Risk-Sharing in Village Economies Revisited

Tessa Bold, Goethe University Frankfurt and IIES, Stockholm University and CEPR
Tobias Broer, IIES, Stockholm University and CEPR*

June 2015

Abstract

The limited commitment model is popular for the analysis of village risk-sharing as it captures both the observed partial character of insurance and the presumption that incomes are well observed but formal contracts are absent in rural communities. Enforcement constraints in insurance contracts, however, typically bind only in case of positive income shocks, when the outside option of leaving the village is attractive. We show how this results in strongly counterfactual asymmetries in the consumption process at usual village sizes and document that this asymmetry arises already at small single-digit group sizes. We endogenize the size of the insurance group by allowing households to renege on informal contracts together with other villagers resulting in groups that are much smaller than the typical village. This brings the predicted consumption process, which is more symmetric in small groups, in line with the data. We thus argue that allowing for endogenous group formation in the dynamic limited commitment model strongly improves its predictive power for analyzing risk-sharing in village economies.

JEL Classification: D11, D12, D52
Keywords: Risk-Sharing, Village Economies, Informal Insurance, Limited Enforcement, Renegotiation-Proofness

*Contact: tobias.broer@iies.su.se, tessa.bold@iies.su.se.

1 Introduction

Low and volatile incomes in agricultural communities of developing countries suggest that targeted policy interventions could significantly improve the lives of the disadvantaged. Both the desirability and the effectiveness of policies to alleviate poverty, however, depend on the ability of communities to share risk among their members, and the mechanisms through which this is achieved. For example, if risk-sharing is strong, transfers to temporarily poor households are less urgent than in a situation where household consumption is intimately tied to their income. And trivially, the ability of policy to improve welfare through transfers, public goods provision, etc. may depend strongly on the incentive effects that these policies have for risk-taking behaviour, provision of effort, etc. More specifically, policy interventions may interact strongly with the frictions at the origin of limited risk-sharing. This is particularly true if risk-sharing is imperfect because households are unable to commit to a village insurance scheme. If the alternative to co-insurance is autarky, it is easy to see how measures of income support could have perverse consequences for the degree of risk-sharing by making this outside option more attractive (Attanasio and Rios-Rull 2000).

The absence of commitment to co-insurance is often seen as a particularly plausible reason for limited risk-sharing in poor villages of developing countries. This is because other impediments of insurance, such as lack of information on households' productive possibilities and effort, should be less pronounced in small communities. At the same time, contract enforcement is difficult in the context of poor villages. Specifically, village members may have the option to look for work in the next bigger city, or to drop out of any village-level risk-sharing arrangement in favour of risk-sharing within a smaller group of neighbours or relatives. Apart from its appeal on a priori grounds, several studies have found the limited commitment model to predict the degree of risk-sharing observed in agricultural village economies better than simple alternatives, such as perfect risk-sharing within villages, or lack of risk-sharing altogether (Townsend 1994, Ligon, Thomas and Worrall 2002, Laczo 2014). This is notwithstanding the fact that the standard model underpredicts observed consumption volatility (Ligon et al. 2002), at least when estimated to replicate the observed cross-sectional distribution of consumption in village economies. This tendency to overpredict the degree of insurance is, infact, a more general feature of the model that also arises in applications to larger economies. There, this deficiency has been shown to improve at lower risk-aversion or when relaxing the assumption of eternal exclusion in favour of a more attractive outside option, where agents can save or store consumption goods, for example (Broer 2013).

All previous studies of village risk-sharing, however, have concentrated on a particularly simple version of the model, where the only alternative that individuals have to the village-level insurance arrangement is complete autarky. It is difficult to see, however, how households that were excluded from village insurance could be prevented from forming a new, smaller risk-sharing group. The first contribution of this paper is to relax the standard assumption of individual deviation, and to study quantitatively the implications of village risk-sharing without commitment when individuals can deviate as sub-groups of the village. The theoretical implications of coalition-proof dynamic limited commitment were explored in Bold (2009), but here we propose a tractable way to compute the key characteristics of this 'coalition-proof' risk-sharing arrangement and estimate the model for the well-known ICRISAT dataset on rural villages in India.[1]

An additional, less well-known reason why the limited commitment model is in general appealing for small communities becomes evident when comparing its implications to those that arise in larger insurance communities. While the model can explain features of risk-sharing also in large western economies, such as higher levels of risk-sharing than predicted by simple 'self-insurance' through borrowing and saving,[2] it does, however, have a strongly counterfactual implication: since agents have incentives to default only in 'good times', when they are called upon to make transfers to others, participation constraints only bind in periods of high income while, after negative income shocks, agents enjoy insurance transfers that optimally smooth their consumption path. This results in a strongly counterfactual asymmetry in the predicted time series of individual consumption growth (which combines strong upward spikes with slow downward drift), as well as the cross-sectional distribution of consumption (Krueger and Perri 2005, Broer 2013).

This asymmetry does not, however, necessarily translate to smaller communities. To see this, consider a village of only two individuals: the feasibility constraint of this economy requires agent 1's consumption change relative to the average consumption to equal the negative of agent 2's. The resulting process is thus completely symmetric. We would expect the asymmetry in consumption processes to emerge, and become more severe, as the village size increases. But

[1]We are only aware of one previous paper that attempts to test empirically whether risk-sharing at the village level is constrained by coalitional deviations. Dubois (2005) shows that in a setting with stationary limited commitment and quadratic utility in which an insurance group's stability is assessed only relative to first-best risk-sharing, limited commitment with individual deviations and limited commitment with coalitional deviations have different implications for how the benefits of insurance change with group size and whether the likelihood of perfect risk-sharing increases or decreases with risk aversion and the degree of income correlation. He tests these hypotheses on a panel of Pakistani villages.

[2]For example, (Krueger and Perri 2005) argue that the small rise in consumption inequality in response to the observed increase in US income inequality since the early 1980s could be explained by an endogenous relaxation of the limited-commitment constraints in response to higher income risk. This is because the outside option of default is punished by exclusion from insurance contracts, which becomes more costly as incomes become more volatile, increasing the threat of market exclusion after default and thus relaxing the incentive constraints that limit insurance transfers.

the degree of asymmetry in the model at realistic village sizes is a quantitative question.

The second contribution of this paper is to document the asymmetry of the consumption process implied by the standard limited commitment model, where households are restricted to deviate into individual autarky, for villages of realistic finite sizes, and compare it to that observed in the data. For this purpose, we use the well-known ICRISAT dataset on rural villages in India to estimate an income process with persistence. We first show how, when assuming a constant interest rate charged by a 'village lender' as in Townsend (1994) and Kinnan (2010), the asymmetries in consumption are, trivially, identical to those in the continuum model. This is because individual consumption processes in a village where the interest rate, and the insurance prices it implies by arbitrage, are constant are observationally equivalent to those that arise in the general equilibrium of a large community where the absence of aggregate risk implies a constant price of insurance. More importantly, we show how, when solving the general equilibrium of the dynamic village economy, the asymmetries at typical village sizes are similarly strong. We document how they arise, in fact, already at small single-digit group sizes.

We conclude from this that, despite its a priori appeal and ability to predict partial consumption risk-sharing, the limited commitment model has strongly counterfactual implications even for village size communities. Importantly, however, this counterfactual implication hinges crucially on the size of the insurance group, as smaller groups are more symmetric. This gives an important role to the standard assumption in models of village risk-sharing, that the insurance group equals the whole village.

Our third contribution is to show how the observable implications of limited commitment-insurance are much less counterfactual once we relax this assumption and determine the equilibrium size of insurance groups as a function of the time-varying costs and benefits of extra members. We achieve this by relaxing the standard assumption that the outside option to insurance is financial autarky. Importantly, rather than making this outside option more appealing though consumption smoothing opportunities such as saving or an option to regain access to the insurance group in the future, we, quite intuively, give agents the opportunity to deviate together with other members of the group.

Crucially, these 'coalitional deviations' endogenously limit the equilibrium size of the insurance group (Genicot and Ray 2003). This is because the costs of making transfers to additional members when they have low income may be too high relative to the outside option of keeping those transfers and deviating with a subgroup of currently high income agents. We show how with these 'coalitional deviations', the equilibrium size of insurance groups is typically very small - between 2 and 5 agents.

Note that the model with coalitional deviations has an additional appealing feature: even for a given size of an insurance group, the opportunity to deviate with other village members makes the outside option to the insurance scheme more attractive, and thus limits the transfers that are sustainable relative to the standard model. To the extent that it predicts too strong risk-sharing this is an additional, welcome implication of the more general version of the model.

Our results show first how, for preference parameters similar to those estimated in previous studies, the alternative specification of the model underpredicts the degree of insurance but lacks the counterfactual asymmetries predicted by the standard specification. We then estimate preference parameters for both specifications, using a simulated method of moments procedure that aims to match both the degree of consumption risk-sharing and the asymmetry in the data. The standard specification is unable to deliver realistic degrees of risk-sharing and symmetry. Specifically, it either can predict the symmetry in the data at the price of counterfactual full insurance against idiosyncratic income shocks; or it predicts realistic degrees of insurance at low values of risk aversion, at the price of counterfactual asymmetries. The alternative specifiction with coalitional deviations, in contrast, predicts both realistic degrees of insurance and symmetric distributions, thanks to small predicted insurance group sizes of between 2 and 4 households and stronger estimated risk-aversion. We think these results should make us reconsider our view of income risk-sharing in poor rural communities. Particularly, we should think about replacing our view of 'village risk-sharing' with one of neighbourhood, or kinship, risk-sharing in smaller groups.

2 The asymmetric nature of risk-sharing in the standard limited commitment model

This section presents the standard model of risk-sharing under limited commitment (Kocherlakota 1996, Ligon et al. 2002, Attanasio and Rios-Rull 2000), where individuals can drop out of a village-level risk-sharing scheme at the price of living in autarky forever. Since this outside option of autarky is more attractive in times of high income, the model incorporates an asymmetry in the process for individual consumption: as participation constraints are more likely to bind in periods of rising income, consumption growth responds strongly to income rises. Individuals whose income declines, however, share equally the fall in resources available to them after their constrained peers are satisfied.

Broer (2013) shows how this leads to strong asymmetries in large economies with a continuum of (infinitely many) agents and no aggregate risk. In villages with a finite number of individuals, however, Broer (2013)'s results no longer hold as the village budget constraint ties consumption decreases of unconstrained agents to the time-varying total increase in consumption of their constrained neighbours. That this can break the asymmetry is trivial to see in the case of a two-person village: as one household's binding participation constraint demands a rise in its consumption share, the other's falls by an equal amount. The degree of asymmetry in villages of more than 2 households, however, was previously unknown. As this section shows, the asymmetry is in fact very strong at usual village sizes. And, importantly, we show how it arises already in small insurance groupd with a single-digit number of households.

2.1 A standard limited commitment village economy

We consider a community with N households. In each period $t = 1, 2, ...\infty$, a household i receives endowment of the only consumption good $y^i(s_t)$, where $s = 1, ..., S$ is the state of nature. The state of nature follows a Markov process with the probability of transition from state s to state r given by π_{sr}.

Households are infinitely lived and discount the future with a common discount factor δ. They have identical and twice continuously differentiable utility functions $u(\cdot)$ defined over consumption $c^i(s_t)$ in period t and state s_t. Households are risk-averse and would therefore find it profitable to enter into a risk-sharing arrangement with other villagers in order to smooth consumption in the face of idiosyncratic income movements.

Households have perfect information about both their own income realizations and those of other villagers, but are not able to write binding contracts. Instead, contracts must be self-enforcing, which requires that at any point in time, and in particular when a household has a high income realization, helping out those with a low income realization must be preferable to reneging. Insurance transfers can be sustained in such a contract because reneging on the contract is punished by being excluded from all future insurance possibilities.

If a household is not part of an insurance arrangement, its consumption equals its (volatile) income in each period and expected life-time utility is

$$V_s = E_s \sum_{t=s}^{\infty} \delta^{t-s} u(y_t). \tag{1}$$

This is a household's outside or autarky option.

An incentive-compatible risk-sharing contract among n households can be interpreted as the equilibrium of an infinitely repeated game sustained by the threat of reversion to autarky. More specifically, an insurance contract is a vector of net transfers $(\tau^i(s_t))_{i=1}^n$ for each state s_t and history of the game consisting of the previous states.

To find the constrained-optimal insurance contract, we can write down the dynamic programme that solves for the Pareto frontier in an insurance group of size n. In particular, we maximise the utility of agent n taking as state variables the promised life-time utilities of the other $n-1$ agents, which summarize the history of the game up to the current period (Abreu, Pearce and Stacchetti 1990, Ligon et al. 2002).

The constrained-optimal contract is the solution to the following Lagrangian:

$$U_s^n(U_s^1, U_s^2, ..., U_s^{n-1}) \quad = \max_{((U_r^i)_{r=1}^S)_{i=1}^{n-1}, (c_s^i)_{i=1}^n} u(c_s^n) + \delta \sum_{r=1}^{S} \pi_{sr} U_r^n \left(U_r^1, ..., U_r^{n-1}\right) \tag{2}$$

subject to a set of promise-keeping constraints

$$\gamma^i \quad : \quad u(c_s^i) + \delta \sum_{r=1}^{S} \pi_{sr} U_r^i \geq U_s^i \quad \forall i \neq n, \tag{3}$$

a set of enforcement constraints for each household i and each state $r = 1, ..., S$ in the next period

$$\delta \gamma^i \pi_{sr} \phi_r \quad : \quad U_r^i \geq u(y_r^i) + \delta V_r, \tag{4}$$

and an aggregate resource constraint

$$\omega \quad : \quad \sum_{i=1}^{n} y_s^i \geq \sum_{i=1}^{n} c_s^i \tag{5}$$

2.2 Asymmetries in the consumption-income distribution

It is easy to see from equation (4) how the limited commitment friction implies asymmetries in the joint process of consumption and income growth: since the enforcement constraint is more likely to bind in periods of rising income - when the right-hand side of equation (4) increases in value - consumption growth reacts more strongly to income rises than to income declines. This leads to a kink in both the conditional mean and variance functions of consumption growth conditional on income growth around zero. Note, also, that the degree of symmetry trivially depends on the degree of insurance. When risk-sharing is either absent or perfect, both the unconditional distribution of consumption growth and its distribution conditional on income growth are trivially symmetric (as consumption growth of all agents equals aggregate and individual income growth respectively). Similarly, when individual income movements are small relative to movements in aggregate village incomes, individual consumption growth will be dominated by movements in aggregate income that move consumption of all villagers in parallel.

