

The Impact of Micro-Banking on Health: Evidence
from Self-Help Group Involvement and Child Nutrition

Madeline McKelway¹

Professor Erica Field, Faculty Advisor

Duke University
Durham, North Carolina

2014

¹In the fall of 2014, I will begin doctoral studies in the PhD program in economics at MIT. I can be reached at mad-die.mckelway@gmail.com.

Acknowledgements

I am deeply indebted to Olivia McKelway, Jaideep Vyas, Jasminder Vyas, and the administration and staff of P.E.D.O. for their assistance in data collection. I would also like to thank my advisor, Professor Erica Field, as well as Professors Duncan Thomas, Michelle Connolly, and Manoj Mohanan for their valuable feedback. And finally, I am grateful for the financial support provided by Duke University Deans' Summer Research Fellowship.

Abstract

Low income is only one financial problem that poor families in developing countries face; impoverished households must also face irregularity of their low incomes. Self-help groups (SHGs) can enhance consumption stability by relaxing savings and credit constraints. In this study, I investigate the extent to which SHGs improve a particular dimension of household wellbeing: child nutrition. I analyze households affiliated with the SHGs started by the People's Education and Development Organization (P.E.D.O.) in rural Rajasthan, India. Children who had greater levels of exposure to household SHG membership at a young age have healthier anthropometric statuses than their siblings who had relatively less. This relationship does not appear to be driven by events coinciding with SHG involvement or by the tendency for certain children, who were also exposed to SHGs, to receive better nutrition than their siblings. These findings suggest that SHGs could improve child nutrition but must be interpreted in light of limitations of and potential biases in the data.

JEL: O12; O15; O16; O22

Keywords: Microeconomic Analyses of Economic Development; Human Resources for Economic Development; Financial Markets in Economic Development; Project Analysis

I. Introduction

The provision of financial services for the poor became a prominent feature of development discourse and policy following the advent of the Grameen Bank in the 1970s. Micro-banking institutions, such as microfinance pools, savings clubs, and self-help groups (SHGs), exist with the goal of reducing poverty by relaxing the saving and credit constraints that many of the world's impoverished households face. Though randomized evaluations of microfinance find modest effects on households' short-term economic circumstances (Banerjee et al., 2013; Crépon et al., 2013; Desai et al., 2011), a growing body of literature suggests that microfinance may improve general household wellbeing through indirect effects on human and social capital (Desai et al., 2011; Feigenberg et al., 2010; Pitt et al., 2003).

In this paper, I consider a long-term and indirect effect of micro-banking services by examining the impact of SHGs on child nutrition. Child nutrition is a particularly important outcome to consider because of its far-reaching economic effects. Malnourishment affects an individual's physical and cognitive development (Barker, 1990), which can limit future educational achievement (Field et al., 2009) and labor productivity (Strauss and Thomas, 1998); low levels of investment in child nutrition may impede human capital growth and can have detrimental long-term impacts on the aggregate economy.

I study children living in the Dungarpur district of rural Rajasthan, India. Understanding the determinants of child nutrition in India is especially important given the country's high rates of malnutrition. According to a UNICEF ranking, India has the highest rate of child stunting of any country in the world; in 2005, 43 percent of Indian children under five were stunted. In the same year, 48 percent of children under five were underweight and 17 percent were wasted (United Nations Children's Fund, 2009).² Rajasthan has particularly low levels of health and nutrition. Banerjee and Duflo (2007) find that the average BMI of poor adults in Udaipur district is 17.8, which is below the standard cutoff of 18.5 for healthy weight. Moreover, only 57 percent of the extremely poor in Udaipur have enough to eat throughout the year and, in 12 percent of these households, children are forced to cut the size of their

²A child is said to be stunted if his or her height has been retarded and said to be wasted if his or her weight is low given his or her height.

meals. Given high rates of poverty, low levels of literacy, and lack of access to modern medicine, standards for child nutrition in Durgapur are especially low (People’s Education and Development Organization, nd). The average child in my sample had a height-for-age Z-score that was 0.817 standard deviations below average in the international reference population (p -value = 0.000).

Household SHG involvement could, conceivably, improve child nutrition. Evidence suggests that the financial tools offered by SHGs could relax the severe savings (Dupas and Robinson, 2013a) and credit (Banerjee and Duflo, 2007; Krishna, 2010) constraints that village households would otherwise face. Member households may be able to rely on savings (Kast et al., 2012) and loans (Gertler et al., 2009; Mohanan, 2013) to smooth consumption in spite of the income fluctuations that are particularly pronounced in rural areas of developing countries. Financial stability could improve child nutrition directly by allowing rural households to maintain their food supply throughout the year; financial tools could improve child nutrition indirectly by facilitating improvements in living conditions (Muralidhar, 2011) or by enabling households to invest in health (Dupas and Robinson, 2013b; Tarozzi et al., 2011) and education (Desai et al., 2011). SHG involvement could also improve child nutrition through non-pecuniary channels if, for instance, members acquired knowledge of health and health behavior by attending meetings.

I evaluate the hypothesis that SHG involvement improves child nutrition by analyzing households affiliated with the SHGs started by the People’s Education and Development Organization (P.E.D.O.) in Durgapur’s tribal villages.³ To avoid selection biases that would distort simple comparisons between children of members and non-members, I choose a nutrition outcome that facilitates within-household comparisons. Specifically, I use child stunting to gauge malnutrition. Wasting and stunting are similar in that both are caused by malnourishment but the two metrics differ in their ability to measure conditions retrospectively; whereas child wasting reflects malnourishment only at or near the time of measurement, I use child stunting because it reflects nutritional deficiencies during early childhood. Thus

³In India, tribal peoples are those groups that are outside of the caste system.

the empirical strategy is to compare the height-for-age Z-scores of children who had high levels of SHG exposure to their siblings who had lower levels. I then check the validity of the causal interpretation by considering two alternate explanations: the possibility that periods of good fortune for a household or village led to SHG involvement and the possibility that the results reflect a tendency for younger children to receive better nutrition than older children. And finally, I use information on meeting attendance to examine the extent to which improvements in anthropometric status are related to knowledge exchanged at group meetings.

SHG involvement and anthropometric status are positively correlated in the data. Further analyses suggest that SHG exposure improves anthropometric status and that this effect intensifies with degree of exposure. These results must, of course, be interpreted in light of various limitations of and potential biases in the data. Though the data prevent a thorough evaluation of each mechanism through which SHG involvement could improve child nutrition, investigations suggest the improvement could be due to a relaxation of financial constraints. More generally, the findings suggest that micro-banking services may improve household well-being.

The remainder of the paper is organized as follows. In Section II, I highlight where my study fits into and attempts to build upon existing literature. I detail the institutional setting, explain the data collection process, and highlight potential sources of bias in the data in Section III. Section IV details the estimation strategy and Section V presents the initial results. I consider threats to validity in Section VI and causal channels in Section VII. In Section VIII, I detail limitations of the study as well as avenues for future research. I conclude in Section IX by summarizing the findings and their implications.

II. Literature Review

In response to the microfinance movement, a number of researchers have attempted to evaluate the impact of micro-banking on health. Fewer studies, however, have investigated the effect of micro-banking on child nutrition. While existing literature suggests that micro-

finance could improve child nutrition, there is a dearth of information on the effect of SHGs and on the extent to which child nutrition varies with degree of exposure.

In a large scale experiment, Banerjee et al. (2013) introduce microfinance to randomly selected communities of Hyderabad, India. Among many outcome variables, the authors examine reports of child illness. These results are relevant to the present study as child nutrition and illness are intimately related. Though microfinance produced some favorable results for treated households, it did not reduce rates of child illness. An important caveat of these results is that outcomes were measured only fifteen to eighteen months after randomization. Because changes in household behavior that would affect child illness evolve gradually, the period between the introduction of microfinance and the follow-up survey may not have allowed sufficient time for rates of reported child illness to improve.

Like improvements in child illness, gains in child nutrition often occur over an extended period of time. As detailed in subsequent sections, current child height reflects investments made in child health over several years; using child stunting as an outcome variable allows researchers to retrospectively track conditions over an extended period of time. DeLoach and Lamanna (2011) exploit this biology to evaluate the impact of microfinance on child nutrition in Indonesia. They use longitudinal data at the level of the individual child to track changes in each child's height over time. In a two-stage least squares model that uses changes in community urbanicity and electricity as instruments, the authors find that change in community microfinance exposure is a strong and positive predictor of change in individual children's standardized heights. The authors conclude that microfinance has a large and positive effect on the anthropometric status of children.

