An Experimental Test of Risk-Sharing Arrangements

Gary Charness and Garance Genicot*

September 7, 2003

EXTREMELY PRELIMINARY. Please do no quote.

Abstract

Our project experimentally tests different implications of risk sharing without commitment. Subjects are matched in pairs. Each period, they receive income with a random component, and after observing their and their partner's income, each person in the pair can decide to make a transfer to the other person, knowing that their relationship may not last. At the end of each period, with a given probability, all pairs are broken and subjects are re-matched. Otherwise, they start a new period with the same partner. At the end of the experiment, one period is randomly drawn to count for real money. Subjects all face the same variance in their income but not necessarily the same mean. The impacts of a higher continuation probability, of changes in the risk aversion, and of inequality on risk-sharing are investigated.

JEL Classification Numbers: C91, D31, D8, D63, O17, Q15

Key Words: Risk-Sharing, Informal Insurance, Experiment

---

* Gary Charness, Department of Economics, UCSB, 2127 North Hall, Santa Barbara, CA 93106-9210; charness@econ.ucsb.edu. Garance Genicot, Department of Economics, Department of Economics, Georgetown University, 559 ICC, 37 & O Streets NW, Washington, DC 20057; gg58@georgetown.edu.

1. Introduction

In most of the developing world, people are exposed to substantial, even catastrophic, risk. Mitigating this risk a central concern. A high degree of dependence on agricultural production, widespread poverty, and the lack of access to formal insurance and credit make the need for consumption-smoothing particularly acute. Most individuals respond to the large fluctuations in their income by engaging in informal risk-sharing by providing each other with help, gifts and transfers, with some reciprocation expected. There is considerable evidence of the presence of some, but limited, insurance in village communities (Deaton, 1992, Townsend, 1994; Udry, 1994, Jalan and Ravallion, 1999; Ligon, Thomas and Worrall, 2002; Grimard, 1997, Gertler and Gruber, 1997; Foster and Rosenzweig, 2002).

The most important limitation appears to arise from the lack of enforceability of risk-sharing agreements. The fact that these agreements must be designed to elicit voluntary participation often seriously limits the extent of insurance they can provide. There is a growing theoretical literature on self-enforcing risk-sharing agreements. Some important contributions are Kimball (1988), Coate and Ravallion (1993), Kocherlakota (1996), Kletzer and Wright (2000), and Ligon, Thomas and Worrall (2002). All of these studies define self-enforcing agreements as those that are proof from non-compliance by individual members of the group. According to the theory, the individual defector is isolated from the community, so that she must self-insure. With this insight in place, the common practice in the literature has been to define self-enforcing risk-sharing agreements as subgame-perfect equilibria of a repeated game (in which self-insurance is always an option), and to characterize the Pareto frontier of such equilibria. Most studies focus on risk-sharing among identical agents. One exception is Genicot (2003), who investigates the effect of wealth inequality on these voluntary risk-sharing agreements. Studying risk-sharing between two agents facing the same income fluctuations and same preferences but differing in their mean income or permanent income, she shows that, surprisingly, inequality helps risk-sharing in a large range of cases.

In this project, we experimentally investigate different implications of these models of risk-sharing without commitment.

We have only found a rather limited experimental literature on risk-sharing. Selten et al. (1998) describe the *Solidarity Game*, a three-person-game where each player has an *ex ante* independent 2/3 chance of winning 10 DM and a 1/3 chance of receiving nothing. Before learning the outcome, each player was asked to decide conditionally how much she was willing to give to one loser or to each of two losers if she actually won the 10 DM; if the event occurred, the indicated allocation was enforced. The great majority of subjects were willing to make substantial conditional gifts. Five types are distinguished, with the most common type (36%) giving the same positive total amount to one loser and to two losers. This behavior does not lend itself to a straightforward interpretation as the result of altruistic utility maximization.[1]

Goeree et al. [2000] examine experimental results for a variety of generalized matching penny games, and show that a structural econometric model which incorporates risk aversion into a quantal response equilibrium explains the data very well. Moreover, they find risk aversion estimates around .5 which are stable across the different games and are close to those obtained from laboratory and field auction data, as well as from individual lottery choice experiments.

[1] Subjects were also asked to estimate the average gifts of others. There is a significant positive correlation between the estimates and the subjects' own gifts. This is similar to the *false-consensus effect* known in the literature (e.g., Ross et al., 1977). Among male subjects, those studying economics show a more "egoistical" behavior than others. Among female subjects, no such "education effect" can be found. Females tend to give more than males.

Bone et al. (2000) report an experiment designed to test whether pairs of individuals are able to exploit efficiency gains in the sharing of a risky financial prospect (taking advantage of their difference in risk aversion, with commitment). The results indicate that fairness is not a significant consideration, but rather that having to choose between prospects diverts partners from allocating the chosen prospect efficiently. The pattern of agreements suggests that, where allocation is the sole issue, partners largely favor *ex ante* efficiency over *ex post* equality. From the transcripts there is little indication that *ex post* fairness is a significant consideration in this.

