Implications of Endogenous Group Formation for Efficient Risk-Sharing

Tessa Bold
New College
University of Oxford
JOB MARKET PAPER

Abstract

The existing literature on sub-game perfect risk-sharing suffers from a basic inconsistency. While a group of size n is able to coordinate on a risk-sharing outcome, it is assumed that deviating subgroups cannot. We relax this assumption and characterize the optimal contract among all coalition-proof history-dependent contracts. This alters the predictions of the standard dynamic limited commitment model. We show that the consumption of constrained agents depends on both the history of shocks and its interaction with the current income of other constrained agents. From this, we derive a formal test for the presence of endogenous group formation under limited commitment.

1 Introduction

With complete financial markets, the neo-classical model of consumption predicts that all idiosyncratic risk will be insured so that the only fluctuations in consumption are due to aggregate shocks. While complete financial markets are notably absent in developing countries, economists have argued that the close-knit structure of communities in developing countries may enable them to successfully replicate full insurance by insuring all idiosyncratic risk at least at the community level. Empirically, however, the full insurance hypothesis is rejected. Individual consumption varies with idiosyncratic shocks to income and the relative return to bearing risk is several orders of magnitudes larger than estimates derived from Arrow-Debreu complete market models (Mace (1991), Deaton (1992), Townsend (1994), Udry (1994), Dercon and Krishnan (2000), Ogaki and Zhang (2001), Ligon et al. (2002), Mehra and Prescott (1985)). Moreover, in developing country contexts, diversification is not only incomplete, but the sharing of idiosyncratic risk is also limited to small groups – such as the extended family and informal reciprocal networks – and does not extend to the entire community – be that the village or larger economic units (Grimard (1997), Morduch (1991), Fafchamps and Lund (2003), De Weerdt and Dercon (2006), Murgai et al. (2002)).

The empirical failure of the complete markets model has inspired a host of theoretical attempts to explain the observed covariation between consumption and income. In the main the literature

has proposed two explanations for incomplete markets: imperfect information and imperfect enforceability – usually with an emphasis on modelling risk-sharing as the sub-game perfect equilibrium of a 2-player game (Wang (1995), Phelan and Townsend (1991), Kehoe and Perri (2002), Kocherlakota (1996), Ligon et al. (2002)). While these models have been extremely successful in explaining some of the qualitative properties of the relationship between consumption and income, a number of important caveats remain: the prevalence of small risk-sharing groups cannot be explained by the above literature as it predicts that in the absence of exogenously imposed barriers, risk-sharing should always take place at the level of the community. Moreover, attempts to calibrate these models to explain the quantitative properties of consumption and income series have failed to generate the large fluctuations and equity premia observed in reality.

This is where this paper hopes to make a contribution. By modelling risk-sharing under imperfect enforceability as the coalition-proof equilibrium of a repeated dynamic stochastic game, we can explain both why the size of the risk pool and consequently the extent of risk-sharing is much smaller than the predictions derived from existing models.[1] In this paper, we derive the theoretical properties of the coalition-proof dynamic risk-sharing contract and derive testable implications on the reduced form that enable us to establish whether coalitional deviations pose a threat to the viability of risk-sharing arrangements and by implication whether the requirement of coalition-proofness can be considered a valid restriction on risk-sharing arrangements that improves the fit of theory and empirics.

Under imperfect enforceability, risk-sharing is limited because contracts are not legally enforceable ex-ante. Therefore, risk-sharing arrangements have to take into account that individuals will renege on the contract ex-post if the benefits from doing so outweigh the costs. In the majority of the extant literature, dynamic risk-sharing under imperfect enforceability is modelled as the equilibrium of an infinitely repeated n-player game, in which cooperation is sustained by assuming that players punish a deviation by permanently reverting to autarky (Kocherlakota (1996), Ligon et al. (2002)). The theoretical properties of the equilibrium lead to relevant predictions on consumption patterns, namely that consumption is positively correlated with current and past income and that consumption allocations depend on the history of the game when enforcement constraints are slack, but display 'amnesia' when enforcement constraints bind. While enormously influential in sharpening our understanding of the limits to risk-sharing imposed by imperfect enforceability, several unresolved issues remain. Empirically, the dynamic limited commitment tends to overpredict the

---

[1] As pointed out by Udry (1994), communities in developing countries can be characterized as high information environments but plagued with a lack of formal enforcement mechanisms. From a theoretical perspective, Seppala (1999) argues that limited commitment allows for richer models of incomplete risk-sharing.

extent of risk-sharing observed in practise (Ligon et al. (2002), Alvarez and Jermann (2001)) while some of its 'sharper' predictions, notably that of 'amnesia' following a binding constraint, are rejected by the data (Kinnan (2008)). Theoretically, one could argue that modelling risk-sharing as the sub-game perfect equilibrium of a repeated game suffers from a conceptual inconsistency. The threat supporting the sub-game perfect equilibrium is not credible because it requires that players consume their own income in each period following a deviation, and therefore mutual benefits from insurance are forgone. Consequently, agents have an incentive to abandon the punishment path and instead renegotiate to an equilibrium that makes all of them better off (see Farrell and Maskin (1989)). While renegotiation-proof punishments that are not subject to this critique are available in a bilateral risk-sharing model (see Asheim and Strand (1991) and Kletzer and Wright (2000)), the problem is not so easily resolved in a multilateral environment and becomes more acute the larger the number of players. This is the case because the punishment path of permanent reversion to autarky is not just vulnerable to collective renegotiation by the grand coalition of all n players, but also to renegotiation by any subset of players, making suspect the viability of the sub-game perfect equilibrium and the predictions on consumption derived from it.

Given these caveats, we argue in this paper that the 'correct' equilibrium concept for an n-player risk-sharing game is that of coalition-proofness – a refinement of sub-game perfection. This means that no subcoalition can ever be required to play a Pareto dominated equilibrium in any sub-game (Bernheim et al. (1987)). By implication, only those risk-sharing arrangements are deemed self-enforcing that are sustained by punishments which are themselves coalition-proof equilibria. Defining and characterizing the coalition-proof equilibria of dynamic risk-sharing games and resultant predictions on consumption is the purpose of this paper. While this paper has a theoretical focus, we will see that the chosen modelling approach can address both some of the conceptual as well as empirical shortcomings of existing models of dynamic risk-sharing under limited commitment.

Genicot and Ray (2003) were the first to point out that sub-game perfection, which requires risk-sharing arrangements to be stable with robust to individual deviations only, is not a satisfactory solution concept for games involving more than two players. It places no bound on group size and predicts that risk-sharing will take place at the level of the community, which is clearly at odds with empirical reality. Genicot and Ray endogenize the group formation process by requiring risk-sharing arrangements to be coalition-proof and prove that coalition-proofness implies an important restriction on the equilibrium set, namely that all – possibly history-dependent and asymmetric – risk-sharing agreements are bounded in group size if they are required to be robust with respect to coalitional deviations.

This paper goes considerably beyond Genicot and Ray (2003) by solving for the efficient dynamic risk-sharing contract in the set of coalition-proof equilibria, characterizing its properties and deriving testable implications that can be used to empirically distinguish exogenous from endogenous group formation in risk-sharing arrangements. In the model presented here, agents play an infinite-horizon game, have identical preferences and discount the future by a common factor. All agents are risk-averse and receive a stochastic endowment stream. Income realizations in each period are i.i.d. and publicly observable.[2] There is no opportunity for saving. Since agents are risk-averse, there are gains from mutual insurance and ex-ante all agents are willing to engage in risk-sharing. We assume that insurance contracts cannot be legally enforced and participation is voluntary. Hence, the only incentive for making a transfer to another agent is the surplus the giver receives from the ongoing risk-sharing arrangement in the future.

To define the coalition-proof equilibrium contract, we must first specify what punishments will be meted out for deviating from the initial equilibrium path. In constructing optimal coalition-proof punishment paths, we take the coalitional minmax payoffs Genicot and Ray define in their paper as a starting point and show how to identify payoffs in these sets that are optimal penal codes in the spirit of Abreu (1988). That is, they completely characterize the set of coalition-proof equilibria. Crucially, we demonstrate that the optimal penal code can deter coalitional deviations to any of the potentially infinite number of strategies and associated payoffs in the minmax set of the deviating coalition by switching to a single element in this set that gives each coalition exactly an allocation on its Pareto frontier. By definition, no further deviations that make all players better off are possible from such an allocation.

The construction of these optimal punishment paths is key to obtaining an analytical characterization of the constrained-efficient dynamic risk-sharing arrangement. As is standard in the literature, we formulate a recursive dynamic programme subject to a set of enforcement constraints that embody the coalition-proof punishment paths (Abreu et al. (1990), Ligon et al. (2002)). Conceptually, the problem of finding the constrained-efficient contract for n players can now be decomposed into two steps. In a first step we find the allocation in the minmax set of each subcoalition for which all members of the subcoalition are made indifferent between cooperation and defection. Having pinned down the punishment payoffs, which form the right-hand side of the enforcement constraints, in this way, the second step solves for consumption and continuation payoffs that can be supported by these punishments. We then characterize the solution and its dynamics and find that many of the attributes that have come to be associated with dynamic risk-sharing under limited commitment

---

are in fact a consequence of modelling risk-sharing as the sub-game perfect equilibrium of a repeated game. That is to say, they are not robust when equilibria are refined to be coalition-proof.

We show that a coalition-proof contract is history-dependent both when enforcement constraints are slack and when they are binding. As long as enforcement constraints are slack, consumption is characterized by Borch's rule. This requires that marginal utility ratios, which are a sufficient statistic for the history of the game, are equalized across all states of the world and time periods. More importantly, a coalition-proof risk-sharing contract – unlike its sub-game perfect counterpart – exhibits history dependence even for those agents who are constrained. The intuition for this result is as follows. When enforcement constraints first bind, consumption and continuation payoffs of constrained agents should be adjusted in a manner that involves the least change in the ratio of marginal utilities relative to the previous period. This implies the smallest departure from first-best risk-sharing according to Borch's rule. In the sub-game perfect contract, minimal departures from first-best are always achieved by equating continuation payoffs with the static autarky payoff regardless of the history of the game. However, in the coalition-proof contract the autarky punishment is not credible. Instead, deterring a deviation of several players requires considering the minmax payoff sets of all the coalitions they could form. Different allocations on the Pareto frontier defined by these sets involve trade offs in the marginal utilities of the players. The constrained optimum will therefore choose the allocation on the frontier that minimizes the movement in the marginal utility ratios and by implication the departure from first-best. Roughly speaking, an individual's payoff in this allocation will be larger the higher his consumption in the previous period. Hence, consumption and continuation payoffs in the coalition-proof contract retain some memory of the past even when enforcement constraints are binding – potentially shedding some light on some of the empirical results listed above.

The fact that the coalition-proof contract no longer displays amnesia when enforcement constraints are binding implies reduced form restrictions on the consumption series generated by the coalition-proof contract. The current consumption of constrained agents depends on the history of shocks, as summarized by the marginal utility ratios in the previous period, of the entire set of constrained agents both directly and interacted with their current income realization. Neither of these properties is exhibited by the consumption series generated by a sub-game perfect risk-sharing contract and this allows us to formulate an empirical test for the presence of endogenous group formation based on comparing agents' relative consumption shares. Dubois (2005) derives a test for the likelihood of endogenous insurance group sizes with heterogenous preferences under the assumption of stationary symmetric risk-sharing. Chaudhuri et al. (2005) and Ahn et al. (2006) present tests in

an experimental setting. In contrast, the test presented here is valid in a non-experimental dynamic setting. It relies on being able to identify constrained individuals and their consumption shares in a risk-sharing arrangement in each period and is thus similar in spirit to Mazzocco (2007) and Krueger et al. (2008). To investigate the performance of the test, we use simulated model solutions to derive the size and power of the test. For low levels of measurement error, the test performs extremely well in distinguishing exogenous and endogenous group formation in a risk-sharing arrangement. As measurement error increases, the test becomes somewhat less reliable, however the version that tests for significance of the current income realization of other constrained agents interacted with their past history still fares reasonably well. Taken together, the results give us confidence that the test is indeed able to identify whether coalitional deviations pose a threat to informal risk-sharing arrangements at least when consumption and income data are well measured.

The paper proceeds as follows: Section 2 outlines the model, its assumptions and defines the coalition-proof equilibrium set. Section 3 presents the analytical results and characterizes the coalition-proof risk-sharing contract and its dynamics. Section 4 derives empirical results and presents a test for the endogeneity of group formation, which is implemented using simulated data based on computed model solutions. Section 5 concludes.

2   The Model

Consider a set of $\mathbb{N} = \{1, ..., n\}$ households in a community. Each period $t = 0, 1, ...\infty$, household i receives an income $y^i(s_t)$, where $s = 1, ..., S$ is the independently and identically distributed state of nature, which occurs with probability π_s. All households have an identical twice continuously differentiable utility function $u(c^i(s_t))$ where $c^i(s_t)$ is consumption. Households are risk-averse, infinitely lived and discount the future with common discount factor β. Perfect information is assumed and there are no opportunities for storage. 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.