To see how, with limited risk-sharing and sufficiently important idiosyncratic risk, the standard limited commitment model of the previous section can lead to asymmetries in the joint process of consumption and income growth, consider the case of log preferences $u = \log(c)$ and income transitions that are i.i.d. over time. With log preferences, the first order condition for consumption of agent i is

$$c_r^i = \gamma^i \frac{1+\phi_r^i}{1+\phi_r^n} c_r^n.$$

where ϕ_r^i is the Lagrange multiplier associated with the participation constraint for agent i in state r in 4.

Together with the village budget constraint, this implies that consumption of individual j is

$$c_r^j \quad = \quad \frac{\gamma^j(1+\phi_r^j)}{\sum_{i=1}^n \gamma^i(1+\phi_r^i)} Y_r \tag{6}$$

where Y_r is village income in state r. Taking first log differences of equation (6), $d\log x_t = \log(x_t) - \log(x_{t-1})$, yields

$$d\log(c_t^j) \quad = \quad d\log(Y_t) + \log(1+\phi_t^j) - \log(1+\frac{\sum_{i=1}^n \gamma^i \phi_r^i}{\sum_{i=1}^n \gamma^i}).$$

Individual consumption growth is thus the sum of three terms: first, it is proportional to output growth $d\log(Y_t)$; second, it has an individual-specific term $\log(1+\phi_t^j) \geq 0$ that is positive when agent j has a binding constraint, and zero otherwise; and finally, there is a negative 'drift-term' $-\log(1+\frac{\sum_{i=1}^n \gamma^i \phi_r^i}{\sum_{i=1}^n \gamma^i}) \leq 0$ that is negative whenever at least one participation constrained is binding in the village.

Equation (7) reveals the strong asymmetry in the consumption effect of income changes between constrained and unconstrained agents: constrained agents' consumption growth is a function of their current income (since $\frac{\partial \phi_t^j}{\partial y_t^j} > 0$). Unconstrained agents, in contrast, share identical consumption growth that is smaller than that of any constrained agent and independent of their current or past individual incomes. This implies a 0 variance in consumption growth across unconstrained agents in any period t, and a 0 correlation of income and consumption growth across unconstrained periods, after controlling for village income. Constrained agents, however, will exhibit a positive correlation between consumption and income growth across time (as $\phi_t^j > 0$ is increasing in income growth), and a positive variance of consumption growth across agents with different income growth in any given period.

Apart from the assumption of log-preferences and i.i.d. transitions, these analytical results require that we can distinguish constrained from unconstrained individuals. Note however that the right-hand side of the participation constraint 4 is increasing in individual endowment y_r^i.

It is thus more likely to bind in periods of rising income.

To illustrate how the asymmetry discussed above translates to the joint distribution of consumption and income growth in a more general setting with persistence, we study two frequently used quantitative specifications of the limited commitment model. First, a simplified version of the general equilibrium environment presented in Ligon et al. (2002). And second, a partial equilibrium version where a village lender has access to outside funds at a constant interest rate, as in Kinnan (2010) or Karaivanov and Townsend (2014).

For this illustrative exercise we choose preference parameters that are representative of previous estimates. Specifically, we assume that period utility has unit-relative risk aversion (log-preferences), and that the annual discount factor equals 0.9, which are values close to the mean of those considered in (Laczo 2014), and within the range of values estimated by Ligon et al. (2002). Later, we will choose preference parameters to match moments of village data explicitly. Note that, relative to Ligon et al. (2002), we make the same simplifying assumption as Laczo (2014) that incomes follow independent realisations of an identical AR(1) process with moderate persistence parameter 0.76 and a variance of shocks equal to 0.23 (Laczo (2014)'s estimates for Aurepalle, one of the villages in the ICRISAT dataset that we look at in the following sections).

In the general equilibrium version we set the number of households equal to 34, corresponding to the number of observations in Aurepalle. We then solve the model with the standard algorithm presented in Ligon et al. (2002)).[3] Figure 2 shows a scatter plot of the resulting joint distribution of consumption and income growth (upper panel) and their residuals from a regression on time dummies (lower panel). Moving from left to right (low to high income growth), the upper panel exhibits clearly increasing conditional variances, and seemingly an increase in the gradient of the conditional mean function.

As equation (7) suggests, both features become more evident when looking at residuals that condition on aggregate income through a regression on time dummies in the lower panel. Consumption growth is approximately constant and independent of income growth during periods of income declines. In periods of rising income, however, consumption growth is more dispersed around a conditional mean that is strongly increasing in income.[4]

In Figure 2, we present similar results for a partial equilibrium specification of the dynamic limited commitment model in a village economy, where an outside lender with access to funds at a constant interest rate $R = 1.06$ acts as an insurer to individual villagers (as in Kinnan (2010) or

[3] As in Laczo (2014), we approximate the individual income process as an AR(1) process. We discretise the process using Rouwenhorst (1995)'s method with 6 support points. Again as in Laczo (2014), income of the rest of the village is estimated as an AR(1) on the basis of 33 independent realisations of individual income, and discretised with 5 support points.

[4] In the appendix we derive the correspondence between income growth and decline and being constrained.

Karaivanov and Townsend (2014).[5] The assumption of outside insurance at a constant interest rate breaks the link between aggregate village income and aggregate consumption and makes individual i's consumption process independent from that of her peers. This can be interpreted as changing the decomposition of consumption growth in equation (7), where the drift term now becomes a constant that is a function of the interest rate only. The asymmetry in the individual consumption process that results from insurance through a village lender is thus the same as the asymmetry in the stationary cross-sectional distribution of consumption in a continuous agent limited-commitment economy derived in Broer (2013). As is easy to see, the asymmetry is even starker in this partial equilibrium version than in the general equilibrium economy whose joint distribution of consumption and income growth was depicted in Figure 1.

In Figures 1 and 2 we concentrate on the joint distribution of consumption and income growth. Ligon et al. (2002) discuss how the general equilibrium version of the limited commitment model fits the joint distributions of both growth rates and levels. Concentrating on growth rates, however, has two advantages: first, income and consumption growth in the data are independent of any fixed effects that may lead to constant differences in the levels of consumption and income across households in the ICRISAT sample. For example, any error in adjusting consumption and income for constant differences in household size would typically affect their levels much more than their growth rates. A second advantage follows from the discussion in section 2.2: the asymmetry implied by the standard model surfaces in a very transparent way through non-linearities in the conditional mean and variance functions of consumption growth conditional on income growth.

Table 1 presents four key moments that summarise both the degree of insurance, and the asymmetry of the joint distributions in Figures 1 and 2. To capture the degree of insurance it shows the regression regression coefficient of consumption growth on income growth β_{dcdy}, as a measure of the average effect of income changes on consumption, and the relative variance of consumption and income growth $\frac{Var_{dc}}{Var_{dy}}$, capturing the volatility of consumption relative to incomes. To capture the asymmetry in the joint distribution, we look at two moments that capture non-linearities in, respectively, the conditional variance and mean of consumption growth as a function of income growth. Specifically, because Figures 1 and 2 suggest non-linearities in the conditional means and variances around zero income growth, we first look at the relative variance of consumption and income growth of households that experience positive income growth divided

[5] The risk-sharing properties of this setup are very similar to that in stationary general equilibrium economies with limited commitment and an infinite number of agents (considered by Broer (2013), Krueger and Perri (2006), and others). Specifically, for an identical exogenous income process that is independent across agents, whenever the exogenous interest rate charged to the village lender R equals the equilibrium interest rate in the continuum economy, the individual consumption processes will be identical in the two setups. However, aggregate income and consumption in the continuum economy are constant due to a law of large numbers, while they fluctuate in the village with a finite number of agents.

Figure 1: Consumption and income growth in general equilibrium

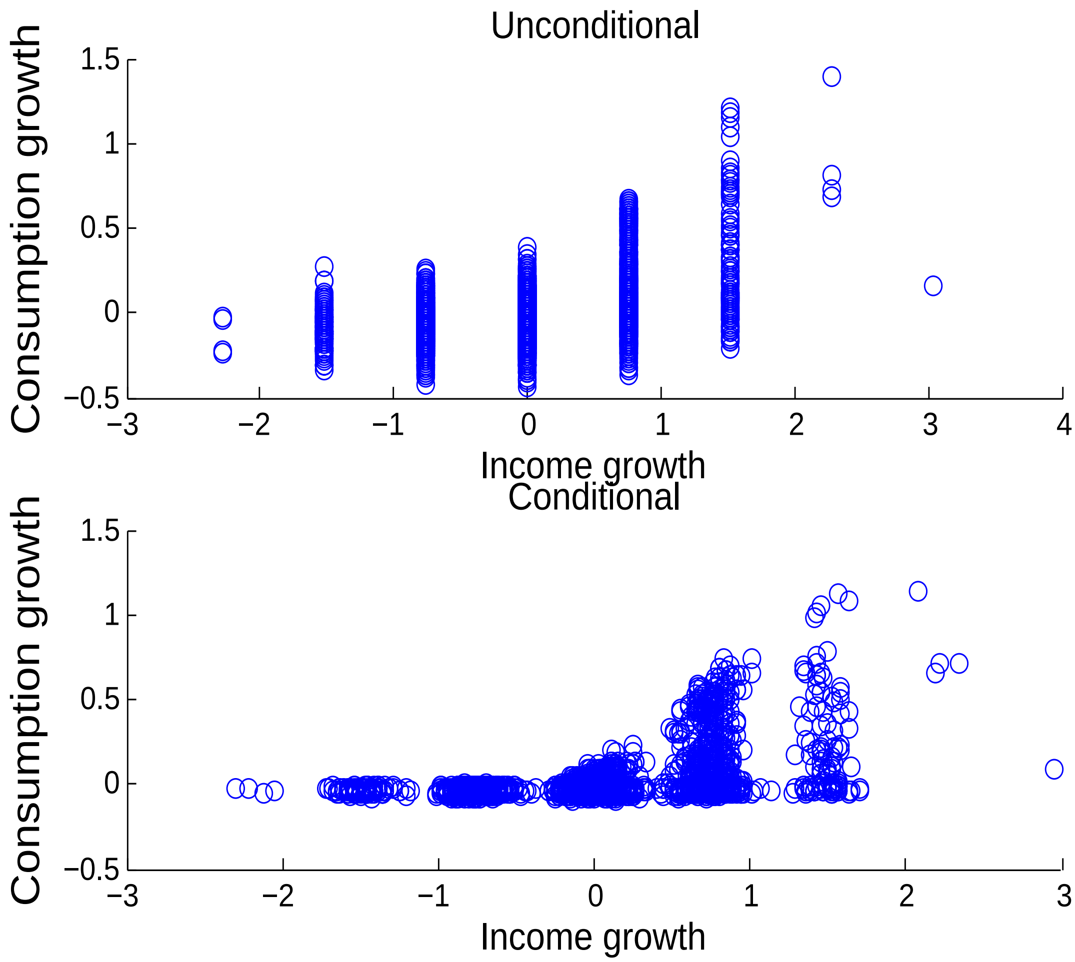

The figure shows a scatter plot of consumption and income growth from a simulation of the standard general equilibrium version of the limited economy. The top panel shows the raw data, the bottom panel the differences in residuals from a regression of both log-income and log-consumption on time dummies, that controls for movements in aggregate resources.