The remainder of our paper is organized as follows. The next Section lays out the basic model of risk-sharing without commitment and describes some important implications. Section 3 then presents the experimental design. In Section 4, the main results of the experiment are presented. Section 5 discusses the implications and some limitations of the paper, and Section 6 concludes.

2. A Model of Risk-Sharing without Commitment

A standard model of risk sharing without commitment goes as follows. Time is discrete and the number of period is infinite. In each period t, two agents, indexed by $i \in \{1,2\}$, receive an income y_i and one of them, randomly chosen, incur a fixed monetary gain h. They each have a probability ½ to receive h but the aggregate income is constant at $Y = y_1 + y_2 + h$ in each period. The following table summarizes the income distribution of the two agents:

individual \ state	state 1 (proba ½)	state 2 (proba ½)
1	$y_1 + h$	y_1
2	y_2	$y_2 + h$

In line with standard practice, let us assume that all agents have additively time-separable Von Neumann-Morgenstern utility functions defined over consumption, such that the expected lifetime utility at time t is given by

$$E_t \sum_{j=0}^{\infty} \delta^j u(c_{t+j}^i) \text{ for all } i \in \{1, 2\}.$$

where $u' > 0$, $u'' \geq 0$, $\lim_{c \to 0} u(c) = -\infty$, and δ is the discount rate. The operator E_t is the expectation conditional on what is known at time t.

Since individuals are risk-averse, optimality would require that the ratio of their marginal utilities remains constant over time and across state of nature. When the aggregate income is constant this implies keeping each individual's consumption at a constant level. The exact levels depend on the welfare weights used but must sum to the aggregate income and satisfy the voluntary participation constraint. If c^* is the optimal level of consumption for individual 1, then it must be that

$$\begin{cases} u(c^*) \geq \frac{1}{2}u(y_1 + h) + \frac{1}{2}u(y_1) \\ u(Y - c^*) \geq \frac{1}{2}u(y_2 + h) + \frac{1}{2}u(y_2) \end{cases}$$

As motivated and discussed in the Introduction, we focus on the theme that insurance arrangements must be self-enforcing, and that this requirement constrains the form of such arrangements. The *enforcement constraint* refers to the possibility that at some date, an individual who is called upon to make transfers to others in the community refuses to make those transfers. To be *self-enforcing*, a risk-sharing agreement must be such that the expected net benefits from participating in the agreement is at any point in time larger than the one time gain from defection. The literature on risk-sharing concentrates on the constrained optimal or "second-best'" self-enforcing schemes. Hence, the constraint is modeled by supposing that the individual is excluded from the insurance pool, so that he must bear stochastic fluctuations on his own. That is, a risk-sharing agreement resulting in a stream of consumptions $\{c_t^i\}_{\forall t}$ for individual i must be such that, at any period t,

$$u(c_t^i) + E_t \sum_{j=1}^{\infty} \delta^j u(c_{t+j}^i) \geq u(z_t^i) + E_t \sum_{j=1}^{\infty} \delta^j u(z_{t+j}^i) \text{ for all } i \in \{1, 2\} \tag{1}$$

where z_t^i is the total income of individual i at time t.

If the power of such punishment is limited, then perfect insurance may not be possible. However, even when full risk-sharing is not possible, individuals may be able to design a risk-sharing agreement by limiting transfers in states for which the enforcement constraint is binding (see Coate and Ravallion (1993) and Kocherlakota (1996) among others).

It is well known that with this simple distribution, when a first best is not incentive compatible, the *constrained optimal agreement* is fully characterized by two values, t^*_1, the transfer made by 1 to 2 when 1 received h and, t^*_2, the transfer made by 2 to 1 when 2 received h. These transfers are such that the incentive constraints (1) hold with equality for both agents, that is $t^* \equiv (t^*_1, t^*_2)$ is defined by

$$\begin{aligned} (1-\frac{\delta}{2})u(y_1 + h - t^*{}_1) + \frac{\delta}{2}u(y_1 + t^*{}_2) &= (1-\frac{\delta}{2})u(y_1 + h) + \frac{\delta}{2}u(y_1) \\ (1-\frac{\delta}{2})u(y_2 + h - t^*{}_2) + \frac{\delta}{2}u(y_2 + t^*{}_1) &= (1-\frac{\delta}{2})u(y_2 + h) + \frac{\delta}{2}u(y_2) \end{aligned} \tag{2}$$

A first implication of this model is that, when full insurance is not achieved, a higher discount rate δ increases the weight put on the long term gain from insurance relative to the short term gain from deviating. Hence, a higher δ raises the transfers that individuals are able to make to each other and so the level of risk-sharing that they can achieve.

A second implication of the model is that an overall increase in the risk aversion exhibited by the agents, by increasing the long term gain from insurance, will increase t^* and the insurance that individuals can provide for each other.