The only equilibrium in a static one-shot game in this environment is the autarky solution, in which each individual consumes his own income, $c^i(s_t) = y^i(s_t)$. In a repeated game, it is possible to sustain insurance contracts that entail some smoothing of consumption by varying future payoffs with current play. In order to characterize these contracts, we must formally describe the game played by the n agents following Asheim and Strand (1991). For each stage t of the infinitely repeated game G, an insurance contract for a group of size n specifies a set of net transfers $(\tau^i(s_t))_{i=1}^n$. If all players make the stipulated transfers, consumption is $c^i(s_t) = y^i(s_t) - \tau^i(s_t)$ for $i = 1, .., n$.

If not, $c^i(s_t) = y^i(s_t)$ for all agents. At stage t, the history of the game consisting of previous transfers and previous states and the current state are common knowledge. At every date $t \geq 0$, G is characterized by a set of t-histories H_t, given by $H_t = \{H_{t-1}, s_t\}$ and $H_0 = \{s_0, (\tau^i(s_0))_{i=1}^n\}$, and an n-tuple of stage t strategies, $(\omega_t^i)_{i=1}^n : H_t \Rightarrow \mathbb{R}^n$, which are mappings from the set of histories to the set of actions. A strategy for player i, ω^i, is a sequence of stage t maps, $\omega_t^i, t = 0, 1, ...\infty$, and we require that strategies are symmetric in the sense that permuting the past history permutes current actions. Player's payoffs are derived from the consumption streams generated by the strategy profile $\omega = \{\omega^1, \omega^2, ..., \omega^n\}$, which consists of the n-tuple of players' strategies. For any history $h_t \in H_t$, contract ω and group size m, the lifetime utility of an individual from time t onwards is

$$U^i(\omega, h_t, m) = u(c^i(s_t)) + E \sum_{\nu=t+1}^{\infty} \beta^{\nu-t} u(c^i(s_\nu)). \tag{1}$$

As shown in Abreu et al. (1990), perfect equilibria of repeated games can be formulated recursively using dynamic programming principles, and therefore the problem of finding an infinite sequence of transfers is equivalent to finding current period actions and next period continuation values such that

$$U^i(\omega, h_t, m) = u(c^i(s_t)) + \beta \sum_{s=1}^{S} \pi_s U^i(\omega, h_t, s_{t+1}, m) \tag{2}$$

for each player i, where the promised values $U^i(\omega, h_t, m)$ summarise the relevant information about each player's history.

To solve for the constrained efficient contract that is self-enforcing, we need to determine which credible threats players can make and what punishments are available to deter deviation from cooperation. In what follows, we briefly revisit the result that reversion to autarky is an optimal penal code (see Abreu (1988)) that can be used to support all efficient sub-game perfect equilibria. We then argue that this punishment strategy is too weak because it does not account for the behaviour of coalitions. For an agreement to be truly self-enforcing, it must be immune to deviations of all coalitions, not just single players. We show how to construct a punishment strategy that is coalition-proof in the sense that neither the grand coalition of n players nor any subcoalition would want to deviate with the added consistency requirement that threatening to deviate is only credible if the deviating coalition is itself robust to further deviations. Since the punishment strategy – if it exists – is an optimal penal code in the spirit of Abreu, it follows that any coalition-proof insurance contract can be supported by it and we will make use of this fact when characterising the constrained-efficient insurance contract and its dynamics in the next section.

Let ω^{aut} denote the strategy in G that for any t and h_t stipulates $\tau(h_t) = 0$. Denote the expected

value of playing the static autarky strategy from period t onwards as

$$\underline{V}_t^i = E \sum_{t=0}^{\infty} \beta^t u(y^i(s_t)) \tag{3}$$

for all players $i = 1, ..., n$. This is the only self-enforcing payoff for an individual and hence the set of stable payoffs for a singleton contains just one element $\mathbb{V}^*(1) = \{\underline{V}\}$.

Since the autarky payoff can be enforced unilaterally by any player i, it follows that $\underline{V}_t^i$ is the minmax payoff for player i. Moreover, the n-tuple $\omega^{aut} = \{\omega^{1,aut}, ..., \omega^{n,aut}\}$ is a sub-game perfect equilibrium after any history h_t. Together, this implies that threatening each player with ω^{aut} following a deviation constitutes an optimal penal code that can hence be used to support all cooperative sub-game perfect equilibria.

Such a contract is not renegotiation-proof, however. While there is no unique definition of renegotiation-proofness, the literature takes as its starting point the notion of Pareto perfection, which requires that players will always negotiate to an equilibrium on the Pareto frontier of the efficient set (see Farrell and Maskin (1989) and Pearce (1988)). Hence, a punishment path of reversion to autarky following a deviation is not credible because it requires that mutual gains from insurance are forgone during the punishment phase. In the context of 2-player games, Asheim and Strand (1991) and Kletzer and Wright (2000) have shown that the punishment path of reversion to permanent autarky can be replaced by a renegotiation-proof punishment path. This path supports exactly the same set of equilibria as the autarky punishment but instead allocates the maximum surplus to the non-deviating player. In what follows, we will use a strategy similar to theirs in the construction of optimal penal codes that are coalition-proof.

Moving from the 2-player to the n-player environment, the concept of renegotiation proofness is not entirely satisfactory, because it only considers the possibility of renegotiation by the grand coalition of n players, but does not allow for renegotiation by any proper subset of $\mathbb{N}$. To resolve this inconsistency, Bernheim et al. (1987) introduce the notion of coalition-proofness, which requires that no coalition ever plays a Pareto dominated strategy in any sub-game. In the context of this paper, this implies that any subcoalition of players can achieve a Pareto improvement by abandoning the original punishment of autarky and instead continuing risk-sharing. Therefore an insurance scheme is defined as stable if no history of states exists, for which a stable subgroup could credibly deviate from the arrangement, consume autarky income in the period of deviation and then continue insurance within the subgroup.[3] Credibility means that a deviating subgroup must itself be immune to further

---

[3]The assumption that a deviating sub-group consumes autarky income during the period of deviation deserves some further discussion. It may be more plausible to assume that agents who deviate continue insurance in the deviating subgroup during the period of deviation. Qualitatively, this makes little difference. It will, however, increase the

deviations (see Bernheim et al. (1987) and Bernheim and Ray (1989)).

To define the set of self-enforcing or stable contracts, we follow Genicot and Ray (2003) and assess stability of a risk-sharing group of size n recursively by first examining the stability of groups of size $1, ..., n-1$ and then checking whether a risk-sharing contract exists for a group of size n that is robust with respect to deviations by stable subgroups. Suppose that we have defined coalition-proof sets of expected payoffs $\mathbb{V}^*(m)$ for all $m = 1, ..., n-1$, where each element in $\mathbb{V}^*(m)$ is a vector of size m with components $V = \{V^1, V^2, ..., V^m\}$ and the requirement of coalition-proofness implies that the vectors in $\mathbb{V}^*(m)$ must be Pareto efficient in the set of coalitionally self-enforcing agreements. Stability of a risk-sharing arrangement of size n is then expressed in the following two conditions stated in Genicot and Ray (2003):

Definition 1 *V is a stable payoff vector for n, if the following two conditions are met:*

[PARTICIPATION] For no history h_t is there a sub-group of individuals $m \leq n$ and a stable payoff vector $V \in \mathbb{V}^*(m)$ such that $V^i(n) < V^i(m)$ for all $i = 1, ..., m$.

[ENFORCEMENT] For no history h_t is there a subset $\mathbb{M}$ of individuals of size $m \leq n$ and a stable payoff vector $V \in \mathbb{V}^*(m)$ such that for all $i \in m$

$$u(y^i(s_t)) + \beta V^i > u(c^i(s_t)) + \beta V^i(h_{t+1}, n). \tag{4}$$

Definition 1 states that in any stable risk-sharing contract for n players, no coalition of size m can be forced to accept a payoff $V^1(n), ..., V^m(n)$, such that there exists a vector in $\mathbb{V}^*(m)$ that makes all of its members better off. In essence, this describes the set of minmax payoffs for each subcoalition of $\mathbb{N}$. While these minmax payoffs are essential ingredients in the definition of the equilibrium set, the participation and enforcement constraints listed above are somewhat too complex to be useful in an analytical characterization of the coalition-proof equilibrium contract – not least because they state conditions that must NOT be true for the equilibrium set and because at each stage t, the minmax threat of a coalition consists of a potentially infinite set rather than a single payoff vector as is the case for the individually rational contract.

We now show how to construct coalition-proof punishment strategies that deter all coalitional deviations to any point in the minmax set by switching play to a single point in this set following a deviation. We need the following notation. Let $\mathbb{J}$ denote the set of proper subsets of $1, ..., n$

---

threat of coalitional deviations meaning that even less risk-sharing can be sustained in equilibrium. Assuming that deviating agents do not engage in risk-sharing during deviation can perhaps be interpreted as pessimism on the part of the deviating sub-group. In any case, the majority of results in this paper do not hinge on this assumption.

and denote an element of $\mathbb{J}$, which is a coalition, as $J \in \mathbb{J}$. Label the members of a coalition J with cardinality k by $j_1, ..., j_k$. Finally, let $m(k)$ denote the largest stable group size not exceeding k.[4] Consider any combination of $n-1$ agents with members $j_1, ..., j_{n-1}$. We will now construct a punishment path that is coalition-proof and deters an individual deviation by agent j_1, a deviation by a coalition consisting of agents j_1 and j_2, a deviation by a coalition consisting of agents j_1, j_2 and j_3 and so on up to a joint deviation of $j_1, ..., j_{n-1}$. Suppose that after any of the deviations listed above, play switches to a punishment path with payoffs V that allocates $V^1 = \underline{V}^1$ to agent j_1, $V^2 = V^2(V^1)$ to agent j_2, $V^3 = V^{m(3)}(V^1, ..., V^{m(3)-1})$ to agent j_3, and so on up to j_{n-1} with the remaining agent receiving $V^n = V^n(V^1, ..., V^{n-1})$. We will now apply the arguments found in Asheim and Strand (1991) and Kletzer and Wright (2000) to show that a punishment path with payoffs V is an optimal penal code in the spirit of Abreu. That is, it is a coalition-proof equilibrium in the set of minmax payoffs.

By the recursive definition of stability, it follows that $V^1 \leq V^2 \leq ... \leq V^{n-1}$. Furthermore, assume that $V^n \geq V^{n-1}$, which is a necessary condition for the existence of a coalition-proof punishment path. First, it is easily seen that the constructed payoffs satisfy the stability conditions in Definition 1. Specifically, $\underline{V}^1$ is the minmax payoff for an individual agent. Collectively, any pair of players can negotiate to any point on the frontier of a group of size 2, which is given by the set $[\underline{V}^1, \bar{V}^1]$ and $[V^2(\underline{V}^1), V^2(\bar{V}^1)]$. Hence, this is the set of minmax payoffs for a group of size 2. Any group of three agents can collectively renegotiate to an outcome in the set $\mathbb{V}^*(3)$, which consists of all combinations of V^1, V^2, V^3 on the Pareto frontier of a 3-agent group (if such a group is itself stable with respect to further deviations). By the same argument, it can be shown that the payoff vectors are elements of the minmax set for any coalition consisting of agents $j_1, ..., j_k$ for $k = 1, ..., n-1$. By construction, there is no permutation of agents such that their joint payoff is dominated by a payoff in their minmax set.

Second, the punishment path is coalition-proof. The punishment path lies on the Pareto frontier of $\mathbb{N}$, and so it is renegotiation-proof in the sense that the grand coalition would not want to abandon it collectively. Moreover, by construction, the payoffs also lie on the Pareto frontier of any stable group that is smaller than n, and there is therefore no subcoalition that would find it profitable to collectively abandon the punishment path. Hence, the punishment path satisfies the requirements of coalition-proofness.

Taken together this implies that V is an optimal penal code and by symmetry, an analogous

---

[4]This notation is necessary, because a cooperative coalition-proof equilibrium does not necessarily exist for all group sizes $k = 3, ..., n-1$ (see Genicot and Ray (2003)). Of course a cooperative equilibrium for groups of size 2 may not exist either if the discount factor is too small. This is, however, not a consequence of the coalitional deviations modelled in this paper, and so we assume that the discount factor is always large enough to support some cooperative play for two agents.

punishment path can be constructed for any combination of agents in $\mathbb{N}$. It follows that all coalition-proof equilibria can be sustained by such punishment paths – when they exist. Moreover, V is simple in the sense that it can be used after any history to deter the deviations described above. Most importantly, V can deter a coalitional deviation to any of the potentially infinite number of elements in the minmax set of this coalition by using a punishment path that switches to a single element in this set. This greatly simplifies the analysis of the constrained-efficient equilibrium contract.

