Risk Sharing in Village Economies Revisited

Tessa Bold, IIES, Stockholm University and CEPR

Tobias Broer, IIES, Stockholm University and CEPR*

October 2016

Abstract

We propose to replace the common notion of 'village' risk sharing by flexible insurance in endogenous risk sharing groups. We model risk-sharing in a quantitative environment with limited commitment where the requirement that contracts be 'renegotiation-proof', or immune to deviations by subcoalitions, makes group size endogenous. Apart from predicting a realistic degree of insurance, the model captures the evidence along two new dimensions: endogenously small insurance groups and symmetric consumption responses to income rises and falls. We argue that this is important, as policy interventions are likely to have different effects in such an environment.

JEL Classification: D11, D12, D52

Keywords: Risk Sharing, Endogenous Groups, Village Economies, Informal Insurance, Dynamic Limited Commitment, Renegotiation-Proofness

*We thank Árpád Ábrahám, Paul Klein, Per Krusell, Nicola Pavoni, Pascual Restrepo, Swapnil Singh, Christian Stoltenberg, Kjetil Storesletten and Nicholas Barton, and seminar participants at Bocconi University, IIES Stockholm University, Normac, the SED annual meeting 2015, and the University of Oslo for comments. Contact: tessa.bold@iies.su.se, tobias.broer@iies.su.se.

1 Introduction

We propose to replace the 'village' in the study of consumption risk sharing in poor agricultural communities by a number of small, endogenous insurance groups. This not only increases the empirical content of the analysis, as group size becomes a testable prediction, but also captures the evidence along two, previously neglected dimensions. Specifically, apart from the observed degree of insurance, the environment we propose predicts, first, equilibrium groups that are substantially smaller than typical villages; and second, the symmetric reaction of consumption to positive and negative income shocks. We think that these results argue in favour of a model with endogenous insurance groups in the context of rural India that we focus on. More generally, they allow researchers to discriminate between models of endogenous vs exogenous group formation that are likely to imply substantially different effects of policy interventions.

The key friction that enables us to study endogenous groups is the absence of commitment to co-insurance. This is often seen as a particularly plausible reason for limited risk sharing in poor villages, where contract enforcement is difficult but other impediments to insurance, such as lack of information on households' productive possibilities and effort, are presumably less pronounced. To study endogenous group formation in a fully dynamic and quantitative model of risk sharing with limited commitment to contracts we assume that households can renege on village insurance not just alone but in subgroups, as in Genicot and Ray (2003). We think this requirement of 'coalition-proofness' is particularly appealing in the context of village economies, where it seems difficult to prevent agents who renege on insurance arrangements to insure each other again in the future.

Our first contribution is to draw attention to group size as an important endogenous determinant of risk sharing, and to propose a tractable way to make insurance groups endogenous outcomes of a dynamic limited commitment risk sharing mechanism. To compute the risk sharing equilibrium quantitatively, we combine the common approximative solution of the standard model, originally proposed by Ligon, Thomas and Worrall (2002) and used, for example, in Laczo (2014) and Dubois, Jullien and Magnac (2008), with the recursive procedure for finding stable group sizes when coalitions can deviate together proposed by Genicot and Ray (2003). We use this to estimate the model for the well-known ICRISAT dataset on agricultural villages in India.

Our second contribution is to show how the model with endogenous groups replicates the degree of risk sharing in those villages well and captures the empirical evidence along two new dimensions: first, it predicts insurance groups of realistic size.[1] And second, it captures the symmetry of empirical consumption-income comovements. This is important because the limited commitment constraint per se is more likely to bind for villagers with high income realisations and therefore attractive outside options. In large insurance groups, such as countries, this feature is known to imply a much stronger response of consumption to positive than to negative income shocks (as the former make the outside option more attractive and thus tighten participation constraints, while the latter do not) that is not seen in the data (Broer, 2013). Beyond pairs, where consumption shares trivially move in symmetry (Kocherlakota, 1996), however, the strength of this asymmetry both in theory and data has so far been unknown for small communities.

