the intervention. While large, maternal deaths are relatively rare, so that an effect of this magnitude corresponds to approximately 3 fewer maternal deaths among treated households in our sample. This effect, on its own, is too small to explain the entire consumption smoothing effect we uncover: A back of the envelope calculation, detailed in Appendix B, indicates that household consumption for these 3 households should have *increased by over 28 times* if the entire effect of the treatment on consumption smoothing was driven by these 3 averted deaths. Studies investigating the consumption growth impacts of adult deaths in developing countries find effects ranging from a 7% drop within

crop loss. We find no evidence that child labour is used to smooth consumption in this context.

5 years of the death in rural Tanzania (Beegle et al., 2008) to an increase documented in Indonesia (Grimm, 2010). This suggests that such a large impact on consumption is highly implausible. The intervention impact on consumption smoothing, thus cannot be entirely explained by the reduction in maternal mortality.

However, the intervention could also have improved adult and child health more generally, thereby leading to fewer health shocks. To investigate this possibility, we first consider intervention impacts on several self-reported indicators of adult health in Figure 6 for males and females respectively. To reduce the chances of detecting a false positive (given the number of outcomes tested), and to improve power, we combine the different measures into a Health Index following the method of Anderson (2008). We report impact estimates for the index and each of the constituent measures. The results provide no evidence that the intervention improved adult health, for both symptoms experienced in a short interval prior to the survey, and indicators relating to activities of daily living, which are better at capturing longer-term health issues.

In the online appendix, we present the intervention impacts on child height and weight for children aged less than 48 months. The intervention would be most likely to improve the health of children in this age range, who would have been less than six months old (or not born) when the group cycles began. Child height and weight are considered to capture long-term and short-term health respectively. We use standardized height-for-age, weight-for-age, and weight-for-height z-scores calculated using the WHO norms. The intervention impacts are displayed in Table 12 in the appendix, which shows that the treatment had no statistically significant effects on child physical growth. Table 13, shown in the appendix, also shows no statistically significant effects of the intervention on maternal reported measures of child morbidity for children aged 5-16 years, whose health was not directly targeted by the intervention.

Thus, the intervention did not have any economically or statistically significant effects on either adult and child morbidity. Moreover, since the improvements in maternal mortality

are unlikely to explain the entire improvement in consumption smoothing, it is likely that other aspects of the groups also had a role to play.

Effects of group strategies Next, we consider whether the strategies implemented by the women's groups directly improved households' consumption smoothing. Though the choice of strategy was endogenous to each group, strategies that generated income, and community gardens growing green leafy vegetables, were very popular across groups. Though the food and income arising from these strategies was primarily targeted at initiatives to improve the health of pregnant women and infants (e.g. to fund transport for pregnant women to get to the hospital or health centre for delivery), they could have been distributed among those affected by crop losses, helping them smooth consumption.

However, our data indicates that this is unlikely: in our data, only 7 households in the treated communities (pooled across both rounds of data), of whom 2 experienced a crop loss, report receiving any loans, gifts or transfers from the women's group. A further 6 households, none of whom reported experiencing a crop loss, indicated that they consumed vegetables grown in the Mai Mwana community garden. Thus, the group strategies are unlikely to explain the improvements in consumption smoothing.

Social Interactions Finally, we consider whether improvements in informal risk sharing were due to improved social interactions. To test this, we focus on a measure of interactions *outside* the participatory women's groups – one-to-one chats with family or friends about a range of health and non-health related topics. In doing so, we capture interactions with 'strong' connections, who are likely to be informal risk sharing partners.³³ Interestingly, a large fraction of respondents reported not chatting with anyone about these topics. This measure provides an indication of changes in the probability of interacting with relatives and friends as a whole. However, we are unable to disentangle

³³Existing literature indicates that the extended family is an important source of risk sharing in rural areas of developing countries (Rosenzweig, 1988; Foster and Rosenzweig, 2001; Angelucci et al., 2018) and in the study context in particular (Fitzsimons et al., 2018).

whether any change comes from more intense use of existing contacts; or from interactions with a broader group of social contacts; or a combination of the two. Furthermore, it is also possible that the intervention increased interaction with new social contacts but crowded out existing ones. However, without detailed information on the identities of who an individual chatted with, it will not be possible to disentangle these effects.

To improve power, we combine indicators of whether a woman chatted with family or friends on a one-to-one basis on each of 6 topics into two indices following the method of Anderson (2008).³⁴ We construct indices for any chats about health related topics (pregnancy and birth delivery, breastfeeding and post-breastfeeding nutrition and family planning), and any chats about non-health topics (local governance, work opportunities and obtaining credit). We estimate regressions of the following form to identify intervention impacts on social interactions outside the group:

$$y_{hvt} = \beta_0 + \beta_1 D_v + X_{hvct} \pi + \eta_t + v_{hvt} \quad (3)$$

where y_{hvt} is the index of chats, D_v is the treatment indicator, X_{hvct} is a vector of observed characteristics including the age and education of the main respondent, and cluster-level baseline controls including average female education, marriage rates and average household size and η_t includes controls for month-year of interview. If the intervention changed social interactions with family and friends outside the groups, $\beta_1 > 0$.