To appreciate these points, it is useful to consider a group of size 3 and suppose that ex-post continuation payoffs in period t both on and off the equilibrium path are

$$U_{s_t}^1 = u(y^1(s_t)) + \beta \underline{V}^1 \tag{5}$$

$$U_{s_t}^2 = u(y^2(s_t)) + \beta V^2(\underline{V}^1) \tag{6}$$

$$U_{s_t}^3 = u(y^3(s_t)) + \beta V^3(\underline{V}^1, V^2(\underline{V}^1)). \tag{7}$$

Given this, there can be no deviation to a different $\widetilde{V} = \{\widetilde{V}^1, \widetilde{V}^2(\widetilde{V}^1), V^3(\widetilde{V}^1, V^2)\}$ in $\mathbb{V}^*(3)$ that would render V unstable, as this would imply

$$u(y^1(s_t)) + \beta \widetilde{V}^1 > u(y^1(s_t)) + \beta \underline{V}^1 \tag{8}$$

$$u(y^2(s_t)) + \beta \widetilde{V}^2 \lessgtr u(y^2(s_t)) + \beta V^2(\underline{V}^1) \tag{9}$$

$$u(y^3(s_t)) + \beta \widetilde{V}^3(\widetilde{V}^1, \widetilde{V}^2) \gtrless u(y^3(s_t)) + \beta V^3(\underline{V}^1, V^2(\underline{V}^1)) \tag{10}$$

by the definition of Pareto efficiency. Since payoffs will only be renegotiated if all members of a deviating subgroup are better off following a joint deviation, threatening to renege and continue with $\widetilde{V}$ is not credible once continuation payoffs have been set equal to $V = \{\underline{V}^1, V^2(\underline{V}^1), V^3(\underline{V}^1, V^2(\underline{V}^1))\}$. Similarly, agent 1 and 2 will not deviate jointly to any other payoff vector $\widetilde{V}$ in their minmax set $\mathbb{V}^*(2)$, since this would imply

$$U_{s_t}^1 \geq u(y^1(s_t)) + \beta \widetilde{V}^1 \tag{11}$$

$$U_{s_t}^2 \leq u(y^2(s_t)) + \beta V^2(\widetilde{V}^1). \tag{12}$$

This point is illustrated in Figure 1, which plots the ex-ante Pareto frontier for $\mathbb{V}^*(2)$. Of course, we could equally well have chosen $\widetilde{V}$ as the punishment payoffs to deter a joint deviation by agent 1 and 2. By construction $\widetilde{V}$ is coalition-proof and an element of the minmax set of coalitions of size 2. It is hence an optimal penal code to deter pairwise deviations (but not individual ones),

11

Figure 1: Ex-ante Pareto frontier in a group of size 2.

even though $\widetilde{V}^1$ is not an element of the minmax set for singletons. This illustrates the fact that an optimal penal code to punish deviations of a coalition of size k does not necessarily have to be constructed such that $V^1, ..., V^{k-1}$ are elements of the minmax sets for coalitions of size $1, ..., k-1$. Finally, note that both V and $\widetilde{V}$ are optimal simple penal codes because they can be used to support cooperation after any history. Moreover, with continuation payoffs given by $U_{s_t}^1$, $U_{s_t}^2$ and $U_{s_t}^3$, there is no other 2-player coalition that could credibly deviate. Hence to deter all coalitional deviations, it is sufficient to satisfy the enforcement constraints for a single point on the Pareto frontier of $\mathbb{V}^*(3)$.

The results on the optimal punishment paths are summarized in Proposition 1.

Proposition 1

1. *In a sub-game perfect efficient contract, the optimal punishment path consists of permanent reversion to autarky.*

2. *In a coalition-proof efficient contract, the optimal punishment path V lies on the Pareto frontier of $\mathbb{V}^*(m)$ for every stable subcoalition of size $m \leq n$. This punishment path can support all*

equilibrium payoffs for which $\exists \widetilde{V} \in \mathbb{V}^*(m)$ *such that*

$$
\begin{aligned}
U_{s_t}^1 &\geq u(y^1(s_t)) + \beta \widetilde{V}^1 \\
U_{s_t}^2 &\geq u(y^2(s_t)) + \beta \widetilde{V}^2 \\
\vdots &\geq \vdots \\
U_{s_t}^m &\geq u(y^m(s_t)) + \beta \widetilde{V}^m.
\end{aligned}
$$

3 Efficient Contracts

Having described behaviour off the equilibrium path, we can now characterize the symmetric coalition-proof equilibrium contract and its dynamics. The constrained-efficient contract is found by solving a dynamic programming problem. In each state and for each date, coalition-proof punishments are used to enforce cooperation. The constrained dynamic programme solves for the Pareto frontier in an insurance group of size n denoted by U_s^n, treating the promised utilities of agent $1, ..., n-1$ as state variables, which summarise the relevant information about each player's history.[5] This problem can be thought of as a social planner having promised agent $1, ..., n-1$ utility $U_s^1, ..., U_s^{n-1}$, which the planner delivers by choosing current consumption and continuation utilities in a way that maximises agent n's payoff and satisfies the self-enforcing constraints.

As a first step toward formulating the constrained dynamic programme, we show how to write the self-enforcing constraints as a set of inequality constraints that embody the optimal penal codes constructed in the previous section. We have shown in Proposition 1 that any point on the Pareto frontier of a stable subgroup of size m can be used as a punishment to deter deviations of a sub-coalition of this size. While this is sufficient to describe the set of coalition-proof equilibria, we must be more precise in order to characterize the properties of the equilibrium analytically. To see this, revisit the example depicted in Figure 1 and suppose that following history h_t,

$$
U_r^1 \geq u(y_r^1) + \beta \widetilde{V}^1 \tag{13}
$$

$$
U_r^2 \geq u(y_r^2) + \beta V^2(\widetilde{V}^1) \tag{14}
$$

are in fact the constrained-efficient payoffs for agent 1 and 2 on the equilibrium path in a group of size n, i.e. those that maximize $U_r^n(U_r^1, U_r^2, ..., U_r^{n-1})$. If these equilibrium payoffs were supported

---

[5]Since the physical environment is the same in every period, we can drop the explicit dependence on time t.

by $\widetilde{V}$, then we can write down a standard Kuhn-Tucker Lagrangean where the multipliers on the enforcement constraints have the usual interpretation. If we use $\{\underline{V}, \bar{V}\}$ to support these payoffs then such a formulation is not possible because (14) may be violated even though all deviations by a pair are deterred. Hence, $\widetilde{V}$ is preferable in terms of analytical tractability. It should be emphasized, however, that this substitution is purely a matter of convenience. It does not affect any of the properties of the solution, it merely helps to characterize it.

Conceptually, we therefore decompose the problem of finding the constrained-efficient contract into two steps. In the first step, the planner solves for the optimal penal code $\widetilde{V}$ after history h_t that maximises the dynamic programme given that each agent's continuation payoff on the equilibrium path will be equal to the discounted punishment payoffs augmented by one period of autarky consumption if enforcement constraints are binding. In the sub-game perfect contract, $\widetilde{V}$ trivially corresponds to the autarky payoff for each agent. In the coalition-proof contract, the planner must choose $2^n - 2$ coalition-proof payoff vectors following each state of the world r: one for each combination of agents of size $k = 1, ..., n - 1$. In the second step, the planner chooses consumption and continuation payoffs that can be supported by the punishment paths determined in the first step.

This discussion gives rise to the following dynamic programme

$$U_s^n(U_s^1, U_s^2, ..., U_s^{n-1}) = \max_{((U_r^i)_{r=1}^S)_{i=1}^{n-1}, (c_s^i)_{i=1}^n, (((\widetilde{V}_r^{jg,J})_{g=1}^{m(k)})_{J \in \mathbb{J}})_{r=1}^S \in \mathbb{V}^*(m(k))} u(c_s^n) + \beta \sum_{r=1}^S \pi_r U_r^n (U_r^1, ..., U_r^{n-1})$$

(15)

subject to a set of promise-keeping constraints

$$u(c_s^i) + \beta \sum_{r=1}^S \pi_r U_r^i \geq U_s^i \quad \forall i \neq n,$$

(16)

an aggregate resource constraint

$$\sum_{i=1}^n y_s^i \geq \sum_{i=1}^n c_s^i$$

and a set of enforcement constraints for each coalition $J \in \mathbb{J}$ and each state $r = 1, ..., S$ in the next

period

$$U_r^{j_1} \geq u(y_r^{j_1}) + \beta \widetilde{V}_r^{1,J}, \tag{17}$$

$$U_r^{j_2} \geq u(y_r^{j_2}) + \beta \widetilde{V}_r^{2,J}, \tag{18}$$

$$... \geq ... \tag{19}$$

$$U_r^{j_{m(k)}} \geq u(y_r^{j_{m(k)}}) + \beta \widetilde{V}_r^{m(k),J} \tag{20}$$

$$... \geq ... \tag{21}$$

$$U_r^{j_k} \geq u(y_r^{j_k}) + \beta \widetilde{V}_r^{k,J}, \tag{22}$$

where the elements of coalition J are labelled such that $\widetilde{V} = \{\widetilde{V}_r^1, ..., \widetilde{V}_r^k\}$ is arranged in increasing order of payoffs and the subscript r indicates that V_r^g is the expected punishment payoff chosen after state r has been realized. Furthermore, $U_r^{j_g} = U_r^{j_g}(U_r^1, ..., U_r^{n-1})$ if $j_g = n$, and it is understood that $V_r^{m(k),J}, ..., V_r^{k,J}$ are set equal to $V^{m(k)}(V^{1,J}, ..., V^{m(k)-1})$.[6]

In order to characterize the solution to this dynamic programming problem analytically, we need to establish differentiability of the value function. For the multilateral sub-game perfect contract which is stable with respect to individual deviations, the results established in Lemma 1 in Thomas and Worrall (1988) apply mutatis mutandis. Since the set of sustainable contracts is not convex in the coalition-proof contract we cannot rely on this property to establish the concavity and differentiability of the value function. Instead, we follow Pavoni (2006) and apply Theorem 1 in Milgrom and Segal (2002) to establish that the value function is differentiable at all optimal points. As a result, the solution of the dynamic programme can be characterized by the usual first-order conditions

---

[6]A possible ambiguity in the above formulation is that the punishment payoffs may specify that a player receives different payoffs depending on which set of coalitional enforcement constraints is considered. This ambiguity is resolved (somewhat informally) in the following manner. Again, consider an example of a 3-player group and suppose we have found punishment payoffs for all possible 2-player coalitions. Further, suppose that following history h_t, player 1 and 2 have a profitable deviation in state r if first-best allocations $U_r^{1,fb}$ and $U_r^{2,fb}$ following this history were actually implemented. That is, $\exists \{V^1, V^2\} \in \mathbb{V}^*(2)$ such that

$$U_r^{1,fb} < u(y_r^1) + \beta V^1$$
$$U_r^{2,fb} < u(y_r^2) + \beta V^2$$
.

Then, we satisfy the enforcement constraints for the 2-player coalition consisting of agent 1 and 2 with equality. That is, we adjust payoffs so that

$$U_r^1 = u(y_r^1) + \beta \widetilde{V}_r^{1,\{12\}} \tag{23}$$

$$U_r^2 = u(y_r^2) + \beta \widetilde{V}_r^{1,\{12\}} \tag{24}$$

Following from the discussion in Section 2, a 3-player group is stable if and only there are no other profitable 2-player deviations for these payoffs. If that is the case, we can ignore the enforcement constraints and punishment payoffs for {13} and {23} when solving for the constrained-efficient allocations. Doing this guarantees that any player can have at most one enforcement constraint that binds with equality in each state r and the enforcement constraints in (17)-(22) are to be interpreted in this way. That is, different sets of coalitional enforcement constraints are applied depending on which subcoalition of players can threaten a profitable deviation if first-best payoffs following h_t were actually implemented.

and the envelope theorem holds. These technical results are summarized in Lemma 1:

Lemma 1

1. *The Pareto frontier $U_s^n(U_s^1, ..., U_s^{n-1})$ is differentiable for any $(U_s^i)_{i=1}^{n-1}$ on the range of the optimal policy correspondence.*

2. *At any point of differentiability, if $U_s^{i'} > U_s^i$, then $\frac{\partial U_s^n}{\partial U_s^i} > \frac{\partial U_s^n}{\partial U_s^{i'}}$.*

3. *The enforcement constraints (17)-(22) satisfy the constraint qualification.*

Proof of Lemma 1

See Appendix.

Having shown that the value function is differentiable at the optimum and that the enforcement constraints satisfy the constraint qualification, we can apply the Kuhn-Tucker theorem. This implies that there exist non-negative multipliers associated with the promise-keeping and enforcement constraints, denoted respectively $(\gamma^i)_{i=1}^{n-1}$ and $(\beta \pi_r \gamma^{j_g} \phi_r^{j_g, J})_{g=1}^k$ for every coalition $J \in \mathbb{J}$ and each $r = 1, ..., S$. The first-order conditions of the dynamic programming problem are

$$\frac{u'(c_s^n)}{u'(c_s^i)} = \gamma^i \quad \forall i \neq n \tag{25}$$

and

$$-\frac{\partial U_r^n}{\partial U_r^i} = \gamma^i \frac{1 + \phi_r^i}{1 + \phi_r^n} \quad \forall r \in S, \quad \forall i \neq n \tag{26}$$

where it is understood that each agent i can have a constraint that binds with equality for at most one deviating subcoalition. The $n - 1$ envelope conditions are given by

$$\gamma^i = -\frac{\partial U_s^n}{\partial U_s^i} \quad \forall i \neq n. \tag{27}$$

Advancing the envelope condition forward by one period, we have

$$\gamma_r^i = -\frac{\partial U_r^n}{\partial U_r^i} = \gamma^i \frac{1 + \phi_r^i}{1 + \phi_r^n} = \frac{u'(c_r^n)}{u'(c_r^i)} \quad \forall r \in S, \quad \forall i \neq n. \tag{28}$$

That is, γ^i measures the trade-off between agent n's and agent i's discounted lifetime utility in the current period, and γ_r^i measures the trade-off in the future period when the state of the world is r.