Two complications in this study's estimation strategy deserve mention. First, it relies on heterogeneity in changes in microfinance exposure at the community level. This strategy assumes that every child in a community is exposed to microfinance once it is introduced in the community. In reality, children in a community will have differing levels of exposure based on whether their households join.

Second, the authors use fixed effects at the level of the individual child to account for

community idiosyncrasies. This involved standardizing the heights of children as old as fifteen. As explained in subsequent sections, child height is most sensitive to environmental factors when children are under the age of four but becomes heavily influenced by genetics at the onset of puberty. It is possible that the observed change in standardized height is related to the timing of adolescent growth spurt in the studied population as opposed to the influence of an environmental factor like microfinance.

In a study that examines the effect of microfinance on child nutrition in Bangladesh, one way in which Pitt et al. (2003) avoid these issues is by comparing siblings. Exploiting variation in microfinance exposure at the household level accounts for variation in microfinance exposure between children of different households in the same community. Because they examine the standardized heights of two groups of comparable children, their empirical design does not track the standardized heights of particular children over time. Results of analyses that account for endogeneity of program participation and placement suggest that access to credit improves child nutrition.

Existing literature suggests that exposure to micro-banking services may improve child nutrition. My study attempts to build upon this research in at least three ways. First, it provides evidence of the impact of micro-banking on child nutrition in rural Rajasthan, India. To my knowledge, this particular question has not been investigated in this setting. Second, I evaluate SHGs instead of microfinance pools. Though the two forms of micro-banking are similar, many SHGs provide loans for health and consumption purposes in addition to entrepreneurial endeavors. SHG members are also able to save money and earn interest on their savings. Thus SHGs could affect child nutrition differently from microfinance pools. And third, my dataset contains information on degrees and dimensions of SHG involvement. This information facilitates analyses that investigate not only the impact of SHGs on child nutrition but also the nature of that effect.

III. Data

III.i. Institutional Setting

Data for this study were collected through P.E.D.O.: an Indian NGO that creates and monitors SHGs in Rajasthan's Dungarpur district. Each of P.E.D.O.'s SHGs is comprised of 15 to 20 female members, all of whom reside in the same village. To ensure that a particular group is small enough to maintain familiarity, but that a large proportion of village households have access to financial services, any given village may have many groups. Women are assigned to their groups based on the order in which they enroll.

To belong to one of P.E.D.O.'s SHGs, a member is required to pay Rs.50 in annual membership fees and to save at least Rs.50 per month.⁴ In exchange, she is eligible to borrow money at a one percent monthly, non-compounded rate, which must be repaid in installments agreed upon by her group. Loans may be used for both consumption purposes, such as school fees or home improvements, and productive endeavors, such as the purchase of livestock or agricultural machinery. Generally, a member cannot take a second loan if she has a first outstanding. One exception is the case of emergency loans; a member may borrow money for acute health care needs even if she has unpaid loans. These loans are granted immediately so that she does not have to wait until her group's next meeting to fund a health emergency. To guard against default, her SHG keeps 10 percent of the profit earned from interest on reserve. She and other group members receive the remaining profit as interest on savings deposits.

Though financial transactions with individual SHG members take place at the group level during monthly meetings, P.E.D.O.'s SHGs depend on the NGO for organizational support and for connections with other SHGs and nearby banks. These connections ensure that each SHG can meet the credit demand of its members, even if its savings are temporarily insufficient.

P.E.D.O.'s SHGs have been successful in many regards. Its network now includes over 2,000 SHGs with roughly 33,000 total members. Nearly two thirds of these groups have had no members leave. Repayment rates are also high; roughly 90 percent of loans have been repaid on time and nearly 100 percent of the loans extended have eventually been repaid

⁴Fifty Indian Rupees (Rs.50) are worth slightly less than one US dollar.

(People's Education and Development Organization, nd).

III.ii. Data Collection

Data were collected over a four week period in May and June of 2013. Data collection occurred in three phases. The first phase involved interviews with P.E.D.O. staff and administrators to select villages for study.⁵ A list of questions for these interviews can be found in Appendix 1. Though administrators were asked to list villages that are representative of those in which P.E.D.O. works, they may have tended to provide villages with SHGs that are well organized and have active members. Hence, a degree of caution is necessary in interpreting the results; members in the sample may be affiliated with SHGs that are closer to an institutional ideal than with groups that are inactive or poorly organized.

In the second phase, I conducted interviews with households in the selected villages. To do so, I created two research teams of NGO fieldworkers and translators, one of which I led and the other an NGO research intern led. Research teams were assigned to villages and tasked with visiting each home in these villages. The teams asked households with children between the ages of four and ten to be interviewed. Both adults and children were likely to be home when the teams visited since the months of May and June are the two that directly follow the local harvest season and that coincide with children's summer recess from school. Though every household that was approached was willing to be interviewed, the teams could only interview households in which both the mother and some of her children were present when the team arrived. This was not the case for the majority of households. In the interest of increasing sample size, the teams did not return to households whose members were absent. The attrition was not necessarily selective; households whose members were present when a team arrived did not appear to differ systematically from their absent counterparts. Absences were generally due to a family member being sent to the general store, to gather firewood, or to collect water at a time that happened to coincide with the teams' visits. Furthermore, the sizes, features, and landholdings of homes seemed ubiquitous across present and absent

⁵Some of these villages are technically hamlets of larger villages but are classified as villages by P.E.D.O.

households. Thus the households in the sample seem representative of households in the selected villages.

During each interview, the research teams obtained information about household and child characteristics by asking women to respond verbally to several questions about their lives. The complete list of questions used in these interviews can be found in Appendix 2.⁶ Following the interview, the teams measured the height of each child between the ages of four and ten who was available during the interview. The heights were standardized using anthropometric reference data from the National Center for Health Statistics (NCHS). This standardization produced height-for-age Z-scores by comparing each child's height to the height of other children of his age and gender in the NCHS reference population.

The major limitation of this process was that many questions required women to recount events that had occurred several years prior. Because height-for-age Z-score would be the study's outcome variable, limited recall of children's birth years was particularly problematic. The teams employed various strategies for obtaining the most accurate age of a child whose mother provided an age range instead of a precise age. One strategy involved asking the woman to chronologize the child's birth year with other memorable community events whose dates the teams knew. Other strategies involved examining the age of the child's siblings, the number of years of schooling the child had completed, the season in which the child was born, and the number of months the child had been breastfed.⁷ Nevertheless, many age-height combinations recorded in the interviews produced extreme values for Z-scores.⁸ For example, some combinations produced Z-scores greater than 5.5 in absolute value or suggested that Z-scores of children within a household differed by more than 3 standard deviations. I deemed such cases implausible and identified them for cleaning. I generally assumed extreme values resulted from child age estimates being one year off; I shifted ages corresponding to these extreme values by one year for analysis. A handful of children had

⁶The protocols established by Duke University's Institutional Review Board (IRB) were adhered to throughout the interview process.

⁷Because a woman cannot breastfeed a child once she becomes pregnant with another, the number of months a child had been breastfed was used to project an implied number of months separating the births of two children. Examining months of breastfeeding was, of course, useful only when the child whose age was in question was a middle or eldest child.

⁸Like any variable with a standard normal distribution, height-for-age Z-scores are generally under 3 in absolute value.

erroneous Z-scores but also had ages provided with certainty. This was particularly true of children aged four or ten as the research teams took particular caution to ensure children surveyed were indeed in the desired age range. Others had implausible Z-scores that would have been made equally extreme in an opposite direction had ages been changed. In such cases, I assumed that the erroneous Z-score was due to an error in measuring, translating, or recording the height. Therefore I shifted heights corresponding to this handful of extreme values for analysis. These strategies may have improved precision of Z-scores but of course cannot guarantee perfect accuracy; participant recall biases must be kept in mind when interpreting the results.

