An Experimental Analysis of Group Size and Risk Sharing

A. Chaudhuri[1], L. Gangadharan[2] and P. Maitra[3]

[1]University of Auckland, [2]Univerisity of Melbourne and [3]Monash University; Pushkar.Maitra@Buseco.monash.edu.au

Keywords: Reciprocity; Risk Sharing; Group Size; Experiments

EXTENDED ABSTRACT

In this paper we examine the relationship between group size and the extent of risk sharing in an insurance game played over a number of periods with random idiosyncratic and aggregate shocks to income in each period. Risk sharing is attained via agents that receive a high endowment in one period making unilateral transfers to agents that receive a low endowment in that period. The Pareto optimal allocation is for all agents to place their endowments in a common pool which is then shared equally among members of the group in every period. Theoretically, the larger the group size, the smaller the per capita dispersion in income and greater is the potential value of insurance. Field evidence however suggests that smaller groups do better than larger groups as far as risk sharing is concerned. These often suffer from differences in institutions and risk sharing arrangements that hinder comparability across groups. Results from our experiments show that the extent of mutual insurance is significantly higher in smaller groups, though contributions to the pool are never close to what efficiency requires.

Typically it has been argued that costs to group formation and other informational problems result in less cohesive behaviour in larger groups. In this laboratory set-up there are no costs to group formation and also there are no informational asymmetries per se. Agents are typically myopic in nature and they fail to realize the full benefits of risk sharing i.e., the fact that contributing to the pool when one receives a high endowment might not generate immediate returns but in the long run the benefits in terms of utility gain can be substantial. In the short run, any amount placed in the group account yields a return of $1/n$ where n is the group size, while amount placed in the private account yields a return of 1. The larger the size of the group, the lower is the short term return from contributing to the group account and this is what appears to be driving the result that contributions to the pool are significantly lower in the larger groups.

1. Introduction

Economic theory suggests that in a homogeneous population, the larger the population the higher is the per capita utility from risk sharing (see for example Genicot and Ray, 2003 and other references cited there). This implies that in the absence of any other impediments to group formation, a Pareto optimal solution to risk sharing would be to form as large a group as possible.

On the other hand, field evidence has shown that smaller groups do better than larger groups with respect to risk sharing. For example, there are a large number of papers that test for full consumption insurance at the village (community) level in developing countries. All of these papers reject complete risk-sharing at the level of the community (or an even larger ethnic group) and find evidence of only partial insurance. Moreover there is evidence suggesting that risk sharing actually occurs within smaller groups rather than at the level of the community as a whole. For example Morduch (1991) and Grimard (1997) find risk sharing within people of the same cast in India and people of the same ethnicity in CoteD'Ivoire. Fafchamps and Lund (2003) find evidence of gifts and transfers among a network of friends and relatives in response to income shocks in rural Philippines. Murgai, Winters, Sadoulet and De Janvry (2002) investigate water transfers among households along a water-course in the Punjab province in Pakistan and find that reciprocal exchanges are localized in units smaller than the entire water course community. It therefore appears that while the larger groups are unable to fully insure households from insuring against income fluctuations, smaller sub-groups are doing a better job of it.

How do we reconcile the theoretical predictions relating to risk sharing and the evidence from the field? One way is to take into account other considerations like informational decay or costs to group formation that increase with the size of the group that ultimately affect large groups. Another issue is that in arguing that smaller groups perform better than larger groups in the field, we are comparing across groups, and in a sense comparing apples and oranges. When we compare a group of size n_1 in community X to a group of size n_2 in community Y with $n_1 < n_2$ we are essentially comparing across different institutions and that might be contaminating the results. To be able to conclude that the extent of risk sharing is greater in smaller groups we need to hold the institutional arrangement fixed and then vary the size of the group within that institution. This is difficult, if not impossible, to do using data from the field. Economic experiments, on the other hand, provide us with a unique opportunity to examine the impact of group size on risk-sharing. Experiments allow us to control for the institution (defined by the experimental design and the parameters) and then vary the size of the group. The relationship between group size and the extent of insurance would no longer be contaminated by variations in institutions.