More generally, the first order conditions show that a constrained-efficient contract can be characterized in terms of the evolution of the vector $\gamma = \{\gamma^1, \gamma^2, ..., \gamma^{n-1}\}$. From the first order condition,

each γ^i measures the ratio of marginal utilities between agent i and n and from the envelope condi-
tion, each γ^i measures the rate at which the discounted utility of agent i can be traded off against
the discounted utility of agent n keeping everybody else's utility constant. Once the state of nature
r is known, the new value for each γ^i_r is determined from equations (26) and (28).

The first-order conditions have a very intuitive interpretation. First, consider a first-best risk-
sharing contract. This must satisfy Borch's rule, which states that the ratio of marginal utilities of
income is constant across all states and time periods. Hence, the initial division of aggregate income
is implemented forever after regardless of how the history of the game unfolds from period $t = 0$
onwards. Now consider the constrained-efficient contract and suppose that enforcement constraints
have been slack up to some period t, so that all $(\phi^i_\nu)^n_{i=1}$ are zero for $\nu = 0, ..., t-1$. Then, (28)
says that the ratio of marginal utilities remains constant and from the envelope condition so does
the rate at which continuation utilities are traded off. Moreover, the marginal utility ratio is equal
to its initial value, which implies that aggregate resources are divided among agents in the same
way in each period $\nu = 0, ..., t-1$. In other words, the first-best risk-sharing contract and the
constrained-efficient contract are equivalent until enforcement constraints bind for the first time.

Suppose this occurs in period t. To fix ideas, assume that agents have identical bargaining
power. In this case, it seems plausible that the risk-sharing arrangement has implemented an equal-
sharing rule. Further, suppose that agent n experiences a negative shock in period t, while agent
i does not. Implementing first-best risk-sharing would require that consumption between the two
agents is equalized. However, insisting on first-best transfers may have the consequence that agent
i finds it in his interest to defect from the arrangement. Suppose that this is indeed the case. That
is, agent i has a binding enforcement constraint in period t and $\phi^i_t > 0$. Then (28) states that
$\gamma^i_t > \gamma^i_{t-1}$ and consequently $\frac{u'(c^n_t)}{u'(c^i_t)} > \frac{u'(c^n_{t-1})}{u'(c^i_{t-1})}$. By the concavity of the utility function, this implies
that consumption growth for agent i is greater than for agent n. Equally, the envelope condition
tells us that $\frac{\partial U^n_t}{\partial U^i_t} < \frac{\partial U^n_{t-1}}{\partial U^i_{t-1}}$ and from Lemma 1, this requires that agent i's continuation payoff is
increased at the expense of agent n in the face of a binding constraint. Together, this implies firstly,
that the transfer from agent i to agent n does not suffice to equalize consumption between the two
agents, and secondly, that the division of aggregate resources is shifted in favour of agent i.

The most interesting property of a dynamic risk-sharing contract – and the one that distinguishes
it from a static limited commitment contract – arises in the period after enforcement constraints
have been binding. Suppose that all enforcement constraints are slack in period $t + 1$, so that all
$\phi^i_{t+1} = 0$. Under these circumstances, a static contract would revert to the initial sharing rule (see
Coate and Ravallion (1993) and Ligon et al. (2002)). In the dynamic contract, (28) states that the

ratio of marginal utilities is to be kept constant even when the previous period has been a constrained optimum. This implies that the increase in agent i's surplus from the risk-sharing arrangement is to be maintained until some period $t + \nu$ in the future when someone else experiences a binding constraint. In other words, the constrained-efficient contract is history-dependent in the sense that what is considered 'first-best' changes as the history of the game unfolds.

How can we understand this result? In contrast to the static contract, the dynamic contract allows agents to trade future claims to consumption in exchange for consumption today. More specifically, agent i is induced to make a payment to agent n in period t in exchange for receiving a larger surplus from the arrangement, not just in the current period but also in the future. This promise of at least partial repayment in the future means that agent i's current consumption can be lower and his transfer to agent n higher than it would be in a stationary setting. In other words, more risk-sharing can be sustained because the dynamic risk-sharing contract allows shocks that cannot be smoothed across states because of limited commitment to be smoothed over time. The fact that this intertemporal smoothing is potentially spread over several periods should not come as a surprise. Of course, it would be possible to reduce agent n's consumption in period $t + 1$ further by forcing him to make a higher repayment to agent i and compensate him with a larger surplus in period $t + 2$. However, this would increase the variability of consumption over time and can therefore not be optimal from the point of view of a risk-averse agent. Intuitively, the constrained-efficient contract introduces an element of quasi-credit (see Fafchamps (1999)): risk-sharing is achieved through informal transfers and loans and future repayment schedules depend on shocks affecting lender and borrower.

Having reviewed the general rationale behind the constrained-efficient contract, we now show what distinguishes the sub-game perfect from the coalition-proof risk-sharing contract. The main difference between the two contracts arises when enforcement constraints are binding. In a sub-game perfect contract, the current income realization of a constrained agent in state r fully determines his utility from period t onwards. The history of income realizations up to time t is therefore irrelevant in determining the payoffs of constrained agents. In contrast, this is not the case in the coalition-proof contract. The current income realizations and history of shocks of all constrained agents together determine their payoffs when enforcement constraints are binding. These results are summarized in the following proposition and proved in the Appendix.

Proposition 2

1. In a sub-game perfect efficient contract, continuation payoffs, U_r^i, are not history-dependent

for agents with binding enforcement constraints.

2. *In a coalition-proof contract, continuation payoffs U_r^i are history-dependent for agents with binding enforcement constraints.*

3. *In a sub-game perfect efficient contract, the continuation payoff U_r^i of a constrained agent depends only on his own current income realization.*

4. *In a coalition-proof contract, the payoff U_r^i of a constrained agent depends on his own current income realization as well as that of other constrained agents. The effect of y_r^j on U_r^i is greater the smaller $|\gamma^i - \gamma^j|$.*

Proof of Proposition 2: See Appendix.

The ideas behind Proposition 2 can be explained using an example of a 3-player game with just two possible income realizations for each player, denoted by y^l and y^h. Suppose agents have implemented an equal-sharing rule prior to period t. In period t, agent 2 and 3 are called upon to help agent 1, who is the only one to experience a low income realization. If the equal-sharing allocation is not self-enforcing, then second-best requires that agent 2 and 3's consumption and continuation payoffs are adjusted so that they are just indifferent between cooperation and defection. The sub-game perfect contract assumes that players will cooperate as long as their payoff under cooperation is at least as large as what they can achieve in isolation. Hence, it is sufficient – and optimal – to equate player 2 and 3's continuation payoffs with their autarky payoffs. That is, $U_t^2 = U_t^3 = u(y^h) + \beta \underline{V}$. Since a player's autarky payoff depends only on his current income realization and not on the history of the game, the same will be true for his continuation payoff when he is constrained in the sub-game perfect contract. In other words, what is considered 'second-best' does not change as the history of the game unfolds in a sub-game perfect contract.

Now consider a coalition-proof contract. Suppose that up to and including period t, the history in the coalition-proof contract has been identical to the one described in the previous paragraph. Then we know from the first-order conditions that the marginal utility ratio between agent 1 and 2 has decreased in period t and agent 1 receives a smaller share of resources than agent 2, so that $\frac{u'(c_t^2)}{u'(c_t^1)} < \frac{u'(c_{t-1}^2)}{u'(c_{t-1}^1)} = 1$. This will remain the case as long as enforcement constraints are slack, because agent 2 has been induced to help agent 1 in period t by being offered a reward in the form of partial repayment in the future. Now suppose agent 1 and 2 both experience a high income realization in period $t+1$ and are called upon to help agent 3. Moreover, assume that the first-best allocation following this history is not enforceable. The coalition-proof contract allows for the

fact that following a joint deviation, agent 1 and 2 can continue to share risk rather than remain in autarky. Hence, agent 1 and 2's consumption must be adjusted so that they are just indifferent between cooperation and jointly defecting to a point on the Pareto frontier of a pair, $\mathbb{V}^*(2)$. Consider the point $\{\underline{V}, \overline{V}\} \in \mathbb{V}^*(2)$ depicted in Figure 1 in Section 2 and suppose that this allocation minimizes the total consumption of agent 1 and 2, $c_r^1 + c_r^2$. This implies that the insurance transfer to agent 3 is maximised for this allocation. By symmetry, the same is true for the allocation in which agent 1 receives $\overline{V}$ and agent 2 receives $\underline{V}$. As a result, the extent of risk-sharing that can be sustained in the current period will not be affected by which allocation is chosen. Looking backwards however, allocating the larger continuation payoff to agent 2 and the smaller continuation payoff to agent 1, so that $U_{t+1}^2 = u(y^h) + \beta\overline{V} > U_{t+1}^1 = u(y^h) + \beta\overline{V}$, has a positive effect on welfare because it increases the scope for positive transfers from agent 2 to agent 1 in period t by promising a higher reward to agent 2 in the future. Hence, following any histories in which agent 2 has been a net giver to agent 1, it is optimal to deter a joint deviation by giving agent 2 a relatively higher payoff than agent 1. Consequently, the coalition-proof contract exhibits history-dependence even when enforcement constraints are binding.[7]

To understand the intuition behind claim 4. in Proposition 1, consider again the three player game described above. We know from the argument made in the previous paragraph that agent 2's continuation payoff exceeds agent 1's continuation payoff in $t + 1$ if agent 1 and 2 both have a high income realization and $\frac{u'(c_t^2)}{u'(c_t^1)} < 1$. Instead, suppose that agent 1 has a low income realization in period $t + 1$. We now argue that the allocation in which agent 1 receives $U_{t+1}^1 = u(y^l) + \beta\underline{V}$ and agent 2 receives $U_{t+1}^2 = u(y^h) + \beta\overline{V}$ is no longer optimal when they are jointly constrained and $\frac{u'(c_t^2)}{u'(c_t^1)}$ is close to 1. To see this, assume that the difference between the income realizations is large enough that $u(y^l) + \beta\overline{V} < u(y^h) + \beta\underline{V}$.[8] Then it follows from part 2. in Lemma 1 that switching to the allocation in which agent 1 receives the highest payoff from risk-sharing in a group of size 2, $\overline{V}$, and agent 2 receives the lowest payoff, $\underline{V}$, increases the payoff of the unconstrained agent. This is the case, because the fall in U_t^3 that results from increasing U_t^1 by $\beta(\overline{V} - \underline{V})$ is smaller than the corresponding increase from reducing U_t^2 by the same amount – a result that is essentially akin to concavity. By the envelope condition in (27), this also applies to agent 3's consumption. Therefore, changing agent 2's allocation in response to a fall in agent 1's income realization allows more risk-sharing to be sustained in the current period. Of course, there is some cost to this in terms of the

---

[7]The argument made here applies to any allocation in which $V^2 \geq V^1$. In the Appendix, we prove that this inequality will always be strict, so that $V^2 > V^1$ when agent 2 has been a net giver to agent 1.

[8]The proof of part 4. Proposition 1 does not depend on this assumption. It is merely introduced here for the purpose of illustration. The proof does, however, use the assumption that agents consume their own income in the period of deviation.

extent of risk-sharing that can be sustained between agent 1 and 2 in period t. However, when $\frac{u'(c_t^2)}{u'(c_t^1)}$ is close to one, which implies that neither agent has had to rely much on the other in the past, the gain in period $t+1$ far outweighs the cost. Taken together, this implies that the continuation payoff of a constrained agent depends on the current income realization of other constrained agents in the coalition-proof contract. Moreover, the size of this effect will depend on the previous risk-sharing history of the constrained agents.

Combining Proposition 2 and the first order conditions (25)-(28), we can see how consumption of a constrained agent is linked to the continuation payoffs and income realizations of other constrained agents. Denote by C_t the set of constrained agents and UC_t the set of unconstrained agents in period t. Define $\underline{U}_r^i$ as the payoff of agent i – not necessarily the autarky payoff – when his enforcement constraint is binding. In a sub-game perfect contract, this payoff depends only on his own income realization. In the coalition-proof contract, it depends on the income realization of other constrained agents as well as the previous history of the contract. This implies that consumption c_r^i of a constrained agent is determined by

$$
\frac{u'(c_r^n)}{u'(c_r^i)} = -\frac{\partial U_r^n}{\partial U_r^i}\Big|_{U_r^i = \underline{U}_r^i(y_r^i), U_r^{j \in C_t} = \underline{U}_r^j(y_r^j), U_r^{k \in UC_t} = U_r^k(h_t)}
\tag{29}
$$

in the sub-game perfect contract. In contrast, in the case of the coalition-proof contract, consumption of a constrained agent c_r^i is given by

$$
\frac{u'(c_r^n)}{u'(c_r^i)} = -\frac{\partial U_r^n}{\partial U_r^i}\Big|_{U_r^i = \underline{U}_r^i(y_r^i, \{y_r^j\}_{j \in C_t}, h_t), U_r^{j \in C_t} = \underline{U}_r^j(y_r^i, \{y_r^j\}_{j \in C_t}, h_t), U_r^{k \in UC_t} = U_r^k(h_t)}.
\tag{30}
$$

As a consequence, consumption depends on the previous history of constrained agents only in the coalition-proof contract. The income realization of other constrained agents enters both in the sub-game perfect and in the coalition-proof contract. However, in the former it enters only indirectly via its effect on U_r^j and since U_r^j does not depend on the previous history of shocks, there is no interaction between the effect of y_r^j and h_t.[9] In the latter, the income realization of constrained agents directly affects U_r^i and the magnitude of this effect depends on the previous history of shocks as outlined in Proposition 2.