Figure 7 displays the results. They indicate that the intervention has positive effects on one-to-one chats on both health and non-health topics, though only the effect on the former is statistically significant. The coefficient for the intervention impact on non-health chats is smaller in magnitude than that on health related chats, and more noisily esti-

³⁴The survey main respondent was asked to give the number of relatives, friends, and health visitors/facilitators she chatted with in the week prior to the survey on a one-to-one basis about pregnancy or birth delivery, breastfeeding or post-breastfeeding nutrition, family planning, local governance, work opportunities and obtaining credit. The second follow-up survey also asked women to provide the number of acquaintances they spoke to about these topics on a one-to-one basis.

mated leading to the statistical insignificance. These increased interactions could have improved the sharing of information, thereby alleviating information related contracting frictions (for example, conversations originally about health related topics could also involve sharing of “gossip” which may spread information about specific households experiencing a shock -and those lying about them-, thereby reducing frictions associated with hidden income); or enforcement constraints (since more frequent interactions might make it more costly for a household to renege on an arrangement); or transaction costs.

To assess whether these increased social interactions (in terms of chats) facilitated better consumption smoothing, we analyze whether the likelihood of chatting about a topic relevant for risk sharing – credit sources – increases more following crop losses in the treatment group. To do so, we estimate a specification similar to Equation (1) with changes in whether or not a household chatted with friends and family about sources of credit as the dependent variable. The results, presented in Figure 8, show that following a crop loss, respondents in control communities either do not change, or even reduce their likelihood of chatting about sources of credit. By contrast, respondents in treated communities that experience a crop loss were more likely to start chatting with friends and family about sources of credit. Such chats could have facilitated access to (informal) credit, allowing households to smooth consumption.

This evidence thus suggests that the increased interactions due to the intervention may have improved informal risk sharing and consumption smoothing.

5.3 Robustness Checks

We consider the robustness of the findings relating to the crop loss intensity measure to alternative ways of normalizing the income lost due to the crop loss. Table 3 presents the findings using three alternatives. Columns 1 and 5 present the coefficients for our preferred measure of crop loss intensity – the crop loss normalized by predicted consump-

tion. Columns 2 and 6 display the findings using changes in income loss itself. Columns 3 and 7 display results taking changes in log of the income loss + 1 (to ensure that it is defined for all households), while columns 4 and 8 use changes in the inverse hyperbolic sine transform of the loss. The latter transformation has the advantage of being insensitive to outliers and being defined for zero values. We obtain similar qualitative findings across the alternative measures of shock intensity: consumption growth falls after a more intense shock in the control clusters. However, there is a positive treatment effect, so that even more intense crop losses are almost completely smoothed in terms of food and total non-durable consumption. Thus, our findings are not sensitive to the proposed normalization.

6 Conclusion

This paper studies the spillover effects of a participatory community based health intervention in rural Malawi on an outcome not directly targeted by the intervention: household consumption smoothing following idiosyncratic crop losses. Group-based community interventions are considered to be a cost-effective way of delivering interventions in low-income settings. However, design features of these programs could also influence dimensions of household and community behavior beyond those targeted by the intervention, generating important spillovers.

To provide causal estimates, we exploit a cluster randomized trial in which a participatory women's group intervention targeting maternal and child health was randomly allocated to some groups of villages (clusters), while a similar number of clusters received no intervention. The intervention encouraged community members, particularly women, to meet on a regular basis to discuss maternal and child health related issues, before mobilizing the community to combat the identified issues. The groups could have improved consumption smoothing through a number of channels, including better health, and by in-

creasing social interactions (and subsequently alleviating barriers to risk sharing) through the formation of groups that met regularly, and that aimed to involve a broad swathe of the community.

We find that consumption smoothing – measured as the response of changes in log household (non-durable) consumption to idiosyncratic crop loss events, net of village level aggregate shocks – worsens following a crop loss for households in control clusters. Those in intervention clusters, however, are able to compensate for this loss and almost perfectly insure their consumption. Crucially, we show that the intervention did not change the distribution of crop loss incidence or severity.

Examining mechanisms underlying this effect, we find that they are not driven by improvements in self-insurance, or adjustments to labour supply. Instead, improvements must come from better informal risk sharing. We find suggestive evidence that these improvements are likely to be a result of the intervention increasing social interactions, which could have alleviated a number of contracting frictions impeding informal risk sharing in this context. Previous research by Feigenberg et al. (2013) showed in the context of microcredit groups that increased social interactions lead to a reduced likelihood of loan default in urban India. Our findings complement this work by showing that participatory community groups, formed to improve health, which increased social interactions outside the groups can also improve consumption smoothing.

