

CSAE WPS/2009-04

Contract Design in Insurance Groups

Tessa Bold*
Stefan Dercon
University of Oxford

Abstract

In many rural settings, informal mutual support networks have evolved into semiformal insurance groups, such as funeral societies. Using detailed panel data for six villages in Ethiopia, we can distinguish two types of contracts, in terms of whether payments are only made at the time of death or savings are accumulated by the group based on premiums paid ex-ante. We characterize these contracts as the coalition-proof equilibria of a symmetric and stationary risk-sharing game, and we show numerically that a contract with savings makes higher demands on enforceability, leading to less cohesive groups finding it in their interest to choose the contract without savings and that coalition-proofness is a necessary condition for the coexistence of both contract types. We show in the data that the type of contract chosen by groups is correlated with the level of trust and other enforcement improving factors. We also predict that among the observed contracts, those with group-based savings and ex-ante payments will attain higher welfare in terms of consumption smoothing than those observed using no group savings. Using panel data, and controlling for household fixed effects and time-varying village level fixed effects, we show that funeral groups are vehicles for risk-sharing and that contract type matters for performance in line with these predictions. The results appear robust to endogeneity of group formation and endogenous selection into contract types.

*This document is an output from research funding by the UK Department for International Development (DFID) as part of the iiG, a research programme to study how to improve institutions for pro-poor growth in Africa and South-Asia. The views expressed are not necessarily those of DFID. Corresponding author: tessa.bold@economics.ox.ac.uk.

1 Introduction

In the absence of complete markets in developing countries, a host of economic activities are carried out by informal groups. These range from labour-sharing and share cropping arrangements to credit and savings associations and even include local insurance clubs that offer protection against a variety of shocks. While these groups differ substantially in their degree of formalization, they all have in common that they cannot rely on external enforcement mechanisms to induce members to comply with their obligations (Ligon et al. (2002), Anderson et al. (2008)). Instead, the only incentive for cooperation is usually provided by the expected gain from continued participation in the arrangement in the future and groups often resort to complex contractual designs and implicit enforcement mechanisms to make this gain from cooperation – or conversely the cost of noncompliance – as large as possible. Although these informal economic arrangements are ubiquitous in the developing world, relatively little – beyond the anecdotal – is known about their institutional design especially in relation to its impact on members' welfare. Understanding the contractual design of informal economic groups, their ability to use implicit enforcement mechanisms and the benefits they provide to their members is particularly important in assessing the scope for scaling up and extending these groups into more formal institutions. In this paper, we examine a particular type of informal arrangement - insurance groups. More specifically, we model the contractual design of funeral insurance groups both theoretically and empirically based on a unique survey of 78 funeral insurance groups in rural Ethiopia. In particular, we focus on how contract design is used to increase the benefits from risk-sharing by alleviating enforcement issues.

Coate and Ravallion (1993) were the first authors to model informal insurance under the requirement that contracts have to be self-enforcing. They focus on a symmetric and stationary 2-player game that allows for sharing risk across agents in a given period, but not for smoothing across time. It is well known, that if the first-best equal-sharing contract is enforceable, then the ratio of marginal utilities is constant across agents and consumption is equalized in every period. However, in the absence of courts, the contract must be designed so that at every point in time, the benefit from sharing risk and remaining in the contract outweighs the benefit from defection and these additional constraints imply that the first-best contract is in general not sustainable.

There is a large literature that examines how the equilibrium of this risk-sharing game can be improved upon by extending the contract space. Both Kocherlakota (1996) and Ligon et al. (2002) have pointed out that conditioning on the previous history of the game can improve the benefits from risk-sharing that can be sustained

in a self-enforcing contract. The self-enforcing contract then becomes dynamic and takes on the form of quasi-credit contracts in which agents are induced to share more of their income in the current period by promises of at least partial repayment in the future. That is, an agent who shares his good fortune now will receive a higher expected surplus from the arrangement in the future. Gauthier et al. (1997) also allow for history dependence and extend the contract space to examine the role of ex-ante transfers in improving the benefits from risk-sharing in the self-enforcing contract. The authors demonstrate that ex-ante payments can be used to relax enforcement constraints and that an agent's constrained efficient ex-ante payment is decreasing in the expected surplus he expects from the arrangement. In other words, the more likely that an agent will want to deviate from the arrangement, the higher will be his ex-ante payment. They show that these payments can be explicitly linked to contract enforcement by interpreting them as a state-independent penalty that is decided upon by the agents at the beginning of the period, is paid only in the case of deviation and is enforceable by courts.

Finally, Gobert and Poitevin (1998) and Ligon et al. (2000) examine how the opportunity for storage affects the benefits from risk-sharing. Of course, under full commitment when the first-best contract is enforceable, such a contract is clearly superior because it allows for the smoothing of consumption both across time and across space. However, as shown in Ligon et al. (2000) the effect of savings is ambiguous when contracts have to be self-enforcing. On the one hand, the opportunity to save allows members of an insurance group to smooth aggregate shocks; on the other hand the ability to self-insure increases the benefits from deviation and this crowds out mutual insurance over time.

An important caveat of the above literature is that risk-sharing arrangements are deemed to be self-enforcing if they are stable with respect to individual deviations. As pointed out by Genicot and Ray (2003), under this assumption any risk-sharing arrangement would comprise the whole community and we can only explain the plethora of small groups found in practice by exogenously imposing a limit on group size. The authors resolve this paradox by pointing out that individuals deviations are not the only threat to the stability of an insurance group. If large groups can form and exploit the benefits from risk-sharing, then there is no reason to rule out that smaller groups can do the same and in turn destabilize larger groups by threatening to deviate from them. In other words, it seems intuitive to require insurance groups to be coalition-proof and this endogenizes the size of these groups Bernheim et al. (1987).

There are only a few papers that have applied contract theory to understand in detail the design of informal groups. Banerjee et al. (1994) study the design of 19th century credit cooperatives. Cabrales et al. (2003) use mechanism design to analyze

the functioning of a mutual fire insurance mechanism in Andorra. Anderson et al. (2008) examine how the organizational structure of Kenyan urban Roscas can be designed so as to address enforcement issues. Similar in spirit to Anderson et al., we compare insurance groups that appear to use different contracts in the context of semi-formal indigenous associations for the provision of funeral insurance in rural Ethiopia. We present a theoretical model that incorporates different institutional features, in terms of the possibility of ex-ante payments and savings. We focus on self-enforcing agreements, and model a new class of contract by extending Gauthier et al. (1997) and Ligon et al. (2000) that does not just allow for ex-ante payments and savings, but is also required to be coalition-proof. To our knowledge, this is the only paper that relates an analysis of institutional design of insurance to welfare outcomes at the level of individual participants in the insurance groups.

We derive a number of key predictions from the model. Under the first best risk-sharing arrangement, contracts allowing for savings are pareto superior to those without savings. Agents will also be indifferent between contracts allowing for ex-post or ex-ante payments. We would predict to only observe groups with savings, but both groups with ex-post and ex-ante payments would co-exist with indifferent performance in terms of consumption smoothing. However, these predictions change when contracts have to be self-enforcing and coalition proof. In particular, not all groups would engage in savings. As in Gauthier et al. (1997), the maximal total ex-ante transfer can be viewed as the total penalty that could be extracted in case of deviations from the contract. By implication, a group will be more cohesive, the higher this total transfer. But high transfers may not be feasible, simply because they cannot be paid for given agents' resources. If this is the case, then our simulations suggest that groups with no savings may be preferred to groups with savings and low ex-ante transfers. Those groups whose constrained efficient contract includes savings would nevertheless perform better than those finding no savings optimal, in terms of welfare via improved consumption smoothing.

The analysis leads to relevant empirical predictions. First, only if the extent of risk-sharing is limited by the threat of coalitional deviations, would we expect to observe heterogeneity in contracts. Thus the requirement of coalition-proofness is not just appealing because it allows for the endogeneity of group formation, but also because it is a necessary requirement for both group types to coexist in equilibrium. In that case, observed groups with savings would outperform groups without savings in terms of consumption smoothing. Furthermore, selection into contracts would be governed by the maximal transfer that is feasible in the group, whereby this maximal transfer is a measure of group cohesion. These transfers were modelled largely as if they were a financial transfer, but we show that it could either be a monetary transfer or a credible threat to a social sanction. As a result, it follows that, *ceteris paribus*,

we expect groups with savings to be groups with a higher ability to raise funds or higher social cohesion (such as in the form of trust or other social connectedness). If the ability to save more is also related to issues of trust (avoiding embezzlement), then this prediction is even stronger.

As our welfare predictions relate to the extent of consumption smoothing and risk-sharing that different contracts may attain, we can nest our paper directly into the existing literature. A vast literature has studied the strategies households use to smooth consumption in the face of large income fluctuations and unforeseeable expenditure needs (see Morduch (1999), and Deaton (1997) for surveys). A particular strand of this literature has focussed on testing whether rural communities manage to replicate Arrow-Debreu economies by employing these risk-coping strategies and achieve full risk-sharing in the absence of complete markets (Deaton (1992), Townsend (1994), Ravallion and Chaudhuri (1997), Udry (1994), Ligon et al. (2002)). Although there is ample evidence that available risk-coping mechanisms serve to provide at least partial insurance at the village level, the hypothesis of full insurance is overwhelmingly rejected (Ravallion and Chaudhuri (1997), Udry (1994), Ligon et al. (2002), Gertler and Gruber (2002)).

Finding partial insurance at the village level has led to a literature exploring risk-sharing at the level of groups or networks. For example, Grimard (1997) studies risk-sharing among ethnic groups in Cote d'Ivoire; Morduch (1991) tests insurance within castes in the ICRISAT data; and Dercon and Krishnan (2000a) find some evidence of full risk-sharing within nuclear households in Ethiopia. While risk-sharing may be incidental to the above groupings, more recently much work has been undertaken to map relevant insurance networks by asking households to identify insurance partners they rely on in times of need (Fafchamps and Lund (2003), De Weerd (2004)). From these mapping exercises it has emerged that individuals in rural communities are connected in insurance networks with a complex structure of direct and indirect links. De Weerd and Dercon (2006) have used such network data for rural Tanzania to test for full insurance through networks by applying the risk-sharing test to groups of directly linked households. Overall, the majority of the above papers continue to reject full insurance but find some evidence that subgroups within the village play a role in consumption smoothing.

Our contribution differs from those examined in the empirical literature so far in important respects, by its focus on different institutional designs in the context of indigenous associations to cope with risk. Firstly, in contrast to insurance networks which are the result of mainly informal bilateral relationships, funeral associations - known as Iddir in Ethiopia - are large and sophisticated institutions, which are formed with the specific purpose of offering insurance for funeral expenses and increasingly other forms of insurance and credit to cope with hardship. Iddir are

based on well-defined rules and obligations in the form of membership regulations and specific contributions and fines relevant to deviant behaviour (Dercon et al. (2006)). Iddir are much larger than informal insurance networks, but still considerably smaller than the community. Groups are also distinguished by the type of contract they offer. There are a large number of groups that demand regular (ex-ante) contributions and accumulate substantial savings and these exist alongside more informal associations that collect contributions only at the time of a funeral (i.e. ex-post) and hold few if any savings; in other words, groups differ in terms of their institutional design.

Secondly, our approach is more structural in nature in the sense that our risk-sharing test is based on a linear approximation of the analytical Euler equation that describes equilibrium consumption in our model of contract design under imperfect enforcement. As such we are estimating a model of incomplete insurance that nests full insurance, but is readily interpretable when the full insurance hypothesis is rejected. This approach is similar in spirit to Foster and Rosenzweig (2001), who estimate a model partial insurance sustained by altruism and Ligon et al. (2002), who estimate a full structural model of dynamic risk-sharing under limited commitment.

Our theoretical analysis makes predictions about the characteristics of the groups choosing different contracts, and has implications for the extent of risk-sharing obtained by different groups. In our econometric analysis, we combine household panel data from rural Ethiopia with the data on funeral insurance groups to estimate the determinants and welfare effects of contract selection. Our findings are directly in line with the theoretical predictions: we observe two types of contracts with groups with more trust and social cohesion more likely to be engaged in contracts with ex-ante payments and savings. Controlling for aggregate insurance group resources and aggregate village resources, we examine whether households are able to smooth consumption in the face of deaths of household members, the likely nature of the enforcement problems they face and the role of different contracts. We find that, at least on average, death shocks are well insured, but that this masks substantial heterogeneity across contract types. The contract including savings delivers full insurance, while those households who are insured in a contract with low ex-ante payments and insignificant savings find their consumption reduced significantly following a death shock. Moreover, the empirical results suggest that the benefits from risk-sharing are indeed limited by imperfect enforceability and more specifically that coalitional deviations pose an important threat to the stability of risk-sharing groups.

In Section 2, the survey data on funeral insurance groups are presented to motivate that analysis. Section 3 presents the theoretical model of contract structure, first focusing on the first best in Section 3.2. In Section 3.3, self-enforcing coalition

proof contracts are introduced. As closed form solutions are intractable, Section 3.4 presents simulations and Section 3.5 highlights the key empirical predictions. Section 4 presents the empirical model including the econometric specification for the risk-sharing tests and the results from the empirical estimation. Section 5 concludes.

2 Funeral Insurance in Rural Ethiopia

The data are based on a sample of funeral societies and their members in six villages in rural Ethiopia in 2003, Funeral Insurance Survey (FIS). In total, detailed data has been collected on 78 funeral societies using interviews of its key members, and the sample covered about half the number of those present in these villages. These data were matched to household panel data, collected as part of the Ethiopian Rural Household Survey, which has been conducted in these (and 9 other) villages since 1994. In each of the villages, the household survey is a random sample. We will restrict our analysis to those households we could match, covering 301 households with detailed information over time (6 rounds) and the funeral societies they belong to.

Iddir are associations that ensure a payout in cash and in kind at the time of a funeral for a deceased relative of a group member (for a detailed description of these groups see Dercon et al. (2006)). These groups do not consist of loose and rapidly changing associations of people. Instead, observed groups have been in existence for an average of 18 years. Iddir have a stable and clearly defined membership, usually based on written lists, and payments are made when members incur costs related to the funerals of a well-defined set of relatives in their household. The actual payout is conditional on the relationship of the member to the deceased: for example, the payment for a spouse is typically different from the payout for a child or for uncles and aunts. The insurance groups have written statutes, bylaws and records of contributions and payouts. The rules define membership procedures, payout schedules, contributions and also a set of fines and other measures for nonpayment of contributions.

Mortality rates are still very high in rural Ethiopia. In our sample, 45% of households report at least one death of a household member in the survey period, despite the fact the households are relatively nuclear (with an average household size of 7.5 in the first round). In each round of the data, on average 18% report a death shock. Virtually all households in our sample - 96.5% - are members of at least one Iddir. Just under 30% report belonging to one Iddir, 33% belong to two, and another 33% belong to three or more. Each village contains at least a dozen Iddir, but the number can rise as high as 30 or 40 Iddir in some communities. With an average size of 85 members (households), these groups are significantly smaller

than the village population. This plethora of relatively small groups is puzzling at first given the benefits of large numbers for risk pooling, providing one motivation for our interest.