Figure 2: Consumption and income growth with a village lender

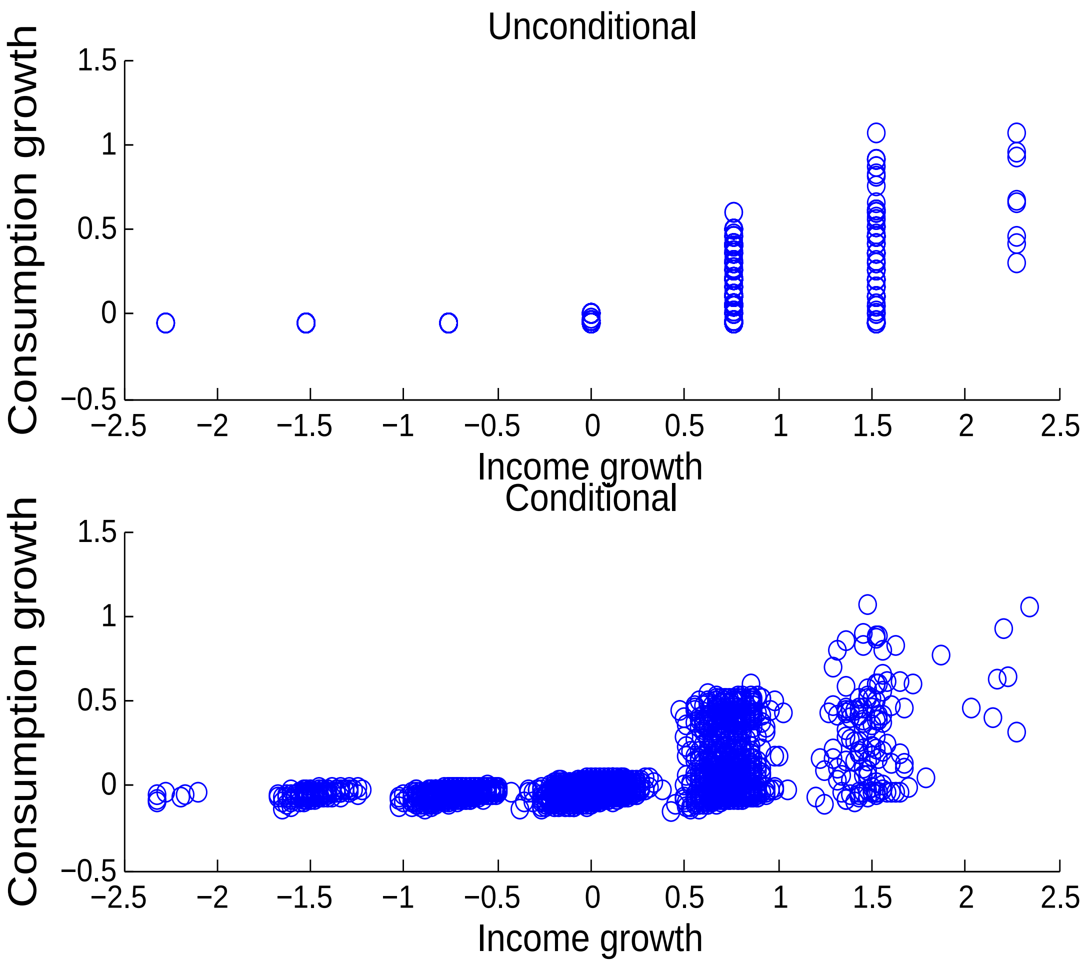

The figure shows a scatter plot of consumption and income growth from a simulation of the partial equilibrium version of the limited economy, for an exogenous interest rate of $R = 1.06$. The top panel shows the raw data, the bottom panel the differences in residuals from a regression of both log-income and log-consumption on time dummies, that controls for movements in aggregate resources.

by those that do not $\frac{Var_{dc|dy>0}}{Var_{dc|dy\leq 0}}$. Second, to capture non-linearities in the conditional mean, we look at the ratio of regression coefficient of consumption growth on income growth for households with rising and non-rising income ($\beta_{dcdy|dy>0}$ and $\beta_{dcdy|dy\leq 0}$). Note that for the calculation of these moments, we lump periods of constant income together with those of falling income. Relative to an alternative procedure that leaves periods of constant income aside in the moment-calculation, this does not change the substance of the results but has the advantage of generating finite moments even in the partial equilibrium model, where the variance of consumption growth across periods of strictly falling incomes is zero.

Table 1 shows how both versions of the standard limited commitment model predict very strong risk-sharing, as summarised by a variance of consumption growth that is only between 4 and 7 percent that of income growth on average, and a regression coefficient of log consumption changes on log income changes between 8 and 12 percent. In the general equilibrium model both measures are reduced somewhat when looking at residuals from a first stage regression on time dummies. This is because that procedure eliminates all variation in average village income from the data. Since village income equals village consumption, that reduces the covariance of the two one for one. And it reduces variation in individual consumption by more than that in individual incomes, which contain a more important idiosyncratic component. For the partial equilibrium model, where there is no aggregate resource constraint that ties aggregate income and consumption together, the coefficients are the same for raw data and residuals.

Importantly, the asymmetry is very strong for all versions of the model. In the general equilibrium version, simple, or 'raw', movements in consumption are five times more volatile in periods of rising income than when incomes fall or are constant. The ratio of regression coefficients for rising and non-rising incomes are of similar magnitude. As Figure 2 suggests, both measures are an order of magnitude higher in the partial equilibrium model. Using residuals strongly increases the asymmetry for the general equilibrium version by eliminating aggregate income and consumption movements that are, by the nature of the income process, symmetric. In the partial equilibrium model, the ratio of regression coefficients is slightly increased, while that of variances falls by about two thirds.[6]

Both standard versions of limited commitment coinsurance in village economies are thus able to predict strong but imperfect risk-sharing. And both predict very strong asymmetries in

[6] As Figure 2 shows, the fall in the relative variance is due to the increase in the conditional variance of consumption, which has a more important effect on the small variance of consumption in periods of income falls in the denominator. The increase in the ratio of regression coefficients results from the fact that observations with zero income growth are numerous in raw data because of strong income persistence and enter the calculation of the moment in the denominator. After a regression on time dummies, as Figure 2 shows, some of these observations become slightly positive, and are thus attributed to the denominator. As Figure 2 suggests, that increases the regression coefficient for income increases.

Table 1: Key moments of the standard dynamic limited commitment model

	GE raw	**PE raw**	**GE condit.**	**PE condit.**
n	33.00	NaN	33.00	NaN
β	0.90	0.90	0.90	0.90
σ	1.00	1.00	1.00	1.00
$\frac{Var_{dc}}{Var_{dy}}$	0.07	0.05	0.04	0.05
β_{dcdy}	0.11	0.12	0.08	0.12
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy\leq 0}}$	4.77	113.65	55.75	34.48
$\frac{\beta dcdy>0}{\beta dcdy\leq 0}$	5.30	14.77	17.78	17.32

The table shows, for the general equilibrium ("GE", in column 1 and 3) and partial equilibrium ("PE", columns 2 and 4) versions of the standard limited commitment model four key moments of the joint distribution of consumption and income growth: the regression coefficient of consumption growth on income growth β_{dcdy}, as a measure of the average effect of income changes on consumption; the relative variance of consumption and income growth $\frac{Var_{dc}}{Var_{dy}}$, capturing the volatility of consumption relative to incomes; the relative variance of consumption and income growth of households that experience positive income growth divided by those that experience negative income growth $\frac{Var_{dc|dy>0}}{Var_{dc|dy\leq 0}}$, to capture non-linearities in the conditional variance function; and finally, the ratio of regression coefficient of consumption growth on income growth for households with rising and non-rising income ($\beta_{dcdy|dy>0}$ and $\beta_{dcdy|dy\leq 0}$). The measures in columns 1 and 2 use simple, or 'raw', changes in log-consumption and income, while columns 3 and 4 are based on residuals from a regression on time dummies.

the joint distribution of income and consumption.

2.3 Asymmetry and village size

This section further investigates the origin of the asymmetry pointed out above. Specifically, we will show that the asymmetry in the joint distribution of consumption and income is increasing in the size of the insurance group and does not arise for small groups. To see this, note that, trivially, the budget constraint imposes symmetry on the joint income and consumption process in a village with only two agents $i = 1, 2$. This is most evident when looking at consumption and income as shares of total resources $Y = C$. Since the changes in both individual consumption shares and income shares are reciprocal to each other, such that $d\frac{c_t^1}{C_t} = -d\frac{c_t^2}{C_t}$ and $d\frac{y_t^2}{Y_t} = -d\frac{y_t^2}{Y_t}$, their marginal and joint distributions are trivially symmetric.

To see how the degree of asymmetry in the growth rates of consumption depends on the size of the insurance group, we use the approximation $dlog(x_t) \approx \frac{dx_t}{x_{t-1}}$ to write equation (7) as

$$dlog(\frac{c_t^j}{Y_t}) \;\; \approx \;\; \phi_t^j - \frac{\sum_{i=1}^n \gamma^i \phi_r^i}{\sum_{i=1}^n \gamma^i} \tag{7}$$

For $n = 2$, this implies that, when only agent 1 has a binding constraint, we have

$$d\log(\frac{c_t^1}{Y_t}) - d\log(\frac{c_t^2}{Y_t}) \approx \frac{(\gamma^2 - \gamma^1)\phi_t^1}{\gamma^2 + \gamma_1} \tag{8}$$

This implies an approximately symmetric distribution of the corresponding log-changes, and of the distribution of consumption and income growth conditional on aggregate income, as long as the shares are not too different.

For $n > 2$, there are additional degrees of freedom in the consumption allocation, so the budget constraint does not imply symmetry of consumption or income shares anymore. From equation (7), the growth in consumption shares is a weighted difference of the individual Lagrange multiplier ϕ_t^j and the "average" multiplier $\frac{\sum_{i=1}^n \gamma^i \phi_r^i}{\sum_{i=1}^n \gamma^i}$. Intuitively, as the village grows in size, the variance in the average declines faster than that of the individual multiplier, leading to an increase in the asymmetry of consumption responses to income rises and income falls. This effect is most pronounced in the limit case of an infinitely large village with a continuum of households because the average multiplier converges to a constant in that case.[7]

While the magnitude of the asymmetry is thus clear for pairs and for infinitely large groups, or equivalently to a village that can borrow at an exogenous interest rate, it is difficult to make precise statements about how the asymmetry evolves when group size gradually increases. We therefore turn to simulations and reestimate the general equilibrium version of the dynamic limited commitment model in Section 2.2, letting group size go from 1 to 34, the size of the economy in the general equilibrium version of the model.[8]

The asymmetry increases strongly and quickly in the size of the insurance group. In Figure 3, we plot the two key ratios describing the asymmetry in Table 1 as a function of village size, for both raw data and the measures based on residuals from a regression on time dummies. Specifically, we first plot the ratio of the variance of consumption growth for periods of rising relative to those of falling incomes (as a solid line, on the left-hand-side axis); and second, the ratio of regression coefficients of consumption on income growth for those with income growth relative to those with income declines (as a dashed line, on the right-hand-side axis). The ratios of these moments are close to 1 only for small group sizes (i.e. $n = 2, 3$), and increase dramatically thereafter. And in line with Table 1, the increase is steeper for the residual-based

[7]With iid incomes, aggregate income is constant with infinitely many agents, and income transitions have a stationary distribution. Krueger and Perri (2011) show how there exists a stationary consumption distribution in this economy at a given interest rate, and Broer (2013) shows the existence of a constant equilibrium interest rate $R^\star$, under some conditions. This implies that the second term in (7) converges to a constant that can be expressed as a function of $R^\star$.

[8]To do this, rather than estimating a coarse AR(1) process on an increasing state space of aggregate income values as the insurance group increases, we use the whole set of aggregate income values as the state space.

measure.

This section has shown the strongly asymmetric nature of risk-sharing in the standard limited commitment when insurance groups are even of moderate size. In small groups, to the contrary, insurance is much more symmetric. Although we have used a simple version of the standard model, we are confident that our results would hold also in a more general context with, for example, heterogeneity in income processes within villages (as in Ligon et al. (2002)) and with heterogeneous risk preferences (Laczo 2014). This is because the asymmetric nature of the participation constraint is independent of the assumption on income processes. And Broer (2013) shows precisely the robustness of the standard model's asymmetries to preference heterogeneity, if in a context of a stationary economy, equivalent to a village lender. The next section looks at the structure of risk-sharing in the village economies that have been used most widely to study models of risk-sharing: the ICRISAT panel. As we will see, we do not find any of the asymmetries implied by the standard limited commitment model there. This presents a challenge to the standard limited commitment model, which predicts that insurance groups should be infinitely large – usually interpreted as comprising all members of an economy. In the remainder of the paper we then take up the challenge and present a model that endogenously determines group size in a risk-sharing community with limited commitment.

3 The Data

Data on consumption and income in large developed economies such as the US typically show a degree of consumption risk-sharing that is high but not perfect, and comovements of consumption and income that are roughly symmetric (in the sense that both conditional mean and variance functions of consumption growth are approximately linear in income growth (see Broer (2013) and also Krueger and Perri (2005)). Risk-sharing has previously also been found to be strong in rural village economies. But less is known about the symmetry of the comovement between consumption and income. This section looks at the village economies that have been used most widely to study models of risk-sharing: the ICRISAT panel. We confirm the strong degree of risk-sharing found in previous studies. We also show how income and consumption comovement is largely symmetric, in contrast to the strong asymmetry predicted by the standard limited commitment model in section 2.2.[9]

[9]The ICRISAT panel data set has been used to test the Pareto-efficient risk-sharing model with homogenous preferences (Townsend 1994), with decreasing relative risk-aversion (Ogaki and Zhang 2001) and with heterogenous risk preferences (Mazzocco and Saini 2012). It has also been used to test the dynamic limited commitment with homogenous preferences (Ligon et al. 2002) and with heterogeneous risk preferences (Laczo 2014).