References

- Abramovsky, L., B. Augsburg, M. Lührmann, F. Oteiza, and J. P. Rud (2019, June). Community matters: heterogeneous impacts of a sanitation intervention. IFS Working Papers W19/11, Institute for Fiscal Studies.
- Aghion, P. and P. Bolton (1997). A theory of trickle-down growth and development. *Review of Economic Studies* 64(2), 151–172.
- Ali, N. and D. Miller (2013). “Enforcing Cooperation in Networked Societies”. *mimeo, University of Michigan*.
- Ambrus, A. and M. Elliott (2018). Investments in social ties, risk sharing and income inequality. *mimeo, Duke University*.
- Ambrus, A., W. Y. Gao, and P. Milan (2019). Informal risk sharing with local information. *mimeo, Duke University*.
- Ambrus, A., M. Mobius, and A. Szeidl (2014, January). Consumption risk-sharing in social networks. *American Economic Review* 104(1), 149–82.
- Anderson, M. (2008). “Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects”. *Journal of the American Statistical Association* 103(484), 1484–1495.
- Angelucci, M. and G. De Giorgi (2009, March). Indirect effects of an aid program: How do cash transfers affect ineligibles’ consumption? *American Economic Review* 99(1), 486–508.
- Angelucci, M., G. De Giorgi, and I. Rasul (2018, November). “Resource Pooling Within Family Networks: Insurance and Investment”. *Economic Journal* 128, 2613–51.
- Banerjee, A. V., R. Banerji, E. Duflo, R. Glennerster, and S. Khemani (2010, February). Pitfalls of Participatory Programs: Evidence from a Randomized Evaluation in Education in India. *American Economic Journal: Economic Policy* 2(1), 1–30.

- Banerjee, A. V. and A. Newman (1993). Occupational choice and the process of development. *Journal of Political Economy* 101(2), 274–98.
- Beegle, K., J. D. Weerdt, and S. Dercon (2008, January). Adult Mortality and Consumption Growth in the Age of HIV / AIDS. *Economic Development and Cultural Change* 56(2), 299–326.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). “How Much Should We Trust Differences-in-Differences Estimates?”. *Quarterly Journal of Economics* 119(1), 249–275.
- Besley, T. (1995). “Non-Market Institutions for Credit and Risk-Sharing in Low-Income Countries”. *Journal of Economic Perspectives* 9(3), 115–127.
- Blattman, C., E. P. Green, J. Jamison, M. C. Lehmann, and J. Annan (2016, April). The returns to microenterprise support among the ultrapoor: A field experiment in postwar uganda. *American Economic Journal: Applied Economics* 8(2), 35–64.
- Bourles, R., Y. Bramouille, and E. Perez-Richet (2020). Altruism and risk sharing in networks. *mimeo, Aix-Marseille School of Economics*.
- Breza, E. and A. Chandrasekhar (2015). “Social Networks, Reputation and Commitment: Evidence from a Savings Monitors Field Experiment”. *mimeo, Columbia University*.
- Cameron, A. C. and D. L. Miller (2015). A practitioner’s guide to cluster-robust inference. *Journal of Human Resources* 50(2), 317–372.
- Cameron, C., J. Gelbach, and D. Miller (2008). “Bootstrap-Based Improvements for Inference with Clustered Errors”. *Review of Economics and Statistics* 90, 414–427.
- Cameron, L., S. Oliva, and M. Shah (2015). "Initial Conditions Matter: Social Capital and Participatory Development". *IZA Working Paper 9563*.
- Casey, K., R. Glennerster, and E. Miguel (2012). Reshaping institutions: Evidence on aid impacts using a preanalysis plan*. *The Quarterly Journal of Economics* 127(4), 1755–1812.

- Chandrasekhar, A., H. Larreguy, and C. Kinnan (2015). "Social Networks as Contract Enforcement: Evidence from a lab experiment in the field". *mimeo, Stanford University*.
- Chetty, R. and A. Looney (2006, December). Consumption smoothing and the welfare consequences of social insurance in developing economies. *Journal of Public Economics* 90(12), 2351–2356.
- Cochrane, J. H. (1991, October). A Simple Test of Consumption Insurance. *Journal of Political Economy* 99(5), 957–76.
- Donald, S. and K. Lang (2007). "Inference with Differences-in-Differences and Other Panel Data". *Review of Economics and Statistics* 89, 221–233.
- Fafchamps, M. and E. L. Ferrara (2012). Self-Help Groups and Mutual Assistance: Evidence from Urban Kenya. *Economic Development and Cultural Change* 60(4), 707 – 733.
- Fafchamps, M., C. Udry, and K. Czukas (1998, April). Drought and saving in West Africa: are livestock a buffer stock? *Journal of Development Economics* 55(2), 273–305.
- Feigenberg, B., E. Field, and R. Pande (2013). "The Economic Returns to Social Interaction: Experimental Evidence from Microfinance". *The Review of Economic Studies* 80(4), 1459–1483.
- Fitzsimons, E., B. Malde, A. Mesnard, and M. Vera-Hernandez (2016). "Nutrition, information and household behaviour: experimental evidence from Malawi". *Journal of Development Economics Forthcoming*.
- Fitzsimons, E., B. Malde, and M. Vera-Hernandez (2018). "Group Size and the Efficiency of Informal Risk Sharing". *Economic Journal* 128(612), F575–F608.
- Foster, A. and M. Rosenzweig (2001). "Imperfect Commitment, Altruism and the Family: Evidence from Transfer Behavior in Low-Income Countries". *The Review of Economics and Statistics* 83(3), 398–407.