Now what would the effect of inequality be? Let's consider different values of y_1 and y_2 keeping the aggregate income Y constant. Clearly, if $y_1 = y_2$ both individuals are ex-ante identical. Now, increasing y_1 and decreasing y_2 to keep Y unchanged would make 1 relatively richer than 2 while keeping the variance of their income constant. To be sure, the set of Pareto optimal allocations is unaffected since the aggregate income is the same. However, the division of wealth affects the autarchic utility and thereby *does* affect the set of self-enforcing allocations. Genicot (2003) shows that for a large range of utility functions such spread-preserving inequality between the two agents increases the likelihood of first-best risk-sharing and increases the transfer that agents make to each other within the constrained optimal agreement.

In what follows, we will describe an experiment in which we replicated the setting of this model of risk-sharing without commitment and test these predictions.

3. Experimental Design

Our experiment was conducted at the CASSEL Laboratory in UCLA. We had six sessions, with either 12, 14, 16, or 18 participants in a session (depending on show-ups). Participants earned an average of about $17, including a $5 show-up fee, for a bit over an hour of their time. Note that participants were never told that this experiment was about risk-sharing nor were the terms risk-sharing or insurance ever used during the experiment. The experimental instructions for the $\delta = .9$ treatment are provided in the Appendix.

We first asked people to complete an investment question. Each person was provisionally endowed with 100 units ($10) and could invest any portion of this amount in a risky asset that had a 50% chance of success and paid 2.5 times the amount invested if successful. The decision-maker retained the units not invested.[2] We told the participants that we would later choose two people at random in each session for actual payoff implementation, and a coin was flipped after the session to determine success or failure for these investors. The objective of this investment question was two-fold. First, it provides us with a measure of risk aversion for each individual. To be sure, the higher the investment the less risk averse the individual is. Second, we use the answer to this question to match individuals with similar degrees of risk aversion.

The body of the session consisted of a number of matches. In a match, each participant was matched with one other person (the same person) for the duration of the match we paired people who had invested more (less) than 67 in the risky investment with other people who had invested more (less) than 67. Each match was comprised of an uncertain number of rounds; the number of these rounds was determined as follows: After each round, the computer determined (for all current matches) whether another round would follow. In three of our sessions (Treatment 1), the continuation probability was 80% and in the other three sessions (Treatment 2), the continuation probability was 90%. In the first case, the expected number of subsequent rounds in a match (at any point in time after the first round) was four; in the latter case, the expected number of subsequent periods was nine; the participants in the corresponding sessions were informed of this mathematical fact.

In this manner, we avoid the unraveling problem resulting when the number of rounds in any match is known in advance. *Ex ante*, we therefore expected each match to last five periods in Treatment 1 or 10 rounds in Treatment 2. When the matches ended, all participants were

[2] This design was used in Charness and Gneezy (2002), who adapted it from the design in Gneezy and Potters (1997).

randomly re-matched for the next match. We had 10 matches in each 80% session and seven matches in each 90% session.

In each round, each person will received income, which was comprised of a fixed portion and an amount that was added to the fixed income for that round for **one** of the people in each match. The person receiving this extra amount was randomly chosen in each pair for every round of the match. The fixed portions did not vary during the match, but did change from match to match. In some matches, both fixed portions were 70 units, while in other matches one fixed portion was 20 and the other was 120. In all cases, the amount randomly assigned and added was 200 units.

In the beginning of each round, each participant learned her fixed income, the fixed income of the person with whom she was paired, and which person received the extra 200 in that round. Everyone then chose a non-negative amount, not to exceed the income received, to transfer to the other person and these designated amounts were then transferred. Participants saw a history of the income and transfers for each previous round in that match, and could also review their previous matches. We also asked participants for some information about their decisions and expectations in the beginning of the first and fourth matches. Specifically, the individuals who received h were asked: "what motivates their choice of transfer?" and the others (who did not receive h) where asked: "what transfer do you expect the other person to give you?" At the end of the session, participants also answered questions concerning their gender and major.

To avoid possible wealth effects and to enhance the element of uncertainty in payoffs, we chose only one round (of the many that were played in the session) for conversion of experimental payoff units to real dollars, at the rate of 17 experimental units to $1.

4. MAIN RESULTS

We first present some summary statistics about our data and then discuss in turn several important questions in relation to our results:

1. Does a higher continuation probability increase the amount of risk-sharing?
2. Does a higher degree of risk aversion increase the amount of risk-sharing during the match?
3. What is the effect of inequality on risk-sharing (matches with equal fixed portions vs. matches with unequal fixed portions)?
4. Does the effect of this inequality differ across continuation probabilities?
5. Do demographics such as gender and major affect the transfer chosen?