Figure 3: Key moments in village economies of different size

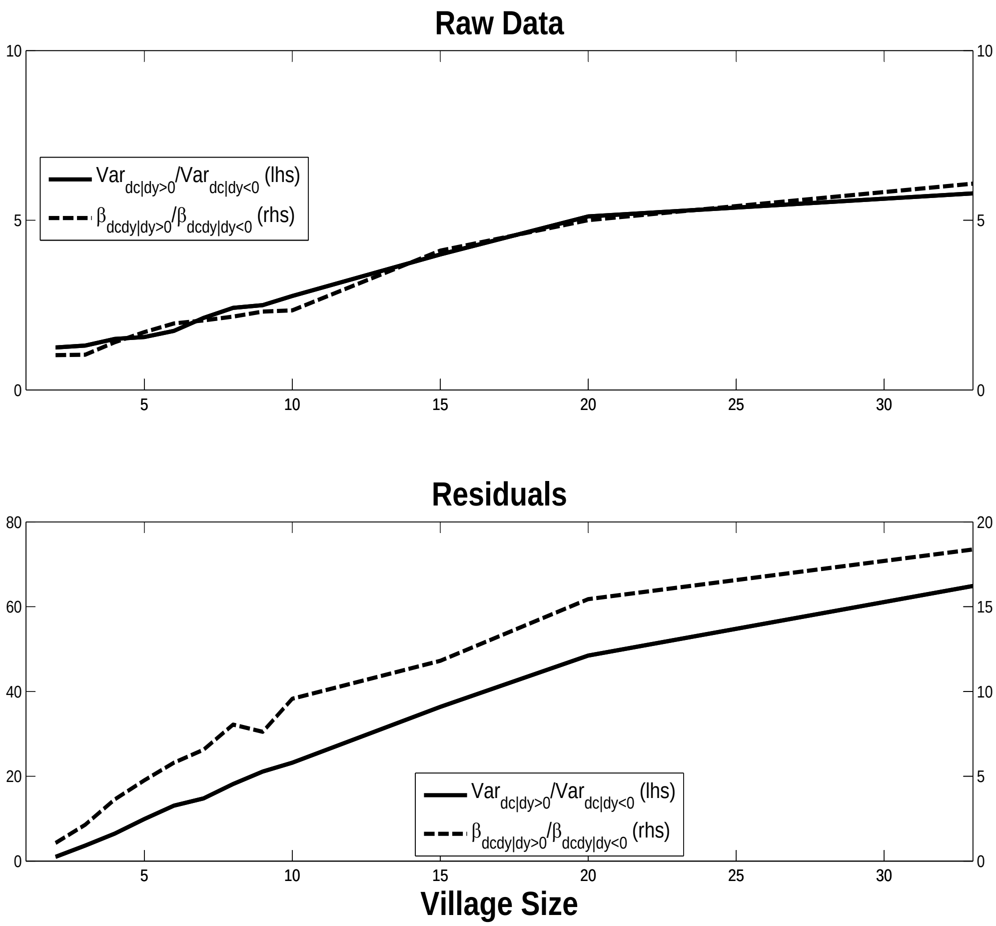

Table 2: Descriptive statistics

	Aurepalle		Kanzara		Shirapur	
Variable	Mean	Sd	Mean	Sd	Mean	Sd
Consumption	1623.10	704.31	2095.43	1113.91	2359.87	1075.85
Consumption (aeq.)	303.47	127.86	400.84	161.42	430.37	170.71
Income	3787.41	3734.31	5623.42	5524.55	4432.26	3490.73
Income (aeq.)	629.58	429.78	984.42	742.54	792.16	577.58
Aeq. household size	5.95	2.70	5.66	2.68	5.85	2.52
No. of observations	204		222		186	
No. of households	34		37		31	

Notes: Monthly consumption and income measured in 1975 Indian rupees per year. In 1975, approximately 8 Indian rupees were worth 1 US dollar.

The ICRISAT data come from three rural and agricultural villages in India: Aurepalle in Andra Pradesh state, and Kanzara and Shirapur (both in Maharashtra State). In each village, 40 households were sampled and detailed expenditure and income data was collected each year between 1975-1984.

For our analysis we need information on both consumption and income aggregates across households and over time. We follow Laczo (2014) and use a consumption aggregate that includes monthly expenditure on food, clothing, services, utilities and narcotics.[10] The income aggregate contains net income from farming and livestock, labour and transfers. All variables are in real and per-adult equivalent units where the same age-gender weights are used as in Townsend (1994). For comparability with other authors, we restrict our analysis to the years 1976-1981 and construct a fully balanced panel.[11]

The ICRISAT villages are poor with the average dweller living well below the $1 dollar a day poverty line (Table 2). On average, daily nondurable consumption per adult equivalent is 0.83, 1.10 and 1.18 in 1975 rupees, which is equivalent to 0.42, 0.55 and 0.59 in 2014 $ US dollars respectively. Income is somewhat higher.

Although villagers are poor on average, there is some evidence of consumption risk-sharing. In Table 3, we report the variance of consumption growth as a proportion of the variance of income growth $\frac{Var_{dc}}{Var_{dy}}$, after partialling out changes in village resources. In all three villages, insurance is present, though far from perfect with volatility of consumption relative to income ranging from 0.3 in Aurepalle to 0.56 in Kanzara.

[10] We thank Sarolta Laczo for making her data and code available to us.

[11] See Morduch (1991) and Ravallion and Chaudhuri (1997) for a detailed discussion of measurement issues in the full ICRISAT panel.

Table 3: Conditional variance of consumption

	Aurepalle	Kanzara	Shirapur
	(1)	(2)	(3)
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.56	0.33
	(0.003)	(.005)	(.003)
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy<=0}}$	0.75	0.64	0.60
	(0.023)	(.031)	(.009)
Obs.	170	185	155
No. of households	34	37	31

Notes: Standard errors in parentheses clustered at household level.

Table 4: Reduced form estimates of the degree of risk-sharing

	Aurepalle	Kanzara	Shirapur
Δ ln of aeq. consumption	(1)	(2)	(3)
Panel A:			
Δ ln of aeq. income	.206	.222	.169
	(.061)***	(.071)***	(.059)***
Time Dummies	YES	YES	YES
Obs.	170	185	155
No. of households	34	37	31
Panel B:			
Δ ln of aeq. income	.383	.450	.124
	(.124)***	(.150)***	(.122)
Δ ln of aeq. income >0	-.236	-.101	.035
	(.186)	(.138)	(.135)
Obs.	166	181	151
No. of households	34	37	31

Notes: Standard errors in parentheses clustered at household level.

In Table 4, we regress the growth rate of adult-equivalent consumption on the growth rate of adult-equivalent income – controlling for changes in village resources by including a full set of year dummies. The coefficients on the growth of adult-equivalent income imply that a 1% change in income leads to roughly a 0.2% change in consumption. The effect is fairly uniform and significant across all three villages.

In Figure 4 we plot the joint distribution of consumption and income growth (left panel) and their residuals from a regression on time dummies (right panel). In contrast to the simulated model solutions in Figure 1 and 2, the data appear strongly symmetric: neither the variance of consumption nor the response of consumption to income look dramatically different as households go from negative to positive income growth.

In the second row of Table 3 and panel B of Table 4 we confirm the impression of symmetry. In Table 3, we report the relative variance of consumption and income growth of households that experience positive income growth divided by those that experience negative or zero income growth. In panel B of Table 4, we regress consumption growth on income growth and the interaction of income growth and a dummy for whether the household saw positive income growth.[12] Both sets of results point in the same direction: The ratio of the variances is in all cases below but close to 1, and there is no significant difference in the amount of smoothing obtained for positive and negative income growth. In both cases – and in stark contrast to the implications of the dynamic limited commitment model in the previous section – the magnitudes imply that, if anything, those with income losses have less insurance.

[12]Panel B of Table 4 differs from Panel A in two ways: first, as can be seen from Figure 4, the data contains outliers (both in consumption and income growth). We therefore exclude the observations falling into the top and bottom 1% of income growth (4 in each village). Inclusion of the outliers affects the results in the two sub-samples (but not the whole sample), making the interaction even more negative in the case of Aurepalle (where it becomes significant) and Kanzara and turning it from positive to negative in Shirapur. Second, rather than including a set of time dummies in the regression, we demean both consumption and income period-by-period before the regression. This is equivalent to including time dummies in a linear regression of consumption growth on income growth as in Panel A of Table 4. It amounts to a slight difference, however, when we allow for non-linearities in the association of income and consumption growth in Panel B. Specifically, inclusion of a full set of time dummies would identify the non-linearity only from within-period differences in (already-demeaned) income growth greater than zero. Since our theoretical model does not allow the non-linearity to differ across time, we opt to retain the between-period variance for the identification of the non-linearities.

Figure 4: Consumption growth as a function of income growth in the Icrisat dataset

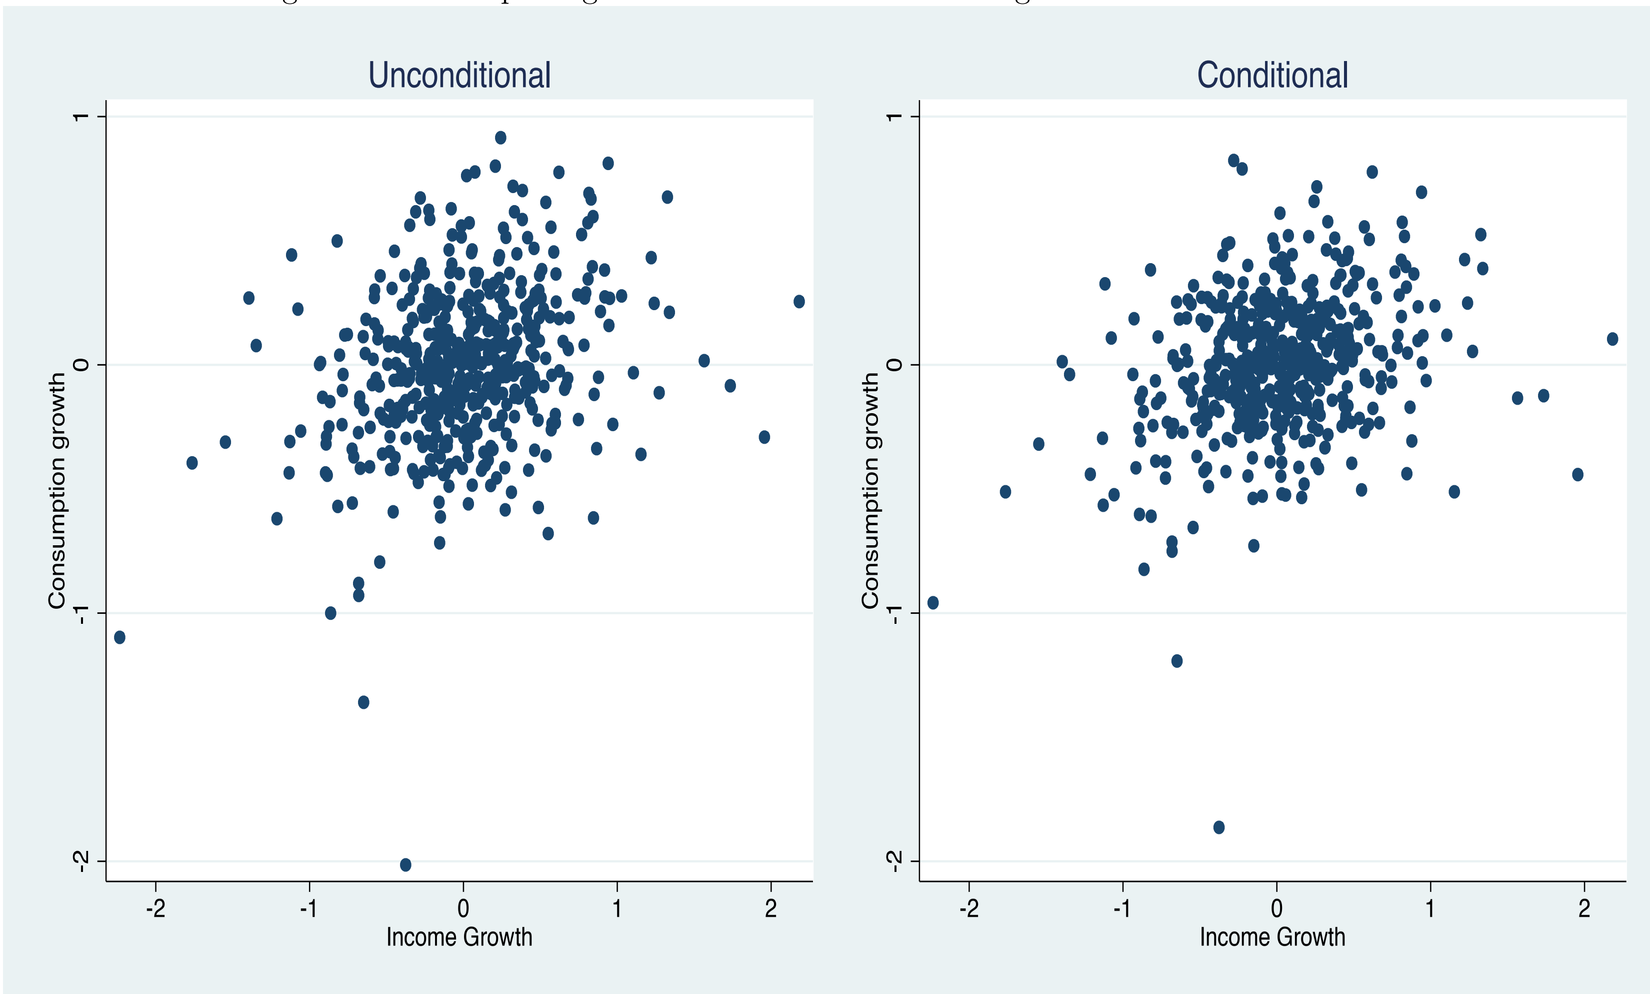

The figure shows a scatter plot of consumption and income growth in Aurepalle, Kanzara and Shirapur. The left panel shows the raw data, the right panel the differences in residuals from a regression of both log-income and log-consumption on time dummies that controls for movements in aggregate resources.

4 Generalising the model: Coalition-Proof Consumption Risk-Sharing

The standard limited commitment model almost always explains observed consumption risk-sharing better than simple alternatives such as autarky or full insurance (Ligon et al. 2002, Laczo 2014). The model has, however, two empirical short-comings. A first, well-known one, relates to its ability to capture the observed volatility of consumption and degree of insurance. Thus, the model often delivers too much insurance (especially when applied to western economies (Ábrahám and Cárceles-Poveda 2009), and struggles to capture the volatility of consumption at the same time as the cross-sectional dispersion of consumption in village economies (Ligon et al. 2002). Similarly, we found in section 2.2 that, for preferences comparable to those estimated in the literature for village economies, the standard model predicts volatility of consumption, and an association of consumption and income movements, that are smaller than in the data. In Section 2.3, we also pointed out a previously unknown problem of the standard model applied to village economies: starting already at medium group sizes risk-sharing is strongly asymmetric, in contrast to the essentially symmetric nature of the data.[13]

We now argue that by addressing a third (theoretical) short-coming of the standard limited commitment model, namely the fact the stipulated contract is in fact not renegotiation-proof, one can also substantially improve its empirical performance. More specifically, the standard dynamic limited commitment model sustains insurance by assuming that any individual that reneges on the contract is punished by being excluded from all future insurance possibilities. Such a punishment may be too harsh to implement, however. Particularly in the context of risk-sharing groups within village economies, it seems difficult to prevent those agents who are excluded from the insurance arrangement in the current period to join forces and insure each other again in the future – albeit in a smaller group.