A striking feature of these associations is their degree of formality and sophistication. There are clear rules, an elected committee, and regular meetings. Iddir are primarily designed for providing insurance in the case of funerals and payouts are mainly made in cash but are also composed of in-kind gifts and labour services. Table 1 offers more details. The average payout per group including in-kind benefits is 201 Ethiopian Birr, which is about 1.2 times average monthly household income. Average contributions are 1.51 Birr and groups retain substantial savings, which average 1921 Birr. The total insurance payouts made by Iddir in the past 12 months average 1116 Birr.

Table 1: Characteristics of funeral insurance groups (mean and standard deviation).

	all types	No Savings	Savings
Regular contribution to Iddir (ETB/month equivalent)	1.51 (1.41)	0.27 (0.05)	2.35 (1.38)
Contributions at funeral (ETB)	1.32 (1.73)	3.67 (1.94)	0.13 (0.08)
Payout in case of member's/spouse death	201.80 (177.56)	143.69 (127.15)	277.00 (205.34)
Number of members	85 (99.34)	60 (40.01)	117 (138.14)
Current funds of Iddir (ETB)	1921.16 (3968.58)	434.58 (771.39)	3488.08 (5473.88)
Amount Iddir paid out in past 12 months (ETB)	1116.0 (3749.58)	564.13 (742.88)	1952.70 (6047.63)
Number of observations	78	41	37

Note: Difference in means statistically significant at 5% level for all Iddir characteristics.

Source: Funeral Insurance Survey.

Ethiopian Birr: 1 USD= 9ETB.

In the villages studied, and elsewhere in Ethiopia, there are two clearly distinguishable types of Iddir, which differ in the extent to which the group uses savings based on contributions. Members of one type of group make (pre-determined) contributions when a funeral occurs and payments are made directly to the member who incurs the funeral expense at the time of the funeral. In contrast, the other type

of Iddir collect regular monthly payments from their members, which are stored in a communal Iddir fund. Members make a claim when a death occurs in their household and are then reimbursed directly out of the Iddir fund. In some cases, these members also make very small additional contributions when they attend the funeral. To put it another way, the second type of group collects contributions before the state of the world is known and saves them, to make payments after the state of the world is known. The first type of group collects contributions only after the state of the world is known, and they are immediately distributed to the members. Hence the two contract types can be distinguished by their use of savings.

As can be seen from Table 1, contract structure has an effect on the amount of contributions collected, payouts and accumulated savings. The average monthly contribution for the iddir with savings is 2.35 birr. In some cases, the first type of Iddir also collect a small regular contribution, but this is significantly smaller at just 0.27 Birr, usually meant for very specific spending related to the running of Iddir. In contrast, the average contribution for this type of Iddir at a funeral is 3.67 Birr, while the second type of Iddir make relatively small direct contributions of about 0.13 Birr at the time of the funeral. As a consequence of collecting regular contributions, the second type of Iddir retain substantial savings. With an average of 3488 Birr, current funds in these groups are almost nine times larger than in the other fund. They also provide a significantly higher payout at the time of a funeral.

Given the formalized nature of the insurance groups, it is also interesting to summarize the procedures groups have in place to enforce payments. Both types of groups waive contributions if a household is not able to pay its contributions in a particular round, but they levy steep fines and threaten exclusion if households refuse to pay and do not give a satisfactory reason for this. The membership of the Iddir is remarkably stable and few people leave either voluntarily or are forced to do so. Equally, reported incidences of existing group breaking up into smaller groups are few and far between. Interestingly, in the Iddir with savings, members who leave the group are entitled to a share of the groups savings which is calculated as total assets at the beginning of the period divided by the number of group members.

Clearly, these groups are rather different from each other. Each group is represented in each village surveyed, so it is of interest to explore why these different contract structures co-exist. One possible explanation for this difference in contracts is of course simply that members of the two group types are quite different. However, comparing the members of the two contract types, there is no difference in risk in terms of instances of illness and mortality.¹

In the next two sections, a theoretical model is developed that examines the features of each contract design with a particular focus on the combinations of savings

¹Details available from authors on request.

and ex-ante transfers and the fact that groups are of limited size. Combined with simulations, we provide predictions on selection into the contract, and its welfare consequences, both under first best conditions and constrained efficient self-enforcing and coalition-proof contracts.

3 The Model

We now formally model the different contract structures observed among informal insurance groups as the coalition-proof equilibria of a stationary and symmetric risk-sharing game with savings and ex-ante transfers. The game is set in an environment in which income realizations and asset holdings are observable and verifiable. We first describe the general set-up of the contract. In the next subsection, we analyze the first-best case in which all contracts are fully enforceable as a benchmark and then proceed to examine which contracts will emerge in an environment of limited commitment in which contracts must be self-enforcing.

3.1 The Environment

Suppose that there is a community of N households. Each period $t = 1, 2, \dots, \infty$, the state of nature s is drawn from a finite set of states $s \in \{1, 2, \dots, S\}$. Each period t is divided into three dates: t_0 , t_1 and t_2 . The state of nature is realized at date t_1 , and we will refer to t_0 as the ex-ante date and t_2 as the ex-post date.

All agents are infinitely lived and risk-averse. An agent's preferences over consumption $c_s^i(t)$ are represented by an increasing and strictly concave utility function $u(c_s^i(t))$. All agents discount the future by the common discount factor β . In each period t , agent i receives an income $y_s^i(t)$ which takes on two values: h with probability p and l with probability $1 - p$. Income is independently and identically distributed across agents and income realizations are observable.

Since all agents are risk-averse, they can profitably enter into a risk-sharing relationship by forming an insurance group of size $n \leq N$. In addition, individuals and insurance groups have access to a storage technology with interest rate r , where $r > 0$. Savings of individual i are denoted by $a_s^i(t)$. Total savings in a group of size n are denoted by $A(t) = \sum_{i=1}^n a_s^i(t)$ and we impose the condition that $A(t) \in [0, \bar{A}]$. Requiring $A(t)$ to be larger than zero implies a no-borrowing constraint. Putting a ceiling on $A(t)$ ensures that we are maximizing over a compact set which guarantees the existence and uniqueness of the maximum. For simplicity, we make no distinction between individual and group savings. That is, we assume savings are held at the group level upon joining an insurance group. Since individual income takes on only one of two values, aggregate income in a group of size n follows a binomial distribution with probability of k high incomes denoted by $p(k, n) =$

$\frac{n!}{k!(n-k)!p^k(1-p)^{n-k}}$. Hence the relevant information about the state of nature s in a risk-sharing group of size n is summarized by the number of agents with a high income realization k .

In each period, an insurance contract for a group of size n specifies the following structure for transfers, consumption and savings in the arrangement:

1. At t_0 each agents pays a positive ex-ante transfer $B(t)$, which is identical across individuals, into the group savings fund.
2. At date t_2 , ex-post transfers $\tau_k^i(t)$ and consumption and savings levels for each agent, $c_k^i(t)$ and $a_k^i(t+1)$ are chosen.

Consumption takes place at t_2 , when the state of nature is known and is defined as

$$c_k^i(t) = y^i(t) - B - \tau_k^i(t) + (1+r)a^i(t) - a^i(t+1)$$

To close the model, let $\sum_{i=1}^n \tau_k^i(t) = -nB$. In addition, we impose the condition that the maximum ex-ante transfer is contained in a compact and convex set, $B \in [0, \bar{B}]$ to ensure that the problem is well-behaved. As we will see, the restriction that $B \geq 0$ is without loss of generality in both the first-best and the second-best environment.

This formulation nests two types of contracts

1. Type I: $\bar{B} = 0, \bar{A} = 0 \Rightarrow$ no ex-ante transfers and no savings.
2. Type II: $\bar{B} \geq 0, \bar{A} > 0 \Rightarrow$ ex-ante transfer and savings.

The contract in which $\bar{A} = 0$ and $\bar{B} > 0$ is not feasible since the ex-ante transfers are stored in the group savings fund at t_0 . In other words, the ability to save is a necessary prerequisite for the use of ex-ante transfers.

3.2 The First-Best Contract

We first examine the benchmark case in which all prescribed transfers are enforceable and solve for the full commitment symmetric insurance contract in a group of size n . Since there are only two states of nature, we can simplify the law of motion for the state variable as

$$a^h(t+1) = (1+r)a^h(t) + h - B - \tau_k^h(t) - c_k^h(t) \quad (1)$$

$$a^l(t+1) = (1+r)a^l(t) + l + \frac{k}{n-k}(\tau_k^h(t) + B) - c_k^l(t) \quad (2)$$

$$(3)$$

where the superscripts h and l indicate the respective individual income realizations and we have used the fact that $k\tau_k^h(t) + (n-k)\tau_k^l(t) = -nB$.

We can now solve for current ex-ante transfers, ex-post transfers in period t and future savings levels in period $t+1$ by maximizing a weighted sum of utilities subject to the aggregate resource constraints. This amounts to maximizing the following dynamic programme, with each agent's Pareto weight set to $\frac{1}{n}$:

$$V^*(A(t), n) = \max_{A(t+1), c_k^h(t), c_k^l(t), B} \sum_{k=0}^n p(k, n) \left[\frac{k}{n} u(c_k^h(t)) + \frac{n-k}{n} u(c_k^l(t)) + \beta V^*(A(t+1), n) \right] \quad (4)$$

subject to the aggregate intertemporal resource constraints in period t and period $t+1$ and the boundary constraints on savings and the ex-ante payments.

$$\begin{aligned} p(k, n) \frac{1}{n} \eta_k(t) & : \quad kc_k^h(t) + (n-k)c_k^l(t) + ka^h(t+1) + (n-k)a^l(t+1) \leq \\ & \quad kh + (n-k)l + (1+r)A(t) \quad \forall k \\ p(k, n)p(j, n) \frac{1}{n} \eta_j(t+1) & : \quad jc_j^h(t+1) + (n-j)c_j^l(t+1) + ja^h(t+2) + (n-j)a^l(t+2) \leq \\ & \quad kh + (n-j)l + (1+r)A_k(t+1) \quad \forall k, j \\ \chi_1(t+1) & : \quad A(t+1) \geq 0 \\ \chi_2(t+1) & : \quad A(t+1) \leq \bar{A} \\ \nu_1(t) & : \quad B(t) \geq 0 \\ \nu_2(t) & : \quad B(t) \leq \bar{B}(t) \end{aligned}$$

where the Lagrange multipliers η in each state and time period are scaled for convenience by the probability of this state occurring.

The first order conditions for an interior solution for consumption and savings prescribe that

$$u'(c_k^h(t)) = u'(c_k^l(t)) \quad (5)$$

and

$$\begin{aligned} u'(c_k^h(t)) & = \beta(1+r) \sum_{j=0}^n p(j, n) \left[\frac{j}{n} u'(c_j^h(t+1)) + \frac{n-j}{n} u'(c_j^l(t+1)) \right] \\ u'(c_k^l(t)) & = \beta(1+r) \sum_{j=0}^n p(j, n) \left[\frac{j}{n} u'(c_r^h(t+1)) + \frac{n-j}{n} u'(c_r^l(t+1)) \right] \end{aligned}$$

This implies that consumption is equalized across all agents regardless of their endowment:

$$c_k^h(t) = c_k^l(t) = \frac{1}{n} \left[kh + (n-k)l + A(t)(1+r) - A(t+1) \right] \quad (6)$$

As a result

$$a^h(t)(1+r) - a^h(t+1) = a^l(t)(1+r) - a^l(t+1) = \frac{1}{n} [A(t)(1+r) - A(t+1)] \quad (7)$$

and

$$\tau_k^h(t) = \frac{n-k}{n}(h-l) - B. \quad (8)$$

Hence, we can normalize $a^h(t) = a^l(t) = \frac{1}{n}A(t)$ and $a^h(t+1) = a^l(t+1) = \frac{1}{n}A(t+1)$ for all $i = 1, \dots, n$.

It follows from this that the optimal contract is stationary in the sense that

$$c_k^h(t) = c_k^h(t+q) \quad \text{and} \quad c_k^l(t) = c_k^l(t+q) \quad (9)$$

and therefore

$$V(A(t)) = V(A(t+q))$$

whenever $A(t) = A(t+q)$. That is, the Pareto optimal contract is stationary conditional on the level of aggregate savings, where the state variable is governed by the following law of motion

$$A(t+1) = (1+r)A(t) + kh + (n-k)l - (kc_k^h(t) + (n-k)c_k^l(t)). \quad (10)$$

Moreover, because we solve for the equal-sharing contract, the expected lifetime utility from risk-sharing $V(A(t))$ is the same for all agents in each period. That is $V^i(A(t)) = V(A(t))$ for all $i = 1, \dots, n$.

To see how the optimal ex-post and ex-ante transfers are determined in the first-best contract, rewrite consumption as

$$c_k^h(t) = h - B - \tau_k^h + \frac{1}{n}((1+r)A(t) - A(t+1)) \quad (11)$$

and

$$c_k^l(t) = l + \frac{k}{n-k}(\tau_k^h + B) + \frac{1}{n}((1+r)A(t) - A(t+1)). \quad (12)$$

Since consumption only depends on net transfers $-B - \tau_k^h$ and $\frac{k}{n-k}(\tau_k^h + B)$, it follows immediately that the ex-ante transfer B is arbitrary. Similarly, the lower and upper bound on the ex-ante transfers are irrelevant, because even when binding, it is always possible to adjust the ex-post transfer so as to achieve the optimal allocation. In other words, under full commitment, ex-ante transfers and ex-post transfers are perfect substitutes in the optimal contract and groups that use ex-ante transfers have no advantage over groups that do not. Therefore, the envelope conditions imply that

$$\frac{\partial V(A(t), n)}{\partial B(t)} = 0. \quad (13)$$

and

$$\frac{\partial V(A(t), n)}{\partial \bar{B}(t)} = 0. \quad (14)$$

in the optimal first-best contract.

To examine the effect of savings on welfare, we note that $B(t)$ and $A(t)$ are both contained in compact and convex sets. It follows then from the assumptions on the utility function (see Stokey et al. (1989)) that

$$\frac{\partial V(A(t), n)}{\partial A(t)} = \frac{1}{n}(1+r) \sum_{k=1}^n p(k, n) \left[\frac{k}{n} u'(c_k^h(t)) + \frac{n-k}{n} u'(c_k^l(t)) \right] > 0. \quad (15)$$

which is the standard envelope condition for savings in the first-best contract. Similarly,

$$\frac{\partial V(A(t), n)}{\partial \bar{A}} = \chi_2(t+1) > 0. \quad (16)$$

That is, the the payoff from being in a risk-sharing group of size n is increasing in the group's capital holding under first best and switching from a type I contract to a contract that uses savings increases welfare.

Under full commitment, the standard symmetric risk-sharing contract is stationary and equal Pareto weights imply equal-sharing across agents in each state and period. More formally, we can now summarize the structure of the optimal contract as follows:

Summary 1

1. At t_0 each agents pays a positive ex-ante transfer $B(t)$, which is identical across individuals, into the group savings fund.
2. At date t_1 , the $nB(t)$ ex-ante transfers are distributed among the $n-k$ agents with a low income realization.
3. At date t_2 , the k agents with a high income realization make a (positive or negative) ex-post transfers τ_k^h to the $n-k$ agents with a low income realization. Each agent saves (or dissaves) $\frac{(1+r)A(t)-A(t+1)}{n}$.