- Gine, X., J. Goldberg, and D. Yang (2012). "Credit Market Consequences of Improved Personal Identification: Field Experimental Evidence from Malawi". *The American Economic Review* 102(6), 2923–54.
- Grimm, M. (2010, April). Mortality Shocks and Survivors' Consumption Growth. *Oxford Bulletin of Economics and Statistics* 72(2), 146 – 171.
- Guiso, L., P. Sapienza, and L. Zingales (2004, June). The Role of Social Capital in Financial Development. *American Economic Review* 94(3), 526–556.
- Guiteras, R. and B. K. Jack (2014). "Incentive, Selection and Productivity in Labor Markets: Evidence from Rural Malawi". *NBER Working Paper* 19825.
- Jack, B. K. (2013, July). Private information and the allocation of land use subsidies in malawi. *American Economic Journal: Applied Economics* 5(3), 113–35.
- Jack, W. and T. Suri (2014, January). Risk sharing and transactions costs: Evidence from kenya's mobile money revolution. *American Economic Review* 104(1), 183–223.
- Kinnan, C. (2014). "Distinguishing Barriers to Insurance in Thai Villages". *mimeo, Northwestern University*.
- Kochar, A. (1999, February). Smoothing Consumption by Smoothing Income: Hours-of-Work Responses to Idiosyncratic Agricultural Shocks in Rural India. *The Review of Economics and Statistics* 81(1), 50–61.
- Kocherlakota, N. (1996). "Implications of Efficient Risk Sharing without Commitment". *Review of Economic Studies* 63(4), 595–609.
- Lewycka, S. (2011). "Reducing Maternal and Neonatal Deaths in Rural Malawi: Evaluating the Impact of a Community-based Women's Group Intervention". Ph. D. thesis, University College London.

- Lewycka, S., C. Mwansambo, P. Kazembe, T. Phiri, A. Mganga, M. Rosato, H. Chapota, F. Malamba, S. Vergnano, M.-L. Newell, D. Osrin, and A. Costello (2010). "A cluster randomised controlled trial of the community effectiveness of two interventions in rural Malawi to improve health care and to reduce maternal, newborn and infant mortality". *Trials* 11:88.
- Lewycka, S., C. Mwansambo, M. Rosato, P. Kazembe, T. Phiri, A. Mganga, H. Chapota, F. Malamba, E. Kainja, M. Newell, G. Greco, A. Brannstrom, J. Skordis-Worrall, S. Vergnano, D. Osrin, and A. Costello (2013). "Effect of women's groups and volunteer peer counselling on rates of mortality, morbidity and health behaviours in mothers and children in rural Malawi (MaiMwana): a factorial, cluster-randomised controlled trial". *The Lancet* 381, 1721–35.
- Ligon, E. (1998). "Risk Sharing and Information in Village Economies". *Review of Economic Studies* 65(4), 847–864.
- Ligon, E., J. Thomas, and T. Worrall (2003). "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies". *The Review of Economic Studies* 69(1), 209–244.
- Maccini, S. and D. Yang (2009, June). Under the weather: Health, schooling, and economic consequences of early-life rainfall. *American Economic Review* 99(3), 1006–26.
- Mansuri, G. and V. Rao (2004). Community-based and -driven development: A critical review. *The World Bank Research Observer* 19(1), 1–39.
- Miller, C. and M. Tsoka (2012). Cash transfers and children's education and labor among malawi's poor. *Development Policy Review* 30(4), 499–522.
- Munthali, A. (2002). "Adaptive Strategies and Coping Mechanisms of Families and Communities Affected by HIV/AIDS in Malawi". *Draft paper prepared for the UNRISD project HIV/AIDS and Development*.