Table 1 shows the average transfer made in each session, and the overall average for each treatment:

Table 1: Average Transfer, by Session and Treatment

	Avg. Transfer	# Observations	Std. Dev.
Session 1	15.83	656	27.91
Session 2	14.97	416	21.19
Session 3	13.42	384	27.12
Treatment 1	*14.95*	*1456*	*25.95*
Session 4	38.84	1162	48.31
Session 5	28.46	648	37.38
Session 6	20.96	846	34.94
Treatment 2	*30.61*	*2656*	*42.54*

We see that the average transfer in Treatment 2 is more than twice as high as the average transfer in Treatment 1. The average transfer in each session in Treatment 1 is lower than the average transfer in Treatment 2. Thus, the Wilcoxon-Mann-Whitney ranksum test (see Siegel and Castellan, 1988) on session-level data, a most conservative test that considers each session as only one observation, finds that transfers are significantly higher in Treatment 2 ($p = 0.050$).[3]

This aggregation ignores the substantial heterogeneity present in the population. Figure 1 shows the frequency with which each range of transfer is made in the two treatments:

Figure 1 - Distribution of Avg. Individual Transfers

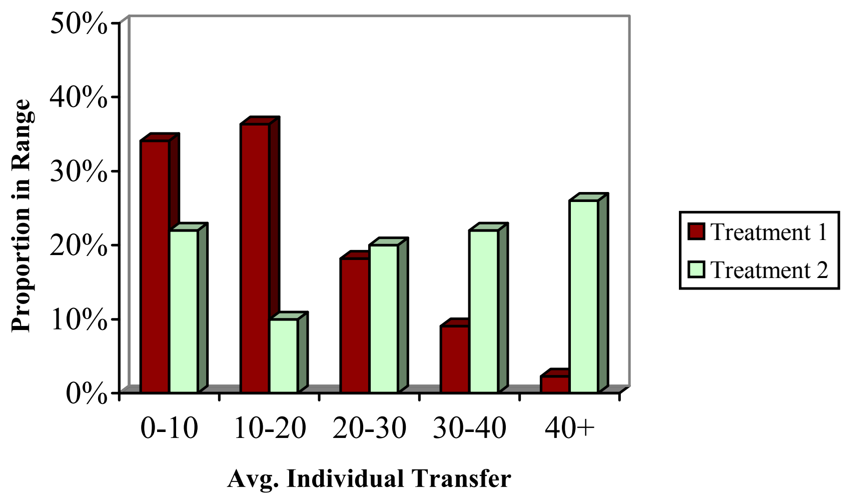

We see a great diversity of average individual transfers, particularly in Treatment 2. The Wilcoxon test on individual average transfers (not completely independent) confirms that these are higher in Treatment 2 ($Z = 3.54$, $p = .000$).

While overall individual transfers are an important metric, it is not clear what to make of positive transfers when a participant has received less income in the round than the other matched

---

[3] Throughout the paper, we round *p*-values to three decimal places.

person. These may have some value through signaling one's cooperative nature, or may simply represent confusion. A better metric may be average transfers made when the chooser has the higher income. Table 2 shows these by session and treatment, and Figure 2 gives the distribution of individual average transfers made when ahead:

Table 2: Average Transfer when Ahead, by Session and Treatment

	Avg. Transfer	# Observations	Std. Dev.
Session 1	26,54	328	34.61
Session 2	22.28	208	26.14
Session 3	20.43	192	34.37
Treatment 1	*23.71*	*728*	*32.42*
Session 4	64.08	581	54.54
Session 5	43.31	324	43.59
Session 6	34.20	423	43.76
Treatment 2	*49.49*	*1328*	*50.48*

Figure 2 - Dist. of Avg. Ind. Transfers when Ahead

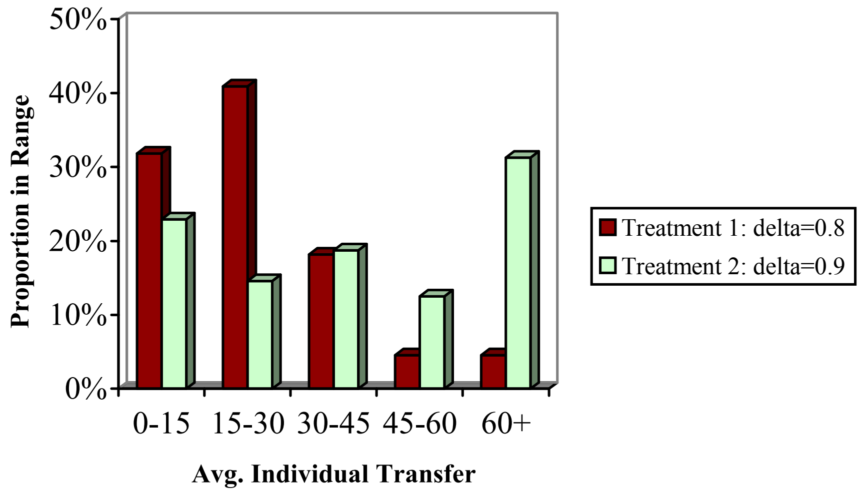

The average transfer made when ahead is higher in every session in Treatment 2, when the continuation probability is 0.9, than it is in Treatment 1 with a continuation probability of 0.8, and this average transfer when ahead is always higher than the corresponding overall average transfer in each session. Figure 2 shows a pattern similar to that seen in Figure 1, but with higher levels of average individual transfers. In fact, the average transfer when ahead was higher than the average transfer when behind for 85 individuals, and this was reversed for eight people.[4] A simple binomial test (See Siegel and Castellan, 1988) finds this to be extremely significant ($Z = 7.98$, $p = 0.000$). It is clear that people are not just randomly and arbitrarily transferring money, but are instead quite sensitive to which matched person receives higher income in the round.