To reiterate, if $\bar{B} \geq 0$ and $\bar{A} > 0$, then this describes a type II contract, in which ex-ante payments are potentially made at date t_0 and ex-post transfers, savings and consumption are decided at date t_2 . Equally, if $\bar{B} = 0$ and $\bar{A} = 0$, then this describes a type I contract. No payments are made at date t_0 , there is no opportunity for storage and ex-post transfers and consumption are decided at date t_2 .

Finally, the welfare properties of the first-best contract derived above are summarized in the following Proposition:

Proposition 1

1. Under first best, the expected utility from participating in a type II contract is independent of the size of \bar{B} . That is, the expected utility from a contract that uses savings and ex-post transfers is equal to the expected utility from a contract that uses savings and ex-ante transfer.
2. Under first best, a household is better off in a type II contract than in a type I contract and welfare is increasing in the savings level of an insurance group.

3.3 The Second-Best Contract for a given group size

We now consider an environment in which contracts are not legally enforceable. Instead risk-sharing contracts must be self-enforcing, which requires that at any point in time the benefit from complying with the contract must outweigh the gain from reneging. In particular, we require risk-sharing groups to be robust not only with respect to individual deviations but also with respect to deviations by subgroups, provided that these subgroups are themselves robust with respect to further deviations. That is, we are looking for the subset of coalition-proof equilibria of the set of sub-game perfect Nash equilibria of the game. This appears to be a natural restriction on the set of enforceable contracts. If a large group can form to reap the benefits from risk-sharing, then a smaller group can do the same and threaten to deviate jointly from the existing arrangement if the payoffs from doing so outweigh the costs.

To define the set of self-enforcing contracts, we follow Genicot and Ray (2003) and assess the stability of a risk-sharing group of given size n recursively by first examining the stability of groups of size $1, \dots, n - 1$ and then checking whether a risk-sharing contract exists for a group of given size n that is robust with respect to deviations by stable subgroups. In the next subsection, we go a step further by examining all the stable groups and self-enforcing contracts that emerge in a community of size N .

The timing of the game is as follows:

1. At date t_0 before the state of nature is realized, every agent makes an ex-ante transfer B that is kept in the group savings fund.
2. At date t_1 , the state of nature is revealed and the nB ex-ante transfers are transferred to the $n - k$ agents with a low income realization.
3. At date t_2 , all agents are asked to save $\frac{1}{n}((1 + r)A(t) - A(t + 1))$ and make ex-post transfers. There are four possibilities:
 - If $\frac{nB}{n-k} < \frac{k(h-l)}{n} + B$, then agents with a high income realization are called upon to make a transfer τ_k^h .

- If the enforcement constraints are satisfied, transfers are made and the group continues.
- If the enforcement constraints are violated, no further transfers are made and the group collapses.
- If $\frac{nB}{n-k} \geq \frac{k(h-l)}{n} + B$, then agents with a low income realization are called upon to make a transfer $\frac{k}{n-k}\tau_k^h$.
 - If the enforcement constraints are satisfied, transfers are made and the group continues.
 - If the enforcement constraints are violated, no further transfers are made and the group collapses.

We postulate that noncompliance with a risk-sharing arrangement has the following effects.

Assumption 1 *If a sub-group of size m reneges from an insurance arrangement of size n at t_2 , the capital stock $A(t)$ of the insurance group is distributed equally among the n agents. The nB ex-ante transfers are distributed among the $n - k$ agents with a low income realization. The deviating sub-group is forever excluded from the existing arrangement, but may continue risk-sharing amongst themselves after breaching the contract.*

Assumption 1 effectively imposes the condition that the group follows the first-best sharing rule outlined in Summary 1 up to the moment of deviation at date t_2 : assets are shared equally and are not expropriated in order to punish deviating agents; ex-ante transfers are distributed among the $n - k$ agents with a low income realization. Consequently, consumption in the case of deviation is written as

$$\hat{c}_k^h(t) = h - B(t) + (1+r)\frac{A(t)}{n} - \hat{a}_k^h(t+1) \quad (17)$$

$$\hat{c}_k^l(t) = l + \frac{k}{n-k}B(t) + (1+r)\frac{A(t)}{n} - \hat{a}_k^l(t+1) \quad (18)$$

where \hat{a} denotes next period's per capita assets in the deviating subgroup.

Even though the ex-ante transfers raised at t_0 are used to pay part of the state-contingent transfers at date t_2 , we should not think of them exclusively as insurance premia. In the second-best contract, they form a crucial part of the enforcement mechanism of the contract. The contract demands that the nB ex-ante transfers are distributed to those agents with a low income realization at t_1 and inspection of equations (17) and (18) shows the impact of this on the enforcement constraints. Under the first-best equal sharing rule, transfers will always flow from agents with a high income realization to agents with a low income realization and so the ex-ante

transfers reduce the payoff from deviating for an agent with a high income realization. However, the scope for using ex-ante transfers to alleviate the enforcement constraints is limited: Because the contract stipulates that agents with a low income realization receive the ex-ante transfers in case the group breaks down they have an incentive to bring about this collapse. As a result, agents with a high income realization may have incentives to overcompensate in some states of the world occasionally leading to those with a low income realization being allocated a higher share of resources in order to keep the group together. Nevertheless, redistributing ex-ante transfers as set out in Assumption 1 will make it easier to enforce an arrangement and we will prove this formally below.²

Our treatment of the ex-ante transfers is similar to Gauthier et al. (1997) who examine the role of ex-ante transfers in a bilateral self-enforcing risk-sharing contract that is non-stationary. The authors demonstrate that ex-ante payments are transferred at t_0 from the agent with the lower expected surplus from the arrangement in period t to the agent with a higher expected surplus. In other words, the more likely that an agent will want to deviate from the arrangement, the higher will be his ex-ante payment. Gauthier and Poitevin show that these payments can be explicitly linked to contract enforcement by interpreting them as a state-independent net penalty that is decided upon by the agents at the beginning of the period, is paid only in case of breach of contract and is fully enforceable. Such a contract is possible in an environment in which it is possible to verify that the contract has been breached, but the guilty party cannot be identified. In the context of an informal insurance arrangement, it is easy to think of examples for this. If an agent is excluded from the arrangement, is that because he deliberately did not pay his contributions, or because the rest of the group decided that his inability to contribute over a certain time period turned him into a liability to the group.

The main difference in our paper is that we are solving for the symmetric and stationary self-enforcing risk-sharing contract. As we will show formally below, in the second-best stationary and symmetric contract, all agents will gain the same expected utility from the risk-sharing arrangement in period t . Consequently, there is no role for using ex-ante payments in the way modelled by Gauthier and Poitevin because all agents are equal ex-ante. Instead, the contract modelled in this paper stipulates that ex-ante payments are stored in the group savings fund until t_1 when

²Alternative assumptions would be that all agents forfeit the ex-ante transfers or that ex-ante transfers are always used for smoothing in case the group collapses. In the latter case, the ex-ante transfers are distributed to those agents with a low income realization only if $\frac{nB}{n-k} \leq \frac{k(h-l)}{n} + B$. If $\frac{nB}{n-k} > \frac{k(h-l)}{n} + B$, then those with a low income realization receive $\frac{k(h-l)}{n} + B$ and those with a high income realization receive $B - \frac{(n-k)(h-l)}{n}$ and full smoothing is achieved by use of the ex-ante transfers alone. In either case, those with a low income realization would never experience a binding constraint. We believe that our assumption is more plausible because it explicitly takes account of the fact that ex-ante transfers are not a 'cost-less' enforcement mechanism.

they are distributed among the $n - k$ agents with a low income realization. This is equivalent to posting a bond that is kept by the group saving fund between period t_0 and t_1 and that is transferred from the high agents to the low agents in case the group collapses. In other words, the net penalty is paid by the agent who has a larger incentive to deviate ex-post.

At first glance, it may seem that this implies more stringent demands on the group's ability to enforce contracts than in the Gauthier and Poitevin environment. On the one hand, this may be a valid criticism. Firstly, we assume that the ex-ante payments can be safely stored in the group saving fund without the risk of embezzlement or other irregularities. Secondly, in line with Gauthier and Poitevin, we assume that ex-ante payments can be enforced, that it is possible to verify that the contract has been breached - but not by whom, and that income is both observable and verifiable. The latter implies that the disposal of the ex-ante transfers can be conditioned on agents' income realizations at date t_1 in case of breach of contract, which is equivalent to adjusting the sign of the net penalty - but not its size - once the state of the world is known. On the other hand, it could be argued that our model is less demanding in terms of the complexity of the second-best contract. Extending Gauthier and Poitevin's model from a bilateral to a multilateral environment would require specifying a unique ex-ante transfer for each of the n agents in the group that is conditional on agent i 's expected surplus from the arrangement. In contrast, our model has no agent specific features and only distinguishes agents by their current income realization, which substantially simplifies the contract structure. This seems to fit well with common practice in insurance groups in developing country contexts: groups for the most part treat their members identically.

Having discussed the one-period effect of deviation, we now turn to the long-run effect of being excluded from an insurance arrangement in case of deviation. To do this, we define stability of an insurance group recursively. Our definition of stability is related to the notion of coalition-proofness and stability derived in Bernheim et al. (1987), Bernheim and Ray (1989) and Genicot and Ray (2003). A group of size n is considered stable if there is no subgroup that can credibly deviate from the arrangement. A credible deviation requires that the subgroup in which insurance is continued, is better off after deviation and that the subgroup itself is stable with respect to further deviations.

We begin with an individual household and define the expected discounted lifetime value starting with per capita assets $A(t) = a(t)$ as

$$V(A(t), 1) = \sum_{k=0}^1 p(k, 1) \left[u(c_k^i(t)) + \beta V(A(t+1), 1) \right] \quad (19)$$

Since there are no possible deviations, stability of an individual is of course guaran-

teed.

This definition of a stable payoff is now extended to an arrangement of size n . Recursively having defined sets of stable payoffs for arrangements of size $m = 1, \dots, n-1$, we consider a coalition of size n . The expected utility for agent i of being in an arrangement of size n for beginning of period total assets $A(t) = \sum_{i=1}^n a^i(t)$ is

$$V^i(A(t), n) = \sum_{k=0}^n p(k, n) \left[u(c_k^i(t)) + \beta V^i(A(t+1), n) \right]. \quad (20)$$

Such an arrangement will be stable if no stable subgroup of size $m \leq n-1$ can credibly deviate from it.

Since the full commitment contract prescribes equal sharing any departure from this rule must arise because deviation payoffs differ across agents. From Assumption 1, deviation payoffs only differ according to an agent's current income realization and not according to their accumulated capital. Similarly, because we restrict ourselves to finding the coalition-proof stationary and symmetric risk-sharing contract, current consumption is not conditioned on the previous history of shocks (or more precisely the previous history of binding constraints) in order to sustain more risk-sharing as it would be in a dynamic contract. Stationarity implies that $V^i(A(t), n) = V^i(A(t+r), n)$ whenever $A(t) = A(t+r)$. Together with symmetry, it follows that $V^i(A(t), n) = V^j(A(t), n) = V(A(t), n)$. To summarize, under the restriction of stationarity, the constrained contract depends on the enforcement constraints in period t , but is not affected by the previous history of binding constraints.

Of course, the presence of savings implies that both period t and period $t+1$ enter the calculation of the optimal contract at time t because beginning of period assets in period $t+1$ are determined by the savings behaviour in period t . Equally, because savings in period t affect the deviation payoff in period $t+1$, enforcement constraints in both periods are potentially relevant for solving for the optimal contract. In other words, in our formulation the stationary contract depends both on current and future binding constraints. An alternative - and more restrictive - interpretation of stationarity would be that future constraints arising out of the savings behaviour of the group are to be ignored.

Because of the introduction of ex-ante transfers and savings, both high income and low income individuals may gain from deviating from the insurance arrangement and this has to be taken into account when defining the enforcement constraints. Any self-enforcing contract $\sigma = \{c, B, a\}$ must at a minimum satisfy the following constraints in period t

$$\begin{aligned} u(c_k^h(t)) + \beta V^*(A(t+1), n) &\geq u(\hat{c}_k^h(t)) + \beta V^*(\hat{A}(t+1), m) \quad \forall \text{ stable } m \leq k \\ u(c_k^l(t)) + \beta V^*(A(t+1), n) &\geq u(\hat{c}_k^l(t)) + \beta V^*(\hat{A}(t+1), m) \quad \forall \text{ stable } m \leq n-k \end{aligned}$$

However, this is not sufficient to guarantee stability. Let $\sigma = \{c, B, a\}$ be a contract that satisfies the above inequalities. It may still be possible that a deviation that involves both high and low income agents may destabilize the group because of the way savings are shared among members. Briefly, suppose the risk-sharing arrangement consists of five members with three high income realizations in the current period and ex-ante transfers low enough such that an agent with a low income realization receives a positive ex-post transfer at t_2 . If savings weren't shared in the group, we would never have to consider a joint deviation of high and low draws at this stage, because the low draws do not stand to gain anything from deviating. However, with shared savings, one of the low draws may find it in his interest to join the three high draws in a deviation, because this deviating subgroup potentially starts with a higher per capita capital stock in the following period than the existing group. The same argument can be made when the high draws stand to receive a positive transfer at date t_2 . Formally, for σ to be a self-enforcing contract, it must not be the case that

$$\begin{aligned} u(c_k^h(t)) + \beta V^*(A(t+1), n) &< u(\hat{c}_k^h(t)) + \beta V^*(\hat{A}(t+1), m) \\ &\quad \forall m = i + j \leq n - 1, \quad i \leq k, j \leq n - k \\ u(c_k^l(t)) + \beta V^*(A(t+1), n) &< u(\hat{c}_k^l(t)) + \beta V^*(\hat{A}(t+1), m) \\ &\quad \forall m = i + j \leq n - 1, \quad i \leq k, j \leq n - k \end{aligned}$$

Denote this set as Ω_t^* . It is then easy to see that any σ , which is not a member of Ω_t^* , cannot be a self-enforcing contract: If both inequalities above hold, then for a given B and a , satisfying the inequality would require increasing both $c_k^h(t)$ and $c_k^l(t)$, which is impossible without violating the aggregate resource constraint.