Moreover, research budget constraints prevented me from using a double blind procedure in data collection and cleaning. Thus results are subject not only to recall biases but also to experimenter biases that were unavoidably introduced in data collection and in the identification and cleaning of extreme values. These issues are further investigated in Appendix 3. If the data collection and cleaning processes introduced marked biases, height-for-age Z-scores would not be correlated with control variables in the predicted direction. Appendix 3 provides reassuring evidence on the effects of these biases by showing that height-for-age Z-scores are correlated in the predicted direction with a number of control variables. Nevertheless, the results must be interpreted in light of subject recall and experimenter biases.

In the third phase, I linked household and child information to P.E.D.O.'s data regarding each woman's SHG involvement. A woman was defined as having been a member in a particular year if she was active in at least one dimension of SHG activity during that year. These dimensions included paying SHG membership fees, attending SHG meetings, saving money with an SHG, and borrowing money from an SHG. A woman was said to be an SHG dropout if she had been an SHG member in the past, is not currently active in any dimension of SHG involvement, and reported having no intention of becoming active in the future. A woman was considered non-affiliated if she had never been a member. A child was said to have had SHG exposure in a particular year if his or her mother was a member in that year.

P.E.D.O.'s records provided dates of membership and hence dates of children's SHG

exposure.⁹ The records also provided the total number of meetings attended, rupees saved, and rupees borrowed in each year of membership. As many older, handwritten files had been lost or destroyed, records were missing for certain households. For these households, dates of membership were gauged using the years women had provided during their interviews and detailed financial behavior was recorded as missing. This source of measurement error would have increased variance but only affected point estimates through the tendency of classical measurement error to bias coefficients to zero.

Interviews were conducted with a total of 212 households from 36 different SHGs and seven different villages. Detailed SHG involvement data was available for 193 of the households. These interviews produced a dataset with information on 424 children. Of the 424, 308 lived in SHG member households, 20 lived in dropout households, and 96 lived in non-affiliated households. Detailed summary statistics will be presented in Section V.

IV. Empirical Specification

I estimate the effect of SHG involvement on child nutrition by using anthropometric status to proxy for nutritional status. Specifically, I consider child stunting.¹⁰ Child wasting and stunting are widely accepted gauges of malnutrition. But despite having similar causes, wasting and stunting do not represent the same information. Because weight responds quickly to nutrition and health inputs, wasting reflects malnutrition at or near the time of measurement. Stunting, on the other hand, reflects the cumulative effect of malnutrition over several years. Low levels of nutrition and health investment early in life can irreversibly stunt an individual's eventual stature, regardless of his or her genotype. Once a child reaches the age of about four, however, the influence of environmental health inputs becomes markedly reduced; by the time the child reaches puberty, genotype becomes a primary determinant of height (Martorell and Habicht, 1986). The height of a child between the ages of four and ten

⁹Child ages and SHG membership were recorded on an annual basis as opposed to a monthly basis due to limitations in recall and documentation. Since the month of a child's birth and the month a household joined an SHG do not perfectly coincide, these limitations may have made the SHG exposure estimates less precise by assigning too many or too few months of exposure. These limitations, however, would not have led to a systematic over- or underestimation of SHG exposure as children were roughly equally likely to have been assigned too many or too few months of exposure.

¹⁰As mentioned in the introduction, a child is said to be stunted if his or her height has been retarded and said to be wasted if his or her weight is low given his or her height.

therefore reflects nutritional and other health inputs the child received in utero and under the age of four.

The empirical strategy seeks to examine the impact of early childhood SHG exposure on child stunting. Child height, h , is a useful descriptor only when contextualized by the heights of other children of the same age, a , and gender, g . Thus I use height-for-age Z-scores, $Z_{a,g}^h$, as the outcome variable.

One might begin the analysis by regressing height-for-age Z-score on early childhood SHG exposure in a least squares model. Such a model would compare children whose households had belonged to SHGs when the children were young to children of non-affiliated households. But because members are fundamentally different from non-members, the model would be subject to selection biases.

I account for endogeneity of household SHG participation by using within-household comparisons. In particular, I compare the standardized height of children who were young when their households belonged to SHGs to the standardized height of those children's siblings who had relatively less SHG exposure at a young age. I estimate the following fixed effects model:

$$Z_{4 \leq a \leq 10, g, i, m}^h = \beta_0 + \beta_1 \frac{\text{yearsSHG}_{-i \leq a \leq 3, i, m}}{5} + \sum_k \beta_{2,k} a_{k, i, m} + \beta_3 g_{i, m} + \sum_n \beta_{4,n} \mathbf{1}_{n, i, m} + \mu_m + \varepsilon_{i, m}. \quad (1)$$

Here μ_m is the fixed effect for each group of siblings, $Z_{4 \leq a \leq 10, g, i, m}^h$ is the height-for-age Z-score, $\frac{\text{yearsSHG}_{-i \leq a \leq 3, i, m}}{5}$ is the fraction of years of early childhood SHG exposure, and $\vec{X}_{i, m}$ is an n-dimensional vector of time-variant, child-level confounds.¹¹ Indicators for child age, $\vec{a}_{i, m}$, and gender, $g_{i, m}$, are included to account for potential biases associated with standardizing child height using international averages.

V. Results

V.i. Initial Results and Summary Statistics

I begin the analysis by regressing height-for-age Z-score on early childhood SHG exposure.

¹¹I use sibling as opposed to household fixed effects because some households in the data represent brothers living together with their wives and children. Though all children in these households have exposure to the same SHG account, each group of siblings has a slightly different genetic makeup. Thus I record SHG exposure at the household level but compare within groups of siblings.

More specifically, I estimate equation (1) without sibling fixed effects and without the vector of control variables. The results are presented in the first column of Table 2. SHG exposure and height-for-age Z -scores are positively correlated; children who were exposed to SHG membership throughout early childhood have height-for-age Z -scores that are roughly 0.758 standard deviations above children who had no exposure (p -value = 0.000).

This model, however, is subject to selection bias as it compares children of members to those of non-members. Many SHG members tend to be wealthier and better educated than non-members. Indeed, when nonmembers were asked why they chose not to join, the most common response was an inability to meet P.E.D.O.'s savings requirements and pay its membership fees. The estimate in column (1) may be biased upwards as members' children would tend to be healthier, regardless of their SHG affiliation.

To correct these selection biases, I use sibling fixed effects. Specifically, I restrict the sample to children of SHG members and dropouts and then compare children who were exposed to household SHG membership during early childhood to their siblings who had no exposure. Sibling fixed effects, however, cannot account for confounds that occur at the level of the individual child. As Table 1 illustrates, children in the implicit comparison groups may be systematically different from one another. By construction, children in the first column were exposed to greater levels of SHG activity during early childhood; households tended to have greater levels of SHG savings, loans, and meeting attendance when these children were young. The remaining mean differences in Table 1 are largely due to the fact that very few members drop out of P.E.D.O.'s SHGs. Of the 162 households in the sample that had ever belonged to an SHG, only 12 had left their group. Most households joined a group at some point in the past decade and remain members. Children that were exposed to household SHG membership during early childhood, and are sorted into the first column, therefore tend to be younger siblings who were born once their households had joined. Older and taller siblings are sorted into the second column. This also explains why, as compared to children in column (1), children in column (2) were born into smaller families.

Another notable difference is that older children in the second column were significantly

less likely to have been born during the monsoon season than were younger children in the first. To understand this discrepancy, it is necessary to understand the weather patterns in Dungarpur district over the past decade. Though Dungarpur's monsoon season has occurred at the same time in each of the past ten years, the annual rainfall has not remained constant over the period. In particular, the region experienced a period of drought around 2000 and 2001 that ended with flooding in 2006. The mean difference in Table 1 can be explained by a reporting bias; although women in the sample were asked to provide the season of their child's birth, they may have been less likely to say their child was born during the monsoon season if they did not remember rainfall around the time of the child's birth. Because older children in the second column were more likely to have been born during a drought, their mothers may have been less likely to report that the children were born during the monsoon season.

Interestingly, height-for-age Z-score does not appear to differ between the comparison groups. This observation seems to contradict the study's hypothesis. Yet, as mentioned above, child age and gender should be controlled for in regressions with height-for-age Z-scores. The statistical insignificance of the relationship in the final row of Table 1 could also be explained by the presence of other confounds or by the implicit assumption of that row's specification that height-for-age Z-scores do not depend on degree of SHG exposure. The results that follow show that height-for-age Z-scores do indeed differ between the comparison groups once these factors are accounted for.