- Paxson, C. H. (1992, March). Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand. *American Economic Review* 82(1), 15–33.
- Peters, P., P. Walker, and D. Kambewa (2008). "Striving for Normality in a Time of AIDS in Malawi". *The Journal of Modern African Studies* 46(4), 659–687.
- Pickering, A. J., H. Djebbari, C. Lopez, M. Coulibaly, and M. L. Alzua (2016). Effect of a community-led sanitation intervention on child diarrhoea and child growth in rural mali: a cluster-randomised controlled trial. *The Lancet Global Health* 3(11), e701–e711.
- Prost, A., T. Colbourn, N. Seward, K. Azad, A. Coomarasamy, A. Copas, T. A. J. Houweling, E. Fottrell, A. Kuddus, S. Lewycka, C. MacArthur, D. Manandhar, J. Morrison, C. Mwansambo, N. Nair, B. Nambiar, D. Osrin, C. Pagel, T. Phiri, A.-M. Pulkki-Brännström, M. Rosato, J. Skordis-Worrall, N. Saville, N. S. More, B. Shrestha, P. Tripathy, A. Wilson, and A. Costello (2013). Women’s groups practising participatory learning and action to improve maternal and newborn health in low-resource settings: a systematic review and meta-analysis. *The Lancet* 381(9879), 1736–1746.
- Roodman, D. (2015). BOOTTEST: Stata module to provide fast execution of the wild bootstrap with null imposed. Statistical Software Components, Boston College Department of Economics.
- Rosato, M. (2012). *How does community mobilisation through MaiMwana women’s groups work? Addressing the social determinants of mother and child health in rural Malawi*. Ph. D. thesis, University College London.
- Rosato, M., S. Lewycka, C. Mwansambo, P. Kazembe, and A. Costello (2009). “Women’s Groups’ Perceptions of Neonatal and Infant Health Problems in Rural Malawi”. *Malawi Medical Journal* 21(4), 168–173.
- Rosato, M., F. Malamba, B. Kunyenge, T. Phiri, C. Mwansambo, P. Kazembe, A. Costello, and S. Lewycka (2012). Strategies developed and implemented by women’s groups to

- improve mother and infant health and reduce mortality in rural malawi. *International Health* 4(3), 176–184.
- Rosato, M., C. W. Mwansambo, P. N. Kazembe, T. Phiri, Q. S. Soko, S. Lewycka, B. E. Kunyenge, S. Vergnano, D. Osrin, M.-L. Newell, and A. M. de L Costello (2006). "Women's groups' perceptions of maternal health issues in rural Malawi". *The Lancet* 368(9542), 1180–1188.
- Rose, E. (2001, April). Ex ante and ex post labor supply response to risk in a low-income area. *Journal of Development Economics* 64(2), 371–388.
- Rosenzweig, M. (1988). "Risk, Private Information and the Family". *The American Economic Review* 78(2), 245–250.
- Rosenzweig, M. R. and K. I. Wolpin (1993, April). Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investment in Bullocks in India. *Journal of Political Economy* 101(2), 223–44.
- Spear, S. E. and S. Srivastava (1987). On repeated moral hazard with discounting. *The Review of Economic Studies* 54(4), 599–617.
- Townsend, R. (1994). "Risk and Insurance in Village India". *Econometrica* 62(3), 539–591.
- Trinitapoli, J., S. Yeatman, and J. Fledderjohann (2014). "Sibling Support and the Educational Prospects of Young Adults in Malawi". *Demographic Research* 30, 547–578.
- Udry, C. (1995, December). Risk and Saving in Northern Nigeria. *American Economic Review* 85(5), 1287–1300.
- White, H., R. Menon, and H. Waddington (2018). Community-driven development: does it build social cohesion or infrastructure? A mixed-method evidence synthesis. Technical report, International Initiative for Impact Evaluation.

Wooldridge, J. (2004). "Cluster-Sample Methods in Applied Econometrics". *The American Economic Review* 93(2), 133–138.