The above results in the following definition of stability

Definition 1 *An arrangement of size n is stable for a contract $\sigma = \{c, B, a\}$ if, and only if, for every possible beginning of period capital $A(t)$ and states k in*

period t and j in period $t + 1$, the following inequality constraints are satisfied

$$\begin{aligned}
(1) \quad & V^*(A(t), n) \geq V^*\left(\frac{m}{n}A(t), m\right) \quad \forall \text{ stable } m \\
(2a) \quad & u(c_k^h(t)) + \beta V^*(A(t+1), n) \geq \\
& u\left(h - B(t) + (1+r)\frac{A(t)}{n} - \hat{a}_k^h(t+1)\right) + \beta V^*(\hat{A}(t+1), m) \quad \forall \text{ stable } m \leq k \\
(2b) \quad & u(c_k^l(t)) + \beta V^*(A(t+1), n) \geq \\
& u\left(l + \frac{k}{n-k}B(t) + (1+r)\frac{A(t)}{n} - \hat{a}_k^l(t+1)\right) + \beta V^*(\hat{A}(t+1), m) \quad \forall \text{ stable } m \leq n-k. \\
(3a) \quad & u(\hat{c}_j^h(t+1)) + \beta V^*(A(t+2), n) \geq u(\hat{c}_j^h(t+1)) + \beta V^*(\hat{A}(t+2), m) \quad \forall \text{ stable } m \leq j \\
(3b) \quad & u(\hat{c}_j^l(t+1)) + \beta V^*(A(t+2), n) \geq u(\hat{c}_j^l(t+1)) + \beta V^*(\hat{A}(t+2), m) \quad \forall \text{ stable } m \leq n-j \\
(4) \quad & \sigma \in \Omega_t^*.
\end{aligned}$$

This definition states that an arrangement is stable if at all times the individual prefers making the contractual transfers and savings to renegeing on the contract and continuing insurance within a stable arrangement of size $m < n$. Constraint (1) is an ex-ante constraint, which must hold at date t_0 . Constraints (2a) – (2b) are the ex-post constraints in period t and must hold at date t_2 when consumption actually takes place. Constraints (3a) – (3b) are the ex-post constraints in period $t + 1$. The enforcement constraints show clearly how ex-ante transfers B can be used to alleviate the ex-post constraints by reducing the utility from deviating for an agent with a high income realization. However, this comes at a price since the ex-ante transfers increase the deviation payoff of those with a low income realization and this limits the scope for using ex-ante transfers to relax ex-post enforcement constraints. Condition (4) states that any self-enforcing contract σ has to be an element of Ω_t^* , which ensures that there are no joint deviations of low and high draws that destabilize the group. Finally, note that the definition of stability nests the subgame perfect contract, which is stable only with respect to individual deviations. This is simply the set of contracts described by setting $m = 1$ in Definition 1, which includes all the coalition-proof contracts.

Having defined the set of self-enforcing contracts, we can now compute the optimal constrained contract as the coalition-proof equilibrium of a stationary and symmetric risk-sharing game with savings and ex-ante transfers. This is done by solving the following dynamic maximization programme:

$$V^*(A(t), n) = \max_{\{a_k^h(t+1), a_k^l(t+1), c_k^h(t), c_k^l(t), B(t)\} \in \Omega_t^*} \sum_{k=0}^n p(k, n) \left[\frac{k}{n} u(c_k^h(t)) + \frac{n-k}{n} u(c_k^l(t)) + \beta V^*(A(t+1), n) \right] \quad (21)$$

subject to the following constraints

$$p(k, n) \frac{k}{n} \mu_k^h(t) : u(c_k^h(t)) + \beta V^*(A(t+1), n) \geq u(\hat{c}_k^h(t)) + \beta V^*(\hat{A}(t+1), m) \\ \forall \text{ stable } m \leq k \quad (22)$$

$$p(k, n) \frac{n-k}{n} \mu_k^l(t) : u(c_k^l(t)) + \beta V^*(A(t+1), n) \geq u(\hat{c}_k^l(t)) + \beta V^*(\hat{A}(t+1), m) \\ \forall \text{ stable } m \leq n-k \quad (23)$$

$$p(k, n) p(j, n) \frac{j}{n} \mu_j^h(t+1) : u(c_j^h(t+1)) + \beta V^*(A(t+2), n) \geq u(\hat{c}_j^h(t+1)) + \beta V^*(\hat{A}(t+1), m) \\ \forall \text{ stable } m \leq j \quad (24)$$

$$p(k, n) p(j, n) \frac{n-j}{n} \mu_j^l(t+1) : u(c_j^l(t+1)) + \beta V^*(A(t+2), n) \geq u(\hat{c}_j^l(t+1)) + \beta V^*(\hat{A}(t+2), m) \\ \forall \text{ stable } m \leq n-j \quad (25)$$

$$p(k, n) \frac{1}{n} \eta_k(t) : kc_k^h(t) + (n-k)c_k^l(t) + A(t+1) \leq (1+r)A(t) + kh + (n-k)l \quad (26)$$

$$p(k, n) p(j, n) \frac{1}{n} \eta_j(t+1) : jc_j^h(t+1) + (n-j)c_j^l(t+1) + A(t+2) \\ \leq (1+r)A(t+1) + jh + (n-j)l \quad (27)$$

$$\nu(t) : B(t) \leq \bar{B} \quad (28)$$

$$\chi(t+1) : A(t+1) \leq \bar{A} \quad (29)$$

where the Lagrange multipliers μ and η in each state and time period are scaled for convenience by the probability of this state occurring. For simplicity, we omit the ex-ante enforcement constraint and the lower bounds on $B(t)$ and $A(t+1)$.

The first-order condition for consumption in period t is given by

$$(1 + \mu_k^h(t))u'(c_k^h(t)) = \eta_k(t) = (1 + \mu_k^l(t))u'(c_k^l(t)) \quad (30)$$

To put it simply, the first-order condition states that if the enforcement constraint is binding for individuals with a high income draw their consumption is to be increased. In the opposite case, their consumption is to be decreased. In contrast to a dynamic contract where the ratio of Pareto weights is time-dependent, the ratio of Pareto weights here depends only on the current state of the system, which in turn depends on the income realization, beginning of period assets, the ex-ante payment, and the size of the maximum stable subgroup that threatens to deviate.

To find the first order condition with respect to $a_k^h(t+1)$ and $a_k^l(t+1)$, we use

the aggregate resource constraints in period t and $t + 1$. This gives

$$\begin{aligned}\eta_k(t) &= \beta(1+r) \sum_{j=0}^n p(j,n) \eta_j(t+1) \\ &- \beta(1+r) \sum_{j=0}^n p(j,n) \left[\mu_j^h(t+1) \frac{j}{n} u'(c_j^h(t+1)) + \mu_j^l(t+1) \frac{n-j}{n} u'(c_j^l(t+1)) \right]\end{aligned}\tag{31}$$

for both $a_k^h(t+1)$ and $a_k^l(t+1)$. Now we can substitute from the first order condition for consumption in period t and $t + 1$ to get the standard conditions for optimal saving in a limited commitment risk-sharing contract (see Ligon et al. (2000)).

$$\begin{aligned}(1 + \mu_h^k(t)) u'(c_k^h(t)) &= \beta(1+r) \sum_{j=0}^n p(j,n) \left[\frac{j}{n} u'(c_j^h(t+1)) + \frac{n-j}{n} u'(c_j^l(t+1)) \right] \\ &+ \beta(1+r) \sum_{j=0}^n p(j,n) \left[\mu_j^h(t+1) \frac{j}{n} (u'(c_j^h(t+1)) - u'(\hat{c}_j^h(t+1))) \right] \\ &+ \beta(1+r) \sum_{j=0}^n p(j,n) \left[\mu_j^l(t+1) \frac{n-j}{n} (u'(c_j^l(t+1)) - u'(\hat{c}_j^l(t+1))) \right].\end{aligned}\tag{32}$$

and

$$\begin{aligned}(1 + \mu_l^k(t)) u'(c_k^l(t)) &= \beta(1+r) \sum_{j=0}^n p(j,n) \left[\frac{j}{n} u'(c_j^h(t+1)) + \frac{n-j}{n} u'(c_j^l(t+1)) \right] \\ &+ \beta(1+r) \sum_{j=0}^n p(j,n) \left[\mu_j^h(t+1) \frac{j}{n} (u'(c_j^h(t+1)) - u'(\hat{c}_j^h(t+1))) \right] \\ &+ \beta(1+r) \sum_{j=0}^n p(j,n) \left[\mu_j^l(t+1) \frac{n-j}{n} (u'(c_j^l(t+1)) - u'(\hat{c}_j^l(t+1))) \right].\end{aligned}\tag{33}$$

The first thing to note is that the right-hand sides of (32) and (33) are identical and since we can write $u'(c_k^l)(1 + \mu_k^l) = u'(c_k^h)(1 + \mu_k^h)$, (32) and (33) have the same solution. Hence it follows that the first-order condition for capital only determines total savings in a symmetric and stationary risk-sharing contract and it is optimal to allocate high and low income individuals an equal share of capital. Therefore, if we implicitly assume that everyone had initial savings of $\frac{A}{n}(0)$ in the first period, there is no loss of generality in assuming that savings are held in a communal fund with each agent having a claim to savings of $\frac{A}{n}(t)$.

The terms $u'(\hat{c}_j^h)$ and $u'(\hat{c}_j^l)$ appear on the right-hand side of (32) and (33), be-

cause a deviating sub-group is allowed to keep its share of assets. If $\mu_j^h = \mu_j^l = 0$, the latter term drops out and we have the usual Euler equation for savings. The sign of the second term in (32) determines the impact of the enforcement constraints on allocative efficiency. The sign of the term depends on the sign of $u'(c_j^h) - u'(\hat{c}_j^h)$ and $u'(c_j^l) - u'(\hat{c}_j^l)$, which in turn depends on whether consumption prescribed by the risk-sharing contract is higher or lower than consumption in the case of deviation. If $u'(c_h^j) > u'(\hat{c}_h^j)$ for example, then additional saving relaxes the enforcement constraint because the marginal gain of additional savings is greater when remaining in the arrangement rather than deviating. The opposite applies when $u'(\hat{c}_h^j) > u'(c_h^j)$. Therefore, if the second term is positive, then current consumption will be lower relative to future consumption than under first-best. Alternatively, if it is negative, a binding enforcement constraint requires that current consumption is raised relative to the first-best outcome (see Ligon et al. (2000) for similar results in the dynamic bilateral contract with savings).

Given the ambiguous effect of increasing storage embodied in the first-order condition, it follows that the effect of storage on welfare is also ambiguous. The effect of increasing savings in a type II insurance group is given by the following envelope condition

$$\begin{aligned} \frac{\partial V(A(t), n)}{\partial A(t)} &= \frac{(1+r)}{n} \sum_{k=1}^n p(k, n) \left[\frac{k}{n} u'(c_k^h(t)) + \frac{n-k}{n} u'(c_k^l(t)) \right] \\ &+ \frac{(1+r)}{n} \sum_{k=0}^n p(k, n) \left[\mu_k^h(t) \frac{k}{n} (u'(c_k^h(t)) - u'(\hat{c}_k^h(t))) \right] \\ &+ \frac{(1+r)}{n} \sum_{k=0}^n p(k, n) \left[\mu_k^l(t+1) \frac{n-k}{n} (u'(c_k^l(t)) - u'(\hat{c}_k^l(t))) \right]. \end{aligned} \quad (34)$$

which is equivalent to the first order condition for capital in (32). The first term is positive and equal to the envelope condition in the first-best contract. The size of the second and third term are only positive if extra storage relaxes the enforcement constraints. If not, then larger savings have a negative effect on the benefits from risk-sharing. Therefore the overall effect of savings is ambiguous. The effect of switching from a type I contract into a contract that uses savings is given by the following envelope condition

$$\frac{\partial V(A(t), n)}{\partial \bar{A}} = \chi_{t+1} \quad (35)$$

This is unambiguously positive and equal to the first-best condition. Note however, that (35) only calculates the effect for a given n . This ignores the fact that increasing \bar{A} would also allow stable sub-groups of n – including individuals – to switch from type I to a Pareto superior contract that uses savings, thereby making the

enforcement constraints harder to satisfy. Hence, under imperfect enforcement, the impact of savings is generally ambiguous.

Finally, the first-order condition with respect to the ex-ante transfer is

$$\sum_{k=0}^n p(k, n) \frac{k}{n} [\mu_k^h(t) u'(\hat{c}_k^h(t)) - \mu_k^l(t) u'(\hat{c}_k^l(t))] - \nu(t) = 0. \quad (36)$$

To show that the optimal ex-ante payment is nonnegative, the first-order condition is evaluated when $B = 0$. In this case, the optimal contract requires that high draw households make a positive transfer to low draw households in each state. If a stable contract exists, then enforcement constraints cannot bind for both high and low draw households and in fact will bind only for the former. By assumption $\nu(t) = 0$ when $B = 0$. Therefore the first-order condition reduces to

$$\sum_{k=0}^n p(k, n) \frac{k}{n} [\mu_k^h(t) u'(\hat{c}_k^h(t))] > 0, \quad (37)$$

which cannot be optimal. Decreasing B will not lead to equality of the first-order condition, because $u'(\hat{c}_k^h) > 0$ for all values of B and μ_k^h is decreasing in B . Therefore the optimal ex-ante payment must be positive, which implies that the lower bound on the ex-ante transfer will never bind. Furthermore, the ex-ante transfers are transferred from the savings fund to those agents whose enforcement constraint is slack - which are of course the agents with a low income realization in our case. This is the case, because optimality in a symmetric and stationary risk-sharing arrangement stipulates that the low income draw agent will never be asked to make a transfer as long as ex-ante transfers are zero and so they will always have a net gain from remaining in the arrangement. Therefore a small redistribution of ex-ante transfers towards agents with a low income realization will not make their enforcement constraint binding but will relax the enforcement constraint of the agents with a high income realization. Of course, if the redistribution becomes too large then the ex-post enforcement constraint for agents with a low income realization may start to bind because it may be in their interest to make the group collapse. It follows from (36) and (37) that it is optimal to increase B until either the first best contract is achieved, low income households have a binding constraint in at least one state or the ex-ante payment reaches its maximum value. Since a household with a binding constraint receives a larger share of resources than a household with a slack constraint, there are states in which consumption of low income draw households exceeds that of high income draw households. This applies to both the sub-game perfect and the coalition-proof contract and ex-ante transfers are therefore generally bounded by the imperfect enforcement constraints.

This entire discussion may suggest that there is no role for the upper bound

\bar{B} in our model. This is certainly true under first best where the upper bound constraint is trivial. Similarly, it is possible to conceive of a second-best contract in which the ex-ante transfer is purely bounded by the enforcement constraints making the upper bound \bar{B} again irrelevant. However, as the ex-ante payment has to be credible introducing a bound is a natural assumption as credibility requires financial feasibility of the payment in all states of the world. Consequently, the maximal upper bound we could consider would be l , the income in the low state of the world. This would ensure that the bond can always be paid regardless of an agent's income realization and the level of savings accumulated by the group. It is therefore credible in every state of the world.³ However, this payment may still be larger than what could ever be credibly enforced in the community, so let us define

$$\bar{B} = \zeta l$$

where ζ , which lies between $0 \leq \zeta \leq 1$, expresses the extent to which a maximal payment can be enforced. For example, agents may be wealthy enough to raise a bond, but don't have enough trust that their ex-ante payments and savings can be kept safely by one of the group members. Defined in this way, it is clear that the maximal \bar{B} is not just related to the ability to make a financial transfer but also to the overall social cohesion of the group including the social pressure it can bring to bear on group members ex-ante to act in the interests of the group.⁴ Similarly, we could think of ζ as a parameter expressing a direct link between social cohesion and a group's ability to agree on credible formal rules. In either case, the more a group can make use of these mechanisms, the more robust it will be with respect to deviations. This leads us naturally to interpret \bar{B} , the maximum size of the bond, as a measure of overall group cohesion which is increasing in social cohesion and the maximum feasible payment which depends on the guaranteed liquid means available to each member:

Definition 2 *The maximum size of the bond \bar{B} is a measure of the cohesion of the insurance group. The cohesion of a group is an increasing function of the social cohesion of the group and the maximum financially feasible ex-ante transfer it can levy:*

$$\bar{B} = \text{Group Cohesion} = \zeta l = f(\text{social cohesion, maximum financially feasible ex-ante transfer}) \quad (38)$$

Introducing the upper bound on the ex-ante transfers not only leads to a defini-

³Of course, the fact that l is the the maximum ex-ante payment that can be levied is a consequence of the no-borrowing constraint.