We show here that in the absence of coalitional deviations, the asymmetry implied by the limited commitment constraint increases quickly with the exogenous group size, and is counterfactually strong not only at usual village sizes, but also for groups equivalent to typical extended families within a village, for example. We also show how this finding is unaffected, if not reinforced, by including a stylised form of heterogeneity in preferences. In contrast, the model with coalitional deviations predicts negligible asymmetry, in line with the data from the ICRISAT villages. We show how two features of the model drive its ability to capture the evidence: endogenously small groups of single digit size; and an outside option that depends on the

[1]The standard punishment assumption of eternal individual autarky implies that there is no limit to the size of insurance groups, as a larger group size increases the benefits of co-insurance but does not affect the outside option. This partly motivates the common focus on 'village-level' risk sharing. Transfers are typically made, however, in groups that are much smaller than the village, giving rise to a small recent literature that focuses on limited commitment to bilateral insurance relationships. The focus on the resulting network structure, however, makes the analysis of truly dynamic constrained risk sharing infeasible, requiring instead a static and, usually, exogenous risk sharing rule (Bloch, Genicot and Ray, 2008; Ambrus, Mobius and Szeidl, 2014).

whole village income distribution and is thus less volatile and, due to the prospect of continued co-insurance in smaller groups, more attractive than without coalitional deviations.[2]

We think that these results are important not only for our understanding of risk sharing in rural India, but also for policymakers in developing countries more generally. This is, first, because the desirability and effectivness of policies, such as income support for the poor, are likely to differ between small groups where most risks are shared and the standard larger groups with more moderate risk sharing. More importantly perhaps, the endogeneity of insurance groups provides an additional and so far neglected channel through which policy interventions affect private risk sharing. This adds a new dimension to the trade-off for policy interventions that arises in the standard limited commitment model, where social safety nets and formal insurance make incomes less volatile and thus the outside option to village insurance more attractive, reducing the need for risk sharing at the same time as making it less powerful by encouraging defections (Attanasio and Rios-Rull, 2000).

We focus on the limited commitment friction for several reasons. First, formal institutions for contract enforcement are typically absent in poor agricultural villages. So lack of commitment captures the reality we are interested in a priori. Second, several previous studies have shown that limited commitment to co-insurance can explain the partial character of risk sharing observed in many agricultural villages (Townsend, 1994; Ligon, Thomas and Worrall, 2002; Laczo, 2014). Finally, and importantly for this paper, when studying the size of insurance groups, it is typically impossible to explain insurance groups of realistic size in a frictionless setting with full commitment. With limited commitment and coalitional deviations alone, groups are, as we show, endogenously small, in fact smaller than most exogenous groups including the village, castes within the village, etc. This does not necessarily contradict the finding of insurance within

[2] Using a simplified, stationary version of the model, Dubois (2006) and Fitzsimons, Malde and Vera-Hernandez (2015) test empirically whether risk sharing at the village level is constrained by coalitional deviations. Both papers find evidence for coalitional deviations, but neither estimates the conditional consumption distribution that would arise if risk sharing was restricted by them. Bold (2009) derives a formal test of the presence of coalitional deviations in a dynamic setting that relies on the finding that groups that are constrained endogenously do not exhibit the amnesia typical of the standard model (Kocherlakota, 1996), but requires an exact identification of constrained households.

kinship groups, castes etc. (Angelucci, de Giorgi and Rasul, 2015; Fitzsimons, Malde and Vera-Hernandez, 2015; Mazzocco and Saini, 2012; Mobarak and Rosenzweig, 2012), to the extent that these may contain several endogenous groups featuring strong risk-sharing. Importantly, small group sizes are not an intrinsic feature of the new model we propose per se. It is therefore possible that our model of risk sharing in endogenous groups predicts larger group sizes in contexts with different income risk, or when including other frictions. For example, information about income and effort of other households might be better in small, closely integrated communities than in larger societies but may not be perfect. While previous studies have looked at limited information and limited commitment together (Broer, Kapička and Klein, 2015), our focus on endogenous group formation makes the inclusion of any additional friction difficult.