V.ii. Results

In response to the selection biases of the first estimation and to the mean differences in Table 1, I estimate equation (1). The estimates presented in the second column of Table 2 account for household selection biases while the estimates in the third column control for these biases as well as potential confounds. In particular, I account for factors that affect nutrition and also differ between siblings. Equation (1) accounts for child age and gender explicitly but the vector of control variables adds birth order, months of breastfeeding,

household size at birth, and number of household income earners at birth. I also include indicators for sickness, death of adult family members, and household financial stresses during each year of early childhood, as well as indicators for season of birth.¹² A remaining concern involves the potential of community-level changes to confound the results. Such changes could include changes in drought conditions or in health infrastructure.¹³ Because siblings live in the same communities, the child age indicators account for each child's exposure to community-level changes. These indicators absorb the net effect of time variant factors that are not controlled for explicitly as well as age-related biases associated with standardizing heights using international averages.

The results presented in the second column of Table 2 suggest that children who had relatively more exposure to SHG membership during early childhood have lower rates of stunting than their siblings who had relatively less. Children who were exposed to SHG membership throughout early childhood have height-for-age Z-scores that are roughly 0.700 standard deviations above those of their siblings who had no exposure (p-value = 0.046). That this estimate is slightly smaller than that of column (1) suggests that household selection may, indeed, have biased the first column's estimate upwards.

To further reduce omitted variable bias, I add the vector of control variables. Controlling for these variables does not change the conclusion of the regression; as evidenced in column (3), children who were exposed to SHG membership throughout early childhood tend to be less stunted than their siblings who had no exposure (p-value = 0.012). The addition of control variables does, however, cause the point estimate to increase to 0.920. This may be due to the addition of household size at birth to the specification. Younger children tended to be exposed to SHG membership during early childhood but were also born into larger families where limited resources were pooled between a larger number of household members.

¹²Dungarpur district experiences four seasons: winter (December – February), summer (March – June), monsoon, (June – September), and post-monsoon (October – November). However, when attempting to recall the seasons in which their children were born, families had difficulty distinguishing between the winter and post-monsoon seasons. My dataset therefore combines these two seasons so that a response to birth season can be one of three options: winter/post-monsoon, summer, or monsoon.

¹³Such health infrastructure changes include the implementation of Janani Suraksha Yojana, the Midday Meal Scheme, and the Mahatma Gandhi National Rural Employment Guarantee Act (MGNREGA), as well as the opening of village health subcenters. The first policy is a conditional cash transfer that aims to reduce maternal and child mortality by incentivizing mothers to deliver in hospitals. As noted earlier, the Midday Meal Scheme provides nutritious meals to children that attend school. MGNREGA guarantees employment for individuals living in rural areas by hiring these individuals to perform manual labor. And finally, village health subcenters are village-level public health facilities that provide primary care.

Household size at birth would produce an effect competing with that of SHG membership; SHG membership would tend to increase height-for-age Z-scores while larger household sizes at birth would pull them down. Though not significant at high confidence levels (p-value = 0.149), the coefficient on household size at birth is indeed negative (-0.557). Failing to control for this variable may therefore have biased the estimate in column (2) downwards.

Because it accounts for household selection biases and controls for potential confounds at the level of the individual child, the specification presented in column (3) can more convincingly be given a causal interpretation. Specifically, the data suggest that exposure to SHGs during early childhood reduces rates of child stunting. Since child stunting is a proxy for nutrition and health, this finding supports the notion that SHG involvement could improve child wellbeing.¹⁴

The analyses presented thus far use an independent variable that allows rates of child stunting to differ based on degree of exposure. A related question is whether or not a corresponding level effect exists. Such an effect would be present if children who were exposed to SHG membership at any point during early childhood had healthier height-for-age Z-scores, regardless of the length of time they were exposed. In the last column of Table 2, I present the estimate of a specification that replaces the independent variable in equation (1) with $SHG_{-1 \leq a \leq 3, i, m}$. This variable takes the value of one if a child was exposed to SHG membership at any point before age four, and zero if otherwise. The data suggest that a level effect does exist; children who had SHG exposure at a young age have height-for-age Z-scores that are, on average, 0.451 standard deviations above those of their siblings who had none (p-value = 0.042). However, replacing the variable that allowed for the effect to differ based on length of exposure with a variable that imposes a level effect on the relationship of interest led the point estimate to decrease from 0.920 to 0.451. This suggests that height-for-age Z-scores improve not just with SHG exposure but also with length of exposure.

VI. Threats to Validity

¹⁴A discussion of how this effect may differ by child gender and by age at exposure can be found in Appendix 4.

VI.i. Do Positive Events Lead to SHG Involvement?

The results of Section V suggest that SHG involvement may improve child health. However, the validity of this interpretation is threatened by the possibility that favorable events for households or villages lead to SHG involvement. One concern is that households join SHGs only when their finances allow them to pay annual membership fees and to meet monthly savings requirements. Similarly, differences in average SHG age between villages may reflect not only the general expansion of P.E.D.O.'s operations over time but also changes in village infrastructure that facilitated the introduction of SHGs. If these changes to village and household wellbeing also improved child nutrition, the child stunting patterns observed in the previous section would be present with or without SHGs.

There are two important dimensions over which these favorable events could be described: predictability and magnitude. Given the enormous amount of uncertainty that villages and village households face with regards to health, income, and employment, one cannot easily predict when such a favorable event will occur. The favorable events can more accurately be thought of as unanticipated, positive shocks. With regards to magnitude, SHG exposure does not seem to be related to minor fluctuations in household or village wellbeing; P.E.D.O. has never scaled back operations in the villages studied, nor have any of the families in the sample joined and left SHGs multiple times. These facts suggest that SHG exposure is not necessarily related to positive events that are minor or unanticipated. However, they do not eliminate the possibility that a household decided to join an SHG following a positive shock to its finances or the possibility that P.E.D.O. introduced SHGs following major improvements in village infrastructure.

To investigate this issue, I suppose that members experienced a positive shock in the years before they joined. This event could have been an improvement in household finances, an improvement in village infrastructure, or some combination of the two. I replace the independent variable in equation (1) with $\frac{yearsshock_{-1 \leq a \leq 3, i, m}}{5}$, which reflects the fraction of years of early childhood in which a child was exposed to a hypothetical shock before his or

her SHG exposure.¹⁵ I consider separately the cases in which the shock occurred one and two years before the household joined. Children who had relatively more early childhood exposure to favorable events would have lower rates of stunting than their siblings who had relatively less. Thus if positive shocks immediately preceding SHG involvement were driving the previous section’s results, the coefficients on the shock exposure variables would be positive and significant. To avoid downward bias on these estimates, I limit the sample to current members; including dropouts would bias the coefficients downwards as children of dropouts may have had early childhood exposure to not only positive shocks preceding SHG involvement but also to negative shocks that caused the household to leave.

The data, however, do not support the hypothesis that positive shocks lead to SHG involvement. The results presented in column (1a) of Table 3 examine the case of a hypothetical shock occurring one year before a household joined its SHG, while those in column (1b) examine a second case of the shock occurring two years prior. The coefficients on the shock exposure variables are negative (-0.076 and -0.139, respectively) and statistically insignificant (p-values = 0.906 and 0.703, respectively). Improvements in child stunting observed in the previous section may have indeed been driven by SHG exposure as opposed to favorable village or household events leading to SHG involvement.

VI.ii. Are Younger Children Better Nourished?

A second threat to validity is the possibility that younger children, who had greater exposure to SHG involvement, tend to receive better nutrition than older children, regardless of their SHG affiliation. This is particularly concerning given the population studied. Traditional cultural practices in rural Rajasthan typically give the responsibility of caring for elderly parents to the youngest son. Parents may have a greater incentive to protect the wellbeing of their youngest son than that of other children. Because younger children in the sample tended to have greater exposure to SHG involvement during early childhood, this bias might have driven the previous section’s results.

¹⁵Because the only possible lengths of time between the hypothetical shocks and SHG activity are zero, one, and two years, $\frac{years_{shock}_{-1 < a < 3, i, m}}{5}$ can only take the values 0, 0.2, or 0.4.