⁴In fact, a number of authors (see Anderson et al. (2008)) argue that social pressure is used to assist in enforcing contracts.

tion of group cohesion, but also allows us to make precise welfare statements about insurance groups that use ex-ante transfers and savings. Writing down the envelope condition for \bar{B} , it follows that increasing the upper bound on the ex-ante transfer unequivocally increases the benefits from risk-sharing for a given group size

$$\frac{\partial V(A(t), n)}{\partial \bar{B}} = \nu_t \geq 0 \quad (39)$$

Quite simply, increasing \bar{B} increases the set from which B can be chosen and this has a positive effect on the expected utility from risk-sharing. In other words, welfare is increasing in group cohesion as defined in (38). However, it should again be noted that (39) ignores the fact that increasing \bar{B} also allows subgroups to use ex-ante transfers. This may increase their benefits from risk-sharing in a type II contract and therefore make it harder to satisfy the enforcement constraints.

The results for the constrained-efficient savings and ex-ante transfers in a self-enforcing contract for a given n are summarized in the following proposition.

Proposition 2

1. *The effect of savings on the benefits from risk-sharing conditional on group size are ambiguous in a self-enforcing type II risk-sharing contract. That is $\frac{\partial V(A(t), n)}{\partial A(t)} \begin{matrix} \leq \\ > \end{matrix} 0$.*
2. *The optimal ex-ante payment in the self-enforcing risk-sharing contract is strictly positive in a type II contract.*
3. *If the first-best contract is not self-enforcing, then either the optimal ex-ante payment is set to its maximum level or there is a state in which households with a low income realization have higher consumption than households with a high income realization.*
4. *The effect of switching from a type I contract to a type II contract is positive. That is $\frac{\partial V(A(t), n)}{\partial A(t)} \geq 0$.*
5. *In a self-enforcing type II risk-sharing contract, the benefits from risk-sharing conditional on group size are increasing in the upper bound on the ex-ante transfer, \bar{B} , $\frac{\partial V(A(t), n)}{\partial \bar{B}} > 0$. Therefore welfare is increasing in group cohesion.*

We can now compare the properties of the second-best contract for a given stable n to the first-best contract. It follows from part (1) of the proposition that the effect of increasing savings in a type II contract is no longer unambiguous under second-best. This is of course the case because an increase in savings both increases the benefits from remaining in the existing arrangement as well as raising the benefits

from deviation. In terms of switching from a type I contract into a type II contract, the predictions for the first-best and the second-best contract are the same, welfare is unambiguously increased (part (4) of the proposition). Note, however, that in the case of second-best this ignores the fact that stable sub-groups can now also resort to savings. Finally, part (5) of the proposition tells us that in contrast to the first-best environment, type II contracts are Pareto ranked by the level of group cohesion as measured by the maximum ex-ante transfer that can be raised. This is the case because the ex-ante transfers are an important part of the enforcement mechanisms of the contract which only come into play when the contract is required to be self-enforcing.

3.4 The second-best contract for all stable group sizes

Having analyzed the constrained-optimal risk-sharing contract for a given group size and shown analytically that type II contracts are Pareto superior to type I contracts, the fact that we observe both type I and type II contracts in the same community still presents something of a puzzle. Of course, the results we have derived are only valid for a given n , which may simply not be stable. Therefore, to understand whether our welfare predictions extend to a community of size $N \geq n$, we must examine all the stable groups and self-enforcing contracts that emerge in a community of size $N \geq n$.

There is no reason to presume that the analytical monotonicity results derived for a given n in Proposition 2 extend to the set of stable sizes in the community. This is the case because the introduction of savings increases the benefits from risk-sharing by allowing for more consumption smoothing over time, but they also increase the benefits from deviation by raising the possibility of self-insurance. In other words, while the introduction of savings allows a given group to switch to a Pareto superior contract, this is also true for all stable subgroups – including individuals – and this potentially destabilizes larger groups. If the latter effect outweighs the former, then the benefits from risk-sharing may be smaller with savings than without because the benefits from risk-sharing are also increasing in the size of the insurance pool. Similarly, the introduction of ex-ante transfers has two conflicting consequences which potentially work in opposite directions. It increases risk-sharing for a given group size and this result of course extends to stable subgroups of larger groups. Because these subgroups derive a larger gain from risk-sharing with ex-ante transfers in a type II contract larger groups that would have been stable in the absence of ex-ante transfers may now be destabilized. Because of this analytical ambiguity, we turn to simulations to investigate how the possibility of ex-ante transfers and savings affects the benefits from risk-sharing over all stable group sizes in a community. We do this both for the sub-game perfect contract and the coalition-proof contract.

As we will see, solving for equilibria of the risk-sharing game that are both sub-game perfect as well as coalition-proof is not just conceptually superior. Unless we require insurance groups to be stable with respect to coalitional deviations, we cannot understand why type I and type II contracts coexist.

To conduct the simulations, we parameterize the model in Section 3.3 and compute the optimal stable contract for a variety of preferences. We calculate the expected utility from risk-sharing over the set of stable sizes and show how different maximum B 's may lead to the selection of different contracts. To reiterate, the different contracts available are:

1. Type I: $\bar{B} = 0, \bar{A} = 0 \Rightarrow$ no ex-ante transfers and no savings.
2. Type II: $\bar{B} \geq 0, \bar{A} > 0 \Rightarrow$ ex-ante transfer and savings.

We maximize the dynamic programme in (21) and solve for ex-ante transfers, ex-post transfers and savings in a risk-sharing arrangement of n individuals in a community of $N = 10$ with individual utility function

$$u(c) = \frac{1}{1-\rho} c^{1-\rho}, \quad (40)$$

where ρ is the coefficient of relative risk aversion. In each period, individual income is l with probability $(1-p)$ and h with probability p , with $h > l$. The discount rate is β and $r = \frac{1}{\beta} - 1$ is the interest rate. For type II contracts, we calculate the optimal contract on a capital grid for individual capital that ranges from $[0, \bar{a}]$ and $\bar{a} = 1$. Furthermore, we compute the upper bound on the ex-ante transfers as laid out in the definition of group cohesion $\bar{B} = \zeta l$ and let $\zeta = \{0, 0.25, 0.5, 0.75, 1\}$ vary from zero to 1. We compare the payoffs from these contracts to a type I contract without savings in which $\bar{A} = 0$. Since ex-ante transfers are transferred into the group savings fund between date t_0 and date t_1 , this of course implies that in the absence of savings $\zeta l = \bar{B} = 0$ in a type I contract. We compute contracts for twenty different combinations of environmental and preference parameters. Specifically, we vary the discount factor and the coefficient of risk aversion and calculate the model for the following parameterizations: $\beta = \{0.81, 0.82, 0.83, 0.84, 0.85, 0.86, 0.87, 0.88, 0.89, 0.9\}$, fixing $\rho = 1.2$, and letting $\rho = \{1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.7, 1.8, 1.9, 2.0\}$, keeping β constant at 0.85. The probability of a high income realization is always set to $p = 0.5$ and $h = 0.75$ and $l = 0.25$.

To measure the payoffs from different contract types, we calculate the relative gain from being in a stable arrangement of size n compared to the maximal per capita payoff that can be achieved, i.e. the first-best payoff $\tilde{V}(a, N)$. To make the notation more transparent, we introduce two more arguments ζ and \bar{a} into the expected utility function to write $V^*(a, n, \zeta, \bar{a})$. Then the payoff from a type I

contract in a group of size n is given by

$$\frac{V^*(a, n, 0, 0) - V^*(a, 1, 0, 0)}{\widetilde{V}(a, N, 1, 1) - V^*(a, 1, 0, 0)} \times 100. \quad (41)$$

The payoff from a type II contract is given by

$$\frac{V^*(a, n, \zeta, 1) - V^*(a, 1, 0, 0)}{\widetilde{V}(a, N, 1, 1) - V^*(a, 1, 0, 0)} \times 100. \quad (42)$$

This gives a measure of the return to insurance in percentage terms (see Genicot and Ray (2003)). Since the largest stable group will typically be smaller than the community and not all groups within a community will be stable, the question arises which group sizes will be observed and how the relevant payoff of being in a particular group can be calculated. Following Genicot and Ray (2003), the population is partitioned into stable groups in a way that maximizes the expected utility of an agent under the assumption that the probability of being in a group is proportional to the size of the group.

Figure 1 plots the largest stable group and Figure 2 plots the expected stable gain for the equilibrium configuration of group sizes for different contract types as ζ varies from zero to one. In each figure, the top two panels plot the results for the sub-game perfect contract and the bottom two panels plot the results for the coalition-proof contract.

We first focus on the sub-game perfect contract. It is notable from Figure 1 that a type I contract is always stable for the whole community. This is the case, because the set of group sizes that is stable with respect to individual deviations is a connected and unbounded set in a stationary risk-sharing contract without savings (Genicot and Ray (2003)). In other words, if we can find a group of size n that is stable, then all $n' > n$ will be stable. As a result, as long as N is large enough, the type I contract is always stable for the whole community. The reason that the type I contract is a connected set is that a group of size n can replicate the autarky outcome and add an ϵ -benefit to it, so enforcement constraints for individual deviations can always be satisfied. This is not the case for the type II contracts. In fact, without the use of ex-ante transfers, type II contracts are not even stable with respect to individual deviations, which is easily seen from the fact that the largest stable group size is one for $\zeta < 0.25$. This is the case because the type II contract requires that savings are shared equally, which implies that the contract cannot replicate the autarky outcome in order to satisfy the enforcement constraints. Once we introduce ex-ante payments, this limitation can be overcome to some extent. However, the set of group sizes that is stable with respect to individual deviations will not be unbounded because as n increases the maximum ex-ante payment that

can be raised is limited by the requirement that those with a low income realization do not have an incentive to deviate.

Moving on to the coalition-proof contract, it is notable from Figure 1 that the largest stable size for both type I and type II contracts is significantly smaller than the community for most parameters and that there appears to be no monotone relationship between the ability of a group to use ex-ante transfers and/or savings and the maximal stable size of the insurance arrangement. In fact, for type II groups the largest stable group size is obtained for $\zeta = 0.5$ in the case of $\beta = 0.88$. Conversely, both for large ζ and small ζ , only smaller groups are stable. This non-monotonicity is a direct consequence of requiring the optimal contract to be coalition-proof. Because the benefit from risk-sharing is increasing in ζ for a given group size, smaller groups have a more destabilizing effect when ζ is large. In contrast, for small but non-zero ζ , there are a few environments in which large groups are stable because the benefits from risk-sharing for small groups are not large enough to destabilize the community. For similar reasons, there are a number of parameterizations in which type I groups are bigger than type II groups. Finally, it is notable that groups which use type II contracts and are required to be robust with respect to coalitional deviations are not much smaller than groups that are only required to be robust with respect to individual deviations. This is the case, because in the case of type II contracts, the impact of the ex-ante transfers on the enforcement constraints, which is independent of the long-term payoff from deviations, drives most of the results.

Since the largest stable group is not increased significantly by the use of ex-ante transfers and savings, any overall welfare effect (if it exists at all) has to work through an increase in the benefits from risk-sharing for a given group size and/or an increase in the number of members of the set of stable sizes. Figure 2 explores these effects by plotting the expected stable gain for type I and II contracts against different values of the discount factor and the coefficient of risk-aversion. In each figure, the top two panels present the results for the sub-game perfect contract and the bottom two panels represent the effects for the coalition-proof contract.

Let us first focus on the comparison of type II contracts for different values of ζ and l . The most important thing to note from Figure 2 is that the higher the value of ζ the higher the expected utility from being in a type II contract and that for the most part the lines for different values of ζ do not cross (there are two exceptions). Therefore the benefit from risk-sharing in a type II group is increasing in the level of group cohesion for a given environment, i.e. conditional on a , β , ρ and l , and the analytical result proved in part (5.) of Proposition 2 for a given group size extends to the set of all stable sizes. This implies that in the second-best environment, type II contracts are Pareto ranked for a given l according to the level of ζl and this




Figure 1: Largest stable group size in Type I and Type II contracts.


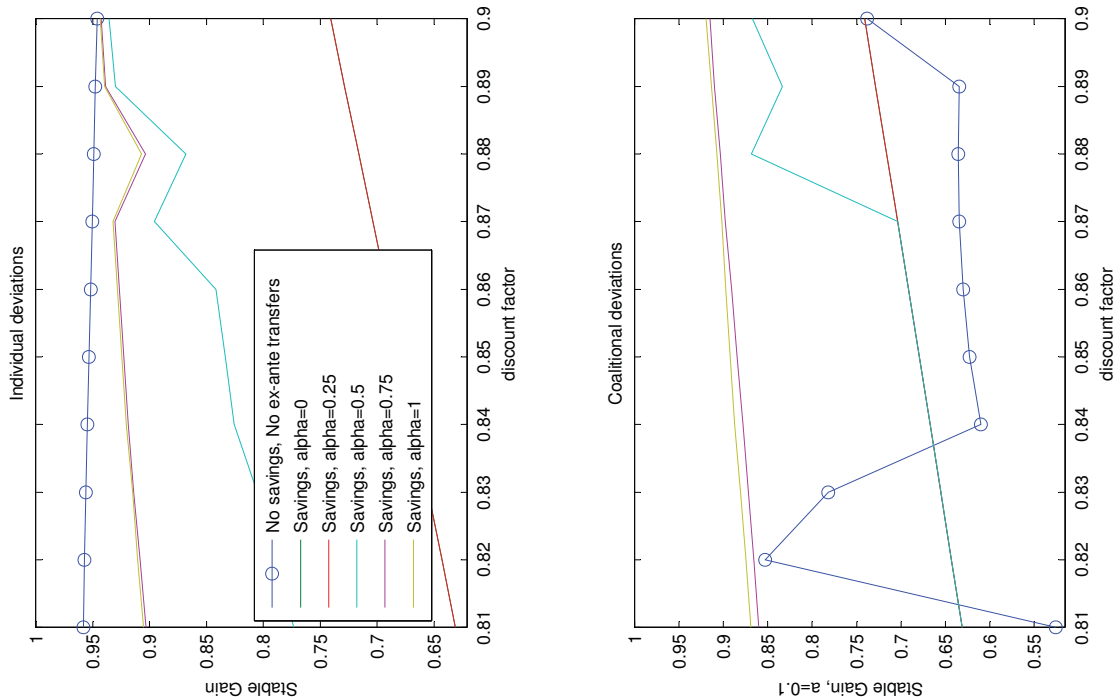


Figure 2: Expected gain from risk-sharing in Type I and Type II contracts.



stands in clear contrast to the first-best environment in which type II contracts that use ex-ante transfers are equivalent in terms of welfare to type II contracts which use ex-post transfers.

Even so, for the sub-game perfect contract, the expected benefit from being in a type I contract always outweighs the benefit from being in a type II contract regardless of the level of group cohesion. This is the case because type I contracts are stable at the level of the community whereas type II contracts are significantly smaller. It also implies that if the benefits from risk-sharing are only limited by contracts being stable with respect to individual deviations, then groups would always select into type I contracts. In other words, the predictions under second-best and sub-game perfection are precisely the opposite to the predictions in a first-best world.