The final step is to write (29) and (30) in terms of the marginal utility ratios in the previous period. This is valid, because from the first-order and envelope conditions, the continuation values are completely determined by the previous history of the game as summarized by the previous

---

[9]Note that (29) conditions on the set C and therefore the number of constrained agents at time t, which will depend on the previous history of the game. Therefore the effect of y_r^j in the sub-game perfect contract is independent of the previous history of marginal utility ratios only for a given number of constrained agents.

period's marginal utility ratios and the current income realizations. Therefore, the evolution of the ratio of marginal utilities over time is characterized by the following Euler equation

$$\frac{u'(c_t^n)}{u'(c_t^i)} = (\gamma_t^i)_{i \in C_t} = g((\gamma_{t-1}^m)_{m \in UC_t}, (y_t^k)_{k \neq i \in C_t}, y_t^i, Y_t) \tag{31}$$

where Y_t denotes aggregate income.

In the coalition-proof contract, the marginal utility ratio of constrained agents as well as their current incomes determine the joint payoff in case of deviation. The implication is that today's marginal utility ratio of agents, who are constrained, will depend on the previous period's marginal utility ratios and current incomes of all agents $k \in C_t$. In particular, the effect of income is larger if the marginal utility ratios of two constrained agents in period $t - 1$ are of similar magnitude:

$$(\gamma_t^i)_{i \in C_t} = g((\gamma_{t-1}^m)_{m \in UC_t}, (\gamma_{t-1}^k)_{k \in C_t}, (y_t^j, (\frac{\gamma_{t-1}^i}{\gamma_{t-1}^j} - 1)^2 \times y_t^j)_{j \neq i \in C_t}, y_t^i, Y_t). \tag{32}$$

Taken together, this implies that the results derived in Proposition 2 are related to the marginal utility ratios as follows.

1. In a sub-game perfect efficient (coalition-proof) contract, the marginal utility ratio of a constrained agent, γ_t^i, is independent of (depends on) the previous history of shocks of other constrained agents as summarized by $(\gamma_{t-1}^k)_{k \in C_t}$.

2. In a sub-game perfect efficient (coalition-proof) contract, the effect of any constrained agent $j's$ income realization, $(y_t^j)_{j \neq i \in C_t}$, on the marginal utility ratio of a constrained agent i, γ_t^i, is independent of (depends on) the previous history of shocks.

To sum up, the analytical properties of the constrained-efficient coalition-proof contract fit well with some of the empirical findings left unexplained by the existing models. For any stable group size, the set of coalition-proof equilibria is contained in the set of sub-game perfect equilibria and the extent of risk-sharing in the former can therefore never be larger than in the latter. Moreover, the set of stable sizes is finite as shown in Genicot and Ray (2003) and this will limit the extent of risk-sharing even further relative to sub-game perfection when the set of stable sizes are compared. Finally, the model shows that 'amnesia' is not actually a direct implication of limited commitment, but rather follows from solving for the set of sub-game perfect equilibria of the dynamic risk-sharing game. Refining equilibria to be coalition-proof can explain the empirical rejection of amnesia while maintaining the assumption of limited commitment.

4   Testing and Simulations

In this section, we exploit the property that history enters both directly and interacted with current income realizations into the first-order condition for consumption of a constrained agent in the coalition-proof contract to derive a test for the endogeneity of insurance group size in an environment of perfect information. We then apply the test to simulated data that is generated by numerically solving a parameterized version of the model derived in Section 3 and summarise the properties of the test.

To implement the test, we must derive an estimable equation from (29) and (30). Since g is an unknown non-linear function, we expand it in order to transform (32) into a linear regression. Define $\mathbf{x} = \{(\gamma_{t-1}^m)_{m \in UC_t}, (\gamma_{t-1}^k)_{k \in C_t}, (y_t^j, (\frac{\gamma_{t-1}^i}{\gamma_{t-1}^j} - 1)^2 \times y_t^j)_{j \neq i \in C_t}, y_t^i, Y_t\}$. This gives the $k'th$ order Taylor expansion around $\mathbf{x}^*$:

$$(\gamma_t^i)_{i \in C_t} = g(\mathbf{x}) = g(\mathbf{x}^*) + Dg(\mathbf{x}^*)(\hat{\mathbf{x}}) + \frac{1}{2!}D^2 g(\mathbf{x}^*)(\hat{\mathbf{x}}, \hat{\mathbf{x}}) + ... + \frac{1}{k!}D^k g(\mathbf{x}^*)(\hat{\mathbf{x}}, \hat{\mathbf{x}}, ..., \hat{\mathbf{x}}) + R_k(\hat{\mathbf{x}}, \mathbf{x}^*),$$

(33)

where $\hat{\mathbf{x}} = \mathbf{x} - \mathbf{x}^*$. Based on this, we can write the following linear estimable equations for a constrained agent i:

$$\gamma_t^i = \beta_0 + \sum_{k \in C_t} \beta_{1,k}\gamma_{t-1}^k + \sum_{k \in C_t} \beta_{2,k}y_t^k + \beta_3 Y_t + \text{higher order terms} \tag{34}$$

$$\gamma_t^i = \beta_0 + \sum_{k \in C_t} \beta_{1,k}\gamma_{t-1}^k + \sum_{k \in C_t} \beta_{2,k}y_t^k + \beta_3 Y_t$$
$$+ \sum_{j \neq i \in C_t} \beta_{4,j}(\frac{\gamma_{t-1}^i}{\gamma_{t-1}^j} - 1)^2 \times y_t^j + \text{higher order terms} \tag{35}$$

There are two more issues that need to be addressed before estimating (34) and (35). The tests above regress marginal utility ratios on their own past and rely on being able to identify constrained and unconstrained agents correctly. Both of these issues pose similar challenges. As proposed in Kocherlakota (1996), constrained and unconstrained agents can be identified by comparing the ratios of marginal utilities for different agents, but of course the shape of the utility function is in general not known. If we are willing to impose a logarithmic or constant relative risk aversion (CRRA) utility function, then constrained and unconstrained agents can be identified by comparing consumption shares, which are readily observable. Similarly, the tests can be implemented by regressing relative consumption shares of any two agents on their own past. However, for more general utility functions, the unknown parameters have to be identified from restrictions implied by the model itself. This

may prove difficult, if the restrictions only apply to constrained agents, but these cannot be identified from the data unless the parameters of the utility function have been estimated. To circumvent these problems, the majority of papers applying tests of this kind assume that agents have a CRRA utility function (see Alvarez and Jermann (2000) and Krueger et al. (2008)). A second, and potentially more serious problem is measurement error. If consumption is measured with error, then unconstrained agents may mistakenly be identified as constrained even when the shape of the utility function is known and past marginal utility ratios may be deemed significant when they are not. It is clear from our discussion that the test for exogenous group formation will be particularly sensitive to measurement error of this kind.

We now use simulated model solutions to show how to discriminate between endogenous and exogenous group formation under imperfect enforceability. To solve the model numerically, we restrict ourselves to 3 households with logarithmic utility functions. Individual income y^i takes on the values 1.5 and 3, with equal probability. Income realisations across households are identically and independently distributed. The discount factor is $\beta = 0.85$. To compute the optimal contract, we follow Marcet and Marimon (1999), Attanasio and Rios-Rull (2000) and Kehoe and Perri (2002) in setting up a Lagrangian for a social planner who seeks to maximize a weighted sum of utilities of a group of n agents by choice of current consumption and future Pareto weights. We implement the algorithm using weighted-residual methods based on finite element approximations to solve for the equilibrium of this model in each income state on a two dimensional grid of Pareto weights, which are normalized such that $\sum_{i=1}^{3} \lambda^i = 1$ (see Judd (1998), Judd (1992), McGrattan (1996), Christiano and Fisher (2000)). We solve for both the efficient sub-game perfect contract, which corresponds to exogenous group formation, as well as for the efficient coalition-proof contract, which corresponds to endogenous group formation.

The simulated data are obtained by starting the simulation at the grid point which gives equal consumption to the three agents and then letting the 3-agent model run for 10,000 periods for both the individual deviations model and the coalitional deviations model drawing a stochastic income realization in each period. In order to assess the effect of measurement error, both the dependent variable and the regressors are perturbed by a random error, $\epsilon \sim N(0, \sigma^2)$, whose standard deviation $\sigma = \frac{k}{10} \times \text{std}(\gamma_t), k = 1, ..., 10$. That is, the standard deviation of the measurement error ranges from 10% to 100% of one standard deviation of the computed marginal utility ratios. In total, this gives 10 simulated economies each for the sub-game perfect and the coalition-proof contract to which the empirical test is applied.

As the test is only valid for constrained agents, we select all time periods in which both agent 1 and

2 are constrained based on the selection rule stated in Kocherlakota (1996): if $\min\left(\frac{\gamma^1_t}{\gamma^1_{t-1}}, \frac{\gamma^2_t}{\gamma^2_{t-1}}\right) > 1$, then both agent 1 and 2 are constrained. Since the data are perturbed by measurement error, we apply a more conservative cut-off rule of 1.1. In the case of log-utility, the expression for the marginal utility ratio between agent 1 and 2, $\frac{u'(c^2_t)}{u'(c^1_t)}$, simplifies to $\frac{c^1_t}{c^2_t}$ and there are no unknown parameters in the utility function. We can therefore run the following two regressions:

$$\frac{c^1_t}{c^2_t} = \beta_0 + \beta_1 \frac{c^1_{t-1}}{c^2_{t-1}} + \beta_2 y^1_t + \beta_3 y^2_t + \beta_4 Y_t + \text{second order terms} \tag{36}$$

$$\frac{c^1_t}{c^2_t} = \beta_0 + \beta_1 \frac{c^1_{t-1}}{c^2_{t-1}} + \beta_2 y^1_t + \beta_3 y^2_t + \beta_4 Y_t + \beta_5 (\frac{c^1_{t-1}}{c^2_{t-1}} - 1)^2 \times y^2_t + \text{second order terms}. \tag{37}$$

These regressions are similar to the 'changes-in-shares' estimator in Ligon et al. (2002). That is, rather than specifying a regression in levels of consumption, we estimate how the relative shares of agent 1 and 2 are determined. The first regression tests whether the past history of constrained agents as summarized by the past marginal utility ratio is significant in explaining own consumption. If group formation is exogenous, the coefficient on $\frac{c^1_{t-1}}{c^2_{t-1}}$ is expected to be zero. If group formation is endogenous, then as c^2_{t-1} increases relative to c^1_{t-1}, we would expect agent 2 to be awarded a higher payoff than agent 1 when enforcement constraints are binding, because this requires a smaller movement in the marginal utility ratios. Hence the coefficient ought to be significant and positive. In (37), we also introduce the interaction term $(\frac{c^1_{t-1}}{c^2_{t-1}} - 1)^2 \times y^2_t$. Under endogenous group formation, the effect of income changes of agent 2 on the current marginal utility ratio depends on the previous period's marginal utility ratio of agent 1 and 2. Broadly speaking, the bigger the difference between initial consumption shares, the less we would expect income to matter in the endogenous group formation case.

The simulation procedure is repeated 500 times. Table 1 reports the average coefficients and p-values for $\sigma = 0.1 \times \text{std}(\gamma_t)$. The results successfully capture the predictions of the two models. The first and third column of Table 1 report the coefficients from regression (36). In the presence of endogenous group formation, the past marginal utility ratio of agent 1 and 2 has significant explanatory power, because it determines the relevant deviation payoff. When c^1_{t-1} is small relative to c^2_{t-1}, it is optimal to deter a deviation by awarding agent 1 the minimum payoff in a group of size 2. When c^1_{t-1} is large relative to c^2_{t-1}, continuation payoffs which give agent 2 the minimum payoff and agent 1 the maximum payoff become optimal. In contrast, when group size is exogenous, the coefficient on $\frac{c^1_{t-1}}{c^2_{t-1}}$ is zero and has no explanatory power because the optimal punishment path always involves reversion to the static autarky equilibrium. It is also noticeable that the coefficient on income is much smaller under endogenous group formation. This does not imply though that own

Relative consumption shares of agent 1 and 2 $\frac{c_t^1}{c_t^2}$	Individual deviation		Group deviation	
	(1)	(2)	(3)	(4)
$\frac{c_{t-1}^1}{c_{t-1}^2}$	-0.002	-0.003	0.233	0.230
	(0.4944)	(0.4911)	(0.000)	(0.000)
y_t^1	0.369	0.368	0.129	0.127
	(0.000)	(0.000)	(0.000)	(0.000)
y_t^2	-0.481	-0.481	-0.163	-0.147
	(0.000)	(0.000)	(0.000)	(0.000)
$(\frac{c_{t-1}^1}{c_{t-1}^2} - 1)^2 \times y_t^2$	-	0.001	-	-0.901
	-	(0.488)	-	(0.000)
Y_t	-0.001	-0.002	-0.028	0.003
	(0.497)	(0.497)	(0.231)	(0.245)
constant	1.367	1.368	0.870	0.790
	(0.000)	(0.000)	(0.000)	(0.000)

Notes: Reported coefficients and p-values in brackets are averaged over 500 repetitions. Data in column (1) and (2) are generated by solving for the equilibrium of a 3-player sub-game perfect limited commitment contract. Data in column (3) and (4) are generated by solving for the equilibrium of a 3-player coalition-proof limited commitment contract. The economy is simulated for 10,000 periods and a random error is added to consumption and income in each period. All time periods in which agent 1 and agent 2 are identified as constrained are included in the regression. The average number of included observations in 500 repetitions is 1600. Second order terms are included, but not reported. The only significant second order term is the interaction $y_t^1 \times y_t^2$.