2 The standard limited commitment model

The standard limited commitment model restricts risk sharing through participation constraints that require households to always prefer co-insurance to the outside option of individual autarky. Since this outside option becomes more attractive as household income rises, participation constraints are more likely to bind in periods of high income. This makes consumption respond more strongly to positive income shocks. The implied asymmetry is mitigated, however, by the group budget constraint, which requires unconstrained households to share efficiently the fall in resources available to them after their constrained peers are satisfied. Previously, the resulting net asymmetry in the consumption response to income shocks was only known for two polar cases: the two-agent economy (Kocherlakota, 1996), where the consumption share of the constrained household rises as much as that of the unconstrained falls, is trivially symmetric; a version of the model with a continuum of (infinitely many) agents (Broer, 2013), in contrast, has strong asymmetries that are significantly larger than those observed in US data. This section shows, for the first time, how the asymmetry is in fact very strong at usual village sizes and arises already in small insurance groups 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 an endowment of the only consumption good $y^i(s_t)$, where $s_t \in \{1, ..., S\}$ is the state of nature in period t. 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 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 realisations 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 realisation, helping out those with a low income realisation 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 $(\tau^i(s_t, h_t))_{i=1}^n$ 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 summarise the history of the game up to the current period (Abreu, Pearce and Stacchetti, 1990; Ligon, Thomas and Worrall, 2002).

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

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

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

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

and an aggregate resource constraint

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

where γ^i, ϕ^i, and ω are the Lagrange multipliers associated with the promise-keeping, enforcement, and resource constraints respectively.

2.2 Asymmetries in the consumption-income distribution

Equation (4) shows 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. Algebraically, this is easily seen with log preferences, where the first order condition for consumption implies that the relative consumption by any two households equals their relative, 'updated' Lagrange multipliers $\frac{\gamma^i(1+\phi_r^i)}{\gamma^j(1+\phi_r^j)}$. Summing across all households i in period t and using the resource constraint $Y_t = \sum_{i=1}^n c_t^i$ to express household j's consumption as a fraction of village income, this implies, after taking first log-differences

$$d\log(c_t^j) \quad = \quad d\log(Y_t) + \log(1+\phi_t^j) - \log\left(1 + \frac{\sum_{i=1}^n \gamma^i \phi_t^i}{\sum_{i=1}^n \gamma^i}\right). \tag{6}$$

where $d\log$ denotes the log difference, we suppress the dependence on state r in period t, and Y_t is village income. 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 the multiplier ϕ_t^j is positive, but zero otherwise; and finally, there is a 'drift-term' $-\log(1 + \frac{\sum_{i=1}^n \gamma^i \phi_t^i}{\sum_{i=1}^n \gamma^i}) \leq 0$ that is common for all group members and strictly negative whenever at least one participation constraint is binding in the village.

Equation (6) immediately yields Lemma 1:[3]

Lemma 1 *For the subsample of unconstrained households, the cross-sectional variance of consumption growth, and thus the covariance of consumption and income growth, are equal to zero.*

Unconstrained agents thus share identical consumption growth, which is smaller than that of any constrained agent in magnitude and independent of their current individual incomes. Constrained agents' consumption growth, in contrast, is an increasing function of their current income (since $\frac{\partial \phi_t^j}{\partial y_t^j} > 0$). Equally, across unconstrained periods, consumption growth is independent of individual income growth, after conditioning on aggregate income growth. Across periods with binding constraints, in contrast, the correlation between consumption and income growth is positive (as $\phi_t^j > 0$ is increasing in income growth).

[3]It is easy to see that Lemma 1 also holds in the more general case ofconstant relative risk aversion.

There are two reasons why it is difficult to test the prediction embodied in equation (6) on observations of consumption and income in actual village economies. First, apart from log-preferences and i.i.d. transitions, the analytical result assumes that we can distinguish constrained from unconstrained individuals, which is typically impossible in the data. And second, the strength of the asymmetry depends, in an obvious way, on the degree of insurance, and thus on the preference parameters used. Specifically, whenever patience and/or risk aversion are such that the allocation is either autarky or perfect risk sharing, the model produces symmetry. Any test thus has to compare jointly the degree of insurance and symmetry in the model and the data. This feature will turn out to be important for the interpretation of the estimation results in Section 5, where preference parameters are chosen to maximise model fit.