Once we require both type I and type II contracts to be coalition-proof, it becomes optimal to switch from a type I contract to a contract that allows for the possibility of savings for sufficiently high levels of ζl . Certainly for low levels of ζl , there are a number of parameterizations in which the benefits from a type I contract outweigh the benefits from a type II contract. This is precisely the nature of second-best solutions. Even though a contract that uses savings is always Pareto superior under first best, the requirement that the contract must be coalition-proof provides strong incentives to select into a contract that would be inferior under full enforceability. Only for large enough levels of group cohesion is it always optimal to switch into a type II contract. Since welfare in a type II contract is increasing in group cohesion, this implies that those choosing type II contracts over type I contracts are always better off in a given environment. Overall then, requiring coalition-proofness is not just conceptually superior. We can only explain why type I and type II contracts coexist – and groups are smaller than the community – when contracts are required to be coalition-proof.

These results have strong implications for the welfare ranking and selection of different contract types that are coalition-proof over the set of all stable group sizes, which are summarized in Conjecture 3:

Conjecture 3 : Welfare Effects of Contract Selection

1. *For low levels of group cohesion, type I contracts sometimes dominate type II contracts.*
2. *For high levels of group cohesion, type II contracts always dominate type I contracts.*
3. *Welfare is increasing in group cohesion so that those for whom a type II contract is observed in a given environment will have higher welfare than those in*

a type I contract.

Comparing these results to part (4) and (5) Proposition 2, it follows that the welfare predictions for \bar{B} extend to the set of stable sizes, whereas an increase in \bar{A} does not lead to a welfare improvement unless \bar{B} is sufficiently high. In other words, while savings are a necessary precondition for raising a bond in practical terms, raising a bond is a necessary precondition for making savings feasible in the presence of coalitional deviations.⁵ By implication, Conjecture 3 predicts that we observe type I and type II contracts side-by-side, but that for a given environment, those able to select into type II contracts will be better off than those in type I contracts.

The relationship between levels of group cohesion and contract selection in the coalition-proof contract is summarized in Conjecture 4.

Conjecture 4 : Group/Social Cohesion and Contract Selection

- 1. Among groups with low to intermediate levels of group cohesion both type I and type II contracts are observed.*
- 2. Groups with high group cohesion always select into type II contracts.*
- 3. Selection into type II contracts is a positively related to group cohesion.*
- 4. For a given l , selection into type II contracts is positively related to social cohesion.*

3.5 The empirical model

The theoretical and numerical results show that contracts are welfare ranked in terms of group cohesion so that those for whom a type II contract is observed in a given environment will have higher welfare than those in a type I contract. Of course, in practise we do not observe utility, but as is well known if agents are risk-averse then – other things being equal – their welfare is increasing in the extent of consumption-smoothing they achieve. In this section, we will show that our welfare predictions can be related to the level of risk-sharing that takes place in insurance groups (and therefore to observable measures of consumption and income). In particular, we show that there is a one-to-one empirical relationship between the

⁵It is perhaps not surprising that our predictions for a given n are a reasonable guide for the effects of increasing \bar{B} over the set of stable sizes, but not for \bar{A} . After all, the results for \bar{B} ignore the effect of better enforcement mechanisms for stable subgroups given that the contract uses savings, whereas the results for \bar{A} ignore the fact that stable subgroups can now switch into a contract that is Pareto superior.

extent of consumption smoothing and the welfare properties of a particular contract in the sense that higher welfare is associated with more smoothing.

As a first step, we show that for the class of CRRA utility functions we can derive a relationship between own consumption and aggregate consumption of the insurance group that is akin to the standard risk-sharing tests derived in the literature but explicitly takes account of the second-best nature of the risk-sharing contract. As above, we assume a CRRA utility function for household i 's consumption in period t and write

$$u(c_t^i) = (1 - \rho)^{-1} (c_t^i)^{1-\rho}. \quad (43)$$

If we apply this to the first-order condition in (30) for any two households i and j , we find

$$\frac{c_t^j}{c_t^i} = \left(\frac{1 + \mu_t^j}{1 + \mu_t^i} \right)^{\frac{1}{\rho}} \quad (44)$$

Now taking logs and adding up the $n - 1$ first order conditions between agent i and the rest of the group, we can write consumption of agent i as a linearly additive function of average group consumption and terms relating to the second-best enforcement constraints

$$\ln c_t^i = \frac{1}{n-1} \sum_{j \neq i} \ln c_t^j + \frac{1}{\rho} \ln(1 + \mu_t^i) - \frac{1}{\rho(n-1)} \sum_{j \neq i} (1 + \mu_t^j) \quad (45)$$

and log-linearizing, we can simplify further to write

$$\ln c_t^i = C_t^g + \frac{1}{\rho} \mu_t^i - \frac{1}{\rho} \mu_t^g \quad (46)$$

where $C_t^g = \frac{1}{n-1} \sum_{j=1}^{n-1} \ln c_t^j$ is average log consumption in the risk-sharing group and $\mu_t^g = \frac{1}{n-1} \sum_{j=1}^{n-1} \mu_t^j$. If we observed the Lagrange multipliers on the enforcement constraints, then we could estimate this equation directly. While this is not possible, the theoretical model tells us that the severity with which the agent is constrained relative to all other agents in the risk-sharing arrangement, $\frac{1}{\rho}(\mu_t^i - \mu_t^g)$, is a function of – among other things – the agent's own income realization and the group's contract structure.⁶ Because we are restricting attention to a stationary and symmetric risk-sharing contract with only two possible individual income realizations, we introduce

⁶Strictly speaking, the theoretical model implies that the Lagrange multiplier is a function of an agent's own income realization, the aggregate income realization, the size of the group, and its savings levels and ability to use ex-ante transfers. That is

$$\mu_t^i = f(y_t^i, Y_t^g, \bar{B}, n, \frac{1}{n} A_t) \quad (47)$$

where $Y_t^g = \frac{1}{n-1} \sum_{j \neq i} y_t^j$. Since we are interested in deriving an empirical specification, we restrict attention to y_t^i and \bar{B} – partly because of data availability and partly because of the high degree of collinearity between aggregate group consumption and aggregate savings and shocks.

a dummy D_t^i that is equal to one if an agent has a low income realization and zero otherwise. Then, we can write

$$\frac{1}{\rho}\mu_t^i - \mu_t^g = f(D_t^i, \bar{B}) \quad (48)$$

and

$$\ln c_t^i = C_t^g + f^i(D_t^i, \bar{B}) \quad (49)$$

Since f is an unknown function, we expand it in order to transform (49) into a linear estimable equation. Define $\hat{\mathbf{x}} = \{D_t^i - D^{i*}, \bar{B} - \bar{B}^*\}$ and $\mathbf{x}^* = \{0, 0\}$. This gives the second order Taylor expansion around \mathbf{x}^* :

$$\rho^{-1}(\mu_t^i - \mu_t^g) = f(\mathbf{x}) = f(\mathbf{x}^*) + Df(\mathbf{x}^*)(\hat{\mathbf{x}}) + \frac{1}{2!}D^2f^i(\mathbf{x}^*)(\hat{\mathbf{x}}, \hat{\mathbf{x}}) + R_2(x, x^*); \quad (50)$$

Based on this, we can write the following linear estimable equation in first differences:

$$\Delta \ln c_t^i = \alpha_0 \Delta C_t + \alpha_1 \Delta \hat{D}_t^i + \alpha_2 \Delta(\hat{D}_t^i \times \bar{B}) + \varepsilon_t^i \quad (51)$$

where α_1 and α_2 are overspecification tests. Finding that either of them is significantly different from zero indicates that full insurance is rejected – for the moment abstracting from preference shifts and measurement error.

However, our theoretical model in conjunction with the simulations presented above can do more. It does not only predict that full risk-sharing will be rejected, it also allows us to give some interpretation to the sign and magnitude of the coefficients on the overidentifying regressors. In particular, we now show that the welfare ranking of contracts we demonstrated above has an empirical analogue in the extent of consumption smoothing that will be observed. Intuitively, we would expect that the superiority of a type II contract in welfare terms translates into $\alpha_1 < 0$ and $\alpha_2 > 0$. That is, while the overall impact of a shock is negative, the fall in consumption is less severe in a type II contract.







This is by no means self-evident, however, because the regression we have formulated will only measure the benefits from smoothing shocks across states, not the benefits that arise from being able to smooth shocks intertemporally, which is only possible in a type II contract. To see this, suppose that type I groups always manage to sustain full risk-sharing while type II groups do not. Then we may well find that $\alpha_1 = 0$ and $\alpha_2 < 0$ even though a type II contract allows for more intertemporal smoothing and therefore gives higher welfare than a type I contract. Similarly, suppose that all group types sustain nearly full insurance and the benefit from using ex-ante transfers and savings arises mainly from larger group sizes. In this case, the risk-sharing test would not pick up the effect of contract structure because the

sensitivity of consumption with respect to income conditional on group resources would be zero regardless of the magnitude of \bar{B} even though expected utility clearly increases with group cohesion as larger group sizes become stable. In other words, the risk-sharing test reflects the superiority of combining ex-ante transfers and savings only if there is some departure from first-best in the type I contract and the extent of cross-sectional smoothing increases as we select into a type II contract.

We will now show that this is indeed the case and that our predictions on the Pareto ranking of contracts can be directly related to the degree of cross-sectional consumption smoothing achieved in different contract types. To this end, we use the computed model solutions from both the individual deviations and the coalitional deviations model from above to estimate (51). For each parameterization (discount factor and level of risk aversion), the community is allocated the contract type that gives the highest payoff among those that are feasible given the level of group cohesion, ζl . We then randomly draw an income realization for each individual in the community of size $N = 10$ for $T = 2$ periods and based on the computed stable insurance contracts calculate consumption for every parameterization. This generates a sample of 1100 observations (five levels of $\zeta \times 10$ agents \times 11 different values of initial capital \times two time periods) for each environment. We then estimate (51) on this sample and record the coefficients and p-values. Finally, the whole process is repeated 100 times.

Figure 3 reports the average coefficients from the Monte-Carlo simulation for the sub-game perfect and the coalition-proof contract. The top panels report the average coefficient on group consumption, α_0 , for all environments, the middle panels report the average coefficient on the shock, α_1 , and the bottom panels report the average coefficient on the interaction, α_2 , of the shock variable and the contract type $\bar{B} = \zeta l$, where \bar{B} has been normalized to lie between zero and one.

Let us first compare the individual deviations model to the coalitional deviations model in the top and middle panels. It is easily seen from the three sets of panels that full risk-sharing will always be achieved in the individual deviations contract. In contrast, the coefficient on group consumption is significantly less than one and the impact of the shock is negative in the contract with coalitional deviations. Hence, this provides another test of whether coalitional deviations present a threat to the stability of risk-sharing arrangements. Of course, the main test of whether coalitional deviations present a threat to stability is provided by contract selection. Only if risk-sharing is limited by coalitional deviations will we observe both contract types (hence the interaction of contract type and shock variable is not plotted for the individual deviations contract). However, even aside from the issue of contract selection, the threat of coalitional deviations has implications for the transfer structure in the constrained-efficient contract.



The bottom panel of Figure 3 shows that the welfare ranking of the different contract types indeed translates into a higher degree of consumption smoothing achieved in a type II contract. For every parameterization, the coefficient on the death shock is negative and significant, while the interaction on the death shock variable and the level of group cohesion is always positive and significant. Taken together, this implies that our theoretical and numerical results can be translated into the following two related hypotheses that will be tested in the next section.

1. H_0 : Risk-sharing is perfectly enforced $\Rightarrow \alpha_0 = 1$ and all other coefficients are zero. H_1 : Risk-sharing is imperfectly enforced \Rightarrow any of the coefficients $\{\alpha_i\}_{i=1}^2$, are significant.
2. H_0 : Contract I outperforms contract II $\Rightarrow \alpha_2 \leq 0$. H_1 : Contract II outperforms contract I $\Rightarrow \alpha_2 > 0$.

4 Empirical Prediction and Estimation

Our theoretical and numerical results derived in the previous section lead to a number of predictions for contract selection and welfare in the different contract types which we now test using data on funeral insurance groups and their members:

Conjecture 5

Contract Selection

1. *Type I and type II contracts coexist only if risk-sharing is limited by coalitional deviations.*
2. *For non-zero levels of group cohesion, it is never optimal to select into a type II contract that does not use ex-ante transfers.*
3. *There is a positive correlation between social cohesion and selection into type II contracts for a given environment and preferences.*

Consumption Smoothing Benefits

1. *Consumption smoothing in both types of contracts is restricted by coalitional deviations.*
2. *The consumption smoothing benefits in a type II contract are larger than the consumption smoothing benefits in a type I contract for a given environment and preferences.*

4.1 The Determinants of Contract Selection

Following our analysis of the different contract types, it is evident that the funeral insurance groups we have described in Section 2 can be categorized as type I contracts and type II contracts with non-zero levels of ex-ante transfers. The FIS contains 38 groups which use ex-post transfers and do not save and 41 groups that use ex-ante transfers and accumulate substantial savings. As predicted, we do not observe any groups that use ex-post transfers and accumulate savings. The co-existence in the same villages of two contracts highlights the presence of coalitional enforcement constraints: as the theory and simulation sections showed, only then can both be observed. We have further hypotheses that can be tested in the data.

Our first hypothesis is that social cohesion will be positively correlated with selection into a type II contract and negatively correlated with selection into a type I contract. The social cohesion of a group clearly depends on a number of individual characteristics of the group's members. As suggested in Anderson et al. (2008), households with stronger ties in the community, stronger ties in the group, and less mobility are more likely to make reliable members of insurance arrangements because they are more susceptible to social and reputational pressure. Conversely, such members may also find it easier to enforce contract rules against others who do not comply. Most importantly, high levels of social interconnectedness and trust make it easier to agree on credible formal rules ex-ante. Taken together, these observations lead to the prediction of a positive correlation between measures of member's reliability, trust in others and social connectedness and the social cohesion component of group cohesion. Consequently, these variables should be positively correlated with selection into type II contracts. Similarly, we predict that the ability to raise ex-ante transfers and savings could be related to the wealth of the group's members because the ex-ante payment required to make group savings profitable may simply exceed the liquidity of poorer households.

We test these hypotheses by estimating a probit regression of contract selection at the household level. In the regression, we capture several dimensions of social cohesion by including information on parental background, whether belonging to a minority ethnicity or religion in the village, whether born in the community and whether ones father has been a member of the Iddir. We also include two measures of the extent to which group members trust each other. The first of these is an 'observed' measure of trust, which records the percentage of people in a households' informal network (defined by common membership in informal labour sharing and credit arrangements) that are also a member of the same Iddir. Intuitively, if this proportion is high, then the household will share Iddir membership with a number of other households who they know well and can rely on in times of need. The second measure is a 'subjective' measure of mutual trust, based on a simple question of

whether others in the community were trusted (on a scale of one to five, with one equal to low). Of course, group selection – which is a prerequisite for contract selection – is a two-way process in which households select a group, but groups can accept or reject members. We recognize this by including all the above variables as group level averages in the regression. Similarly, we include correlates of wealth related to landholding – again averaged at the group level – to determine whether higher liquidity of group members explain selection into type II contracts.

Of course, our empirical predictions are strictly speaking only valid for a given environment and preferences of group members. To take account of this, we include village dummies and time dummies as well as a large number of other household and group level characteristics as explanatory variables. These include household size, age, sex and education levels of the household head, as well as information on the risk profile of groups as measured by the average number of days lost to illness and average number of deaths in the network in the past 5 years. Aside from controlling for preferences, the level of education may in itself be positively correlated with selection into type II Iddir, if these require more sophistication in terms of the day-to-day running of the group.