Table 1: Predicted relative consumption shares when agent 1 and 2 are constrained.

shocks are better insured under endogenous group formation. The coefficient on income is smaller, because the current relative consumption share of agent i is determined not only by his income, but also by the history of the game.

In column (2) and (4) of the table we add the interaction effect of $(\frac{c_{t-1}^1}{c_{t-1}^2} - 1)^2 \times y_t^2$ to the specification. Here, the differential effect of y_t^2 in the case of endogenous versus exogenous group formation becomes very pronounced. In the former scenario, both the direct effect of agent 2's income and the past consumption shares of agent 1 and 2 as well as the interaction effect of these two variables are large and significant. To illustrate the magnitude of the effects, first suppose $\frac{c_{t-1}^1}{c_{t-1}^2} = 1$. Then raising agent 2's income from 1.5 to 3 decreases the marginal utility ratio in the current period $\frac{c_{t-1}^1}{c_{t-1}^2}$ by 0.15. If $\frac{c_{t-1}^1}{c_{t-1}^2} - 1$ is set to its maximum of 0.5, then increasing y_t^2 decreases the marginal utility ratio by 0.36. This is the case because for marginal utility ratios close to 1 it is optimal to increase agent 1's punishment path relative to the case when income realizations of constrained agents are equal in order to compensate for the increased inequality between agents due to agent 2's higher income realization. That is, consumption of agent 1 is raised relative to agent 2's income (even though in absolute terms, agent 2's consumption may exceed agent 1's consumption).

However, as $(\frac{c_{t-1}^1}{c_{t-1}^2} - 1)^2$ increases, the impact of agent 2's income realization diminishes. In the case of exogenous group formation, none of these effects matter.

Figure 2 plots the size and power of these tests as a function of measurement error. The size calculations in the three panels on the left-hand side are based on solving the model under exogenous group formation and recording the fraction of times in 500 repetitions the null hypothesis is rejected for different significance levels $\alpha = 0.01, 0.05, 0.1$. The first panel shows the size of the test for $\beta_1 = 0$ in regression (36). The test performs well for low to intermediate levels of measurement error, however as measurement error increases, the test becomes less robust. For $\sigma \geq 0.5 \times \text{std.}(\gamma_t)$, the null is rejected 80 % of the time when it is in fact true. One might expect that this happens because constrained agents are not identified correctly as measurement error increases. However, this does not appear to be the case. The conservative cut-off rule implies that even for the highest level of measurement error, only an average of 4 in 1600 observations are classified incorrectly. Rather, the null of $\beta_1 = 0$ is sensitive to even moderate amounts of measurement error. The second and third panel on the left-hand side report the size of the test for $\beta_5 = 0$ and the joint test of $\beta_1 = 0$ and $\beta_5 = 0$ in regression (37). Both these tests perform well and even for large levels of measurement error, the null is rejected less than half the time at a significance level of 5%. The three panels on the right-hand side report the power of the tests. The calculations are based on solving the model under endogenous group formation and recording the fraction of times in 500 repetitions we accept the null for different significance levels. The first panel shows the power of the test for $\beta_1 \neq 0$ in regression (36). The test is extremely strong and the null is never rejected regardless of the degree of measurement error. This is the case because the impact of the previous period marginal utility ratio both directly and interacted with income is large and because measurement error at least to some extent biases the results in favour of finding endogenous group formation.

Taken together, the results give us some confidence that the tests are able to distinguish between exogenous and endogenous group formation at least when data are well measured. Comparing the size and power of the test, it is clear that the main challenge lies in being able to reject endogenous group formation when it is not present. Of course, the testing strategy relies on a number of underlying assumptions. As discussed above, the power of this test depends crucially on the ability to identify constrained and unconstrained agents correctly. Moreover, even when constrained agents are identified correctly, the fact that the past marginal utility ratio matters may be due to other factors. An obvious weakness is that the model excludes the possibility of savings and assumes perfect information. If these assumptions do not adequately describe the environment, then the tests described here may not be able to distinguish between endogenous and exogenous group formation.
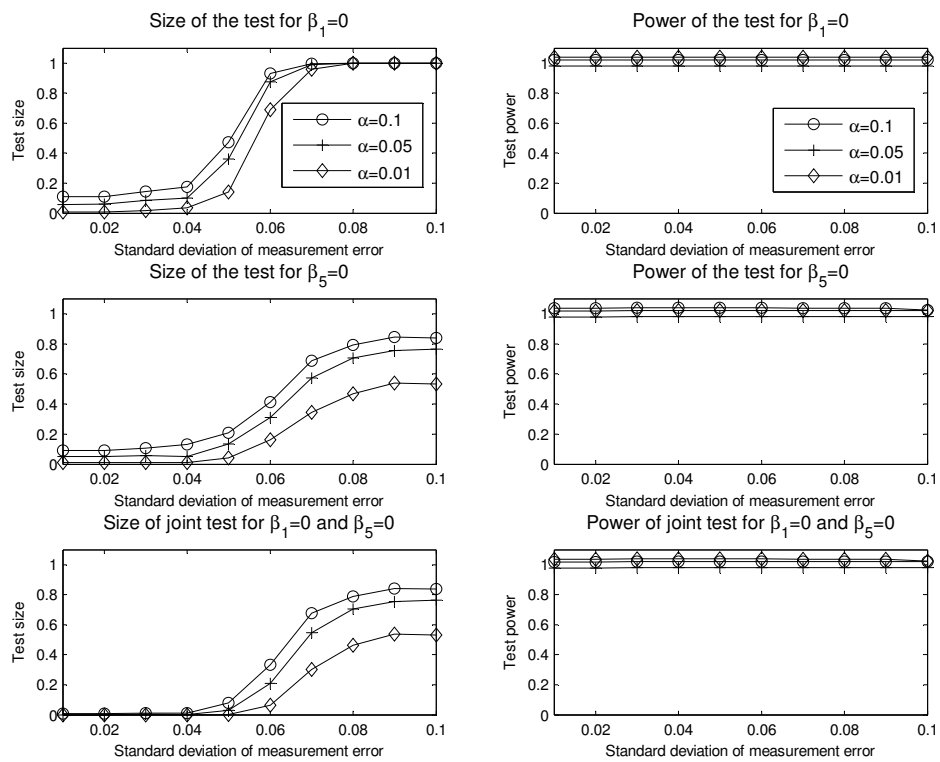
Figure 2: Size and power of the tests

In a recent paper, Dubois et al. (2008) show that the marginal utility ratio retains information about the past when enforcement constraints are binding even when group formation is exogenous in an environment that combines formal short-term contracts and informal long-term contracts. However, this is manifested by the fact that lagged consumption explains current income, which is an endogenous variable in the presence of formal contracts. This is not a feature of the pure limited commitment model under either exogenous or endogenous group formation and therefore the way in which history matters is quite different in their model and ours.

Secondly, the model of recursive stability of groups employed in this paper is partial in the sense that it does not consider deviations with agents outside the group. Group formation is also likely to be determined by other factors outside of the model such as proximity, covariance of incomes and heterogeneity of agents. Nevertheless, the constraints on group formation derived in this paper can be considered a subset of the factors that explain group formation – a first hurdle any group must pass. Given well measured income and consumption data on a census of mutually exclusive risk-sharing groups over time, the above tests can be applied to examine whether coalitional deviations form part of the constraints that limit group formation and as such may be considered a necessary condition for the stability of a risk-sharing group. Alternatively, the algorithm derived above enables us to estimate a full structural model of efficient risk-sharing when group formation is endogenous and compare it to a number of alternatives including exogenous group formation, imperfect information and models that include savings (see Ligon et al. (2002) and Karaivanov and Townsend (2008)). To be sure, this is computationally cumbersome for large risk-sharing arrangements. But there is increasing evidence that risk-sharing is confined to small networks of family and friends. Estimating a structural risk-sharing model for such groups could be a fruitful next step for further research.

5   Conclusion

This paper has shown how to model an efficient risk-sharing contract in the presence of endogenous group formation. Requiring groups to be stable with respect to deviations of subgroups as well as individuals alters the predictions of the dynamic limited commitment risk-sharing model substantially. In particular, 'amnesia' in the wake of binding constraints is no longer a property of the model. Instead, the consumption of constrained agents depends on both the history of shocks and its interaction with the current income of other constrained agents. More generally, the analytical properties derived can explain larger deviations from first-best insurance both for a given stable size and even more so over the entire set of stable sizes than predicted by existing models of sub-game

perfect dynamic risk-sharing.

The theoretical properties derived lead to relevant restrictions on consumption and income that allow the researcher to empirically test whether the threat of coalitional deviations is relevant in explaining the observed departure from full insurance. On a more practical level, establishing the presence or absence of endogenous group formation is an important question in developing country contexts, because the impact of policies such as outside financial intermediation and the provision of external safety nets in terms of crowding out of existing arrangements are likely to be very different if insurance group sizes are limited endogenously rather than exogenously.

Finally, the model derived in this paper provides a blue-print for structurally estimating a model of dynamic risk-sharing with coalitional deviations. In a companion paper (Bold (2008)), we test this model using data on intergenerational risk-sharing among families from the ICRISAT data set against the alternatives of full insurance, autarky and dynamic risk-sharing with individual deviations (see Ligon et al. (2002)). Moreover, we decentralize the economy to derive quantitative implications on the relative price of risk generated by our model of risk-sharing. Preliminary results suggest that dynamic coalition-proof risk-sharing does significantly better at explaining the data than the alternatives suggesting that our model can potentially play an important role in enhancing our understanding of consumption behaviour, risk-sharing and the returns to bearing risk.

6   Appendix A: Proofs

Proof of Lemma 1

1. The enforcement constraints in (17)-(22) are not convex. Hence, we cannot prove differentiability using standard arguments. Instead, we follow the strategy in Pavoni (2006) to establish differentiability of the value function at all points of interest.

The proof proceeds by defining a series of dynamic programmes of maximizing agent n's utility subject to agent $1, ... n-1$ receiving at least $U_s^1, ..., U_s^{n-1}$ and the enforcement constraints being satisfied conditional on fixed punishment payoffs for all potential subcoalitions. We then show that the conditional functions are differentiable in $U_s^1, ..., U_s^{n-1}$, that $U^n(U_s^1, ..., U_s^{n-1})$ is indeed the upper envelope of these conditional functions, and that $U^n(U_s^1, ..., U_s^{n-1})$ has left- and right-hand derivatives everywhere. Hence, the conditions of Theorem 1 in Milgrom and Segal (2002) are satisfied. Combining this with Theorem 2 in Cotter and Park (2006), it follows that the value function is differentiable on the range of the optimal policy correspondence.

Let $\mathcal{V}$ denote the space of infinite sequences of all possible punishment payoffs $(((\widetilde{V}_r^{j_g, J})_{g=1}^{m(k)})_{J \in \mathbb{J}})_{r=1}^{S}$ $\in \mathbb{V}^*(m(k))$ for a group of size n and let $\mathbf{V}$ denote a particular sequence in this space. For a given

$\mathbf{V}$, write the conditional maximisation problem as

$$U_s^n(U_s^1, U_s^2, ..., U_s^{n-1}; \mathbf{V}) = \max_{((U_r^i)_{r=1}^S)_{i=1}^{n-1}, (c_s^i)_{i=1}^n} u(c_s^n) + \beta \sum_{r=1}^S \pi_r U_r^n(U_r^1, ..., U_r^{n-1}; {}_1\mathbf{V}) \quad (38)$$
$$\text{s.t.} \quad (17) - (22).$$

The conditional function $U_s^n(U^1(s), U^2(s), ..., U_s^{n-1}; \mathbf{V})$ represents agent n's maximised discounted lifetime utility conditional on a given sequence of punishment payoffs, where ${}_1\mathbf{V}$ stands for the one step ahead continuation of $\mathbf{V}$. To proceed, we assume that a group of n is stable. Then it follows from adapting Proposition 3 in Pavoni (2006) that (38) exists and is unique.

We now establish that the conditional function is well-behaved and in particular, that it is differentiable. We then show that the value function is the upper envelope of the collection of conditional functions using an inductive argument. The set of feasible promised utilities $U_s^1, ..., U_s^{n-1}$ is compact, which can be shown by adapting the argument in Lemma 1 of Thomas and Worrall (1988). Secondly, note that $V^*(1)$ and $V^*(2)$ are compact by standard arguments (Thomas and Worrall (1988) and Kocherlakota (1996)). $\mathcal{V}$ is therefore compact by Tychonoff's Theorem for $n \leq 3$. Suppose moreover that compactness of $\mathcal{V}$ has been established for all $1, ..., n-1$ (ignoring for the moment that not all of these group sizes may be stable). The utility function is continuous and bounded. Substituting from the promise-keeping constraints when they bind with equality, it is easy to see that the choice set for consumption and continuation utilities is monotone and that the utility function is decreasing in each $(U_s^i)_{i=1}^{n-1}$. Therefore, Theorem 4.6 and 4.7 in Stokey et al. (1989) apply and the value function of the conditional problem is bounded, continuous, unique and strictly decreasing in its first $n-1$ arguments. In addition, the constraint set is convex conditional on $\mathbf{V}$ and it follows that the value function is concave and the set of maximizers is continuous and single valued. Finally, convexity of the constraint set and concavity of the utility function allows us to apply Lemma 2 in Benveniste and Scheinkman (1979) to show that the value function is continuously differentiable at any interior point of its domain. The usual envelope condition then applies and

$$\frac{\partial U_s^n(U_s^1, ..., U_s^{n-1}; \mathbf{V})}{\partial U_s^i} = -\frac{u'(c_s^n)}{u'(c_s^i)} \quad \forall i. \quad (39)$$

Before proceeding, we show that the maximization of the conditional function with respect to $\mathbf{V}$

$$U_s^n(U_s^1, ..., U_s^{n-1}) = \max_{\mathbf{V} \in \mathcal{V}} U_s^n(U_s^1, ..., U_s^{n-1}; \mathbf{V}) \quad (40)$$

is well-defined. By the inductive hypothesis, $\mathcal{V}$ is compact and the conditional function is continuous.