To investigate this threat, I add an indicator to equation (1) that takes the value of one if the child is a youngest son, and zero if otherwise. If lower rates of stunting amongst children that were exposed to SHG membership during early childhood reflect a tendency to protect the health of the youngest son, the addition of this indicator would reduce the predictive power of the SHG exposure variable. The results presented in column (2), however, reject this hypothesis. The coefficient on the indicator is positive (0.298) but not statistically significant at generally accepted confidence levels (p -value = 0.182). Moreover, the coefficient on SHG exposure remains large (0.972) and significant (p -value = 0.008) after the addition to the model. The findings of this analysis, as well as those of the one presented earlier in this section, provide further support for the hypothesis that SHG involvement improves child nutrition.

VII. Causal Channels

VII.i. Pecuniary Channels

The results presented thus far suggest that SHG involvement improves child nutrition but do not specify how the improvement occurs. The improvement could, conceptually, have been caused by pecuniary or non-pecuniary facets of SHG exposure.

The pecuniary services offered by SHGs allow households to smooth consumption. Agricultural incomes in Rajasthan are not only low but also unpredictable due to their dependence on weather and climate. Further compounding the issue is an inability to save and borrow money. Because they lack collateralizable assets, have limited credit histories, and tend towards higher default rates, the poor are largely ignored by the formal banking system. Unfortunately, informal lending sources suffer from financial shallowness, unpredictable availability, uncertain repayment structures, and exorbitant fees (Banerjee and Duflo, 2007; Krishna, 2010). The poor also face formidable savings constraints (Dupas and Robinson, 2013a) and are forced to save at home. Money saved at home represents forgone interest earnings and may be lost, stolen, or spent trivially (Collins et al., 2009). By operating as small and local institutions, SHGs reduce informational asymmetries and capitalize on

social collateral to an extent that allows them to provide financial tools at relatively low cost to arguably high-risk individuals. These pecuniary services could facilitate consumption smoothing by providing financial streams to draw upon in the face of adverse shocks to household finances (Gertler et al., 2009; Mohanan, 2013) and by offering a means of storing money during periods of high income (Kast et al., 2012).

The stabilization that these pecuniary services allow could improve child nutrition directly by improving food supply. Household consumption regularity can provide children with constant nourishment in spite of unforeseen expenditures. SHG members could further improve household food supply by using savings and loans to fund larger investments in livestock and farmland.

Consumption regularity may also reduce child malnutrition indirectly. Micro-banking services may allow households to maintain investments in health (Dupas and Robinson, 2013b; Tarozzi et al., 2011), education (Desai et al., 2011), and living conditions (Muralidhar, 2011). Certain programs implemented by the Indian government provide nutritional benefits for children who attend school and receive appropriate medical care.¹⁶ Investments in medical care that treat infectious diseases can also improve child nutrition since such diseases would otherwise drain calories from the body. Similarly, investments in living conditions that improve households' access to sanitation and drinking water can improve child nutrition by reducing rates of infectious disease.

Consumption stability could also improve child nutrition through its impact on household psychology. Preliminary findings from a study of factory workers in Mexico (Atkin, 2009) suggest that income stability, and not just income itself, secured by long-term employment is a significant predictor of child height. Poor individuals who believe they will have a stable money supply in both current and future periods may feel empowered and enthused to plan for the future. Since investments in health can be seen as plans for the future, child nutrition could benefit from income stability.

¹⁶Two of these programs include the Midday Meal Scheme and the Anganwadi. Through the first program, the government provides nutritious lunches for children that attend school. Anganwadis, on the other hand, are village-level public health facilities that provide nutrition supplementation, health behavior education, vaccinations, and medications for young children.

It is important to note that child stunting is particularly responsive to consumption stability. Though child weight can recover from periods of relative scarcity, child height is restricted irreversibly. The pecuniary services that SHGs offer may be important in explaining the relationship between SHG exposure and height-for-age Z-scores since they allow households to maintain both direct and indirect investments in child health.

VII.ii. Non-Pecuniary Channels

The improvement could occur through non-pecuniary channels as well. An important difference between commercial banks and SHGs is that customers of the latter are required to meet periodically. Officially, meetings convene so that savings can be collected and loans dispersed. But evidence suggests that group meetings could affect members' lives independently of financial activity (Feigenberg et al., 2010).

Women may gain knowledge at group meetings that would improve their children's health independently of the pecuniary channels discussed above. This knowledge might involve information regarding health and health behavior. One member explained that her son had been ill several months prior. She explained his symptoms to other members during an SHG meeting and learned that the child of another member had recently overcome a similar affliction. This member provided instructions for seeking care from the local health clinic. The mother took this advice and brought her son to the clinic where his ailment was subsequently diagnosed and treated. In this case, and perhaps in many others, knowledge exchanged at an SHG meeting had a direct impact on child health and nutrition.

Meeting attendance may also give women financial knowledge that allows them to take a more active role in household decision-making. Data from developing countries suggest that, as compared to men, women are more apt to make financial decisions with the wellbeing of their children in mind (Duflo, 2003; Pitt and Khandker, 1998; Pitt et al., 2003; Qian, 2008; Thomas, 1990). SHGs may improve child nutrition by giving women the ability to exercise greater intrahousehold bargaining power.

VII.iii. Evidence

I begin by considering non-pecuniary channels. I use information from P.E.D.O.'s records regarding each woman's meeting attendance to examine the possibility that knowledge exchanged at group meetings improves child nutrition.

In this context, information on meeting attendance is valuable for three reasons. First, at the margin, meeting attendance does not necessarily reflect financial activity: a woman could attend a meeting to save or borrow relatively large sums, or attend the meeting without saving or borrowing any money. Thus the information allows for the isolation of non-pecuniary channels. Second, the predicted direction of the relationship between child nutrition and meeting attendance is unambiguous. To the extent that SHG meetings spread accurate information about health behavior, encourage women to exercise greater intrahousehold bargaining power, and proxy for other non-pecuniary child health benefits, meeting attendance could be positively related to height-for-age Z-scores. But because they require only a few hours per month and no monetary cost, group meeting attendance is unlikely to divert resources away from child health. Third, meeting attendance may be exogenous to anthropometric status. Women would not attend meetings with the goal of improving child health; if women needed advice on health or health behavior, they could always consult members outside of group meetings. Since meetings demand only a few hours of time per month and since women could rely on family and neighbors for childcare, it is unlikely that members would choose not to attend a meeting because their children were sick. Improvements in anthropometric status seem unlikely to lead to higher levels of meeting attendance as women are unlikely to make decisions about meeting attendance based on the health of their children.

In hopes of quantifying the benefits of meeting attendance, I first consider the possibility that the knowledge a woman gained by attending meetings when her child was young improved the child's health. To do so I replace the independent variable in equation (1) with $meetings_{-i \leq a \leq 3, i, m}$, which represents the number of SHG meetings a child's mother attended between the time that child was conceived and age four. Because knowledge may

accumulate over time, I estimate a second specification using as an independent variable the number of meetings a child’s mother attended before the child was born: $meetings_{a<0,i,m}$. To control for a level effect of SHG exposure on child stunting that is independent of meeting attendance, I include $SHG_{-1\leq a\leq 3,i,m}$ in both specifications.

The results are presented in Table 4. Columns (1a) and (1b) examine the possibility that children who were young during periods of regular meeting attendance have higher height-for-age Z-scores than those of their siblings who were exposed to lower levels of meeting attendance when young. The first of these specifications assumes a linear relationship while the second allows the relationship to be quadratic. Though the estimate in column (1a) is positive, it is insignificant (p-value = 0.250). Likewise, the coefficients of the meetings variables in column (1b) do not differ significantly from zero (p-values = 0.385 and 0.158, respectively). Children do not appear to have experienced marked gains in anthropometric status from meetings their mothers attended when the children were young.

Columns (2a) and (2b) investigate the possibility that children who were born when their mothers had attended many SHG meetings have higher height-for-age Z-scores than those of their siblings who were born when the mother had attended fewer meetings. However, neither the meetings variable in column (2a) (p-value = 0.665) nor the meetings variables in column (2b) (p-values = 0.790 and 0.569, respectively) differ significantly from zero. The data do not suggest that knowledge gained at meetings accumulates over time and improves child health. These analyses provide little evidence for the hypothesis that non-pecuniary benefits of SHG involvement improved child health. The benefits of SHG involvement for child nutrition may instead occur through a relaxation of savings and credit constraints.