Table 2 offers the results from a probit analysis of the difference between the two types of groups.⁷ It is striking that the only variables that are significant and consistent are related to social cohesion and wealth. Being born in the village, higher levels of trust, a large proportion of the informal network relationships within the group and ones father having been a member of the Iddir, all increase the likelihood of selecting a type II insurance group. The variables related to wealth are all positive, and strongly significant. Taken together, these results provide strong evidence that social cohesion and liquidity constraints are positively correlated with selecting into a type II contract as predicted. The variables related to the risk profile of members are not significant, suggesting that heterogeneity in contract choice is not driven by these ‘demand-side’ factors.

4.2 Testing the Consumption Smoothing Benefits of Contract Selection

In this section we test the effects of contract design under imperfect enforcement on consumption smoothing empirically. The empirical analogue of (51) is given by the following first-difference specification:

$$\Delta \ln \left(\frac{c_t^i}{n_t^i} \right) = \alpha_0 \Delta C_t^h + \alpha_1 \Delta D_t^i + \alpha_2 \Delta (D_t^i \times B_t^h) + \delta_1 \Delta B_t^h + \delta_2 V_t + \delta_3 \Delta \ln H_{it} + \varepsilon_{it}. \quad (52)$$

⁷The pooled probit regression is based on the 266 of the 301 sampled households for whom all explanatory variables are available in at least one round. This gives a total of 1106 observations in six rounds.

Table 2: Determinants of contract selection

	(1)	
	contract type	
	1=ex-ante transfers, 0=ex-post	
Father was member of Iddir	0.093 ^{**}	(0.038)
% of network born in village	0.308 [*]	(0.071)
% of network that belongs to ethnic minority in village	-0.172	(0.517)
% of network that belongs to religious minority in village	0.178	(0.445)
Proportion of informal network in same Iddir	0.367 ^{**}	(0.045)
Average measure of trustworthiness of others in network (1=Low, 5=High)	0.133 ^{***}	(0.004)
Second land-owning quartile†	0.123 ^{***}	(0.003)
Third land-owning quartile†	0.280 ^{***}	(0.000)
Fourth land-owning quartile†	0.220 ^{***}	(0.000)
Ln age of head	-0.180 ^{***}	(0.004)
Ln of household size	-0.006	(0.858)
Head has schooling†	-0.036	(0.354)
Female-headed household†	0.003	(0.944)
Number of days lost to illness in past 4 weeks	0.008	(0.648)
Average number of deaths in network in past 5 years	0.115	(0.190)
Observations	1106	

Source: ERHS and Funeral Insurance Survey.

Notes: Marginal effects. P-values in parentheses. Village and time dummies included, but not reported. Robust standard errors. (†) for discrete change of dummy variable from 0 to 1. * significant at 10% level, ** significant at 5% level, *** significant at 1% level.

D_t^i is defined as the number of deaths in household i in period t . We focus on death shocks as an overspecification test because they are large, idiosyncratic and to some extent unpredictable. More importantly, the groups under study provide insurance to cover the costs associated with the ceremonial expenditure on the funeral of the deceased. Hence, the impact of death shocks is of immediate interest to test our hypotheses related to contractual design and enforcement. $C_t^h = \frac{1}{N_t^i} \sum_{j \neq i \in N_t^i} \ln \left(\frac{C_t^j}{n_t^j} \right)$ measures the average resources and $D_t^h = \frac{1}{N_t^i} \sum_{j \neq i \in N_t^i} \ln \left(\frac{D_t^j}{n_t^j} \right)$ the average number of shocks of risk-sharing partners N_i of household i in several (possibly overlapping) groups. B_t^h is a dummy that takes on the value 1 if group h uses primarily ex-ante transfers and saves and zero otherwise.⁸ Finally, V_t is a set of village-time dummies and H_{it} denotes a set of time-varying household characteristics, which are included to capture preference shifts.

There are a number of econometric issues that have to be addressed when estimating (52). If there is full insurance at the village level, then the time-specific village fixed effect should explain all changes in consumption growth and α_0 will be zero. If risk-sharing is confined to insurance groups within the community, then changes in household consumption are at least partly determined by changes in average group consumption and α_0 will be significant and measures the extent to which own consumption covaries with group consumption.

However, the finding of a significant partial effect of changes in group consumption on changes in own consumption does not necessarily provide evidence that Iddir are the vehicles for insurance, because the OLS estimator mechanically fits the mean regardless of whether the data is generated by a model of risk-sharing or not (see Suri (2005)). This is known as the reflection problem in the peer effects literature (see Manski (1993) and Brock and Durlauf (2001)). Of course, our regression does not actually include the mean of group consumption, but the mean of group consumption excluding own consumption. However, this does not solve the issue and as pointed out by Boozer and Cacciola (2001) and Suri (2005), the magnitude of α_0 depends on the inclusion or exclusion of covariates. The authors show that α_0 can be written as $1 - \frac{WSS}{(N-1)BSS}$ where WSS and BSS are the conditional within and between sum of squares respectively. Including covariates that vary only at the group level will bias the estimate of α_0 towards zero, while covariates that explain within group variation will make α_0 closer to one.

While the identification of α_0 as a true measure of risk-sharing is not of primary interest in our estimation, failing to estimate it correctly may present cause for

⁸Because of multiple group memberships, B_t^h strictly speaking measures the percentage of Iddir of which a household is a member that use ex-ante transfers. In practice, most households only join one type of group, so B_t^h remains a dummy even when averaged over multiple groups. Contract type is included in levels, because a household's membership varies at least to some extent over the 6 rounds of data.

concern if the remaining coefficients in the regression are biased as a result. We therefore present a number of alternatives to the first-difference estimator. One solution suggested by Townsend (1994) is to simply impose $\alpha_0 = 1$ and subtract the mean of group consumption from both sides before estimating the regression.

$$\begin{aligned} \Delta \left(\ln \left(\frac{c_t^i}{n_t^i} \right) - C_t^h \right) &= \alpha_1 \Delta D_t^i + \alpha_2 \Delta (D_t^i \times \hat{B}_t^h) \\ &+ \delta_1 \Delta B_t^h + \delta_2 V_t + \delta_3 \Delta \ln H_{it} + \varepsilon_{it}. \end{aligned} \quad (53)$$

An alternative specification proceeds in terms of village and group level fixed effects (see Ravallion and Chaudhuri (1997), Deaton (1992), Morduch (1991) among others).

$$\begin{aligned} \Delta \ln \left(\frac{c_t^i}{n_t^i} \right) &= \sum_{h=1}^M \vartheta_h \Delta G_t^h + \alpha_1 \Delta D_t^i + \alpha_2 \Delta (D_t^i \times B_t^h) \\ &+ \delta_1 \Delta B_t^h + \delta_2 V_t + \delta_3 \Delta \ln H_{it} + \varepsilon_{it} \end{aligned} \quad (54)$$

where G_t^h is a dummy equal to 1 if household i is a member of group h in period t . Thirdly, we use an instrumental variable strategy. As shown in Brock and Durlauf (2001), IV is an appropriate solution to the reflection problem if there is at least one individual variable whose group average is not an element of the individual risk-sharing equation. It is standard in the risk-sharing literature to include individual level preference shifters in the individual equation, but to assume that group level averages of preference shifters do not affect own consumption once group consumption has been controlled for (see Mace (1991) and Townsend (1994)). If we are willing to maintain this assumption, then all the group level averages of H_t^i are appropriate instruments for group consumption.

In addition to the reflection problem, we face another potentially serious problem. In the regression equation we have specified it is implicitly assumed that network formation is exogenous. However, this is unlikely to be the case, particularly in the case of funeral insurance groups, where a household can choose between a number of groups in a particular village and this choice is unlikely to be random. There may then be (time-varying) unobservables that may be correlated with both own consumption and group consumption. The effect that this will have on the bias in $\hat{\alpha}_0$ is not a priori clear. On the one hand, households may choose network partners with whom they are connected through ethnic or geographic ties and this proximity may imply that income and therefore consumption streams within the network are positively correlated even in the absence of risk-sharing. This would lead to an upward bias in $\hat{\alpha}_0$. On the other hand, households may choose network partners whose income streams are negatively correlated with their own in order to

increase the benefits from insurance. This would lead to a negative bias in $\hat{\alpha}_0$ if endogenous network formation is not adequately accounted for.

Appropriate instruments for changes in average network consumption in (52) are required for identification: exogenous factors that are correlated with changes in network consumption and only affect changes in own consumption via network consumption. The number of deaths in the network, beyond the deaths experienced by household i , is a plausible identifying instrument for network consumption. The argument is that, given an insurance group, changes in deaths of other network members only affect own consumption via the insurance group. Own family mortality shocks are being controlled for by the design of the risk-sharing test, and (unsurprisingly) deaths within the community and network are far from perfectly correlated with own family mortality ($\rho=-.143$). Network remittances to network members from outside the network also have the potential to be reasonable instruments. As argued in De Weerd and Dercon (2006), the motivation for this assumption is that outside remittances are less likely to be highly correlated among households than other components of consumption that are wholly determined inside the village. Overall, this suggests that we may have a priori confidence in the validity of our instrument.

A final challenge in the estimation of (52) is that contract choice B_t^h may be driven by unobservables also affecting consumption outcomes c_t^i . For example, the theoretical model has suggested that the extent of trust and cohesion in the community or group one belongs to may affect contract choice (as the better contract requires more trust and cohesion). Trust may also affect how well one is doing in terms of consumption, c_t^i . However, this particular case is not a problem: if trust and cohesion are fixed, time-invariant factors, then estimating the difference model in (52) would purge the equation of these time-invariant factors. A more problematic case would be if variables such as trust and cohesion are affecting the growth in consumption, as if there are externalities from living in high trust environments (as suggested by social capital models). In that case, $\Delta(D_t^i \times B_t^h)$ may well be correlated with trust and cohesion and endogeneity linked to a missing variable problem would give inconsistent estimates of β_2 .

In the data, we have however a large number of variables that can reflect the underlying trust and cohesion in the group and community and these have been shown to be significant determinants of contract selection. We follow Wooldridge (2002) and include the propensity scores from the contract selection probit regression interacted with D_t^i to control for selection on observables in the risk-sharing equation.⁹ The result is an estimate of β_2 offering the potential differential impact

⁹Technically, this requires that the difference between the unobserved errors of those in a type II contract and those in a type I contract has zero mean conditional on the regressors in the risk-sharing regression. Moreover, it requires that the difference in outcomes for those in the two

of the contract itself, controlling for the possibly convoluted impact of general trust and cohesion present in the group.

We can now estimate (52) using data on six villages and households with complete matched data on the funeral groups they are members of. From the survey data, it is possible to construct an estimate of monthly household consumption for each round of the survey that will be used in the augmented risk-sharing test. The consumption data are obtained by summing over all sources of food and non-food consumption, deflated by a consumer price index using average prices in the first round as a base (Dercon and Krishnan (2000b)). The data are purged of any additional expenditure, linked to exceptional items, such as health treatment or ceremonial expenditure, so any increases in consumption directly related to the mortality episode in the household are excluded. Total consumption is expressed in per capita terms and in Ethiopian Birr of 1994 (1 USD = 9 Birr). Average per capita consumption is low at about 86 Birr per month or just US\$0.32 per day.

To estimate the impact of death shocks on household consumption we use an aggregate measure of deaths that have occurred in a household, which has been constructed from detailed ERHS household composition and mortality data. The measure is heavily weighted to deaths of dependents (78% of recorded deaths), which trigger the need for ceremonial expenditure. About 45% of households reported at least one death during the survey period. On average, this corresponds to 18% of households reporting a death in each round. Iddir only offer insurance against the funeral expenditure related to these deaths, but these can be very substantial: average payouts amount to the equivalent of three months per adult consumption.

To control for the effect of Iddir membership in the risk-sharing equation, we need to identify the different funeral insurance groups and their membership in the sample. Triangulating the ERHS and FIS data, we have detailed data on 78 funeral groups and their members. As the ERHS sampled a high percentage of the population of each village, about a quarter of the members of each group were found to be included in the ERHS, allowing us to use an estimate of the mean consumption level across the group; the minimum number of observations for estimating group consumption is five.¹⁰ The contract design variable is based on the interviews with

different contracts in the risk-sharing regression is not correlated with the variance of selecting into a type II contract conditional on observables (see Wooldridge (2002)).

¹⁰This has been necessary for two reasons: first of all, the fewer households that report membership and therefore information about a group, the less confident one can be that a group has actually been identified correctly. Secondly, the fewer members of a group that are observed in the sample, the bigger the approximation error in computing mean group consumption and therefore the attenuation bias in estimating the network effect on changes in household consumption when insurance is incomplete. Comparison of a number of socio-economic characteristics shows that included and excluded households are not systematically different. Similarly, the characteristics of insurance groups whose membership has been successfully mapped do not differ significantly from groups for which matching of members has not been possible. Given the first difference specification and the fact that some households are not observed in all periods, the effective number of

the Iddir leadership, and as Table 1 showed the two types are well identified in the data.¹¹

We now turn to a discussion of the regression results. Table 2 offers the result of the augmented risk-sharing test in (52). The model controls for a full set of interactions of village and time dummies, capturing in each period total resources available in the village, as well a full set of taste shifters, using time-varying demographic characteristics (males and females aged 0 to 5, 6 to 15 and 16+). We present the OLS estimation in column (1), the Townsend specification in column (2), a specification in terms of group dummies in column (3) and the instrumented regression using outside remittances and aggregate death shocks as identifying instruments for group consumption in column (4).

The results are consistent with an environment of imperfect enforcement, which is indicated by the fact that the overidentifying regressors are strongly significant. The estimated coefficients are also remarkably similar across the four specifications implying that estimation is unlikely to be affected by the identification problems surrounding the group consumption variable. As discussed in the previous section, the coexistence of both contract types together with the departure from first-best risk-sharing suggests that risk-sharing is not only constrained by individual deviations but also by coalitional deviations.

In terms of the impact of a death shock, groups that use both ex-ante transfers and savings clearly offer superior consumption smoothing to those that do not. On average, a household that is insured in a Type I group suffers a drop in consumption of 12% following a death shock in column (4).¹² If a household is insured in a type II group, the 12% drop in consumption following a death shock is reversed completely and consumption even increases slightly. The total effect of $\alpha_1 + \beta_2 = 0.06$ is not significantly different from zero, indicating that on average type II groups achieve full insurance.¹³ Based on the numerical results presented in the previous section, this indicates that households who are members of type II groups have higher welfare

groups included in the regression is 248. We also attempt a fixed effects specification which does not require households to be observed in subsequent periods and retains 301 observations. The results are virtually identical.

¹¹For a very small number of groups, based on the data of the Iddir, arguments could be made that they need to be included in another type than what the leadership had reported in the survey. Using an alternative classification was found to make little difference in the results and their significance.

¹²The consumption aggregates used in the regression have been purged of any increases in expenditure necessitated by the occurrence of a shock (such as ceremonial expenditure in case of a funeral or hospital bills in case of illness). Therefore our insurance test will measure whether a death shock has a lasting negative impact on consumption either because of an income short fall or because the added expenditure at the time of a funeral makes it necessary to decrease consumption subsequently.