Given our matching structure, it may also be useful to consider the average transfer made in each match. The distribution of these average transfers is shown in Figure 3:

[4] The remaining person always chose a transfer of 0.

Figure 3 - Distribution of Avg. Transfers, by Match

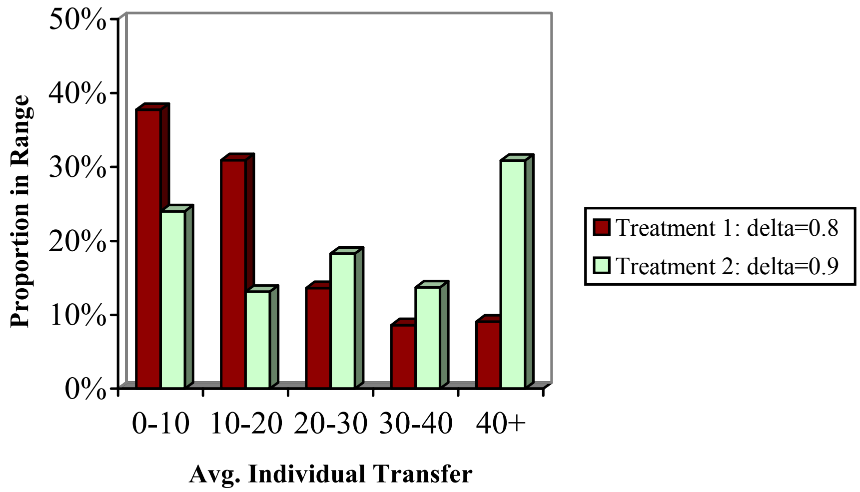

The difference between treatments in Figure 3 is perhaps even stronger than in Figures 1 and 2. A Wilcoxon test confirms the difference is highly significant ($Z = 5.96$, $p = 0.000$).

Another way to see that individuals in a match are insuring each other is to consider the standard deviation of realized consumption for a pair of individuals in a match.[5] In the absence of transfers, they would face a standard deviation of about 100. Table 3 illustrates clearly that, by making transfers to each other, individuals are sharing risk. This is particularly true in Treatment 2 when the continuation probability is high and when the fixed income of the individuals is the same ($y_1 = y_2 = 70$).

Table 3: Average Std. Dev. of Consumption, by fixed income (y_i) and treatment (δ)

		Avg. Std. Dev. of Consumption	# Observations
	$y_i = 20$	90.65	446
Treatment 1 (δ=0.8)	$y_i = 70$	85.07	564

[5] Note that in this simple model the standard deviation of consumption of individuals matched with each other will be the same.

	y_i =120	90.65	446
Treatment 2 (δ=0.9)	y_i =20	84.18	643
	y_i =70	75.39	1,370
	y_i =120	84.18	643

Thus far we have established that we see significant transfers and that these transfers are highly dependent on whether the chooser has a higher or lower endowment. Hence, individuals are sharing risk. Moreover, the transfers and risk-sharing are definitely higher when the continuation probability increases, and when individuals have the same fixed-income (notice that the standard deviation is always lowest when y_i = 70).

To address the remaining questions mentioned at the beginning of this section, we supplement our non-parametric statistical analysis with some regression analysis.

Since we are interested in insurance, as a first cut it is interesting to look at the determinant of the standard deviation of consumption within matches. In Table 4, we can see that the risk aversion of the least risk averse agent (Rho1) significantly improve risk sharing.[6] This effect is slightly decreased by the interaction of the individuals' risk aversions (Rho1*Rho2). A higher continuation probability substantially reduces the standard deviation of consumption achieved. Equality in the fixed income has also a strongly positive effect on risk-sharing. Finally, we observe a weakly significant gender effect, as having at least one female in the group *increases* the standard deviation of consumption.

[note that the standard deviations are biased since we are ignoring the fact that each individual is in many pairs.]

[6] Assuming a constant relative risk aversion utility, their risk aversion would be given by

$$\rho \equiv \begin{cases} \dfrac{\ln(1.5)}{\ln(\text{inv} * 2.5 + 100 - \text{inv}) - \ln(100 - \text{inv})} & \text{if inv} < 100 \\ 0 & \text{if inv} = 100 \end{cases}.$$

Table 4 – Cross-Match Regressions

	Dependent Variable	
Independent variables	Std. Dev. of Consumption (1)	Std. Dev. of Consumption (2)
Rho1 (min rho)	**-114.0467** **[19.2027]**	**-115.8031** **[19.2612]**
Rho2 (max rho)	0.2577 [6.1168]	-0.2667 [6.1096]
Rho1*Rho2	**43.5896** **[14.8133]**	**44.6190** **[14.7540]**
Equal	**-12.3819** **[3.8215]**	**-12.0676** **[3.8074]**
Delta	**-273.7349** **[39.1720]**	**-282.3923** **[39.1680]**
Sex1		-3.4412 [4.3972]
Sex2		**12.0011** **[5.0872]**
Constant	**350.8338** **[33.8998]**	**349.9789** **[33.7482]**
Observations	395	395
R-squared	0.18	0.19

Standard errors in brackets
Bold indicates significance at 1% and **Gray** indicates significance at *5%*. Rho1 is the lowest risk aversion of the pair and Rho2 the highest; Sex1 = 0 if at least one men and 1 otherwise; Sex2 = 1 if at least one female and 0 otherwise; Equal = 1 if both have the same fixed income and is 0 otherwise.