As proposed in Genicot and Ray (2003) and analyzed in Bold (2009), the standard dynamic limited commitment model is made renegotiation-proof by abandoning the sub-game perfect Nash equilibrium solution concept in favor of coalition-proofness, which allows sub-groups of agents to depart from an insurance group and enter a new insurance mechanism thereafter. Importantly, this alternative equilibrium concept has the potential to improve the empirical performance of the standard model in two ways: first, by making deviations – to a new insurance group rather than individual autarky – more attractive, it should reduce the size of sustainable

[13]This strong asymmetry was shown by Broer (2013) for the case of large economies with a continuum of consumers. But given the symmetric nature of the two-agent version (Kocherlakota 1996) of the standard limited commitment model, the degree of asymmetry predicted by the model for small villages had remained an open question.

transfers, and thus the degree of insurance. In this sense, the alternative equilibrium concept could work like other, simpler extensions to the standard model that reduce insurance by making the outside option more attractive, such as allowing some participation in financial markets after autarky (Krueger and Perri 2006), or let agents return to an insurance mechanism with a certain probability (Broer 2013). Second, because the maximum sustainable size of an insurance group is known to be bounded in the alternative model, it might also help to reduce counterfactual asymmetries, which we showed arise everywhere except in very small groups. This is important, because the limited commitment model with individual deviations has no mechanism to generate small group sizes. Instead, the model says that group size should be increased until the marginal benefit of adding members to the risk-sharing group is zero (Murgai, Winters, Sadoulet and de Janvry 2002, Genicot and Ray 2003). This effectively implies that the risk-sharing group should comprise the entire economy, which is naturally interpreted as the village in a developing country context.[14]

We will now endeavor to derive a somewhat more realistic version of the standard dynamic limited commitment model in village economies – by allowing sub-groups who have reneged on a larger group to provide insurance to each other in the future – that is also computationally tractable. The model is set in the same environment as the standard limited commitment model presented in Section 2.1 with one exception: we now require equilibria to be coalition-proof, which requires risk-sharing groups to be robust not only with respect to individual deviations but also with respect to deviations by subgroups, provided that these subgroups are themselves robust with respect to further deviations (Bernheim, Peleg and Whinston 1987, Genicot and Ray 2003).

As shown theoretically by Genicot and Ray (2003), requiring groups to be coalition-proof, has two implications for the size of observed insurance groups (across a large class of contracts): (i) for a given group of size n, the coalition-proof equilibrium risk-sharing contract may fail to exist and (ii) the largest possible insurance group size is bounded. Both of these results open up the possibility that the equilibrium group size of insurance groups may be small once one allows for coalitional deviations. This would reduce the degree of risk-sharing as well as the asymmetry, both of which are increasing in group size. We will examine whether this is the case in practice by solving the model numerically.

Finding the constrained-efficient risk-sharing mechanism in a group of size n whose members can deviate by entering alternative risk-sharing arrangements with a subset of other group members is an order of magnitude more complicated than the standard model. In the standard

[14]Of course, one could consider other (finer) partitions to improve the match between model and data. Still, any of these partitions – families, neighbourhoods, kinship groups, networks, castes, religious or ethnic groups either inside the village or even stretching across village borders – are exogenous to the model.

limited commitment model, the set of sub-coalitions that can threaten deviation equals n. In the coalition-proof limited commitment model, the size of this set swills to $2^n - 2$. Moreover, the outside option is completely determined by current income for any member of the insurance scheme in the case of individual deviations. In the case of coalitional deviations, however, a deviating sub-group has potentially an infinite number of ways to divide the surplus from the new insurance scheme among its members.

We now present the general problem of finding the coalition-proof dynamic risk-sharing contract and then describe how we amend the standard approximation procedure, designed for the case of individual deviations, to fit the context of our model and examine its quantitative implications.

We follow Genicot and Ray (2003) by defining coalition-proofness for a group of size n recursively. Denote the set of all sub-coalitions of n by $\mathbf{J} = \{J_1, J_2, ...J_m\}$, and let $j_{i,1}, ...j_{i,|J_i|}$ denote the members of J_i. Suppose that we have defined for each sub-coalition J_i in $\mathbf{J}$ a set of credible and feasible threats they can make, $\mathbb{V}^{J_i}$. A member of this set corresponds to a $|J_i|$-dimensional vector $\mathbf{V}^{J_i} = \{V^{1,J_i}, V^{2,J_i}, ..., V^{|J_i|,J_i,}\}$ of expected life-time utilities from a particular consumption allocation for the members of J_i, where the first entry of $\mathbf{V}^{J_i}$ corresponds to agent $j_{i,1}$, the second corresponds to $j_{i,2}$ and so forth. The consumption allocation has to be feasible in the sense that it respects the aggregate resource constraint of the sub-coalition and credible in the sense that there are no further sub-coalitions that could profitably deviate from it.

Having defined the set of credible and feasible threats each sub-coalition can make, the Pareto frontier of the coalition-proof contract

$$U_s^n(U_s^1, U_s^2, ..., U_s^{n-1}) = \max_{((U_r^i)_{r=1}^S)_{i=1}^{n-1}, (c_s^i)_{i=1}^n} u(c_s^n) + \delta \sum_{r=1}^{S} \pi_{sr} U_r^n \left(U_r^1, ..., U_r^{n-1}\right) \tag{9}$$

must then satisfy the following enforcement constraints: For each coalition J_i in $\mathbf{J}$, there must be no $\mathbf{V}^{J_i} \in \mathbb{V}^{J_i}$, such that

$$\begin{aligned} U_r^{j_{i,1}} &< u(y_r^{j_{i,1}}) + V_r^1(J_i) \\ U_r^{j_{i,2}} &< u(y_r^{j_{i,2}}) + V_r^2(J_i) \\ &..... \\ U_r^{j_{i,|J_i|}} &< u(y_r^{j_{i,|J_i|}}) + V_r^{|J_i|}(J_i) \end{aligned}$$

where the superscript r on V_r indicates that it is the life-time expected utility over the probability distribution for income that ensues if r was the income state in the previous period. The usual aggregate resource and promise-keeping constraints must also hold.

The coalitional enforcement constraints describe all contracts that are stable with respect to deviations by stable sub-groups. To find the constrained-optimal contract within this set is not trivial, however, since the constraints in their form above cannot be used to define a constrained maximization problem. In particular, note that in contrast to the enforcement constraint for individual deviations, which gives an individual one threat-point (conditional on current income), the coalitional enforcement constraints leave a deviating subcoalition with a large (in fact infinite) number of possible threats since the set $\mathbb{V}^{J_i}$ may contain infinitely many elements.

Bold (2009) shows how to reduce the dimensionality of the set of possible deviations, thus allowing the problem to be cast as a constrained dynamic social planner program. Intuitively, the argument relies on the following observations: (i) any coalition-proof contract must lie on the constrained-optimal Pareto frontier of the risk-sharing group since those are the only contracts that are renegotiation-proof, (ii) therefore one only needs to consider threats by sub-groups of size m that lie on the constrained-optimal Pareto frontier of the deviating sub-group, (iii) to deter a threat it suffices to make a sub-group indifferent between exactly one point on their Pareto frontier and continuing to share risk within the group of size n as this deters all other possible threats from this sub-group, because by definition any other threat along the Pareto-frontier (or within it), would make at least one member of the deviating sub-coalition worse off.

Hence, with coalitional deviations, the outside option of any individual is optimally chosen by the social planner and will typically depend on the past utility promises and current incomes of all members of the insurance group. Intuitively, the planner can punish members more harshly by allocating them to a less preferred location on the Pareto frontier. And optimally, she promises to punish those members less harshly whom she has promised high utility under the insurance arrangement.

As the previous discussion shows, the requirement of coalition-proofness is appealing because it makes risk-sharing allocations immune against any collective deviation of any subgroup. The price of this conceptual appeal, however, is complexity. While, with individual deviations, the set of potential punishments is equal to that of actual punishments, coalition-proofness makes finding the optimal punishment code a complicated problem of its own: for a group of size n the planner must first find the optimal threat one can feasibly deter for each stable subcoalition, where feasible means that there is no other coalition that has a profitable deviation from it if the utilities corresponding to the optimal threat were actually implemented. Presuming one can find an optimal threat that is feasible to deter for each sub-coalition (if not, the group of size

n is not stable), the planner must then decide for which sub-coalition she wants to deter the optimal threat to arrive at the constrained-optimal allocation of consumption and continuation utilities in the original group.

While computation of the optimal coalition-proof allocation is feasible for small insurance groups (see Bold (2009)), it is not feasible for the size of the villages under study. We therefore adapt and extend the common approximation to the solution of the standard limited commitment model with individual deviations to the case of coalition-proof risk-sharing. The approximation to the standard model, originally proposed by Ligon et al. (2002)) and used, for example, in Laczo (2014) and Dubois, Jullien and Magnac (2008), reduces the dimensionality of finding the constrained-efficient risk-sharing allocation in a village of n members by considering the simpler problem of an individual that shares risk with an agent who represents the rest of the village of $n-1$ individuals. The use of this representative agent, assumed to have the same preferences as all village members and to receive an endowment equal to the average across $n-1$ villagers, implicitly assumes that the rest of the village can share risk in a first-best, unconstrained way. The vector of outside options of both agents equals the consumption values of individual and average incomes respectively. By reducing the number of agents in the problem to two, and by considering, typically, a coarse discretisation of both individual and aggregate income processes, this procedure significantly reduces the complexity of solving the standard model.

To adapt this standard approximation to the case of coalitional deviations, we combine the 'one-against-the-rest-of-village' strategy with a recursive identification of stable coalitions of increasing size. Our aim is to define an outside option of an individual i sharing risk with a rest group of size n that captures the idea of coalition-proofness but does not require us to optimally choose which coalitional threats to deter. To this end, we consider coalition-proof insurance arrangements between an individual and an agent representing the rest of an insurance group whose total size rises from $n=1$ to $n=N$. As in the standard model, for every n this reduces the dimension of the Pareto frontier, given by the solution to (9), from $n \leq N$ to 2.

For a given n, the outside option of the rest of the village is unchanged relative to the standard model. Rather than individual autarky, however, the individual now has the option to deviate by entering any sustainable sub-coalition of size $n-1$ or smaller. In line with the assumption of equal treatment for rest-of-village members in the standard model, we assume that, after one period of autarky, a sub-coalition promises all members the same expected value or the minimum required for sustainability, whatever is larger. This reduces the set of values that the individual can obtain in any given sub-coalition J_i from the potentially infinite set of feasible values $\mathbf{V}^{J_i} \in \mathbb{V}^{J_i}$ to a single point, as in the standard model. Since, with persistent incomes, higher income individuals are more attractive coalition partners even when deviation starts with a period of individual autarky, we restrict our attention to sub-coalitions with those village members with highest current income. In turn, this reduces, for every n the dimension of the set of possible coalitions to consider from $2^n - 2$ to $n - 1$.

Our recursive solution procedure starts with the case $n = 2$, for which the standard and coalition-proof models are identical, and uses its solution to define the value of deviation for the individual in the case of $n = 3$ as follows: for any realisation of the three-dimensional income vector, the outside option of the individual equals the value of entering, after one period of autarky, a deviation-proof risk-sharing arrangement with the richer of the other remaining villagers that promises both of its members the same utility, or their value of deviation to autarky, whatever is larger. The solution for the case of $n = 3$, in turn, defines the value to the individual of deviating to a risk-sharing arrangement with the two richest rest-of-village members in the case of $n = 4$, and so forth. Note that, typically, not all risk-sharing arrangements of size $n = 2, 3, ..., N$ will be sustainable. If this is true for $n = k$, we derive the coalition-proof risk-sharing arrangement for an insurance group of size $k+1$ with outside options based on sub-coalitions of size $k - 1$, rather than k. This recursive sustainability criterion implies also that, for every $n > 2$, we only have to consider deviations to the next smaller sustainable coalition. For every n, this further reduces the number of coalitions to consider from $n - 1$ to 1.

These successive reductions in complexity, derived on the basis of the standard one-against-the-rest-of-the-village approximation, and by applying its simplifying assumption of equal treatment to deviating sub-coalitions, reduces the complexity of the coalition-proof model considerably. There remains one dimension, however, along which this model remains considerably more complex than the standard one. In the latter, only aggregate income of the rest of the village matters, which determines outside option and resources for the insurance arrangement. With coalitional deviations, however, the whole distribution of incomes matters, as individuals deviate with a selection of the richest village members. Rather than using a coarse discretisation of rest-of-village income, we therefore have to use the whole income distribution as state variable. This puts an upper bound on the maximum village size N we can consider.

5 Results

This section presents the main results of this paper. It shows how both versions of the limited commitment model, with individual deviations on the one hand and coalitional deviations on the other, are consistent with realistic degrees of risk-sharing in small villages – if we accept estimates of risk-aversion that are in the standard model lower and in the model with coalitional deviations higher than those most common in the literature. The standard model, however, robustly predicts strong counterfactual asymmetries in the consumption process. The coalitional deviations model, on the other hand, generates a symmetric process for consumption and income growth just as observed in the ICRISAT villages. In sum, we show that the latter is more consistent with the data.

5.1 Quantitative model evaluation

The main aim of this section is to compare the two models to data moments from the three ICRISAT villages. This comparison can be performed in different ways, for example on the basis of scatter plots like those in section 2.2 (and presented, for example in Ligon et al. (2002)), or by making assumptions about the distribution of consumption and income conditional on model parameters to derive the likelihood of different model specifications given the data as in Laczo (2014).