VIII. Discussion

VIII.i. Causal Channels

Though the results suggest that SHG involvement improves child nutrition and that this improvement does not occur through meeting attendance, remaining ambiguities exist. The first involves a precise identification of causal channels. The previous section suggests that

marginal meeting attendance may not produce marked effects on child stunting but does not address the possibility that non-pecuniary benefits follow a level effect. Such effects, however, could not be disentangled from pecuniary benefits since indicators that represent a level effect of SHG exposure may represent both financial and non-financial benefits.

Furthermore, the data do not facilitate analyses that distinguish between pecuniary channels. A lack of detailed information on consumption patterns prevents an examination of the extent to which the improvements in financial stability, food expenditure, and human capital investments that SHG membership enables improve in child nutrition. In the previous section, I also note that SHG involvement could benefit child nutrition if it enabled households to improve living conditions, to invest in clean water supplies, to purchase livestock, or to increase the productivity of farmland. Problems of recall and documentation prevented me from determining precisely when such investments were made and inhibit analyses of children's exposure to these investments. Though the previous sections' results suggest that pecuniary channels are more likely drivers than non-pecuniary channels, a precise identification of the aspects of SHG involvement that improve child nutrition would require further research.

VIII.ii. Spillovers

A second ambiguity relates to the presence of positive spillovers whereby children who had no exposure to SHG involvement during early childhood benefitted from other children having exposure. Because children of non-members were not included in the primary analyses, the presence of spillovers from member children to non-member children would not have affected the results. Positive spillovers within households could, however, have affected the fixed effects models. An assumption of the estimation strategy is that child height does not respond to environmental factors after age four. While the influence of environmental factors may be markedly reduced after this age, height in older children may not be driven exclusively by genotype. The presence of catch-up growth could have affected the results if older children who had no SHG exposure before age four became less stunted due to expo-

sure at a later age. But since any such spillover would have led to an underestimation of the relationship of interest, catch-up growth would not change the main conclusion of the study.

Though positive spillovers may not have affected the study's conclusion, their presence is an important consideration when evaluating the overall impact of SHGs on child nutrition. Unfortunately, I am not able to quantify spillovers due to the small number of non-members in my dataset and the challenge of disentangling inputs to child height that occurred before age four from those that occurred after. Further research is needed to test for spillovers to children who had no exposure to SHG involvement during early childhood.

IX. Conclusion and Implications

This paper contributes to the literature relating micro-banking services to health by investigating the impact of SHGs on child nutrition. I use data on households affiliated with P.E.D.O.'s SHGs in rural Rajasthan, India. Children in the sample who were exposed to household SHG involvement at a young age have lower rates of stunting than children who had relatively less exposure. This result provides evidence of a positive correlation between SHG exposure and anthropometric status. To account for household selection biases and other confounds, I examine the relationship of interest in a sibling fixed effects model. The data suggest that children whose households were involved in SHGs when the children were young have lower rates of stunting than their siblings who were born during periods of inactivity. This effect seems to intensify with degree of exposure. The findings do not appear to be driven by favorable events that led to SHG involvement or by the tendency of certain children, who also happened to have had exposure to SHG involvement, to receive better nutrition than their siblings. Though analyses suggest that improvements occur through pecuniary pathways, data limitations prevent a thorough analysis of causal channels. Further research is needed to identify these channels, to isolate spillovers, and to correct potential biases introduced in data collection.

This study suggests that SHG involvement improves children's nutrition. The findings provide evidence that micro-banking services could improve dimensions of household well-

being that are not purely financial. They motivate further research into the potential for financial tools to improve human capital stock and necessitate a consideration of these effects when evaluating the overall impact of micro-banking services on poverty reduction.

References

- Atkin, D. (2009). Working for the future: Female factory work and child height in Mexico. Working Paper, Yale University.
- Banerjee, A. and Duflo, E. (2007). The economic lives of the poor. *Journal of Economic Perspectives*, 21(1):141–167.
- Banerjee, A., Duflo, E., Glennerster, R., and Kinnan, C. (2013). The miracle of microfinance? Evidence from a randomized evaluation. NBER Working Paper No. 18950.
- Barker, D. J. (1990). The fetal and infant origins of adult disease. *British Medical Journal*, 301(6761):1111.
- Collins, D., Morduch, J., Rutherford, S., and Ruthven, O. (2009). *Portfolios of the poor: How the world's poor live on \$2 a day*. Princeton, NJ: Princeton University Press.
- Crépon, B., Devoto, F., Duflo, E., and Parienté, W. (2013). Estimating the impact of microcredit on those who take it up: Evidence from a randomized experiment in Morocco. Working Paper, Massachusetts Institute of Technology.
- DeLoach, S. and Lamanna, E. (2011). Measuring the impact of microfinance on child health outcomes in Indonesia. *World Development*, 39(10):1808–1819.
- Desai, J., Johnson, K., and Tarozzi, A. (2011). On the impact of microcredit: Evidence from a randomized intervention in rural Ethiopia. Working Paper, Duke University.
- Duflo, E. (2003). Grandmothers and granddaughters: Old-age pensions and intrahousehold allocation in South Africa. *The World Bank Economic Review*, 17(1):1–25.
- Dupas, P. and Robinson, J. (2013a). Savings constraints and microenterprise development: Evidence from a field experiment in Kenya. *American Economic Journal: Applied Economics*, 5(1):163–192.
- Dupas, P. and Robinson, J. (2013b). Why don't the poor save more? Evidence from health savings experiments. *American Economic Review*, 103(4):1138–1171.

- Feigenberg, B., Field, E., and Pande, R. (2010). Building social capital through microfinance. NBER Working Paper No. 16018.
- Field, E., Robles, O., and Torero, M. (2009). Iodine deficiency and schooling attainment in Tanzania. *American Economic Journal: Applied Economics*, 1(4):140–169.
- Gertler, P., Levine, D., and Moretti, E. (2009). Do microfinance programs help families insure consumption against illness? *Health Economics*, 18(3):257–273.
- Iyengar, S., Iyengar, K., and Gupta, V. (2009). Maternal health: A case study of Rajasthan. *Journal of Health, Population, and Nutrition*, 27(2):271–292.
- Kast, F., Meier, S., and Pomeranz, D. (2012). Under-savers anonymous: Evidence on self-help groups and peer pressure as a savings commitment device. NBER Working Paper No. 18417.
- Krishna, A. (2010). *One illness away: Why people become poor and how they escape poverty*. Oxford: Oxford University Press.
- Martorell, R. and Habicht, J. (1986). Growth in early childhood in developing countries. In Falkner, F., and Tanner, J.M. (Eds.), *Human Growth*. New York, NY: Plenum Press.
- Mohanan, M. (2013). Causal effects of health shocks on consumption and debt: Quasi-experimental evidence from bus accident injuries. *Review of Economics and Statistics*, 95(2):673–681.
- Muralidhar, G. (2011). Impact of SHGs: A study of P.E.D.O.’s intervention in microfinance and livelihoods under sakh-se-vikas initiative. Impact Assessment, Centre for MicroFinance.
- People’s Education and Development Organization (n.d.). Sakhi: Sustainability and institution building of SHGs. National Printers, Dungarpur, India.
- Pitt, M. and Khandker, S. (1998). The impact of group-based credit programs on poor

- households in Bangladesh: Does the gender of participants matter? *Journal of Political Economy*, 105(5):958–996.
- Pitt, M., Khandker, S., Chowdhury, O., and M., D. (2003). Credit programs for the poor and the health status of children in rural Bangladesh. *International Economic Review*, 44(1):87–118.
- Qian, N. (2008). Missing women and the price of tea in China: The effect of sex-specific earnings on sex imbalance. *The Quarterly Journal of Economics*, 123(3):1251–1285.
- Strauss, J. and Thomas, D. (1998). Health, nutrition, and economic development. *Journal of Economic Literature*, 36(2):766–817.
- Tarozzi, A., Mahajan, A., Blackburn, B., Kopf, D., Krishnan, L., and Yoong, J. (2011). Micro-loans, insecticide-treated bednets and malaria: Evidence from a randomized controlled trial in Orissa (India). ERID Working Paper No. 104.
- Thomas, D. (1990). Intra-household resource allocation: An inferential approach. *The Journal of Human Resources*, 25(4):635–664.
- United Nations Children’s Fund (2009). Tracking progress on child and maternal nutrition: A survival and development priority. Retrieved from http://www.unicef.org/publications/index_51656.html.