In the absence of information about binding constraints, we exploit the fact that the right-hand side of the participation constraint (4) is increasing in the individual's endowment y_t^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 therefore study a simplified version of the general equilibrium environment presented in Ligon, Thomas and Worrall (2002).[4,5] Figure 1 shows a scatter plot of the resulting joint distribution of consumption and income growth, after controlling for aggregate income movements.[6] Moving from left to right (low to high income growth), consumption growth

[4]In a previous version of this paper, we performed the same analysis for a partial equilibrium version where a village lender has access to outside funds at a constant interest rate, as in Kinnan (2014) or Karaivanov and Townsend (2014). Apart from a different interest rate and income process, this version of the model implies the same comovement of consumption and income as the general equlibrium continuum economy analysed in Broer (2013), and thus has even stronger asymmetry.

[5]For this illustrative exercise we choose preference parameters that are representative of previous estimates, with log-preferences, and an annual discount factor equal to 0.9, close to the mean of values considered in Laczo (2014), and within the range of values estimated by Ligon, Thomas and Worrall (2002). Later, we will choose preference parameters to match moments of village data explicitly. Relative to Ligon, Thomas and Worrall (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, her estimates for Aurepalle, one of the villages in the ICRISAT dataset that we look at in the following sections. We set the number of households equal to that in Aurepalle (34), and solve the model with the standard algorithm presented in Ligon, Thomas and Worrall (2002). As in Laczo (2014), we approximate the income process of the individual (rest of the village) as an AR(1) process, discretised using Rouwenhorst's (1995) method with 6 (5) support points.

[6]We control for aggregate income movements for consistency with later sections and the literature. As Section A.4 in the Online Appendix shows, the asymmetry in the raw simulated data is equally very

Figure 1: Consumption and income growth in general equilibrium

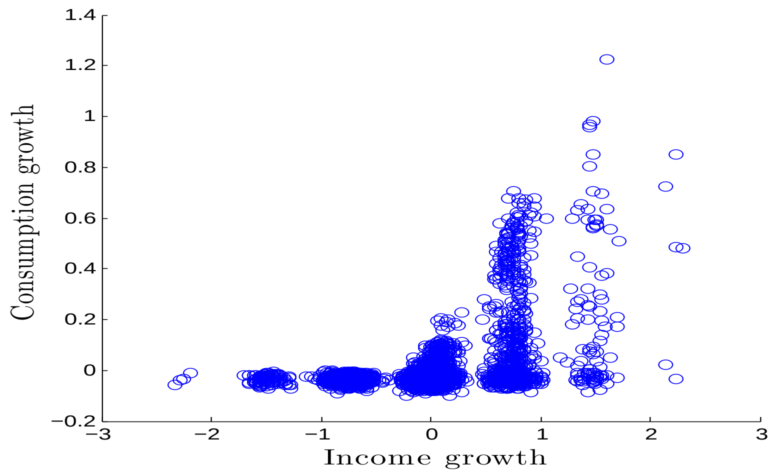

Notes: The figure shows a scatter plot of consumption and income growth from a simulation of the standard general equilibrium version of the limited commitment economy. The figure plots residuals from a regression of both log-income and log-consumption on time dummies, that controls for movements in aggregate resources.

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. In other words, there is a kink in the conditional mean and variance functions at zero income growth.

Note that we focus on the joint distribution of consumption and income growth. Ligon, Thomas and Worrall (2002) discuss how the general equilibrium version of the limited commitment model fits the joint distributions of both growth rates and levels. Focussing 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 Figure 1: 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.

---

strong.

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

n	34
δ	0.90
σ	1.00
$\frac{Var_{dc}}{Var_{dy}}$	0.04
β_{dcdy}	0.08
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy\leq 0}}$	55.75
$\frac{\beta dcdy>0}{\beta dcdy\leq 0}$	17.78

Notes: Based on simulated data from the standard limited commitment model, the table shows the regression coefficient of consumption growth on income growth β_{dcdy}; the relative variance of consumption and income growth $\frac{Var_{dc}}{Var_{dy}}$; the variance of consumption of households that experience positive income growth divided by that of those that experience non-positive income growth $\frac{Var_{dc|dy>0}}{Var_{dc|dy\leq 0}}$; and the ratio of regression coefficients 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 table uses residuals from a regression on time dummies to control for movements in aggregate village income.