In this paper, we concentrate on a small number of moments that describe the key risk-sharing properties of the models, as well as the asymmetries in the joint distribution of consumption and income. This approach has the disadvantage of using a limited amount of information contained in that joint distribution. Moreover, it is conditional on the choice of moments. In our view, these drawbacks of our approach are more than made up for by the advantage of summarising a complicated distribution through some key moments that have a close link to intuitive features such as the degree of risk-sharing in the models and the degree of asymmetry in their implied reaction of consumption to positive and negative income shocks. Before we can calculate these moments, however, we first have to choose a vector of parameters that allows us to solve the models. We then use the solution to simulate panel data for artificial villages, and finally calculate a vector of moments and compare them to their data equivalents.

5.2 Parameter choice

To solve the model, we need to determine the size of the insurance group, the income process and preferences.

As discussed above, in the individual deviations model, group size is unbounded – or as large as the economy. In developing countries, researchers have equated the economy with the village as the largest – albeit bounded – group a household could interact with. In the ICRISAT data set, this amounts to setting $N = 34$ when simulating and comparing the standard model to the data in Aurepalle, $N = 37$ when doing the same for Kanzara, and $N = 31$ in the case of Shirapur.

In the case of coalitional deviations, we also assume that households can only form risk-sharing arrangements with other households inside the village. Hence, the largest group that can form has the same size as in the individual deviations model, but the largest stable group

Table 5: Estimated income processes

	Aurepalle	**Kanzara**	**Shirapur**
ρ	0.70	0.77	0.61
Var_ϵ	0.29	0.15	0.22

Notes: The table presents the point estimates for the persistence parameter ρ and the shock variance Var_ϵ for the AR(1) process (11).

will typically be much smaller. While the set of stable sizes may contain several different group sizes, we assume that only the largest stable one m^{max} is formed. This implies that the village typically contains $k > 1$ insurance groups. To maintain comparability across models and with the data, we calculate moments of interest for the model with coalitional deviations based on a simulation of the smallest number of groups that comprise at the least the number of villagers in the data. Or more formally, we find the smallest k such that $N' = k \times m^{max} >= N$, where N is again the sample size in the three ICRISAT villages.

We identify a separate income process for each of the three villages. For this, we abstract from heterogeneity in the income process by assuming that log-incomes of all village members y_{it} follow an identical AR(1) process with persistence parameter ρ

$$y_{it} = \alpha + \rho y_{it-1} + \epsilon_{it} \tag{11}$$

where ϵ_{it} are mean zero shocks that are identically and independently normally distributed across households. We identify ρ and the variance of shocks Var_ϵ from the autocovariance and variance of household incomes Var_y as

$$\begin{aligned} \rho &= \frac{Cov_y}{Var_y} \\ Var_\epsilon &= Var_y * (1 - \rho^2) \end{aligned} \tag{12}$$

Table 5 presents the estimates for the AR(1) parameter ρ and the shock variance Var_ϵ for the three villages. Income shocks in Kanzara have the highest persistence (0.77) and lowest variance, a feature to bear in mind for our estimates of risk-sharing there. Income persistence is lower, but the variance higher in Aurepalle and Shirapur. For the quantitative solution of our model, given ρ and Var_ϵ we approximate y_{it} as a Markov process with 3 support points using Rouwenhorst (1995)'s method.

Following the discussion in Section 4, the cross-sectional distribution of individual incomes, which determines the value of reneging on the insurance contract by entering a coalition with the most attractive sub-group, is a state-variable of the model with coalitional deviations. We therefore cannot simply approximate the income process for the rest of the village as a discretised version of an aggregate autoregressive process for the sum of n independent realisations of (11), as in Laczo (2014). Instead, for any given size of an insurance group m, we use the whole cross-sectional distribution of incomes across the $m-1$ members of the the rest of the village as an exogenous state variable and compute the transition probabilities on the basis of the individual income process in (11).

The remaining parameters to be determined are those that govern preferences. For this, we assume that period-utility is of the constant relative risk-aversion type

$$u(c_t) = \frac{c_t^{1-\sigma}}{1-\sigma} \tag{13}$$

Below we present results corresponding to different ways of choosing the risk-aversion parameter σ and the discount factor δ.

5.3 Simulation and calculation of moments

Given solutions to the standard model and the model with coalitional deviations, we draw a vector of income realizations and then simulate consumption on the basis of the models' policy functions for N and N' households in $T = 1,000$ periods. From this simulated sample, we then calculate moments that summarize the key features of the risk-sharing mechanism. For this, we concentrate again on the joint distribution of consumption and income growth and look at the moments familiar from table 1: the regression coefficient of consumption growth on income growth β_{dcdy}, as a measure of the average effect of income changes on consumption; the variance of consumption growth as a proportion of the variance of income growth $\frac{Var_{dc}}{Var_{dy}}$, capturing the volatility of consumption relative to incomes; the relative variance of consumption and income growth of households that experience positive income growth divided by those that experience negative or zero income growth, to capture non-linearities in the conditional variance function; and finally, the difference of regression coefficients of consumption growth on income growth for households with rising and non-rising income $(\beta_{dcdy|dy>0} - \beta_{dcdy|dy\leq 0})$.

We calculate all moments after subtracting period-specific village-averages from the individual data (observed and simulated) to make our results more robust to any correlation in individual incomes not captured by the assumption of independent individual incomes.[15]

[15] As Laczo (2014) points out, the correlation of incomes across individuals in the three villages is positive, but small. Nevertheless, we decide to be conservative and condition on movements in aggregate village income.

Table 6: Risk-sharing moments for common preferences

	Aurepalle			Kanzara			Shirapur		
	Data	**ID**	**CD**	**Data**	**ID**	**CD**	**Data**	**ID**	**CD**
n	NaN	34.00	2.00	NaN	37.00	2.00	NaN	31.00	2.00
δ	NaN	0.90	0.90	NaN	0.90	0.90	NaN	0.90	0.90
σ	NaN	1.00	1.00	NaN	1.00	1.00	NaN	1.00	1.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.03	0.56	0.56	0.13	0.60	0.33	0.03	0.53
β_{dcdy}	0.21	0.07	0.53	0.22	0.20	0.63	0.17	0.08	0.52
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy\leq 0}}$	0.75	61.62	1.16	0.64	55.55	1.19	0.60	55.39	1.08
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.24	0.14	0.01	-0.10	0.36	0.00	0.03	0.15	0.00

The table shows for the data and for the individual deviations "ID" and coalitional deviations "CD" versions of the dynamic limited commitment model the four moments of interest of the joint distribution of consumption and income growth. For the simulated model solutions, the table also presents the size of the insurance groups and the preference parameters.

5.4 Model performance with common preference parameters

To understand the different predictions of the two model specifications, Table 6 presents the four moments of interest for the income process estimated in each of the three villages and the same standard preference parameters as in section 2.2.

As we saw in section 3, consumption risk-sharing in the three ICRISAT villages is strong but far from perfect, with a relative variance of individual consumption and income growth of between 0.3 and 0.56, and a regression coefficient of consumption on income growth between 0.17 and 0.22. In line with the results reported in Section 2.2, the simple version of the standard model with homogeneous incomes and preferences predicts more insurance than observed in all three villages. Specifically, the regression coefficient β_{dcdy} is predicted to equal about one third of the coefficient estimated from the data in Aurepalle and one half in Shirapur, but is closer to the data in Kanzara. More starkly, the relative consumption volatility $\frac{Var_{dc}}{Var_{dy}}$ is an order of magnitude smaller than in the data.[16]

At the same level of risk-aversion and impatience, the alternative model with coalitional deviations predicts that any coalition with more than two people (up to the size of the village) is unsustainable and all three villages are thus predicted to consist of a multitude of two-person risk-sharing groups. As a result, the model predicts less insurance than the data, but gets closer than the standard model to matching observed consumption volatility. The alternative model does not, however, give a better prediction of the observed average association between income

[16]One should also note that the degree of risk-sharing in the standard model is somewhat smaller than that in section 2.2. This is due to the different approximation of both the individual income process (where we choose 3 rather than 5 points) and the village income process (which takes on over 2000 values corresponding to all different combinations of $N = 34, 37, 31$ in our solution, rather than just 5 values).

and consumption changes: the regression coefficient β_{dcdy} is estimated to be two to three times larger than in the data.

In line with the results in section 2.3, with 2-household insurance groups, the alternative model predicts a symmetric distribution of consumption and income changes: there is no difference in the variances of consumption growth or in the linear association of income and consumption growth between households with income increases and those with decreases. In the standard model, in contrast, the asymmetry is again extreme. The alternative model thus fits better this dimension of the data, which, as already noted in Section 3, is largely symmetric and – if anything – has asymmetries in the other direction. Note that, in contrast to what the point estimates of the data moments indicate in several cases, none of the two models delivers asymmetries that correspond to declines in the slope of the conditional mean of consumption growth and in the conditional variance as income grows. This is due to the very nature of limited commitment, which makes consumption sensitive to income growth only in periods of rising income.

While these results help to understand the economic mechanism at work in the models, they are conditional on a parameterisation of preferences that is close to estimates found in previous studies but ultimately arbitrary. The next section presents results for all three villages when preference parameters are estimated on the basis of a standard criterion.

5.5 Estimating the model

We now estimate the preference parameters using a simulated method of moments approach. The goal here is to choose the discount factor δ and coefficient of risk aversion σ to minimize the squared weighted distance between the selected moments from the ICRISAT villages and moments from simulated samples generated by the standard dynamic limited commitment model and the alternative with coalitional deviations.

In order to understand the estimates, it is useful to recall the role of discount factor and risk-aversion for risk-sharing in both models. Higher risk-aversion trivially makes risk-sharing more desirable, and, even for given values of deviating (to individual autarky or a smaller risk-sharing group in the standard and alternative models, respectively), increases the constrained optimal degree of risk-sharing. In addition, it also makes the increase in consumption volatility that would follow from a deviation more costly, and thus increases the sustainable size of transfers and therefore the degree of risk-sharing in equilibrium. Since deviation delivers higher mean consumption in earlier periods at the price of eternally higher consumption volatility, higher discount factors also deter deviation and increase risk-sharing. Finally, in the coalitional deviations model, a rise in in risk-aversion also increases the benefits from increasing the size of the insurance community, leading to larger risk-sharing groups in equilibrium.

For the estimation of the models, the criterion to be minimized is

$$\Lambda(\delta,\sigma) \quad = \quad [f-g(\delta,\sigma)]'W^{-1}[f-g(\delta,\sigma)]$$

where f is the vector of moments calculated from the ICRISAT data and $g(\delta,\sigma)$ is the vector of simulated moments. For our estimation we use a diagonally weighted minimum distance procedure, corresponding to a weighting matrix W that has the variances of the moments on the diagonal and is zero everywhere else.[17] The variances are obtained by bootstrapping the data 1,000 times.

To minimize the criterion we conduct a grid search on a discrete grid of $\delta \in [0.1, 0.99]$ and $\sigma \in [0.1, 7]$. Because the criterion is not necessarily a smooth function of the preference parameters, we do not use gradient methods: in the model with coalitional deviations, small changes in preferences lead to discrete jumps in equilibrium group size and consequently the estimated moments.

We now present two sets of estimates. First, we estimate the preference parameters by targeting only the 'standard' moments related to the average degree of risk-sharing, namely the relative volatility of consumption and income growth and the average linear association of income and consumption growth. Second, we estimate the preference parameters by targeting all four moments, including those that capture the asymmetry, where we split the sample into households with income gains and households with income losses.

Table 7 presents the moments of interest when the preference parameters σ and δ are estimated to target $\frac{Var_{dc}}{Var_{dy}}$ and β_{dcdy}, the two moments that summarize the extent of insurance in the whole sample.

Allowing the preference parameters to vary, the two models' estimates depart from the common preferences of the previous section in opposite directions. In order to decrease the high degree of risk-sharing found with the common parameterisation, the standard version with individual deviations estimates markedly lower risk-aversion, on the lower bound of the parameter space and thus close to risk-neutrality in the case of Aurepalle.[18] By way of illustration, an estimated coefficient of risk-aversion of 0.3 implies an arguably implausibly low risk premium of 1.6% among Indian farmers.[19] This low risk-aversion in the standard model is counteracted by

[17]For a further description and application of this procedure see for example Blundell, Pistaferri and Preston (2008). We conduct robustness checks with different weighting matrices and different functional forms for the asymmetric moments. This makes no difference to the results.

[18]Estimates of the coefficient of relative risk aversion tend to be clustered around 1, though values as small as 0.3 and as large as 2.8 have been measured.

[19]The risk premium is defined here as the difference between the expected value of a bet and the

Table 7: Preferences estimated to target degree of risk-sharing - 2 moments

	Aurepalle			**Kanzara**			**Shirapur**		
	Data	**ID**	**CD**	**Data**	**ID**	**CD**	**Data**	**ID**	**CD**
n	NaN	34.00	4.00	NaN	37.00	4.00	NaN	31.00	4.00
δ	NaN	0.97	0.85	NaN	0.93	0.89	NaN	0.94	0.80
σ	NaN	0.10	4.00	NaN	0.50	5.00	NaN	0.30	7.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.23	0.26	0.56	0.20	0.26	0.33	0.18	0.25
β_{dcdy}	0.21	0.31	0.23	0.22	0.27	0.23	0.17	0.26	0.23
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy\leq 0}}$	0.75	34.95	1.02	0.64	38.50	0.99	0.60	37.02	1.04
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.24	0.51	0.01	-0.10	0.45	0.00	0.03	0.43	-0.01
Goodness of fit	NaN	4.48	0.77	NaN	4.64	2.85	NaN	6.12	2.10

The table presents the size of the insurance groups, the preference parameter σ and δ and the four moments of interest when σ and δ are chosen to minimise the sum of differences between the first two moments ($\frac{Var_{dc}}{Var_{dy}}$ and β_{dcdy}) predicted by the model and those observed in the data, weighted by the inverse variance of the data moments calculated using a bootstrap procedure. The estimated preference parameters are those that minimize this criterion on a grid of $\delta \in [0.1, 0.99]$ and $\sigma \in [0.1, 7]$.