Hence, Berge's Theorem of the Maximum can be applied to this problem and the value function is continuous and unique – and its range is therefore a compact set – and the optimal policy correspondence $\mathcal{V}^*$ is non-empty, compact valued and upper-hemicontinuous. Since each V^i is a continuous function of $(U_r^i)_{r=1}^S$, it follows that $\mathcal{V}$ is compact for n. By strong induction, the properties derived for the conditional functions and the value function $U^n(U_s^1, ..., U_s^{n-1})$ hold for all n. Moreover, both the conditional and the value function are bounded. The equivalence between the sequential and the recursive formulation then follows easily by applying standard arguments found in Stokey et al. (1989).

To apply Theorem 1 in Milgrom and Segal (2002), we make the following assumption

Assumption 1 *The set of maximizers* $\mathcal{V}^*(U_s^1, ..., U_s^{n-1})$ *is finite.*[10]

Together with differentiability of the conditional functions, this implies that the value function is differentiable almost everywhere and has left- and right-hand derivatives everywhere. Hence, Theorem 1 in Milgrom and Segal (2002) tells us that

$$\frac{\partial U_s^{n,+}(U_s^1, ..., U_s^{n-1})}{\partial U_s^i} \geq \frac{\partial U_s^n(U_s^1, ..., U_s^{n-1}; \mathbf{V})}{\partial U_s^i} \geq \frac{\partial U_s^{n,-}(U_s^1, ..., U_s^{n-1})}{\partial U_s^i} \quad \forall i = 1, ..., n-1. \quad (41)$$

If $\frac{\partial U_s^{n,+}}{\partial U_s^i} = \frac{\partial U_s^{n,-}}{\partial U_s^i}$, the value function is differentiable at U_s^i and its derivative equals $-\frac{\partial u'(c_s^n)}{\partial u'(c_s^i)}$.

We now show that the value function is differentiable at all optimal points and the first order conditions are therefore necessary to characterize the maximum. Note that if the value function is not differentiable at U_s^i, then $\frac{\partial U_s^{n,+}(U_s^1, ..., U_s^{n-1})}{\partial U_s^i} \geq \frac{\partial U_s^{n,-}(U_s^1, ..., U_s^{n-1})}{\partial U_s^i}$ describes a downward kink of the value function and so U_s^i cannot be a maximizer. To see this, note that we can always turn (15) into a free maximization problem by substituting in the promise-keeping and binding enforcement constraints. At any point of optimality, it must then be the case that

$$\frac{\partial U_r^{n,+}}{\partial U_r^i} \leq \frac{\partial U_r^{n,-}}{\partial U_r^i} \quad (42)$$

Combining this with (41) above, it follows that the value function is differentiable at any point

---

[10]While we have not been able to prove that this will in general be the case, the simulation results show that this assumption is not implausible. The total utility to be allocated to the set of constrained agents and the allocations making up this total utility are usually unique. The only non-uniqueness arises from the order in which the chosen payoffs are allocated among the constrained agents. Of course, finiteness of the optimal policy correspondence in the recursive formulation is not the same as finiteness of the set of infinite sequences that maximize the programme in (40). It is merely a necessary condition. We would argue that assuming finiteness of the set of optimal sequences for the sake of analytical tractability is not too controversial for the following reason: The main task of the paper is to show that the sub-game perfect and the coalition-proof risk-sharing contract behave differently when constraints are binding. It is well-known that the sub-game perfect contract can be characterized by a state-dependent, but history independent updating rule. Hence, the optimal choice set in the sub-game perfect contract is finite and moreover, it is single-valued. Therefore, if the assumption of finiteness was violated, then this in itself would imply that the coalition-proof contract behaves differently from the sub-game perfect contract when constraints are binding.

on the optimal policy correspondence and the first-order conditions are necessary (see Theorem 2, Cotter and Park (2006)).

2. Solving for c_s^i from the promise-keeping constraints, we write

$$\frac{\partial U_s^n}{\partial U_s^i} = -\frac{u'(\sum_{i=1}^n y_s^i - \sum_{i=1}^{n-1} u^{-1}(U_s^i - \beta \sum_{r=1}^S U_r^i))}{u'(u^{-1}(U_s^i - \beta \sum_{r=1}^S U_r^i)}. \tag{43}$$

It follows from the concavity of the utility function that the right-hand side of this expression is decreasing in U_s^i.

3. Having shown that the problem is differentiable at the optimum, we can use the Kuhn-Tucker Theorem to show that the multipliers associated with the enforcement constraints are positive. To see that the Kuhn-Tucker constraint qualification holds, note that at the optimum each agent can have a binding constraint in at most one deviating subcoalition and at most $n-1$ agents can have a binding enforcement constraint. Then it can be shown that the Jacobian matrix of first derivatives of the binding enforcement constraints has maximal rank (see Rincon-Zapatero and Santos (2007)). Hence, the constraint qualification holds and there are nonnegative multipliers associated with the enforcement constraints.

∎

Proof of Proposition 2:

1) Suppose $\phi_r^i > 0$ for $i = 1, ..., n-1$ and $\phi_r^n = 0$. As a first step, note that (15) is equivalent to a social planner maximizing a weighted sum of expected discounted utilities of n agents. The planner takes as given the planning weights $\{\gamma^i\}_{i=1}^{n-1}$, which are equivalent to the Lagrange multipliers on the promise keeping constraints in (16). He then chooses consumption c_s^i, promised continuation values U_r^i and the punishment payoffs $(((\widetilde{V}_r^{j_g,J})_{g=1}^k)_{J \in \mathbb{J}})_{r=1}^S$ subject to satisfying the enforcement constraints corresponding to the equilibrium concept that is applied:

$$U_s^P = \max_{((U_r^i)_{r=1}^S)_{i=1}^{n-1}, (c_s^i)_{i=1}^n} \sum_{i=1}^{n-1} \gamma^i \big(u(c_s^i) + \beta \sum_{r=1}^S \pi_r U_r^i\big)$$
$$+ u(c_s^n) + \beta \sum_{r=1}^S \pi_r U_r^n \big(U_r^1, ..., U_r^{n-1}\big). \tag{44}$$

If $\phi_r^i > 0$ for $i = 1, ..., n - 1$ and $\phi_r^n = 0$, the first order condition with respect to U_r^i is

$$-\frac{\partial U_r^n}{\partial U_r^i} = \gamma^i(1 + \phi_r^i)$$

$$\Leftrightarrow \quad -\frac{\partial U_r^n}{\partial U_r^i} > \gamma^i \tag{45}$$

i.e the marginal gain from increasing U_r^i to satisfy the enforcement constraints is less than the marginal loss in utility due to decreasing $U_r^n(U_r^1, ..., U_r^{n-1})$. Therefore, the planner will increase $\{U_r^i\}_{i=1}^{n-1}$ by as little as possible to satisfy the enforcement constraints. In the sub-game perfect contract, this implies setting $U_r^i = u(y_r^i) + \beta V^*(1)$ and hence choosing the punishment path $\widetilde{V}^i = V^*(1)$, since this is the payoff that agents can achieve by unilateral deviation. Because the autarky payoff is fully determined by the current income realization of agent i, it corresponds to a static equilibrium and is not history-dependent.

2) Next, we show that the optimal punishment payoffs are history-dependent in the coalition-proof contract. As they are chosen with the fact in mind that continuation payoffs will be equal to the discounted punishment payoffs augmented by one period of autarky consumption when enforcement constraints are binding, it then follows that the continuation payoffs on the equilibrium path are history-dependent when enforcement constraints are binding. In what follows, a superscript fb indicates the first-best allocation given the planning weights after history h_t, whereas a superscript c denotes the constrained contract. Suppose that in the next period in state r, $y_r^1 = y_r^2$, and agent 1 and 2 would both find an individual deviation profitable if first-best transfers were implemented, and can therefore also threaten to deviate jointly. All other enforcement constraints are satisfied. The total loss in utility U_s^P when enforcement constraints are binding can be approximated by

$$
\begin{aligned}
dU_s^P = U_s^{P,fb} - U_s^{P,c} &\approx \beta \pi_r \sum_{i=1}^{n-1} \left(\gamma^i + \frac{\partial U_r^n}{\partial U_r^i}\right)\Big|_{U_r^{i,c}} (U_r^{i,fb} - U_r^{i,c}) \\
&= \beta \pi_r \sum_{i=1}^{n-1} (\gamma^i - \gamma_r^i)(U_r^{i,fb} - U_r^{i,c}) \\
&= -\beta \pi_r [\gamma^1 \phi_r^{1,\{12\}} (U_r^{1,fb} - U_r^{1,c}) + \gamma^2 \phi_r^{2,\{12\}} (U_r^{2,fb} - U_r^{2,c})].
\end{aligned}
\tag{46}
$$

where the last equality follows from the fact that $\gamma^i = \gamma_r^i$ for all unconstrained agents. By assumption, the first-best allocation following this history would give agent 1 and 2 less than autarky utility and therefore $(U_r^{i,fb} - U_r^{i,c}) < 0$ for $i = 1, 2$. Hence dU_s^P is positive and the planner will choose punishment paths to support the continuation payoffs $U_r^{1,c}$ and $U_r^{2,c}$, such that $U_s^{P,fb} - U_s^{P,c}$ is as close to zero as possible. Suppose, the planner has chosen a particular punishment path $\widetilde{V}_r^{1,\{12\}}$ for

agent 1 with $V^2(\widetilde{V}_r^{1,\{12\}})$ on the Pareto frontier of a risk-sharing group of size 2. By symmetry of the stable payoff vectors, for a given total utility $\widetilde{V}_r^{1,\{12\}} + V^2(\widetilde{V}_r^{1,\{12\}})$, the planner can either pick a punishment path, such that $\widetilde{V}_r^{1,\{12\}} = \widetilde{V}^h \geq \widetilde{V}^l = V^2(\widetilde{V}_r^{1,\{12\}})$ or vice versa, where h and l denote high and low. Also write $U_r^h = u(y_r^1) + \beta\widetilde{V}^h$, $U_r^l = u(y_r^1) + \beta\widetilde{V}^l$. Then if $\gamma^2 > \gamma^1$, which implies that $U_s^1 < U_s^2 \Rightarrow dU^P(\widetilde{V}^h, \widetilde{V}^l) \geq dU^P(\widetilde{V}^l, \widetilde{V}^h)$. That is, the fall in total utility is less when $U_r^1 = U_r^l$ and $U_r^2 = U_r^h$. Towards a contradiction, suppose

$$
\begin{aligned}
dU_s^P(\widetilde{V}^h, \widetilde{V}^l) &\approx \beta\pi_r\Big[(\gamma^1 + \frac{\partial U_r^n}{\partial U_r^h})dU_r^1 + (\gamma^2 + \frac{\partial U_r^n}{\partial U_r^l})dU_r^2 + \sum_{i=3}^{n-1}(\gamma^i + \frac{\partial U_r^n}{\partial U_r^i})dU_r^i\Big] < \\
dU_s^P(\widetilde{V}^l, \widetilde{V}^h) &\approx \beta\pi_r\Big[(\gamma^1 + \frac{\partial U_r^n}{\partial U_r^l})dU_r^1 + (\gamma^2 + \frac{\partial U_r^n}{\partial U_r^h})dU_r^2 + \sum_{i=3}^{n-1}(\gamma^i + \frac{\partial U_r^n}{\partial U_r^i})dU_r^i\Big].
\end{aligned}
\tag{47}
$$

From (47), it follows that

$$
\begin{aligned}
(\gamma^1 + \tfrac{\partial U_r^n}{\partial U_r^h})(U_r^{1,fb} - U_r^h) + (\gamma^2 + \tfrac{\partial U_r^n}{\partial U_r^l})(U_r^{2,fb} - U_r^l) \quad &< \\
(\gamma^1 + \tfrac{\partial U_r^n}{\partial U_r^l})(U_r^{1,fb} - U_r^l) + (\gamma^2 + \tfrac{\partial U_r^n}{\partial U_r^h})(U_r^{2,fb} - U_r^h) & \\
\Leftrightarrow \quad (\gamma^1 - \gamma^2)(U_r^l - U_r^h) + (\tfrac{\partial U_r^n}{\partial U_r^h} - \tfrac{\partial U_r^n}{\partial U_r^l})(U_r^{1,fb} - U_r^{2,fb}) \quad &< \quad 0,
\end{aligned}
\tag{48}
$$

which is a contradiction, because the first term is positive by assumption, and the second term is positive because $\frac{\partial U_r^n}{\partial U_r^h} < \frac{\partial U_r^n}{\partial U_r^l}$ from Lemma 1 and the fact that $U_r^{1,fb} < U_r^{2,fb}$ when $\gamma^1 < \gamma^2$.