Now let's look at the determinant of the transfers that individuals make when receiving a good draw as well as the net transfer that they give. These results are shown in Table 5. To account for unobserved individual characteristics, we use individual effects. Table 5 presents the results of the fixed effect, between-effect and random-effect regressions.

Table 5 - Transfer Regressions

Independent variables	Transfer when h fe (1)	Transfer when h be (2)	Transfer when h re (3)	Net Transf. when h fe (4)	Net Transf. when h be (5)	Net Transf. when h re (6)
Investment		-0.1976 [0.1271]	-0.2644 [0.1068]*		-0.2018 [0.1186]	**-0.2737 [0.1000]****
Partner's invest	-0.0803 [0.0484]	-0.2154 [0.1805]	-0.0920 [0.0468]*	-0.0953 [0.0501]	-0.2481 [0.1685]	-0.1107 [0.0481]*
Fixed_income	**0.1696 [0.0237]****	0.1305 [0.1633]	**0.1699 [0.0234]****	**0.2811 [0.0245]****	0.2436 [0.1524]	**0.2817 [0.0242]****
Equal	**7.7885 [1.6365]****	53.4916 [21.8162]*	**8.0533 [1.6333]****	**7.2987 [1.6933]****	50.9089 [20.3620]*	**7.6086 [1.6892]****
Discount rate		**161.3488 [54.5428]****	**211.1056 [49.6890]****		114.6531 [50.9073]*	**161.9128 [46.3856]****
Female dummy		-11.3957 [5.1219]*	-10.8036 [5.0572]*		**-13.0080 [4.7805]****	**-12.5715 [4.7188]****
Constant	**29.4328 [3.3164]****	-106.0718 [45.0035]*	**-134.5081 [42.8046]****	**12.9951 [3.4314]****	-78.9313 [42.0038]	**-106.4329 [39.9849]****
Observations	2056	2056	2056	2056	2056	2056
# of indiv	94	94	94	94	94	94
R-squared	0.04	0.30		0.07	0.31	

Standard errors in brackets

Bold indicates significance at 1% and Gray indicates significance at 5%. Invest and partner's invest are the answer to initial investment question for individual and his partner; Equal = 1 if both have the same fixed income and is 0 otherwise.

In the fixed-effect model, the hypothesis that all individual constant are the same is strongly rejected. A Hausman test rejects the random effect hypothesis that the individual effects are uncorrelated with the other regressors.

The coefficient on Investment is significantly negative in, and highly so when looking at the net transfer. As more investment means less risk aversion, we conclude that a higher degree of risk aversion increases the transfer that one chooses. On average, we would expect an individual who has the higher income and who chose to invest 20 in the risky asset to transfer 24 units more than a similar individual who chose to invest 100 in the risky asset.

A higher continuation probability clearly results in higher transfers. According to the random effect specification, increasing the continuation probability from 0.8 to 0.9 increase the transfers when high by 20 and increases the net transfer by 16.

Risk-sharing requires reciprocity. The amount invested by the individual's partner in a match should be negatively related to his risk aversion and so the transfers he makes. Hence, the negative relationship found between an individual's transfer and the investment made by his partner (which is not present when looking at only the first transfers made in a match) shows that reciprocity is an important factor.

When looking at individual transfers we cannot isolate the effect of equality and the income level effect, since changes to an individual's fixed income are concurrent with changes in his partner's income. The estimated coefficients of Fixed_income and Equal show that these two forces result in a strong non-linear effect of y_i, first increasing then decreasing. A fixed income of 70 for both individuals in a match results in substantially higher transfers upon receiving 200 then a fixed income of 20 or 120 in an unequal match. In unequal matches, individuals with a fixed income of 120 give slightly higher transfers than individuals with a fixed income of 20. Note that we did not find any significant or substantial difference in the effect of ex ante equality across continuation probabilities.

We also observe a modest but significant gender effect, as females make smaller transfers than males do. Net consumption is also significantly higher for females ($Z = 2.68$, $p = 0.007$, using a Wilcoxon test on individual average consumption.[7]

[7] Note that this is not due to females having better draws; females comprised 56.4% of the population and females had the larger endowment 55.6% of the time.

5. DISCUSSION, IMPLICATION AND LIMITATIONS

1. Implications of the model:

Our experiment provides some strong support for the model of risk-sharing without commitment. First, there is strong evidence that individuals are providing some but limited insurance to each other. Net positive transfers are going from individuals receiving a high shock to the other and these transfers substantially reduces the standard deviation of consumption.