Table 8: Preferences estimated to target degrees of risk-sharing and asymmetry

	Aurepalle			**Kanzara**			**Shirapur**		
	Data	**ID**	**CD**	**Data**	**ID**	**CD**	**Data**	**ID**	**CD**
n	NaN	34.00	4.00	NaN	37.00	4.00	NaN	31.00	4.00
δ	NaN	0.99	0.85	NaN	0.99	0.89	NaN	0.99	0.85
σ	NaN	5.00	4.00	NaN	5.00	5.00	NaN	3.00	4.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.00	0.26	0.56	0.00	0.26	0.33	0.00	0.26
β_{dcdy}	0.21	0.00	0.23	0.22	0.00	0.23	0.17	0.00	0.25
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy\leq 0}}$	0.75	3.66	1.02	0.64	2.86	0.99	0.60	4.35	0.96
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.24	0.00	0.01	-0.10	0.00	0.00	0.03	0.00	0.00
Goodness of fit	NaN	288.34	4.56	NaN	100.18	5.48	NaN	597.22	7.94

The table presents the size of the insurance groups, the preference parameter σ and δ and the four moments of interest when σ and δ are chosen to minimise the sum of differences between all four moments predicted by the model and those observed in the data, weighted by the inverse variance of the data moments calculated using a bootstrap procedure. The estimated preference parameters are those that minimise this criterion on a grid of $\delta \in [0.1, 0.99]$ and $\sigma \in [0.1, 7]$.

an increase in the discount factor.

In contrast, the alternative model with coalitional deviations estimates very high levels of risk aversion, counteracted partly by a lower discount factor. Notwithstanding the high estimates for σ (ranging up to 7), the implied risk premium (25%) is well in line with magnitudes for low and middle income countries measured in the literature.[20] This increased risk-aversion increases the maximum sustainable insurance group to 4 in all three villages, and increases the degree of risk-sharing to be close to that observed in the data (apart from the relative variance of consumption and income growth in Kanzara, which continues to be higher than that predicted by the model).[21]

Interestingly, the individual deviations model performs worse than the alternative, as it underpredicts the relative consumption volatility $\frac{Var_{dc}}{Var_{dy}}$ and overpredicts the linear association of consumption and income growth β_{dcdy} more strongly in all three villages. The reason for this is that, with average income fluctuations smaller in the whole village than in small insurance groups, a given β_{dcdy} delivers smaller consumption volatility in the standard model than the alternative with coalitional deviations. To reduce the deviation of relative consumption volatility $\frac{Var_{dc}}{Var_{dy}}$ from that observed in the data, the standard model thus overpredicts β_{dcdy} more strongly than the alternative.

Although the standard model performs somewhat worse in predicting the observed degree of insurance, both models arguably do a reasonable job once preferences are estimated. This implies that the inability of the standard model to generate small group sizes is not a serious handicap when one estimates 'standard' moments of risk-sharing. It does become a handicap, however, when one considers the asymmetry in the joint distribution of consumption and income: just as with common preferences, the asymmetry is extreme and of an order of magnitude higher than in the data. Interestingly, the alternative specification with coalitional deviations continues to predict a symmetric distribution despite the increase in the predicted size of the insurance groups. The alternative specification is thus closer to the data, although, as expected, it continues not to deliver the decrease in the slope of the conditional mean function and in the conditional

certainty equivalent, i.e. the amount of money and agent needs to receive for certain to be indifferent between taking the bet and the certain outcome. Here we calculate the risk premium for a 50-50 bet between \$1 and \$2 and – for comparability with the literature – divide the risk premium by the expected value of the bet.

[20]Based on author calculations, Barr and Packard (2000) estimates a risk premium of 13.45% using Holt-Laury risk experiments among a sample of wage-employed and self-employed workers in Chile, Henrich and McElreath (2000) estimate a risk premium of 22% among Chilean farmers using a similar method and Ayenew et al. (2014) estimate a risk premium of 18.27% using observed production decisions among Ethiopian farmers.

[21]Of course, one could argue that these estimates for risk aversion are somewhat extreme. As a robustness exercise, we repeat the simulated method of moments estimation after restricting the range of admissible coefficients of relative risk aversion in Section 8 and find very similar results.

variances as income growth moves from positive to negative implied by the point estimates of the data moments.

There are two reasons why the model with coalitional deviations can better account for the symmetry in the data: first, asymmetries are less pronounced within smaller insurance groups, as section 2.3 has shown. Second, as explained in section 5.1, we keep the number of village members in the alternative model approximately equal to that in the standard model by simulating multiple insurance groups. The resulting average village income and consumption, which we control for by regressing household incomes on time dummies, is less than perfectly correlated with average incomes in the insurance group. Thus, with multiple risk-sharing groups, conditioning on village variables, or equivalently time dummies, leaves a group-component in household-level variables that makes the observed data more symmetric. Since the treatment of the simulated data follows directly from the standard conditioning we apply to the empirical data, this differential effect does not imply, in our view, any inconsistency. As a robustness exercise, however, we repeat the simulated method of moments estimation on the unconditional moments in Section 8, where we find very similar results.

In sum, although the standard dynamic limited commitment model performs somewhat worse than the model with coalitional deviations in predicting the average degree of risk-sharing seen in the data, both arguably do reasonably well – if at the cost of estimating rather extreme levels of risk aversion. The model with coalitional deviations is far more successful, however, in matching additional features of the data, namely, the degree of risk-sharing for households with income increases relative to those with decreases. Of course, so far, we have estimated preferences by targeting only the moments related to the extent of insurance in the whole sample. It is therefore as yet unclear, whether the standard limited commitment model is able to match the average degree of risk-sharing in the data without generating unwanted asymmetries. To answer this question, we now add the moments calculated on the two sub-samples to the estimation criterion.

Table 8 presents the results. For the coalitional deviations model, neither the moments nor the estimated preferences are changed much compared to targeting only two moments. Again, the model estimates high degrees of risk-aversion and successfully matches the degree of risk-sharing, both on average and for those with income gains and losses.

The standard dynamic limited commitment, however, performs quite disastrously: including the additional moments in the estimation criterion leads to a dramatic deterioration in model fit and wild swings in the estimated preference parameters. This is the case because the individual deviations model cannot match the symmetry in the data by generating smaller group sizes. Therefore, it has to exploit the fact that the consumption and income distribution is symmetric for any group size only if insurance is perfect. As a consequence, the model now gets closer to matching the moments on the sub-samples of income winner and losers, but it does so at the cost of predicting full insurance on the whole sample (hence high levels of risk-aversion and patience), which is clearly at odds with the data. In other words, using the standard limited commitment model to explain the symmetry observed in the data, only reinforces its tendency to predict levels of insurance that are too high.

In sum, we have shown that the standard specification of the dynamic limited model is able to predict realistic degrees of insurance (with low risk-aversion in Table 7), or realistic degrees of asymmetry (with high risk-aversion in Table 8), but not at the same time. The model with coalitional deviations, in contrast, predicts reasonably accurately both the degree of risk-sharing and the relative symmetry in the model.

6 Comparing the models for a common group size

All the results so far point to the fact that the model with coalitional deviations performs better than the standard model because it makes different predictions about the size of the insurance group. In particular, while the coalitional deviations model restricts group size endogenously to levels that are consistent both with the average degree of risk-sharing and the observed symmetry in the consumption distribution of income winners and losers, the standard model predicts that group sizes are infinitely large.

We have shown that this prediction is inconsistent with the data even when one interprets 'infinitely large' as the village economy. While the village is the exogenous bound on group size most in the spirit of the 'infinitely' large group size prediction in the standard model, there is of course a large literature in development economics that has pointed out that risk-sharing is often confined to smaller groups (or rather that has established that perfect insurance takes place not at the village level but within smaller units, such as clans, castes, households, extended families, networks).[22]

We therefore now ask whether the ICRISAT data contain a plausible exogenous partition of the village that leads to a better fit between the standard model and the data and by implication, whether the ability to predict small group sizes is really the only advantage of the model with coalitional deviations. We show that for given group size and preferences, the coalitional model

[22]For example, Grimard (1997) studies risk-sharing among ethnic groups in Cote d'Ivoire; Morduch (1991) tests insurance within castes in the ICRISAT data; and Dercon and Krishnan (2000) find some evidence of full risk-sharing within nuclear households in Ethiopia. While risk-sharing may be incidental to the above groupings, much work has been undertaken to map relevant insurance networks by asking households to identify insurance partners they rely on in times of need ((Fafchamps and Lund 2003, De Weerdt 2004).

delivers (often) less insurance which we think is an additional attractive feature of the model over and above the endogenous group size restriction.

The ICRISAT data contains a couple of variables that allow one to partition the village into mutually exlusive subsets, namely the caste and the extended family a household belongs to. On average, the ICRISAT villages contain 10 members of the same caste and the average extended family size is 2.8 (with a minimum of 1 and a maximum of 5). Given that the dynamic limited commitment model performs best for small to medium sized groups, the extended family size partition would seem most appropriate for our purposes.[23]

To implement the estimation for a given group size, we set $n = 3$ in the standard model (i.e. average extended family size rounded to the next integer) and incorporate the same restriction in the model with coalitional deviations. Although this restricts the parameter space of the alternative to those combinations that endogenously generate $n = 3$, the fit of the two models is very similar (see Table 9). It thus seems that, even when we give the standard model an additional degree of freedom along the dimension that we have shown generates the asymmetry, it doesn't do better than the alternative.

The preference estimates still differ, however. This is, on the one hand, because of the restricted parameter space in the model with coalitional deviations. On the other hand, this arises because the alternative model bases its predictions on an outside option of a two-agent coalition, while the standard model relies on individual autarky. We would thus expect the participation constraint to restrict insurance more in the coalitional deviations model. And this could potentially improve its fit with the data relative to the standard model, which tends to predict too strong insurance.

To see the effect of the different outside options on the degree of risk-sharing, we now compare the degree of insurance in the standard model and the alternative at identical group sizes across a variety of parameter specifications on our estimation grid. As our measure of the relative degree of insurance we use the ratio of the β coefficients in the two models, that is the sensitivity of consumption to income in the model with coalitional deviations relative to the sensitivity of consumption to income in the standard model.

Since we are interested in cases that deliver realistic degrees of partial insurance, we also restrict attention to stable insurance groups where the degree of insurance lies between 0.1 and 0.35, which we deem to be appropriate given that the true degree of insurance lies around 0.2 in all three villages. In addition, we only consider cases where insurance within groups is not

[23]Information on the extended family network was obtained as follows: in 2001, the ICRISAT villages were revisited and researchers tried to identify all households that split off from the households contained in the original panel used here. We think of a set of households connected in this way as an extended family.

Figure 5: The degree of insurance is smaller in the coalitional deviations model even for given group size

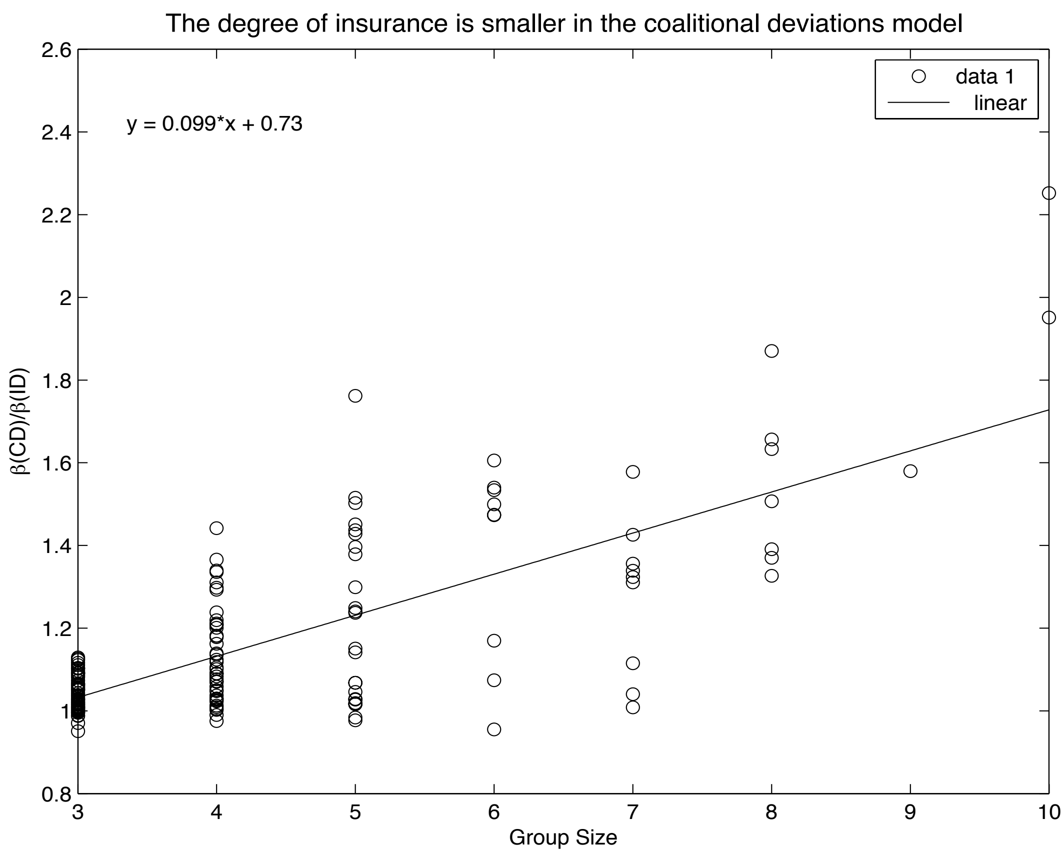

Table 9: Preferences and group size estimated to target all 4 moments

	Aurepalle			Kanzara			Shirapur		
	Data	**ID**	**CD**	**Data**	**ID**	**CD**	**Data**	**ID**	**CD**
n	NaN	3.00	3.00	NaN	3.00	3.00	NaN	3.00	3.00
δ	NaN	0.93	0.95	NaN	0.97	0.95	NaN	0.93	0.87
σ	NaN	5.00	1.50	NaN	2.50	4.00	NaN	7.00	4.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.35	0.35	0.56	0.34	0.34	0.33	0.34	0.34
β_{dcdy}	0.21	0.30	0.30	0.22	0.31	0.31	0.17	0.30	0.31
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy\leq 0}}$	0.75	0.97	0.97	0.64	1.00	0.99	0.60	0.98	0.96
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.24	-0.01	-0.01	-0.10	-0.01	-0.01	0.03	-0.00	-0.01
$\frac{Var_{d\epsilon}}{Var_{dc}}$	NaN	6.34	6.48	NaN	5.55	5.44	NaN	11.06	10.97

The table presents the size of the insurance groups, the preference parameter σ and δ and the four moments of interest when group size is restricted to equal 3 and σ and δ are chosen to minimise the sum of differences between all four moments predicted by the model and those observed in the data, weighted by the inverse variance of the data moments calculated using a bootstrap procedure. The estimated preference parameters are those that minimise this criterion on a grid of $\delta \in [0.1, 0.99]$ and $\sigma \in [0.1, 7]$.

perfect, i.e. where there is idiosyncratic variation in consumption even after controlling for group resources.