Finally, we show that $\widetilde{V}^h$ is strictly greater than $\widetilde{V}^l$ when $\gamma^1 < \gamma^2$. The difference in total utility between choosing $\widetilde{V}_r^{1,\{12\}} = V^2(\widetilde{V}_r^{1,\{12\}}) = V$ and $V_r^{1,\{12\}} = V - \epsilon$ and $V^2(V - \epsilon)$ is given by

$$
\begin{aligned}
&U_s^P(V - \epsilon, V^2(V - \epsilon)) - U_s^P(V, V^2(V)) \approx \\
&- \beta\pi_r\epsilon\Big[\Big(\gamma^1 + \frac{\partial U_r^n}{\partial V_r^{1,\{12\}}}\Big|_{V_r^{1,\{12\}}=V}\Big) + \frac{\partial V_r^{2,\{12\}}}{\partial V_r^{1,\{12\}}}\Big|_{V_r^{1,\{12\}}=V}\Big(\gamma^2 + \frac{\partial U_r^n}{\partial V_r^{2,\{12\}}}\Big|_{V_r^{2,\{12\}}=V}\Big)\Big]
\end{aligned}
\tag{49}
$$

Since $\frac{\partial V_r^{2,\{12\}}}{\partial V_r^{1,\{12\}}}\Big|_{V_r^{1,\{12\}}=V} = -1$, the above expression is greater than zero when $\gamma^1 < \gamma^2$.[11] Therefore, $\gamma^1 < \gamma^2$ and $U_s^1 < U_s^2$ imply that $\widetilde{V}_r^{1,\{12\}} < V^2(\widetilde{V}_r^{1,\{12\}})$. If $y_r^1 = y_r^2$, this necessarily implies that $U_r^2 > U_r^1$ in the coalition-proof contract. An analogous argument can be made for any stable subgroup $m = 1, ..., n-1$. Therefore the coalition-proof efficient contract is history-dependent when enforcement constraints are binding.

3) This follows trivially from the fact that the punishment path in a sub-game perfect efficient contract involves reversion to the autarky payoff, which in state r is given by $u(y_r^i) + \beta V^*(1)$.

---

[11]The argument made in this proof depends on the fact that the ex-ante value function $V^2(V_r^{1,\{12\}})$ is differentiable at V. As shown in Koeppl (2003), for this to be the case, the constraint qualification must be satisfied at V.

4) Again, we consider the potential deviation of a coalition consisting of two agents in a group of n agents. We compare the case where both agents have income realization $y_r^1 = y_r^2 = y_r''$ to the case in which agent 1 has a lower income realization, $y_r^1 = y_r' < y_r^2 = y_r''$. Again, let $\gamma^2 > \gamma^1$. We know from 2) that the optimal choice in the first case implies that $U_r^1 = u(y_r'') + \beta \widetilde{V}^l < U_r^2 = u(y_r'') + \beta \widetilde{V}^h$, where $\widetilde{V}^l < \widetilde{V}^h$ implies that agent 2's continuation payoff from next period onwards is raised relative to agent 1's payoff. Because the set of punishment paths is symmetric, $\{V_r^{1,\{12\}}, V^2(V_r^{1,\{12\}})\} = \{\widetilde{V}^h, \widetilde{V}^l\}$ is also a feasible punishment path. We now show that the allocation $U_r^1 = u(y_r') + \beta \widetilde{V}^h$, $U_r^2 = u(y_r'') + \beta \widetilde{V}^l$ increases welfare relative to $U_r^1 = u(y_r') + \beta \widetilde{V}^l < U_r^2 = u(y_r'') + \beta \widetilde{V}^h$, which means that the latter cannot be in the set of maximizers. To see this, write the change in the planner's value function as

$$
\begin{aligned}
U_s^P(\widetilde{V}^h, \widetilde{V}^l) - U_s^P(\widetilde{V}^l, \widetilde{V}^h) \quad \approx \quad & \beta^2 \pi_r \big[\gamma^1 (\widetilde{V}^h - \widetilde{V}^l) + \gamma^2 (\widetilde{V}^l - \widetilde{V}^h) \\
& + \big[\frac{\partial U_r^n}{\partial U_r^1} \Big|_{u(y_r') + \beta \widetilde{V}^l} - \frac{\partial U_r^n}{\partial U_r^2} \Big|_{u(y_r'') + \beta \widetilde{V}^h} \big] (\widetilde{V}^h - \widetilde{V}^l) \big]
\end{aligned}
$$

Suppose that $u(y_r') + \beta \widetilde{V}^h < u(y_r'') + \beta \widetilde{V}^l$. If $\gamma^1 = \gamma^2$, then $U_s^P(\widetilde{V}^h, \widetilde{V}^l) > U_s^P(\widetilde{V}^l, \widetilde{V}^h)$ because the gain in U_r^n as U_r^2 decreases by $\widetilde{V}^h - \widetilde{V}^l$ exceeds the fall in U_r^n as U_r^1 increases by the same amount (see part 2. of Lemma 1). Hence it follows from the continuity of the value function that $\exists \epsilon > 0$, such that $\gamma^2 - \gamma^1 = \epsilon$ and $U_s^P(\widetilde{V}^h, \widetilde{V}^l) > U_s^P(\widetilde{V}^l, \widetilde{V}^h)$, but then $\{U_r^1, U_r^2\} = \{u(y_r') + \beta \widetilde{V}^l, u(y_r'') + \beta \widetilde{V}^h\}$ cannot be an optimal choice on the range of deviation payoffs induced by $y_r^2 = y_r'' > y_r^1 = y_r'$. If $u(y_r') + \beta \widetilde{V}^h > u(y_r'') + \beta \widetilde{V}^l$, then an analogous argument can be made by comparing $U_r^n(u(y_r'') + \beta \widetilde{V}^h, u(y_r') + \beta \widetilde{V}^l, ...)$, which by symmetry equals $U_r^n(u(y_r') + \beta \widetilde{V}^l, u(y_r'') + \beta \widetilde{V}^h, ...)$, to $U_r^n(u(y_r') + \beta \widetilde{V}^h, u(y_r'') + \beta \widetilde{V}^l, ...)$. ∎

References

Abreu, D., 1988. 'On the theory of infinitely repeated games with discounting'. *Econometrica* vol. 56, pp. 383–396.

Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica* 58, 1041–1063.

Ahn, T. K., Isaac, R. M., Salmon, T. C., 2006. Endogenous Group Formation, Mimeo, Florida State University.

Alvarez, F., Jermann, U. J., 2000. 'Efficiency, equilibrium, and asset pricing with risk of default'. *Econometrica* vol. 68, pp. 775–797.

Alvarez, F., Jermann, U. J., 2001. 'Quantitative Asset Pricing Implications of Endogenous Solvency Constraints'. *Journal of Fiscal Studies* vol. 14, pp. 1117–1151.

Asheim, G., Strand, J., 1991. 'Long-term union-firm contracts'. *Journal of Economics* vol. 56, pp. 161–184.

Attanasio, O., Rios-Rull, J.-V., 2000. Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare? European Economic Review 44, 1225–1258.

Benveniste, L. M., Scheinkman, J. A., 1979. On the Differentiability of the Value Function in Dynamic Models of Economics. *Econometrica* 47, 727–732.

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.

Bold, T., 2008. Quantitative Implications of Coalition-Proof Dynamic Risk-Sharing, Draft Paper, University of Oxford.

Chaudhuri, A., Gangadharan, L., Maitra, P., 2005. An Experimental Analysis of Group Size and Risk Sharing, The University of Melbourne, Department of Economics, Research Paper 955.

Christiano, L. J., Fisher, J. D. M., 2000. Algorithms for Solving Dynamic Models with Occasionally Binding Constraints. Journal of Economic Dynamics and Control 24, 1179–1232.

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

Cotter, K., Park, J.-H., 2006. Non Concave Dynamic Programming. Economic Letters 90, 141–146.

Deaton, A., 1992. Household savings in LDCs: Credit Markets, Insurance and Welfare. Scandinavian Journal of Economics 96, 253–273.

De Weerdt, J., Dercon, S., 2006. Risk-Sharing Networks and Insurance against Illness. Journal of Development Economics 81, 337–356.

Dercon, S., Krishnan, P., 2000. In Sickness and in Health: Risk Sharing within Households in Rural Ethiopia. Journal of Political Economy 108, 688–727.

Dubois, P., 2005. Heterogeneity of Preferences, Limited Commitment and Coalitions: Empirical Evidence on the Limits to Risk Sharing in Rural Pakistan, Mimeo, University of Toulouse.

Dubois, P., Jullien, B., Magnac, T., 2008. 'Formal and informal risk sharing in LDCs: theory and empirical evidence'. *Econometrica* vol. 76, pp. 679–725.

Fafchamps, M., 1999. Risk Sharing and Quasi-Credit. Journal of International Trade and Economic Development 8, 257–278.

Fafchamps, M., Lund, S., 2003. Risk-Sharing Networks in Rural Philippines. Journal of Development Economics 71, 261–287.

Farrell, J., Maskin, E., 1989. Renegotiation in Repeated Games. Games and Economic Behavior 1, 327–360.

Genicot, G., Ray, D., 2003. Group Formation in Risk-Sharing Arrangements. *Review of Economic Studies* 70, 87–113.

Grimard, F., 1997. Household Consumption Smoothing through Ethnic Ties: Evidence from Cote d'Ivoire. Journal of Development Economics 53, 391–422.

Judd, K., 1992. Projection Methods for Solving Aggregate Growth Models. *Journal of Economic Theory* 58, 410–452.

Judd, K. L., 1998. Numerical Methods in Economics. MIT Press, Cambridge, Massachusetts.

Karaivanov, A., Townsend, R. M., 2008. Enterprise dynamics and finance: distinguishing mechanism design from exogenously incomplete market models, Mimeo, University of Chicago.

Kehoe, P., Perri, F., 2002. International Business Cycles with Endogenously Incomplete Markets. *Econometrica* 58, 1041–1063.

Kinnan, C., 2008. Distinguishing Barriers to Insurance in Thai Villages, mIT, Mimeo.

Kletzer, K. M., Wright, B. D., 2000. Sovereign Debt as Intertemporal Barter. American Economic Review 90, 621–639.

Kocherlakota, N. R., 1996. Implications of Efficient Risk Sharing without Commitment. *Review of Economic Studies* 63, 595–609.

Koeppl, T. V., 2003. Differentiability of the Efficient Frontier when Commitment to Risk Sharing is Limited, Mimeo, European Central Bank, Frankfurt, Germany.

Krueger, D., Lustig, H., Perri, F., 2008. 'Evaluating asset pricing models with limited commitment using household consumption data'. *Journal of the European Economic Association* vol. 6, pp. 715–726.

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.

Mace, B. J., 1991. Full Insurance in the Presence of Aggregate Uncertainty. Journal of Political Economy 99, 928–956.

Marcet, A., Marimon, R., 1999. Recursive Contracts, Mimeo, Universitat Pompeu Fabra.

Mazzocco, M., 2007. Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment. *Review of Economic Studies* 74, 857895.

McGrattan, E. R., 1996. Solving the Stochastic Growth Model with a Finite Element Method. Journal of Economic Dynamics and Control 20, 19–42.

Mehra, R., Prescott, E. C., 1985. The Equity Premium – A Puzzle. *Journal of Monetary Economics* 15, 145–161.

Milgrom, P., Segal, I., 2002. Envelope Theorems For Arbitrary Choice Sets. *Econometrica* 70, 583–601.

Morduch, J., 1991. Consumption Smoothing Across Space: Tests for Village-Level Responses to Risk, Mimeo, Harvard University.

Murgai, R., Winters, P., Sadoulet, E., de Janvry, A., 2002. Localized and Incomplete Mutual Insurance. Journal of Development Economics 67, 245–274.

Ogaki, M., Zhang, Q., 2001. 'Decreasing relative risk aversion and tests of risk sharing'. *Econometrica* vol. 69, pp. 515–526.

Pavoni, N., 2006. Optimal Unemployment Insurance, with Human Capital Depreciation, and Duration Dependence, Mimeo, University College London and IFS.

Pearce, D., 1988. Renegotiation-proof equilibria: Collective rationality and intertemporal cooperation, Mimeo, University of Yale.

Phelan, C., Townsend, R. M., 1991. Computing Multi-period, Information Constrained Optima. *Review of Economic Studies* 58, 853–881.

Rincon-Zapatero, J. P., Santos, M. S., 2007. Differentiability of the value function without interiority assumptions, Mimeo, Universidad Carlos III de Madrid.

Seppala, J., 1999. Asset Prices and Business Cycles under Limited Commitment, University of Chicago, Mimeo.

Stokey, N., Lucas, R., Prescott, E. C., 1989. Recursive Methods in Economic Dynamics. Harvard University Press, Cambridge, Massachusetts.

Thomas, J., Worrall, T., 1988. Self-Enforcing Wage Contracts. *Review of Economic Studies* 55, 541–554.

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.

Wang, C., 1995. Dynamic Insurance with Private Information and Balanced Budgets. *Review of Economic Studies* 62, 577–595.