7 Tables

Table 1: Sample Balance

	Sample Drawn			Sample Interviewed			Analysis Sample		
	Control	Diff: T-C	p-value	Control	Diff: T-C	p-value	Control	Diff: T-C	p-value
Main floor material: dirt, sand or dung	0.914	-0.015	0.498	0.916	-0.007	0.765	0.922	-0.016	0.492
Main roof material: natural	0.853	-0.002	0.937	0.857	-0.008	0.819	0.868	-0.021	0.482
Agricultural HH (yes=1)	0.995	-0.009	0.492	0.995	0.000	0.921	0.994	0.000	0.991
Traditional pit toilet (yes=1)	0.772	0.017	0.717	0.791	-0.009	0.855	0.787	-0.011	0.797
Piped water (yes=1)	0.011	0.110	0.132	0.009	0.108	0.171	0.011	0.103	0.248
Number of HH members	5.772	-0.129	0.446	5.848	-0.144	0.293	5.911	-0.262	0.124
Number of sleeping rooms	2.116	-0.003	0.963	2.152	-0.007	0.931	2.177	-0.041	0.757
HH has electricity (yes=1)	0.002	0.002	0.278	0.002	-0.001	0.603	0.003	-0.002	0.561
HH has radio (yes=1)	0.629	-0.015	0.643	0.641	-0.027	0.509	0.672	-0.052	0.116
HH has bicycle (yes=1)	0.508	-0.001	0.965	0.512	0.003	0.927	0.531	-0.007	0.901
HH has motorcycle (yes=1)	0.008	-0.002	0.723	0.007	-0.001	0.841	0.006	0.002	0.827
HH has car (yes=1)	0.006	-0.002	0.400	0.007	-0.006	0.030**	0.006	-0.005	0.144
HH has paraffin lamp (yes=1)	0.926	0.026	0.427	0.926	0.025	0.354	0.952	-0.001	0.981
HH has oxcart (yes=1)	0.058	-0.008	0.577	0.059	-0.006	0.713	0.062	-0.014	0.411
Main respondent education: none (yes=1)	0.294	-0.029	0.462	0.318	-0.035	0.420	0.342	-0.046	0.404
Main respondent education: primary (yes=1)	0.641	0.026	0.462	0.622	0.040	0.278	0.598	0.053	0.298
Main respondent education: secondary or more (yes=1)	0.066	0.003	0.801	0.060	-0.003	0.813	0.060	-0.006	0.386
N	1249	1247		846	790		629	620	

Notes: * significant at the 10% level, ** at the 5% level. p-values computed using the wild cluster bootstrap-t method recommended by Cameron et al. (2008). The left (middle) [right] panel displays the control group mean, difference in means for the treated and control groups, and the p-value of a t-test of this difference for the sample drawn for the first follow-up in 2008-09 (sample interviewed) [final analysis sample].

Table 2: Intervention Effects on Crop Loss Incidence and Intensity

	[1]	[2]	[3]
	crop=0 or 1	$\Delta crop$ = crop - income lost	crop = loss/pred cons
D_v	-0.005 [0.111] {0.999}	-0.107 [1.906] {0.991}	-0.0702 [0.138] {0.663}
Observations	1,249	1,227	1,225
R-squared	0.000	0.000	0.001
Intraclass Correlation	0.167	0.024	0.062
Mean (wave 1)	0.353	5.588	0.398
Std dev (wave 1)	0.478	20.858	1.069
Mean (wave 2)	0.215	2.100	0.176
Std dev (wave 2)	0.411	6.390	0.526

Notes to Table: Sample includes households resident in the same village in both follow-up surveys. *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered at the cluster level. Wild cluster bootstrap p-values reported in curly braces. Income loss is measured in 1,000's of Malawi Kwacha. Predicted consumption is calculated as a function of pre-intervention education.

Table 3: Alternative measures for crop loss intensity

	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
	$\Delta \log(C_{tot})$	$\Delta \log(F_{tot})$						
	crop = loss/pre-shock cons	crop = loss/pre-shock cons	crop = loss/pre-shock cons	crop = loss/pre-shock cons	crop = loss/pre-shock cons	crop = loss/pre-shock cons	crop = loss/pre-shock cons	crop = loss/pre-shock cons
$\Delta crop(\alpha)$	-0.0592** [0.0220] {0.014}	-0.0048** [0.0017] {0.016}	-0.0101* [0.0057] {0.082}	-0.0392** [0.0180] {0.044}	-0.0485* [0.0278] {0.050}	-0.0039* [0.0022] {0.032}	-0.0102 [0.0063] {0.124}	-0.0359* [0.0203] {0.094}
$\Delta crop * D_v(\gamma)$	0.0411** [0.0227] {0.020}	0.0033** [0.0018] {0.026}	0.0146* [0.0083] {0.080}	0.0447* [0.0260] {0.080}	0.0308 [0.0288] {0.287}	0.0026 [0.0023] {0.242}	0.0154* [0.0089] {0.070}	0.0449* [0.0291] {0.094}
$\alpha + \gamma$	-0.0180* [0.088]	-0.001 [0.148]	0.004 [0.392]	0.006 [0.721]	-0.0177 [0.281]	-0.001 [0.306]	0.005 [0.380]	0.009 [0.615]
Observations	1221	1221	1221	1221	1219	1219	1219	1219
R-Squared	0.4121	0.4125	0.4104	0.4107	0.4016	0.4018	0.4015	0.4012

Notes: Sample includes households resident in the same village in both follow-up surveys. *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered at the cluster level reported in brackets. Wild cluster bootstrap p-values reported in curly braces. D_v is the treatment indicator.