In Figure 5, we plot the ratio of β coefficients for each group size, where each point corresponds to a combination of the discount factor and the coefficient of relative risk aversion for which the insurance group is stable in the model with coalitional deviations. As one would expect, for small groups the predicted degree of risk-sharing is about the same in both models. As group size increases, the range of the ratio of β coefficients increases: at larger group sizes, as the difference in the outside option becomes more pertinent, the difference in the predicted degree of risk-sharing in the two models becomes larger on average (while the number of stable insurance groups decreases strongly). For example, for groups that comprise ten members, the consumption is on average twice as sensitive to income in the coalitional deviations model than in the standard model.

We can also relate the ratio of β coefficients to group size in a linear regression. We plot the regression line in Figure 5 and find that for every additional member of an insurance group, the ratio of the two coefficients increases by 10%.

As we argued above, a main advantage of the coalitional deviations model is that it reduces asymmetries by endogenously limiting the size of insurance groups. As 5 has shown, there are additional differences between the two models that tend to decrease the predicted degree of insurance in the coalitional model relative to the standard specification even at given group sizes. To the extent that the standard model tends to predict insurance that is too high, we think this may be an additional advantage of the alternative specification we have explored in this papers.

7 Conclusion

The dynamic limited commitment model has long been popular for the analysis of village risk-sharing because it can explain observed consumption risk-sharing better than simple alternatives such as autarky or full insurance. In this paper, we have, however, shown that it generates strongly counterfactual asymmetries in the consumption process at usual village sizes, because enforcement constraints in insurance contracts typically bind only in case of positive income shocks, when the outside option of leaving the village is attractive. We have then shown that by making the dynamic limited commitment renegotiation-proof, one can substantially improve the model's empirical performance. In particular, when households can renege on informal contracts together with other villagers, the size of the insurance group becomes endogenous and is usually much smaller than typical villages. Moreover, even for a given group size, the model with coalitional deviations predicts less insurance than the standard model because its outside option of joining a new risk-sharing coalition is more attractive. This is potentially another advantage of allowing for coalitional deviations in the dynamic limited commitment model, because the standard version tends to predict too much insurance. We also note that despite our prediction that risk-sharing takes place within small groups within the village, our approach does not require knowledge of the membership of these groups or even their size. We think that this is an advantage of our approach as it limits data requirements. In sum, we show that the limited commitment model's strongly counterfactual implication in large economies are no reason to reject it for the modeling of village (or any other small-group) risk-sharing. Rather, we show how a plausible restriction of the set of equilibria one considers, namely those that are coalition-proof, can successfully reconcile model and data.

8 Online Appendix

In the online appendix, we discuss the results and their robustness along two dimensions. First, we will estimate both models on the basis of unconditional moments, rather than the residuals from a regression on time dummies. Second, we will explore how the results change when estimating the models on a restricted grid for risk-aversion, since both the standard model estimates rather small levels of risk-aversion, while the coalitional deviations model's estimates appear rather large.

Table 10: Preferences estimated to target all 4 moments (unconditional)

	Aurepalle			**Kanzara**			**Shirapur**		
	Data	**ID**	**CD**	**Data**	**ID**	**CD**	**Data**	**ID**	**CD**
n	NaN	34.00	4.00	NaN	37.00	5.00	NaN	31.00	4.00
δ	NaN	0.94	0.85	NaN	0.93	0.89	NaN	0.94	0.90
σ	NaN	4.00	4.00	NaN	1.50	7.00	NaN	1.50	3.00
$\frac{Var_{dc}}{Var_{dy}}$	0.33	0.03	0.29	0.61	0.03	0.22	0.37	0.03	0.27
β_{dcdy}	0.22	0.03	0.25	0.18	0.03	0.20	0.15	0.03	0.25
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy\leq 0}}$	0.96	0.98	1.01	0.74	1.04	0.98	0.68	0.92	0.93
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.22	-0.00	-0.03	-0.03	0.01	-0.00	0.09	-0.01	-0.00
Goodness of fit	NaN	47.05	2.64	NaN	17.47	5.34	NaN	23.85	5.66

The table presents the size of the insurance groups, the preference parameter σ and δ and the four moments of interest when σ and δ are chosen to minimise the sum of differences between all four moments predicted by the model and those observed in the data, weighted by the inverse variance of the data moments calculated using a bootstrap procedure. The estimated preference parameters are those that minimise this criterion on a grid of $\delta \in [0.1, 0.99]$ and $\sigma \in [0.1, 7]$.

8.1 Estimating the model on the unconditional distribution of consumption and income

It is standard practice to evaluate the performance of risk-sharing models by conditioning both the data and model simulations on movements in aggregate resources, using residuals from a regression on time dummies. As we discussed in section 5.5, this procedure has somewhat different effects, however, in the two models we analyse. In the standard model, with individual deviations, the only risk-sharing group coincides with the village. Conditioning on village-level aggregate income (equal to aggregate consumption) thus isolates the idiosyncratic movements in income and consumption. The alternative model, with coalitional deviations, however, predicted a village to consist of several insurance groups. A regression on time dummies, therefore, does not eliminate fluctuations in group-level incomes, but only in village-level incomes. Since the remaining fluctuations in group-level income are symmetric and translate to individual consumption fluctuations, this may increase the symmetry in the alternative model.

In Table 10 we therefore repeat our benchmark estimation based on 'raw' data, rather than residuals from a regression on time dummies. This makes very little change to the degree of insurance measured in the three villages and it increases the observed symmetry, apart from the difference of the regression coefficients β_{dcdy} for Shirapur. Also, as expected from the previous discussion, considering unconditional moments leaves the predicted group sizes in the alternative model unchanged, and does not qualitatively affect its estimated preference parameters or predicted moments.

Table 11: Preferences estimated on restricted parameter space to target all 4 moments

	Aurepalle			**Kanzara**			**Shirapur**		
	Data	**ID**	**CD**	**Data**	**ID**	**CD**	**Data**	**ID**	**CD**
n	NaN	34.00	4.00	NaN	37.00	4.00	NaN	31.00	5.00
δ	NaN	0.99	0.87	NaN	0.99	0.92	NaN	0.99	0.92
σ	NaN	3.00	3.00	NaN	3.00	3.00	NaN	3.00	3.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.00	0.26	0.56	0.00	0.26	0.33	0.00	0.20
β_{dcdy}	0.21	0.00	0.24	0.22	0.00	0.25	0.17	0.00	0.19
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy\leq 0}}$	0.75	4.51	1.03	0.64	3.41	1.01	0.60	4.35	0.98
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.24	0.00	0.01	-0.10	0.00	0.01	0.03	0.00	0.01
Goodness of fit	NaN	450.39	4.95	NaN	145.14	6.02	NaN	597.22	8.91

The table presents the size of the insurance groups, the preference parameter σ and δ and the four moments of interest when σ and δ are chosen to minimise the sum of differences between all four moments predicted by the model and those observed in the data, weighted by the inverse variance of the data moments calculated using a bootstrap procedure. The estimated preference parameters are those that minimise this criterion on a grid of $\delta \in [0.1, 0.99]$ and $\sigma \in [0.7, 3]$.

More interestingly, the performance of the standard model is equally changed very little: while both the discount factor and the coefficient of relative risk aversion are estimated to be somewhat lower than in the benchmark estimation, risk-sharing is still predicted to be, essentially, perfect. Without prior conditioning on village-level variables, however, the very small remaining idiosyncrative movements in individual consumption and income are now dominated by those in village income and consumption. Since the latter are, by construction, symmetric, this combination of using raw data and (almost) perfect insurance against idiosyncratic income shocks yields a symmetric joint distribution of consumption and income changes, similar to that in the alternative model with coalitional deviations. The standard model nevertheless performs overall worse than the alternative because the relatively small idiosyncratic component of consumption movements – a necessary condition for the predicted symmetry – is achieved through counterfactually high insurance. In other words, even in the raw data, idiosyncratic movements in individual consumption and income are sufficiently important that the standard model can only generate symmetry by predicting essentially perfect insurance.

8.2 Estimating the model on a restricted parameter space

Since one might argue that the superior performance of the model with coalitional deviations is driven by estimates of the coefficient of relative risk aversion that are implausibly large, we now compare the two models after restricting the admissible parameter range for risk-aversion to lie between 0.7 and 3.[24]

The results are presented in Table 11. Restricting the range of values the risk aversion parameter can take barely affects the performance of the coalitional deviations model. The estimates for consumption volatility and its sensitivity to income growth on the whole sample are virtually unchanged, because the model can adjust for the slightly lower level of risk aversion by increasing its estimate of the discount factor. This does not affect its estimates of the asymmetry, because these are pinned down by the equilibrium group size predictions – which are not overly sensitive to the preference parameters in this range.

While not the focus here, it is worth noting that the standard model suffers a lot more from the restricted parameter space. This is the case because it did not only estimate very high risk aversion, but also very high levels of patience, so it cannot trade off along these dimensions. In fact, the results show that even small numerical departures from full insurance, which occur for the slightly lower levels of risk aversion, lead to large increases in the asymmetry when the risk-sharing group is assumed to be the village.

In sum, the results in Table 11 give us some confidence that the model with coalitional deviations provides a better fit to the data 'globally' and not just for a rather extreme combination of preference parameters.

References

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

Árpád Ábrahám and Eva Cárceles-Poveda, "Competitive Equilibria with Production and Limited Commitment," *mimeo*, 2009.

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

Bernheim, B. D., B. Peleg, and M. Whinston, "Coalition Proof Nash Equilibria. I: Concepts," *Journal of Economic Theory*, 1987, *42* (1), 1–12.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption inequality and partial insurance," *The American Economic Review*, 2008, pp. 1887–1921.

[24]This restriction of the grid is based on a literature search looking for the most common vales of risk-aversion estimated empirically.

Bold, Tessa, "Implications of endogenous group formation for efficient risk-sharing," *Economic Journal*, 2009, *119*, 562–591.

Broer, Tobias, "The Wrong Shape of Insurance? What Cross-Sectional Distributions Tell Us about Models of Consumption Smoothing," *American Economic Journal: Macroeconomics*, 2013, *5* (4), 107–40.

De Weerdt, J., *Insurance Against Poverty*, Oxford University Press, 2004.

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

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

Dubois, Pierre, Bruno Jullien, and Thierry Magnac, "Formal and Informal Risk Sharing in LDCs: Theory and Empirical Evidence," *Econometrica*, 07 2008, *76* (4), 679–725.

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

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

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

Karaivanov, Alexander and Robert M Townsend, "Dynamic financial constraints: Distinguishing mechanism design from exogenously incomplete regimes," *Econometrica*, 2014, *82* (3), 887–959.

Kinnan, Cynthia, "Distinguishing barriers to insurance in Thai villages," 2010.

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

Krueger, Dirk and Fabrizio Perri, "Understanding Consumption Smoothing: Evidence from the U.S. Consumer Expenditure Data," *Journal of the European Economic Association*, 2005, *3* (2-3), 340–349.

____ **and** ____ , "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *Review of Economic Studies*, 2006, *73* (1), 163–193.

—— **and** —— , "Public versus private risk sharing," *Journal of Economic Theory*, 2011, *146* (3), 920–956.

Laczo, Sarolta, "Risk Sharing With Limited Commitment and Preference Heterogeneity: Structural Estimation and Testing," *Journal of the European Economic Association*, 2014.

Ligon, E., J.P. Thomas, and T. Worrall, "Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies," *Review of Economic Studies*, 2002, *69*, 209–244.

Mazzocco, Maurizio and Shiv Saini, "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," *American Economic Review*, 2012, *102* (1).

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

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

Ogaki, Masao and Qiang Zhang, "Decreasing Relative Risk Aversion and Tests of Risk Sharing," *Econometrica*, 2001, *69* (2), 515–526.

Ravallion, M. and S. Chaudhuri, "Risk and Insurance in Village India: Comment," *Econometrica*, 1997, *65*, 171–184.

Rouwenhorst, K. Geert, "Asset pricing implications of equilibrium business cycle models," in Thomas F. Cooley, ed., *Frontiers of business cycle research*, Princeton University Press, 1995, p. 294330.

Townsend, R. M., "Risk and Insurance in Village India," *Econometrica*, 1994, *62*, 539–591.