Table 1: Summary Statistics

	(1)	(2)	(3)	(4)
	Early Childhood SHG exposure	No Early Childhood SHG exposure	<i>n</i>	<i>t-static</i>
	<i>Mean</i>	<i>Mean</i>		
Child male	0.520	0.467	270	0.40
Child age	6.191	9.418	270	-6.14***
Child height (cm)	112.999	132.013	270	-5.52***
Born in monsoon season (Jul - Sept)	0.253	-0.023 ^a	269	2.75***
Born in summer season (Mar - Jun)	0.310	0.474	270	-1.40
Birth order	3.334	1.863	270	5.82***
Birth complications	0.047	0.081	270	-0.70
Sick while young:				
Sick in infancy	0.069	0.075	270	-0.09
Sick at age 1	0.036	0.059	270	-0.56
Sick at age 2	0.033	0.056	270	-0.53
Sick at Age 3	0.037	0.037	270	0.00
Months breastfed	13.652	13.515	270	0.17
Household size at birth	5.387	3.951	270	5.70***
Household income earners at birth	1.216	1.216	270	0.00
Death of adult family member:				
Family death under age 1	0.032	-0.002 ^a	270	0.87
Family death at age 1	0.021	0.021	270	0.00
Family death at age 2	0.027	0.027	270	0.00
Family death at age 3	0.032	-0.012 ^a	270	1.37
Household financial stress: ^b				
Financial stress under age 1	0.015	0.015	270	0.00
Financial stress at age 1	0.015	0.015	270	0.00
Financial stress at age 2	0.015	0.015	270	0.00
Financial stress at age 3	0.015	0.015	270	0.00

Table 1 (continued): Summary Statistics

	(1)	(2)	(3)	(4)
	Early Childhood SHG exposure	No Early Childhood SHG exposure		
	<i>Mean</i>	<i>Mean</i>	<i>n</i>	<i>t-static</i>
Value household SHG savings (INR):				
Value SHG savings in utero	405.211	383.557	237	0.20
Value SHG savings in infancy	491.340	491.340	237	0.00
Value SHG savings at age 1	681.553	299.531	237	2.23***
Value SHG savings at age 2	958.398	-268.898 ^a	237	7.72***
Value SHG savings at age 3	1242.902	-732.163 ^a	237	14.81***
Value household SHG loans (INR):				
Value SHG loans in utero	2337.931	2337.931	237	0.00
Value SHG loans in infancy	2540.206	2540.206	237	0.00
Value SHG loans at age 1	2438.496	1910.937	237	0.38
Value SHG loans at age 2	3286.841	806.526	237	1.23
Value SHG loans at age 3	3702.030	-1635.24 ^a	237	2.91***
Number SHG meetings mother attended:				
Number SHG meetings in utero	5.108	4.675	237	0.37
Number SHG meetings in infancy	5.897	5.897	237	0.00
Number SHG meetings at age 1	6.957	3.918	237	2.67***
Number SHG meetings at age 2	8.153	1.589	237	6.77***
Number SHG meetings at age 3	9.377	-0.568 ^a	237	13.71***
Height-for-age Z-score	-0.570 ^c	-0.763 ^c	270	1.00

Notes - Each mean difference was estimated using sibling fixed effects. Only children of members and dropouts were included. A child was said to have been exposed to SHGs if his or her household belonged to an SHG at any point between when the child was conceived and age four.

a. While the true means are positive, the means predicted by the fixed effects regression may be slightly negative.

b. Financial stresses included property divisions, job losses, large medical bills, marriages, and family separations.

c. Both height-for-age Z-score averages are slightly larger than the overall average of -0.817 mentioned in the introduction. This is likely due to the fact that averages presented here do not include non-affiliated households, which tend to be poorer than members and dropouts.

Table 2: Results

	(1)	(2)	(3)	(4)
	Height-for-age Z-score	Height-for-age Z-score	Height-for-age Z-score	Height-for-age Z-score
	<i>Coefficient</i>	<i>Coefficient</i>	<i>Coefficient</i>	<i>Coefficient</i>
	<i>(Standard Error)</i>	<i>(Standard Error)</i>	<i>(Standard Error)</i>	<i>(Standard Error)</i>
Fraction years SHG_{-1 ≤ a ≤ 3}	0.758*** (0.150)	0.700*** (0.347)	0.920*** (0.361)	- -
SHG_{-1 ≤ a ≤ 3}	-	-	-	0.451** (0.220)
Controls	Child age Child gender	Child age Child gender	Child age Child gender	Child age Child gender
Fixed Effects	- -	- Sibling	X Sibling	X Sibling
Other Sample Restrictions	-	Member & dropout children	Member & dropout children	Member & dropout children
Observations	422	270	269	269

Table 3: Threats to Validity

	(1a)	(1b)	(2)
	Height-for-age Z-score Coefficient (Standard Error)	Height-for-age Z-score Coefficient (Standard Error)	Height-for-age Z-score Coefficient (Standard Error)
Fraction years shock $_{-i \leq a \leq 3}$ (one year prior to joining)	-0.076 (0.646)	-	-
Fraction years shock $_{-i \leq a \leq 3}$ (two years prior to joining)	-	-0.139 (0.364)	-
Fraction years SHG $_{-i \leq a \leq 3}$	-	-	0.972*** (0.362)
Youngest son	-	-	0.298 (0.222)
<i>Controls</i>	Child age Child gender	Child age Child gender	Child age Child gender
<i>Fixed Effects</i>	X Sibling	X Sibling	X Sibling
<i>Other Sample Restrictions</i>	Member children	Member children	Member & dropout children
<i>Observations</i>	256	256	269

Table 4: Causal Channels

	(1a)	(1b)	(2a)	(2b)
	Height-for-age Z-score Coefficient (Standard Error)	Height-for-age Z-score Coefficient (Standard Error)	Height-for-age Z-score Coefficient (Standard Error)	Height-for-age Z-score Coefficient (Standard Error)
SHG _{-1 ≤ a ≤ 3}	0.339 (0.256)	0.693* (0.356)	0.081 (0.448)	0.166 (0.474)
Meetings _{-1 ≤ a ≤ 3}	0.009 (0.008)	-0.017 (0.020)	- -	- -
Meetings ² _{-1 ≤ a ≤ 3}	-	0.000 (0.000)	- -	- -
Meetings _{a < 0}	-	-	-0.007 (0.015)	0.008 (0.029)
Meetings ² _{a < 0}	-	-	-	0.000 (0.001)
Controls	Child age Child gender	Child age Child gender	Child age Child gender	Child age Child gender
Fixed Effects	X Sibling	X Sibling	X Sibling	X Sibling
Other Sample Restrictions	Member & dropout children Meeting information available	Member & dropout children Meeting information available	Member & dropout children Meeting information available	Member & dropout children Meeting information available
Observations	237	237	237	237

Appendix 1: P.E.D.O. Questionnaire

1. List 5-10 villages that are representative of all of the villages in which P.E.D.O. works in terms of demographics, socioeconomic conditions, and the degree of NGO involvement.
2. Are you aware of any changes to the health infrastructure and health environment of the village over the past decade? (Such changes might include droughts, the opening of a new health center, the expansion of an existing health center, the installment of water purification systems, or the implementation of the Midday Meal Scheme and other health-related policies by the government or another local NGO.)

Appendix 2: Household Questionnaire

For each woman:

1. Describe your current and historical affiliation with P.E.D.O.'s SHGs. If you have ever been a member, what year did you join? If you dropped out, what year did you leave?
2. What were your primary motivations for joining (or not joining) an SHG?
3. What were your primary motivations for continuing (or not continuing) to be a member?
4. On a scale of 1-5, with 1 being low and 5 being high, how strong a sense of community do you feel in your village? Is your sense of community related to your SHG membership or non-membership?
5. How many children do you have?
6. How many of those children are between the ages of four and ten, and are available to be interviewed?