8 Figures

Figure 1: Village level proportion of crop losses among all sampled households

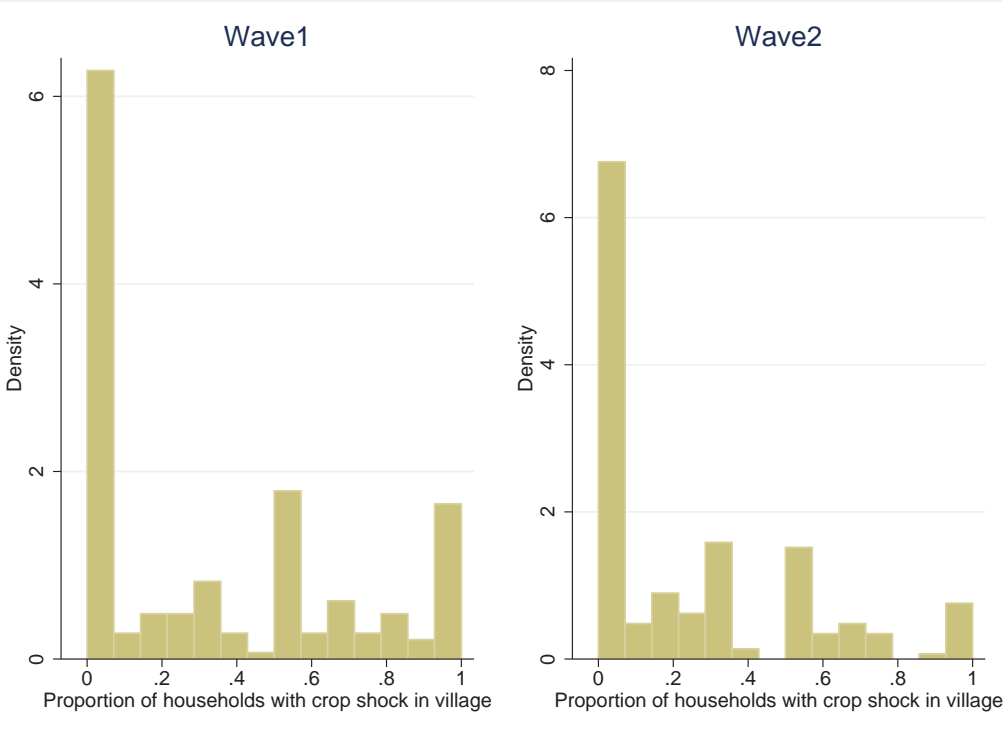
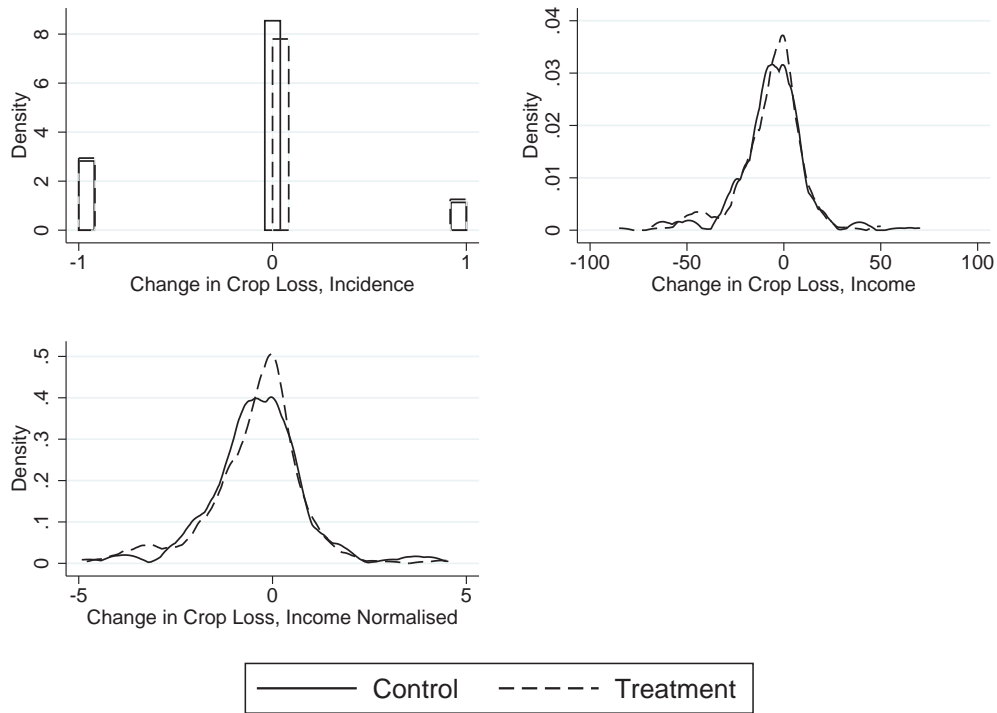



Figure 2: Distributions of the changes in crop loss, by treatment group




Notes: This Figure plots the distributions of the change in incidence of the crop loss (top left panel), change in income lost from the crop loss (top right panel) and change in income lost from the crop loss normalized by predicted consumption (bottom left panel) separately for households in the control (solid line) and treatment (dashed line) groups. The combined Kolmogorov-Smirnov test statistic values (p-value) for the equality of distributions are respectively 0.0194 (1.000), 0.0533 (0.348) and 0.0570 (0.272) respectively.