In this paper, we use an insurance game to compare the behaviour of small groups (with 5 members) with that of large groups (with 25 members). We implement a multi-period game, in which subjects in both small and large groups get either a high or a low endowment with equal probability. Apart from this individual level risk, subjects also experience an aggregate uncertainty in the sense that the number of people with a high or low endowment varies from one period to the next depending on a random draw. Subjects can fully insure their earnings against the individual uncertainty by placing their entire endowment into a group account in each period with the total amount in the group account being distributed equally among all group members.[1]

2. Theoretical Framework

Consider a community of n identical agents engaged in the production and consumption of a perishable good at each time period t. Each agent receives a random income that takes on two values h (with probability p) and l (with probability $1-p$) with $h > l > 0$. Income realizations are independent and identical both over individuals and also over time periods. There is full information – all agents know the realization of the shock. Each agent has the same utility function that is increasing, smooth and strictly concave in consumption. This is therefore a classical group insurance problem. The (symmetric) Pareto optimal allocation is attained by dividing the aggregate resource available at each period equally among all members of the community. It follows

[1] They cannot insure against aggregate uncertainty.

immediately that the larger the group size the smaller the dispersion of per capita output and the larger the potential value of insurance.

The (symmetric) Pareto optimal allocation is attained by dividing the aggregate resource available at each period equally among all members of the community. What this implies is that the different members of the community pool all the available resources and the total amount of pooled resources are then divided among the different members of the community. Mutual insurance requires that once the shock is realized, agents that receive a high endowment make unilateral transfers to agents that receive a low endowment. Risk sharing is obtained because of reciprocal behaviour on the part of agents. We can call this the *insurance game*.

How does risk sharing work within the group? Once the shock is realized agents contribute a share of their income to a common pool. The pool is then distributed among members of the group according to some pre-determined rule. The pre-determined rule that we will consider in this paper is equal sharing of the common pool, so that the amount collected in the common pool is divided equally among the members. What we have therefore is essentially a transfer of resources from those with a high draw to those with a low draw. All members of the group can contribute to the pool and the pool is divided equally among all members of the group, irrespective of whether the member receives a high endowment or not and whether she contributed or not. Each player knows his own endowment and the total number of high (or low) draws in the community. We can then compare the actual extent of risk sharing with what should happen under Pareto optimality. Bear in mind that once the shock is realized there is an incentive for a subject with a high draw to deviate and contribute nothing to the group account in that particular period. But the expected utility from contributing is always higher than the expected utility under autarky. Note that once the shock is realized, there is always an incentive for the individuals with a high draw to deviate (and contribute 0),

The type of risk sharing mechanism being implemented here (and by extension those in many field situations) is essentially based on mutual obligation and reciprocity, even though there is no commitment or enforceability.[2] There is no punishment for "cheating" via non-contribution in the event of receiving a high draw. The only potential consequence is the loss of faith by other group members who in turn might respond by not putting anything into the group account when they get high draws in turn. Given that each member of the group knows the distribution of high and low draws in each period, it is trivial for them to detect cheating on the part of a group member.

We test several hypotheses relating to individual behaviour in terms of risk sharing. The first hypothesis relates group size and contributions to the pool.

Hypothesis 1:

Agents in large groups contribute more to the pool (as a proportion of their endowment in the period) compared to agents in small groups

Mutual obligation and reciprocity implies that agents with a high endowment would voluntarily transfer some of their resources to those who are less fortunate (agents that receive low endowment). Of course it must be noted that while this reciprocal behaviour is consistent with risk sharing and risk aversion on the part of economic agents (endowments are uncertain and those with high endowment contribute more to the pool today, hoping that they would in turn be the beneficiary of voluntary contribution by some others when they have a low endowment), it could also be the result of altruistic behaviour on the part of agents or inequality aversion on the part of agents. Becker (1974) models a utility function that is comprised of two elements: the agent's own wealth and the wealth of other members of the group. Utility increases as the agent's wealth increases and as the wealth of the other group members increases. This model predicts that individuals with a high endowment will contribute more, in an absolute sense, to the pool than individuals with a low endowment. The second type of model is inequity-aversion. Fehr and Schmidt (1999) use a utility function where utility decreases (asymmetrically) when an agent earns either more or less than the average group payoff. Bolton and Ockenfels (2000) use a utility function that combines self-interest with a concern for relative standing. These models