For each of the woman's children mentioned in response to question 6:

1. What is this child's gender?
2. How old is this child?
3. How many older siblings does this child have?
4. How many people lived in your household when this child was born? How many of those household members were wage earners?
5. In what season was this child born?
6. For how many months was this child breastfed?
7. Did this child suffer from a particularly severe ailment when young? If so, how old was the child when this occurred?
8. Did any unusually adverse event affecting your household's finances occur when this child was young? If so, how old was the child when this occurred? (Such events may include the death, illness, or injury of a family member)
9. How many years of schooling has this child completed? Does this child still attend school?

Appendix 3: Assessing Potential Biases in the Data

Here I investigate concerns raised in Section III regarding biases in data collection and cleaning of height-for-age Z-scores. I examine correlations between height-for-age Z-scores and five control variables: household size at birth, number of household income earners at birth, years of schooling, and indicators for being sick while young and being a youngest son. I choose these six because they are the control variables that are unlikely to have been recalled with error. To investigate this issue, I regress height-for-age Z-scores on each control variable. For reasons detailed above, the regressions only analyze children of members and dropouts and also include controls for child age, gender, and sibling fixed effects. If the data collection and cleaning processes introduced marked biases, height-for-age Z-scores would not be correlated with control variables in the predicted direction.

One would expect height-for-age Z-scores to be negatively correlated with household size at birth as larger households would have to split a limited number of household resources between a greater number of household members. Likewise, conditional on household size, the number of income earners at birth should be positively correlated with height-for-age Z-scores. Years of schooling should also correlate positively with Z-scores since, in informal village schools, taller children may be assumed to be older and placed in higher grades. Moreover, height represents general nutrition, which may be positively correlated with scholastic aptitude. Illnesses at early ages may represent decreases in net calorie intake and hence should be negatively correlated with height-for-age Z-scores. And finally, for reasons detailed in Section VI.ii, youngest sons in this setting are expected to have higher height-for-age Z-scores than their siblings.

Table A.3 presents the results. Though not all precisely estimated, each of the relevant point estimates is relatively large in magnitude and is correlated with height-for-age Z-scores in the predicted direction. Column (1) suggests that household size at birth is indeed negatively correlated with height-for-age Z-score; the marginal household member at birth is associated with a 0.353 decrease in Z-score (p -value = 0.005). Conditional on household size at birth, the number of income earners at birth is positively correlated with Z-score.

This point estimate is large in magnitude (0.386) but not statistically significant (p-value = 0.543), which is likely to do multilinearity between household size at birth and income earners at birth. Column (3) suggests that an additional year of schooling corresponds to a 0.222 increase in Z-score (p-value = 0.001). The point estimates presented in columns (4) and (5) are not statistically significant (p-values = 0.563 and 0.259, respectively) but are relatively large in magnitude (-0.137 and 0.224) and suggest that being ill during childhood and being a youngest son are correlated in the predicted direction with Z-scores. Overall, this evidence is reassuring in light of concerns in data collection and cleaning of height-for-age Z-scores.

Table A.3: Assessing Potential Biases in the Data

	(1)	(2)	(3)	(4)	(5)
	Height-for-age Z-score				
	<i>Coefficient</i>				
	<i>(Standard Error)</i>				
Household size at birth	-0.353*** (0.124)	-0.352*** (0.125)	- -	- -	- -
Income earners at birth	- -	0.386 (0.633)	- -	- -	- -
Years of schooling	- -	- -	0.222*** (0.064)	- -	- -
Sick while young	- -	- -	- -	-0.137 (0.236)	- -
Youngest son	- -	- -	- -	- -	0.224 (0.198)
<i>Controls</i>	Child age and gender				
<i>Fixed Effects</i>	Sibling				
<i>Other Sample Restrictions</i>	Member & dropout children				
<i>Observations</i>	269				

Appendix 4: Additional Analyses

Two additional questions are whether the effect of SHG exposure depends on the age at which the child was exposed or upon the gender of the child. Equation (1) assumes that the effect on anthropometric status is the same in every period between conception and age four. However, evidence suggests that health investments made in utero may be particularly important in determining health-related outcomes (Barker, 1990; Field et al., 2009). To distinguish between the effect of SHG membership on household conditions in utero from the effect on conditions after birth and before age four, I replace the independent variable in equation (1) with two indicators: $SHG_{a=-1,i,m}$ and $SHG_{0 \leq a \leq 3,i,m}$. The first represents SHG exposure in utero and the second represents SHG exposure between birth and age four. As evidenced below in the first column of Table A.3, the coefficient on the first indicator is smaller than that of the second (coefficient difference of -0.266) and the difference is insignificant (p-value = 0.319); the data do not suggest that the effect of SHG membership on child nutrition depends on the age at exposure.

This analysis, however, is complicated by the fact that household SHG membership and child age are measured annually rather than monthly. As mentioned in Section III, this introduces a degree of imprecision in estimates of SHG exposure at every age. But the imprecision is more severe for in utero exposure and could produce a biased estimate. SHG exposure lasts for twelve or fewer months of the year in which a household joins but the in utero period lasts for a maximum of only nine months in the calendar year preceding a child's birth year. Hence $SHG_{a=-1,i,m}$ may have systematically overestimated exposure for children who were in utero in the calendar year during which their household joined. This would bias the estimate in column (1) downwards and could prevent the data from providing evidence of the fetal origins hypothesis.

Next, I examine how the effect of SHG involvement on anthropometric status differs based on child gender. To do so, I add a variable to equation (1) that interacts the SHG exposure variable with child gender. The results are presented in the second column of Table A.3. While the positive point estimate (0.155) suggests that male children may receive greater

benefits from SHG involvement than female children, this value is not significant (p -value = 0.633). The impact of SHG involvement on child nutrition does not appear to differ by child gender.

These findings seem to contradict the results of studies that suggest female children receive greater benefits than male children from policies that grant women greater intrahousehold bargaining power (Duflo, 2003; Qian, 2008; Thomas, 1990). There are at least three factors that could explain the apparent inconsistency. The first relates to differences in the populations and programs that were studied. Because female SHG members may only execute financial decisions their husbands make, SHGs may not truly alter intrahousehold decision-making. Moreover, any amount of female empowerment that SHGs do produce may not be large enough to overcome Dungarpur's patriarchal culture.

The second factor is differential attrition rates by child gender. One complication in data collection was that many older children in the age range could not be interviewed because they had been sent to tend the fields or to collect firewood. These responsibilities are typically given to female children, which means there are disproportionately fewer older girls in the sample. Even if girls did receive greater benefits than boys, this differential would be hard to identify due to the lack of older girls in the sample.

A third factor that could explain the inconsistency relates to the practice of sex selective abortion. Although it is illegal to do so, many households in rural Rajasthan choose to abort female fetuses (Iyengar et al., 2009). It is possible that female children in the sample tended to belong to larger households that chose to continue having children until they had a son rather than selectively aborting. Because determining child gender through ultrasound can be costly, these girls may also have belonged to poorer families. Children in larger, poorer households may gain less from SHG involvement for two reasons: first, household benefits would be split between a larger number of household members, and, second, household benefits would be smaller due to the inability of poorer households to save and borrow large amounts of money. Since these households may have more daughters, any tendency for female children to receive greater benefits from policies that empower their mothers may be

partially offset in this context by the practice of sex selective abortion. Due to the sensitive nature of the subject, however, it was not possible to obtain reliable information on the practice through the brief interviews the research teams conducted. Moreover, examining household gender ratios is not possible because the dataset contains information only on the children between the ages of four and ten that were available to be interviewed. Since I cannot identify households that selectively aborted female fetuses, I cannot truly isolate differential impacts of SHG involvement by child gender.

Table A.4: Additional Analyses

	(1)	(2)
	Height-for-age Z-score	Height-for-age Z-score
	<i>Coefficient Difference</i>	<i>Coefficient</i>
	<i>(F Statistic)</i>	<i>(Standard Error)</i>
SHG_{a=-1} - SHG_{0≤a≤3}	-0.226	-
	(1.000)	-
Fraction years SHG_{-1≤a≤3}	-	0.872***
	-	(0.375)
Fraction years SHG_{-1≤a≤3}x child male	-	0.155
	-	(0.325)
<i>Controls</i>	Child age	Child age
	Child gender	Child gender
	X	X
<i>Fixed Effects</i>	Sibling	Sibling
<i>Other Sample Restrictions</i>	Member & dropout children	Member & dropout children
<i>Observations</i>	269	269