Second, transfers are limited in a way that is consistent with the relatively low level of risk aversion exhibited by the participants.[8] The strong positive effect of higher continuation probability and risk aversion on the transfers and on the level of risk sharing provides strong support for the limited-commitment story.

The effect of inequality is harder to reconcile with the model presented in Section 2. For utility functions of the HARA class (hyperbolic absolute risk aversion), inequality should improve risk-sharing and not decrease it. Different explanations for this result are possible. First, inequality could make it harder for individual to coordinate. This seems to be part of the explanation. Looking at the pattern of transfer when high over time, we find strikingly different patterns when equal or not. Transfers are increasing and concave when equal and decreasing and convex when unequal (the critical points being high enough that usually not reached). However, this is probably not the only explanation.

Second, individual preferences may be very different from HARA utility functions. Under the model presented in Section 2, if individuals have HARA utility functions with decreasing risk aversion (as traditionally assumed) we should observe that individuals with a fixed

[8] These levels of risk aversion are similar to the ones found in other experiments.

income of 20 are making higher transfer when receiving 200 than individuals with a fixed income of 120 when receiving 200. The poorest agent is trading some mean consumption in exchange of more insurance. In our experiment, we observe that individuals with 120 are actually transferring more than individuals with 20, such that overall there is a small but positive net transfer from the individuals whose fixed income is 120 to the individuals whose fixed income is 20.

Third, preferences may not be defined only over consumption. However, it is important to note that in a model of risk aversion without commitment with utility functions such as Fehr and Schmidt (1999), Bolton and Ockenfels (2000), or Charness and Rabin (2002) inequality should improve risk sharing, not decrease it. Fourth, beliefs of reciprocity could be lower when heterogeneity is higher. More tests are clearly needed to better understand these results.

2. Risk aversion: Naturally inducing risk aversion on the participants rather than relying on their preferences would allow us to have a more accurate measure of risk aversion. However, there is little evidence that the method used to induce risk aversion, the binary lottery procedure, works and a couple of studies actually showing that it does not work (Camerer and Ho 1994, Selten et al. 1999).

6. CONCLUSION

7. APPENDIX

First transfer made in a match:

Table 6 – First Transfer Regressions

	Independent Variable					
Dependent Variables	(1) 1st tr when h, fe	(2) 1st tr when h, be	(3) 1st tr when 0, re	(4) 1st tr when 0, fe	(5) 1st tr when 0, be	(6) 1st tr when 0, re
Fixed_income	**0.1950 [0.0380]****	0.0403 [0.1737]	**0.1898 [0.0371]****	**0.1038 [0.0185]****	**0.2313 [0.0767]****	**0.1113 [0.0180]****
Equal	3.6726 [2.4683]	18.5340 [13.6715]	4.0693 [2.4336]	1.5841 [1.2067]	-5.5831 [5.8795]	1.2302 [1.1886]
Invest		**-0.3906 [0.1246]****	**-0.3781 [0.1246]****		-0.0349 [0.0570]	-0.0305 [0.0566]
Continuation Probability		**200.6754 [59.2359]****	**209.8432 [58.1746]****		**61.6849 [27.3570]***	50.5147 [26.7582]
Female dummy		-8.8300 [6.1019]	-9.4099 [5.9590]		-2.8130 [2.7533]	-2.5080 [2.7256]
Constant	**23.7547 [3.0338]****	**-117.5331 [50.9948]***	**-127.7270 [50.0443]***	1.9309 [1.6294]	**-51.6372 [24.0374]***	-37.6795 [22.9044]
Observations	395	395	395	395	395	395
# of indiv	94	94	94	94	94	94
R-squared	0.09	0.22		0.10	0.15	
Standard errors in brackets						
* significant at 5%; ** significant at 1%						

A Hausman test accepts the random effect hypothesis that the individual effects are uncorrelated with the other regressors for the first transfer when receiving 200 but rejects it for the first transfer when not receiving 200. In the fixed-effect model, the hypothesis that all individual constant are the same is strongly rejected.

Instructions:

Welcome to our experiment. For showing up on time, we will pay you a $5 show-up fee. In addition, you may receive additional earnings as the result of the outcomes in the experimental session. Today's session will take about an hour.

To begin, we ask you to complete a brief questionnaire. The body of the session will be comprised of a number of segments. In each of these segments, each participant will be matched with one other person. Each segment is comprised of an uncertain number of periods. The number of periods in a segment is determined as follows: After each period, the computer will 'roll a die' (for the entire room) to see whether another period will follow, with an 90% chance that another period will follow, and a 10% chance that the segment ends immediately. The computer will 'roll the die' after every period. With this continuation probability, the expected number of subsequent periods in a segment, **at any point in time**, is 9.

When the segment ends (10% chance after each period), **all participants will be randomly re-matched with other participants for the next segment**. We anticipate that there will be approximately 7 segments in the session, but this will vary according to how many periods there are in the segments – we aim to complete the session in about an hour.