¹³While at first somewhat surprising, the positive (albeit insignificant) impact of a death shock on consumption in a type II contract is consistent with our model of a contract that distributes ex-ante transfers to agents with a low draw and occasionally overcompensates these agents to induce them to remain in the insurance arrangement.

Table 3: Testing the welfare effects of contract selection

	(1) FD $\Delta \ln$ consumption	(2) FD Townsend $\Delta \ln$ consumption	(3) FD Dummies $\Delta \ln$ consumption	(4) FD-IV $\Delta \ln$ consumption
$\Delta \ln$ netw. consumption	-0.273** (0.050)			0.701** (0.020)
Δ death shock	-0.095 (0.128)	-0.123* (0.069)	-0.093 (0.130)	-0.124* (0.062)
Δ contract	-0.564* (0.081)	-0.776** (0.042)	-0.407 (0.122)	-0.719** (0.048)
Δ (death \times contract)	0.151** (0.042)	0.198** (0.013)	0.179** (0.014)	0.184** (0.002)
p-value of joint significance of Iddir dummies (χ^2 -test(40))			0.000	
Number of households	248	248	273	248
F-statistic of overall significance	9.96	4.46	34.68	9.71
Hansen statistic				0.178
Kleibergen-Paap statistic				12.73
p-value of Cragg-Donald underidentification test				0.000

Source: ERHS and Funeral Insurance Survey.

Notes: Household fixed effects. P-values in parentheses. Village-time dummies and individual preference shifts included, but not reported. Robust standard errors. * significant at 10% level, ** significant at 5% level, *** significant at 1% level.

than households insured via type I groups. Averaging over both contract types, the impact of a death shock is 0.02, which is also not significantly different from zero. Hence, on average risk-sharing groups sustain perfect risk-sharing. However, this masks substantial variation in the extent of insurance across contract types.

Finally, considering α_0 , we see that network consumption and the group dummies are strongly significant in column (1),(3) and (4) after controlling for period-specific village effects. This suggests that risk-sharing occurs via the network and not just the village. The negative sign in the OLS regression is possible, if insurance is partial and the covariance between group resources and village resources is large. Given that village level resources vary between groups but not within groups, we expect this to reduce $\alpha_0 = 1 - \frac{WSS}{(N-1)BSS}$ and a stepwise regression approach reveals that the negative coefficient is indeed driven by the inclusion of the village level fixed effect. The coefficient on group consumption is large, significant and positive in the IV-regression. This is consistent with full insurance of ceremonial expenditure via Iddir in the sense that the added expenditure at the time of a funeral does not require households to subsequently cut back on consumption. The first stage regression is well behaved. The identifying instruments are overall significant and network remittances are positively correlated with network consumption and the diagnostics reported in Table 4 suggest that we can have some confidence in the instruments used. The Cragg-Donald test of underidentification is strongly rejected, the Kleibergen-Paap weak identification statistic exceeds the critical values computed by Stock and Yogo (2002) and the Hansen statistic suggests that the overall set of instruments is exogenous.¹⁴

So far, our analysis points to the importance of insurance groups in sharing risk related to funeral expenditures. Specifically, the results suggest that funeral expenditure is insured well enough that households do not suffer any lasting negative impact following a death, i.e. they do not find it necessary to decrease consumption subsequently. Moreover, there is evidence that groups carefully design contracts to overcome the impediments of imperfect enforcement and threats of coalitional deviations. In particular, groups that combine ex-ante transfers and savings seem to be remarkably successful at approximating full insurance. Given this superiority, it is surprising that not all groups use ex-ante transfers and accumulate savings. However, as we have shown in the modelling section, sufficient trust among group members is an important prerequisite for such an organizational structure, and our analysis in the previous section has demonstrated that the level of trust is significantly higher among members of Type II groups than members of Type I groups. Of course, it is also possible that this higher level of trust affects not just contract choice, but also how well shocks are insured once a group has decided on a

¹⁴Detailed first-stage results are available from the authors on request.

particular institutional set up. If this is the case and we fail to control for the effect of trust and cohesion, then the coefficient on the contract variable may suffer from omitted variable bias. We investigate this in Table 4, where we add the propensity score from the first-stage contract selection regression both in levels and interacted with the death shock to the instrumented risk-sharing regression (column (4) in Table 3) in order to control for any bias that arises from selection on observables.¹⁵

Column (1) of the table reports how inclusion of the propensity score affects the predicted average impact of a death shock in a type II contract. The first thing to note is that the predicted effect of a death shock in a type I contract remains negative and significant, while the interaction of the death shock variable with the contract dummy is positive and significant. A death shock in a type I group is predicted to result in a 12% drop in consumption, which is reasonably close to the range of values reported in Table 3. The interaction of the propensity score with the death shock variable is negative and significant indicating that selection on observables is important. However, the total impact of a death shock in a type II contract is predicted to be $-0.12 + 0.31 - 0.26 \times (E(\text{propensity score}|\text{contract type=II}) - E(\text{propensity score}|\text{contract type=I})) = 0.08$, which is virtually the same as in Table 3 and not significantly different from zero.

In column (2), we investigate an alternative specification in which the variables relating to social cohesion from the contract selection regression are included directly in the regression instead of using the propensity score as a summary measure. The results from the two specifications are more or less identical. The impact of a death shock in a type I contract is -0.167 , while the impact of a death shock in a type II contract is predicted to be 0.07 , which is not significantly different from zero. The coefficients on the interaction of the average trust in the network and a household's death shock and the interaction of proportion of a household's informal network and the death shock are the only ones to be significant but cancel each other out in magnitude. Finally, the effect of (instrumented) changes in group consumption is large and significant for both specifications in Table 4. Taken together, these results can give us some confidence that the coefficient on the contract variable does indeed measure the impact of the organizational structure of the Iddir itself rather than merely the effect of a group's membership composition.

¹⁵Following the method in Wooldridge (2002), the propensity score is estimated using all the variables in Table 2 as well as the explanatory variables in (52). The effects of the variables from the contract selection regression are virtually unchanged, while exogenous regressors from the regression in (52) mostly fail to be significant in this 'first-stage' regression. The propensity score variable is demeaned as follows: $\text{propensity score} - E(\text{propensity score}|\text{contract type=I})$. This ensures that the coefficient on the death shock measures the impact of a shock in a type I group.

Table 4: Testing the welfare effects of contract selection controlling for selection

	(1) (FD-IV) $\Delta \ln$ consumption	(2) (FD-IV) $\Delta \ln$ consumption
$\Delta \ln$ netw. consumption	0.710** (0.030)	0.709* (0.019)
Δ death shock	-0.151** (0.018)	-0.167** (0.014)
Δ contract	-0.688 (0.156)	-0.858** (0.021)
Δ (death \times contract)	0.314*** (0.003)	0.237*** (0.004)
Δ propensity score	0.422 (0.392)	
Δ (death \times prop. score)	-0.260 (0.132)	
Δ (death \times netw. trust)		-0.205** (0.035)
Δ (death \times % of netw. born in village)		-0.002 (0.991)
Δ (contract \times % of netw. minority religion)		0.097 (0.688)
Δ (death \times % of netw. minority religion)		-0.216 (0.364)
Δ (death \times % of informal netw. in Iddir)		0.401** (0.042)
Number of households	208	248
F-statistic of joint significance	7.95	8.29
p-value of Hansen J-statistic	0.234	0.530

Source: ERHS and Funeral Insurance Survey.

Notes: Household fixed effects. Propensity scores derived from contract selection regression in Table 2. P-values in parentheses. Village-time dummies and individual preference shifts included, but not reported. Robust standard errors. * significant at 10% level, ** significant at 5% level, *** significant at 1% level.

5 Conclusion

This paper has explored theoretically and empirically contract choice in funeral insurance groups and its consequences for welfare. We focused on two types of contracts found in rural Ethiopia: one with considerable savings at the level of the group and regular payments unrelated to shocks, and one without savings and payments to its members after they experienced a mortality shock. Consistent with theoretical predictions and simulation results, we found that both contracts co-exist in the same villages and that contract choice is related to social cohesion. In general, we found evidence that on average death shocks are well insured but that there are catastrophic shocks, for which insurance is incomplete. Furthermore, we could confirm that groups with contracts including group savings resulted in welfare gains in the form of higher consumption smoothing among the members of groups engaging in this type of contract, compared to those without group savings and only ex-post payments. Despite this heterogeneity, our results suggest nevertheless that funeral associations in rural Ethiopia are remarkably successful in providing insurance against death shocks.

In a broader context, our paper also contributes to the extensive and continuously growing literature on social capital and its welfare impacts. While there is no unique definition of social capital, the instrumental view of social capital holds that in and of itself it is a neutral force (see Coleman (1988)). That is, it has no intrinsic benefits unless it is put to some productive use. The results in this paper mirror this view. While high levels of trust and reciprocity – which are suitable proxies for social capital according to Putnam (2000) – are a necessary prerequisite for the enforcement of more sophisticated contractual arrangements, they are neither sufficient nor are they substitutes for the benefits derive from more formalized insurance provision. In this respect, the paper also sheds some light on the important question of how social capital can be turned into a "resource for action" in practice (Coleman (1988)).

Taking a longer term view, there are even more fundamental reasons, why semi-formal insurance groups deserve our attention. While some may argue that the prevalence of these institutions in developing countries signifies economic backwardness, and that they will rightfully become obsolete if growth takes off, there is an alternative view which posits that these institutions are an integral part of a country's growth strategy especially in terms of achieving financial proliferation. After all, there are numerous examples of humble mutual and friendly societies in developed economies that have turned into large and sophisticated financial institutions over the last century.¹⁶ This provides some hope that informal financial institutions

¹⁶The German Raiffeisen cooperatives started as localized self-help organizations in the 19th century and today have the densest network of bank branches in Germany. Similarly, friendly

can be stepping stones – rather than stumbling blocks – on the path to economic development.

References

- Anderson, S., Baland, J.-M., Moene, K. O., 2008. Enforcement in Informal Savings Groups, forthcoming, *Journal of Development Economics*.
- Banerjee, A. V., Besley, T., Guinnane, T. W., 1994. Thy Neighbor's Keeper: The Design of a Credit Cooperative with Theory and a Test. *Quarterly Journal of Economics* 109, 491–515.
- Bernheim, B. D., Peleg, B., Whinston, M., 1987. Coalition Proof Nash Equilibria. I: Concepts. *Journal of Economic Theory* 42, 1–12.
- Bernheim, B. D., Ray, D., 1989. Collective Dynamic Consistency in Repeated Games. *Games and Economic Behavior* 1, 295–326.
- Boozer, M., Cacciola, S., 2001. "Inside the 'Black Box' of Project Star: Estimation of Peer Effects Using Experimental Data", Economic Growth Center Discussion Paper 832, Yale University.
- Brock, W., Durlauf, S., 2001. *Handbook of Econometrics*. Elsevier Science B.V, Ch. Interactions-Based Models, pp. 3299–3380.
- Cabrales, A., Calvo-Armengol, A., Jackson, M. O., 2003. La Crema: A Case Study of Mutual Fire Insurance 111, 425–456.
- Coate, S., Ravallion, M., 1993. Reciprocity without Commitment: Characterization and Performance of Informal Insurance Arrangements. *Journal of Development Economics* 40, 1–24.
- Coleman, J., 1988. Social Capital in the Creation of Human Capital. *American Journal of Sociology* 94, S95–S121.
- Deaton, A., 1992. Household savings in LDCs: Credit Markets, Insurance and Welfare. *Scandinavian Journal of Economics* 96, 253–273.
- Deaton, A., 1997. *The Analysis of Household Surveys*. Johns Hopkins University Press, Baltimore, Maryland.
- De Weerd, J., 2004. *Insurance Against Poverty*. Oxford University Press, Ch. Risk-Sharing and Endogenous Group Formation.

societies in the UK cater for as many as six million customers.

- De Weerd, J., Dercon, S., 2006. Risk-Sharing Networks and Insurance against Illness. *Journal of Development Economics* 81, 337–356.
- Dercon, S., Bold, T., De Weerd, J., Pankhurst, A., 2006. Group-based Funeral Insurance in Ethiopia and Tanzania. *World Development* 34, 685–703.
- Dercon, S., Krishnan, P., 2000a. In Sickness and in Health: Risk Sharing within Households in Rural Ethiopia. *Journal of Political Economy* 108, 688–727.
- Dercon, S., Krishnan, P., 2000b. Vulnerability, Seasonality and Poverty in Ethiopia. *Journal of Development Studies* 36, 25–53.
- Fafchamps, M., Lund, S., 2003. Risk-Sharing Networks in Rural Philippines. *Journal of Development Economics* 71, 261–287.
- Foster, A., Rosenzweig, M., 2001. Imperfect Commitment, Altruism, and the Family: Evidence from Transfer Behaviour in Low-Income Rural Areas. *Review of Economics and Statistics* 83, 389–407.
- Gauthier, C., Poitevin, M., Gonzalez, P., 1997. Ex Ante Payments in Self-Enforcing Risk-sharing Contracts. *Journal of Economic Theory* 76, 106–144.
- Genicot, G., Ray, D., 2003. Group Formation in Risk-Sharing Arrangements. *Review of Economic Studies* 70, 87–113.
- Gertler, P., Gruber, J., 2002. Insuring Consumption against Illness. *American Economic Review* 92, 51–70.
- Gobert, K., Poitevin, M., 1998. Non-Commitment and Savings in Dynamic Risk-Sharing Contracts, Mimeo, University of Montreal.
- Grimard, F., 1997. Household Consumption Smoothing through Ethnic Ties: Evidence from Cote d’Ivoire. *Journal of Development Economics* 53, 391–422.
- Kocherlakota, N. R., 1996. Implications of Efficient Risk Sharing without Commitment. *Review of Economic Studies* 63, 595–609.
- Ligon, E., Thomas, J., Worrall, T., 2002. Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies. *Review of Economic Studies* 69, 209–244.
- Ligon, E., Thomas, J. P., Worrall, T., 2000. Mutual Insurance, Individual Savings and Limited Commitment. *Review of Economic Dynamics* 3, 216–246.
- Mace, B. J., 1991. Full Insurance in the Presence of Aggregate Uncertainty. *Journal of Political Economy* 99, 928–956.

- Manski, C., 1993. Identification of Endogenous Social Effects: The Reflection Problem. *Review of Economic Studies* 60, 531–542.
- Morduch, J., 1991. Consumption Smoothing Across Space: Tests for Village-Level Responses to Risk, Mimeo, Harvard University.
- Morduch, J., 1999. Between the State and the Market. *World Bank Research Observer* 14, 187–207.
- Putnam, R. D., 2000. *Bowling Alone: The Collaps and Revival of American Community*. Simon and Schuster, New York.
- Ravallion, M., Chaudhuri, S., 1997. Risk and Insurance in Village India: Comment. *Econometrica* 65, 171–184.
- Stock, J., Yogo, M., 2002. Testing for Weak Instruments in Linear IV regression, National Bureau of Economic Research, Technical Working Paper 284.
- Stokey, N., Lucas, R., Prescott, E. C., 1989. *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, Massachusetts.
- Suri, T., 2005. "Spillovers in Village Consumption: Testing the Extent of Partial Insurance", mimeo, MIT.
- Townsend, R. M., 1994. Risk and Insurance in Village India. *Econometrica* 62, 539–591.
- Udry, C., 1994. Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria. *Review of Economic Studies* 61, 495–526.
- Wooldridge, J. M., 2002. *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, Massachusetts.