Table 1 presents four key moments that summarise both the degree of insurance, and the asymmetry of the joint distribution in Figure 1. To capture the degree of insurance it shows the 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 the conditional variance and mean of consumption growth as a function of income growth, respectively. Specifically, because Figure 1 suggests non-linearities in the conditional means and variances around zero income growth, we first look at the variance of consumption of households that experience positive income growth divided by that of those that do not, $Var_{dc|dy>0}/Var_{dc|dy\leq 0}$. Second, to capture non-linearities in the conditional mean, we look at the ratio of the regression coefficient of consumption growth on income growth for households with rising and non-rising income, $\beta_{dcdy|dy>0}$ and $\beta_{dcdy|dy\leq 0}$.[7]

Table 1 shows how the standard limited commitment model predicts strong, but not perfect risk sharing, as summarised by a variance of consumption growth that is only 4 percent that of

[7]Note that for the calculation of these moments, we group 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.

income growth on average, and a regression coefficient of log consumption changes on log income changes of 8 percent. Importantly, the asymmetry is very strong: movements in consumption are fifty times more volatile in periods of rising income than when incomes fall or are constant; and the ratio of the regression coefficients for rising and non-rising incomes exceeds 15.

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.

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 (6), the growth in consumption is a weighted difference of the individual Lagrange multiplier ϕ_t^j and the "average" multiplier $\frac{\sum_{i=1}^n \gamma^i \phi_t^i}{\sum_{i=1}^n \gamma^i}$. 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.[8]

While the magnitude of the asymmetry is thus known for pairs and for infinitely large groups, or equivalently for a village that can borrow at an exogenous interest rate, we simply do not know 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 number of households in the sample

[8]When income realisations are independent across households, 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 third term in (6) converges to a constant that can be expressed as a function of $R^\star$.

Figure 2: Key moments in village economies of different size

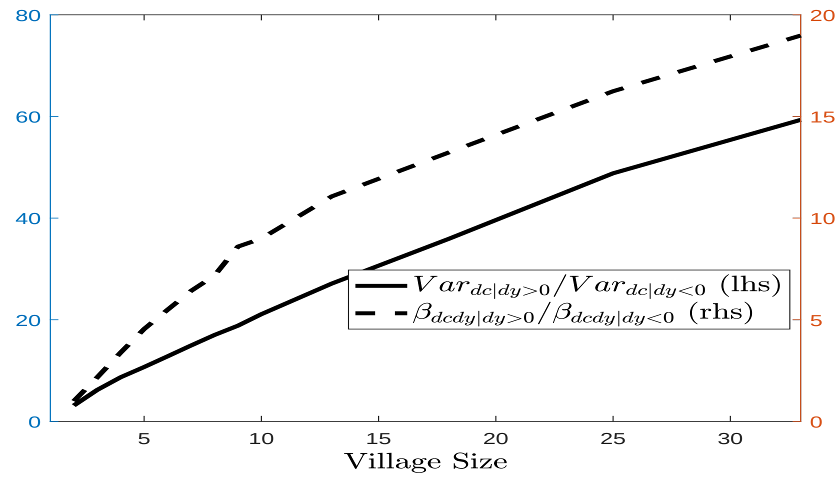

Notes: The figure plots $\frac{Var_{dc|dy>0}}{Var_{dc|dy\leq 0}}$ (the relative variance of consumption growth of households that experience positive income growth divided by that of those that do not, as a solid line) and $\frac{\beta_{dcdy|dy>0}}{\beta_{dcdy|dy\leq 0}}$ (the ratio of regression coefficients of consumption growth on income growth for households with rising and non-rising income, as a dashed line) from simulations of the model with increasing village size (reported along the bottom axis). The figure uses the full set of aggregate incomes in the rest of the village (rather than the approximation of the aggregate income process as in section 2.2, explaining small differences in the moments for $n = 34$ relative to table 1) and is based on residuals from a regression on time dummies to control for movements in aggregate village income.

for Aurepalle.[9]

The asymmetry increases strongly and quickly in the size of the insurance group. In Figure 2, we plot the two key ratios describing the asymmetry in Table 1 as a function of village size. 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 positive 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, single-digit group sizes, and increase dramatically thereafter.

[9]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.