In each segment, you and the person with whom you are matched will receive income. This income is composed of a fixed portion and an amount (200) that is added to the fixed income for that period for **one** of the people in each match; the person receiving this extra amount is randomly chosen in each pair for every period of the segment. The fixed portions will not vary during the segment, but will change from segment to segment; these fixed portions may or not be the same for the two people matched. **In all cases, this fixed portion will be considerably smaller than the 200 units that are randomly assigned.**

In the beginning of the period, you will learn your fixed income, the fixed income of the person with whom you are paired, and who received the extra 200 in the period. At this point, you choose to **transfer money** to the other person. **This amount must be non-negative and no more than the total income you received in that period.** The other person in your match simultaneously chooses to transfer money to you, subject to the same restrictions on the amount to be transferred. The designated amounts are then transferred, and the computer then determines whether another period follows in this segment. You will see a history of the income and transfers for each previous period in that segment.

Thus, you will be involved in many periods. We wish to make it clear that **only one of these periods will be chosen at random for conversion to real dollars, at the rate of 17 experimental units to one cash dollar.**

Let's take an example. Assume that your fixed income be 50 and that you are matched with someone whose fixed income is 90. In each round, either you get an additional 200 (50% chance) or the person with whom you are matched gets an additional 200 (50% chance). If in this round the other person receives this 200, your total income is 50 while his or her total income is 290. Now, you and the person with whom you are matched decide on transfers. Suppose you transfer x to the other person while he or she transfers y to you; then your "income net of transfer" or "consumption" for this round is $50-x+y$ while his "income net of transfer" or "consumption" for this round is $290-y+x$. If this round happens to be the one selected to count for actual payoffs, these are your and your match's payoffs for the experiment. For instance if x = 1 and y = 61 then your payoff would be 110 and your match's payoff would be 230.

At some points along the way, you will be asked for some information about your decisions and/or your expectations.

The history of income and transfers for the current match appears on your screen. By pressing the "full view" button you can also review the history of past matches.

At the end of the experiment, one period will be chosen at random for payment. The screen will state your earnings. When you have completed a short questionnaire on your demographics, we will distribute receipts forms for participants to sign, and will pay people individually and privately.

We highly encourage clarifying questions. Thank you for your participation.

References

Bone, J., J. Hey and J. Suckling (2000), "A Simple Risk-Sharing Experiment," mimeo.

Camerer, C. and T. Ho (1994), "Isolation Effects and Violation of Compound Lottery Reductions," mimeo.

Charness, G. and U. Gneezy (2002), "Portfolio Choice and Risk Attitudes," mimeo.

Coate, S. and M. Ravallion (1993), "Reciprocity without Commitment: Characterization and Performance of Informal Insurance Arrangements," *Journal of Development Economics* **40**, 1--24.

Deaton, A. (1992), *Understanding Consumption*. Oxford: Clarendon Press.

Foster, A. and M. Rosenzweig, "Imperfect Commitment, Altruism, and the Family: Evidence from Transfer Behavior in Low-Income Rural Areas," *Review of Economics and Statistics* **83**, 389--407.

Genicot (2003), "Inequality and Informal Insurance," mimeo.

Gertler, P. and J. Gruber (2002), "Insuring Consumption Against Illness," *American Economic Review* **92**, 51-70.

Goeree, J., C. Holt and T. Palfrey (2000), "Risk Averse Behavior in Generalized Matching Pennies Games," *Games and Economic Behavior*, forthcoming

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

Jalan, J. and M. Ravallion (1999), "Are the Poor Less Well Insured? Evidence on Vulnerability to Income Risk in Rural China," *Journal of Development Economics* **58**, 61--81.

Kimball, M. (1988), "Farmer Cooperatives as Behavior Toward Risk," *American Economic Review* **78**, 224--232.

Kletzer, K. and B. Wright (2000), "Sovereign Debt as Intertemporal Barter," *American Economic Review* **90**, 621--639.

Kocherlakota, N. (1996), "Implications of Efficient Risk Sharing without Commitment," *Review of Economic Studies* **63**(4), 595-609.

Ligon, E., J. Thomas and T. Worrall (2002), "Mutual Insurance and Limited Commitment: Theory and Evidence in Village Economies," *Review of Economic Studies* **69**, 209-44.

Selten, R., and A. Ockenfels (1998), "An Experimental Solidarity Game," *Journal of Economic Behavior and Organization* **34**, 517-519.

Selten, R., A. Sadrieh, and K. Abbink (1999), "Money Does Not Induce Risk Neutral Behavior, But Binary Lotteries Do Even Worse," *Theory and Decision* **46**, 211-49.

Townsend, R. (1994), "Risk and Insurance in Village India," *Econometrica* **62**(3), 539-91.

Townsend, R. (1995), "Consumption Insurance: An Evaluation of Risk-Bearing Systems in Low-Income Economies," *Journal of Economic Perspectives* **9**(3), 83-102.

Udry, C. (1994), "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria," *Review of Economic Studies* **61**(3), 495-526.

Udry, C. (1995), "Risk and Saving in Northern Nigeria," *American Economic Review* **85**(5), 1287 1300.