Figure 3: Intervention impacts on consumption smoothing




Notes to Figure: Sample includes households resident in the same village in both follow-up surveys. 95% confidence intervals calculated using a wild cluster bootstrap-t procedure reported. 'Incidence' is measured as the change in whether or not a household experienced a crop loss; 'Intensity' is the change in the income lost from the crop loss normalised by predicted baseline consumption.

Figure 4: Savings as a consumption smoothing device




Notes: Sample includes households resident in the same village in both follow-up surveys. 95% confidence intervals calculated using a wild cluster bootstrap-t procedure reported. 'Savings' capture cash savings in Malawi Kwacha; 'Small Animals' include livestock such as poultry and rabbits while 'Large Animals' include cattle, goats and sheep.

Figure 5: Labour supply as a consumption smoothing device
 (a) Males




(b) Females



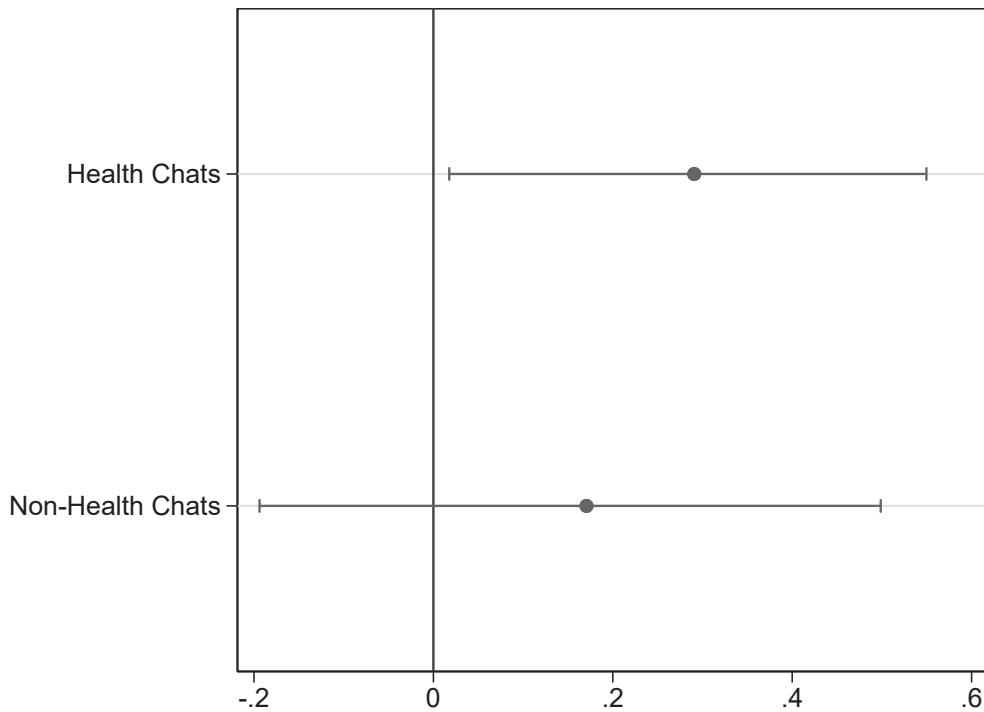
Notes: Sample includes households resident in the same village in both follow-up surveys. 95% confidence intervals calculated using a wild cluster bootstrap-t procedure reported.

Figure 6: Impacts on adult health




Notes: Sample in the left (right) panel includes male (female) individuals aged 15-64 living in households that were resident in the same village in both follow-up surveys.. 95% confidence intervals calculated using a wild cluster bootstrap-t procedure reported. The indicators for 'No diarrhoea', 'No fever', 'No cough', 'No chills' and 'No vomiting' refer to the 30 days prior to the survey. The Health Index aggregates the other variables using the method of Anderson (2008).

Figure 7: Impacts on chats



Notes: Sample includes households resident in the same village in both follow-up surveys. 95% confidence intervals calculated using a wild cluster bootstrap-t procedure. 'Health Chats' is an index aggregating binary variables indicating whether the main respondent chatted with a friend or family member on a one-to-one basis about pregnancy and birth, breastfeeding and post-breastfeeding nutrition, and family planning in the week prior to the survey. 'Non-health chats' constructs a similar index for chats related to credit, jobs and local politics. All indices are constructed according to the method of Anderson (2008). All regressions include controls for age, age-squared and education for the main respondent, month of survey dummies, household demographics, and baseline cluster-level variables including proportion of women with secondary schooling, average female marriage rate and average household size.

Figure 8: Changes in probability of chatting about credit by crop loss



Notes: Sample includes households resident in the same village in both follow-up surveys. 95% confidence intervals calculated using a wild cluster bootstrap-t procedure reported.