[2] Actually these models have been categorized under the broad heading of models of risk sharing without commitment.

have the same predictions in our setting. Hence irrespective of which behavioural pattern (reciprocity, altruism or inequity aversion) motivates them, individuals with higher income or wealth should contribute a larger share of their income to the pool (i.e., do more risk sharing) than those with lower income or wealth. So the second hypothesis we test is:

Hypothesis 2:

Agents with high endowment contribute more to the pool compared to agents with low endowments both as a proportion of their endowment in the period and in an absolute sense.

3. Experimental Design

160 subjects participated in the experiment. These are undergraduate and post-graduate students from Monash University and the University of Melbourne.

Each session consists of two parts. In the first part subjects fill out a questionnaire designed to elicit their risk preferences. For this part, participants were presented with ten lotteries (referred to as choice games in the instructions given to the subjects). Holt and Laury (2002) use a similar set up to experimentally elicit risk preference from agents. Each lottery involved a choice between two options. Option A always yielded $5.00 with certainty. Option B was risky and paid either $10.00 or $0.00 (see Appendix A) with the probability of winning $10.00 changing (in 10% increments) from 10% on the first lottery to 100% on the tenth lottery. Lottery 10 (where Option B paid $10.00 with certainty) was included to ensure consistency (ideally every agent should choose Option B in the 10th game). The participant's pattern of choices provided an ordinal measure of their risk attitude in this context. Risk aversion is represented by the convexity or concavity of an individual's utility function when faced with the choice between an uncertain payoff and a safe bet. One way to assess the convexity or concavity of this function is to find the bet at which the participant is indifferent between the safe and risky option. In the present context this point is represented by the lottery at which the participant switches from choosing Option A to Option B. Individuals who switched from Option A to Option B after Game 5 are coded as being risk averse, those that switched before Game 5 are risk lovers and those who switch at Game 5 are risk neutral.

Once the lottery choices had been made subjects moved on to the insurance game which was conducted using the ZTREE software (Fischbacher, 1999). Each group played the game for at least 20 periods and the end period was randomly determined by throwing a six sided die. After the 20th period, the experiment continued for an additional period with a probability of $\frac{5}{6}$ and the experiment stopped as soon a "6" was rolled. At the beginning of each period the subjects were informed about their endowment for that period, which could be either high or low. A high endowment was 100 tokens and a low endowment was 20 tokens in all treatments except the increased inequality treatment where the high endowment was 200 tokens. Subjects did not know the exact endowments of the other members of the group but they were told how many players in the group received a high endowment in that period. The endowments were randomly determined using a random number generator. The subjects then had to decide how many tokens to contribute to the pool. Tokens placed in the group account were added up and divided equally among the group members. At the end of each round the players received the following feedback: the number of tokens contributed to the group account. and their earnings for that round. The subjects could track their earnings on a personal record sheet. Each session lasted around 1 hour (including the lottery game) and the average payoff for the insurance game was AUD 14.10.

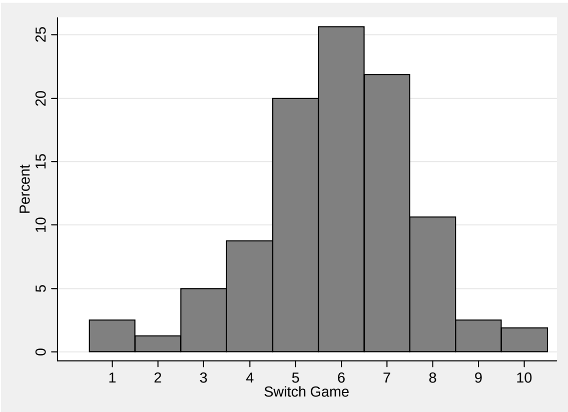

Figure 1: Choice in the Risk Assessment Game

4. Results

Figure 1 presents the histogram of the choice where the participants switched from the risk free Option A to the risky Option B. It is clear the majority (62.5%) of the subjects are risk averse in the sense that they switch from Option A to Option B in Game 6 or later, 20% of the subjects are risk neutral (switch from Option A to Option B in Game 5) and the remaining (17.5%) of the subjects are risk lovers (switch prior to Game 5).[3]

The majority of the subjects contribute 0 and the percentage contributing 0 varies from 73 percent in the large group sessions to 49 percent in the increased inequality sessions. Likewise the proportion contributing the maximum (10) varies from 9 percent in the baseline sessions to around 2 percent in the high probability of bad shock sessions.

Our primary interest is the proportion of their endowment that individuals contribute to the group account. We normalize the contributions by the endowment of the individuals: define p_{it} as the proportion of his/her endowment contributed by player i in period t. This is estimated using a Random Effects Tobit Model.

The following results are worth noting.

1. Contributions fall over time.
2. the proportion contributed to the public pool is significantly lower in the larger groups. *Hypothesis 1* is therefore not supported by the data.
3. Players receiving low endowment contribute to the pool significantly more in terms of the proportion of their endowment compared to players who receive high endowment. *Hypothesis 2* is therefore not supported by the data. Our results also do not support the inequity aversion model (Fehr and Schmidt, 1999 and Bolton and Ockenfels, 2000) where agents are concerned with relative standing and predict that individuals with higher income or wealth should contribute a larger share of their income to the public good than those with lower income or wealth. Our results show that the opposite is true – individuals with lower income actually place a larger share of their income to the pool.
4. Does the aggregate state in the period have an effect? The regression results show that the proportion contributed is significantly lower when there are a larger number of subjects with low endowment in the group.
5. The Random Effect Tobit regressions show that the proportion contributed by the subject in the previous period does not have a statistically significant effect on the proportion contributed in this period, but interestingly the proportion of the total endowment contributed by the group in this period has a positive and statistically significant effect on contributions in this period. The proportion of the total endowment placed in the pool could therefore be viewed as an indicator of the reciprocity of the other members of the group and players use this information to determine the extent of reciprocity in the group and hence their level of contributions.
6. It is interesting to note that contribution levels are significantly lower for risk-averse agents.
7. Contribution levels are not significantly different for males or for Business/Commerce majors.

[3] 21 of the 160 subjects either did not switch or kept switching between Options A and B. We defined their switch as the game where they switched from Option A to B for the first time. For those that always chose Option A, we coded their switch at Game 10.

References:

Becker, G. (1974): "A Theory of Social Interactions," *Journal of Political Economy*, **82** (6), 1063 – 1093.

Bolton, G. E. and A. Ockenfels (2000): "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, **26**(1), 51 – 74.

Charness, G. and G. Genicot (2004): "An Experimental Test of Risk-Sharing Arrangements", *Mimeo University of California, Santa Barbara*.

Fafchamps, M. and S. Lund (2003): "Risk-Sharing Networks in Rural Philippines", *Journal of Development Economics*, **71**(2), 261 – 287.

Fehr, E. and K. Schmidt (1999): "The Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, **114**(3), 817 – 868

Fischbacher, U. (1999): "Z-Tree. Toolbox for Readymade Economic Experiments", *IEW Working paper 21, University of Zurich.*

Holt, C. and S. Laury (2002): "Risk Aversion and Incentive Effects", *American Economic Review*, **92(**5), 1644 – 55.

Genicot, G. and D. Ray (2003): "Group Formation in Risk-Sharing Arrangements", *Review of Economic Studies*, **70**(1), 87 – 113.

Grimard, F. (1997): "Household Consumption Smoothing through Ethnic Ties: Evidence from Cote D'Ivoire", *Journal of Development Economics*, **53**(2), 391 – 422.

Morduch, J. (1991): Consumption Smoothing Across Space: Tests for Village Level Response to Risk", *Mimeo, Harvard University*.

Murgai, R., P. Winters, E. Sadoulet and A. De Janvry (2002): "Localized and incomplete mutual insurance", *Journal of Development Economics*, **67**(2), 245 – 274.

Ray, D. (1998): "Development Economics", *Princeton University Press*.