3 The Data

Risk sharing has previously been found to be strong in rural village economies. Less is known about the symmetry of the comovement between consumption and income, however. 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 asymmetries in income and consumption comovement are small and insignificant for the most part (and of the wrong sign in the few instances where they are) in contrast to the strong asymmetry predicted by the standard limited commitment model in Section 2.2.[10] Note that the ICRISAT data do not include information on transfers between households, which would be needed to identify risk sharing units. This is why our approach to model evaluation is based on key moments of consumption risk sharing that can be compared to data without specific knowledge of the potentially complex transfer structures underlying them.

The data come from the village level studies conducted by the International Crop Research Institute for the Semi-Arid Tropics (ICRISAT) in India from 1975-1984. In particular, we focus on three rural and agricultural villages surveyed, Aurepalle in Andra Pradesh state, and Kanzara and Shirapur (both in Maharashtra State). In each village, detailed expenditure and income data were collected for 40 randomly sampled households on an annual basis.

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 intoxicants, such as paan, alcohol and tobacco.[11] The income aggregate contains net income from farming and livestock, labour and transfers from outside the village. 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

[10]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 model with homogenous preferences (Ligon, Thomas and Worrall, 2002) and with heterogeneous risk preferences (Laczo, 2014).

[11]We thank Sarolta Laczo for making her version of the data available to us.

Table 2: Conditional variance of consumption

	Aurepalle	Kanzara	Shirapur
	(1)	(2)	(3)
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.56	0.33
	(0.052)	(.179)	(.077)
$\frac{Var_{dc\|dy>0} - Var_{dc\|dy<=0}}{Var_{dy}}$	-0.08	-0.25	-0.16
	(0.063)	(.205)	(.078)
Obs.	170	185	155
No. of households	34	37	31

Notes: The table shows the variance of consumption growth divided by that of income growth and the difference in the variance of consumption growth for those experiencing income growth and those experiencing income losses, scaled by the variance of income growth. Both measures are conditional on changes in aggregate resources. Standard errors in parentheses are clustered at the household level.

other authors, we restrict our analysis to the years 1976-1981 and construct a fully balanced panel.[12]

The ICRISAT villages are poor with the average dweller living well below the $1 dollar a day poverty line (Table 13 in Section A.7 of the Online Appendix). On average, daily nondurable consumption per adult equivalent is 0.83, 1.10 and 1.18 in 1975 rupees, which is equivalent to 0.48, 0.63 and 0.68 in 2016 US dollars respectively. Income is about twice as high, and the difference between income and consumption might be accounted for by durables consumption, investment in livestock and housing, but also measurement error.

Although villagers are poor on average, there is strong evidence of consumption smoothing. In Table 2, 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, consumption smoothing is strong, though far from perfect with the variance of consumption relative to income ranging from 0.3 in Aurepalle to 0.56 in Kanzara.

In Table 3, 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%

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

Table 3: 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	.441	.384	.176
	(.080)***	(.137)***	(.092)*
Δ ln of aeq. income > 0	-.413	-.141	-.053
	(.137)***	(.174)	(.095)
Obs.	170	185	155
No. of households	34	37	31

Notes: Panel A shows the results from a regression of consumption growth on income growth in Aurepalle, Kanzara and Shirapur and includes a full set of time fixed effects in each village. Panel B estimates the coefficient separately for those with positive and negative income growth by including a dummy for households with rising income, and its interaction with income growth. Both consumption and income are demeaned period-by-period before the regression. Standard errors in parentheses are clustered at the household level.

change in income leads to roughly a 0.2% change in consumption.[13] We take this effect, which is fairly uniform and significant across all three villages, and the smoothness of consumption growth relative to income growth as strong evidence for substantial consumption risk sharing.[14] Nevertheless, we note that the degree of smoothing is smaller than that predicted by the standard limited commitment model with standard preferences in the previous section.

In Figure 3 we plot the joint distribution of the residual consumption and income growth (after a regression on time dummies). In contrast to the simulated model solutions in Figure 1, the data appear symmetric: neither the variance of consumption nor the response of consumption to income look dramatically different as households move from negative to positive income growth.

In the second row of Table 2 and panel B of Table 3 we test the impression of symmetry more formally. As can be seen from Table 2, consumption growth of households that experience positive income growth is actually less volatile than that of those who experience negative or

[13]For comparability, we follow Laczo (2014) in the implementation of the risk sharing test and selection of consumption and income aggregates. Of course the estimates of the degree of risk sharing in Table 3 need to be interpreted with caution because of concerns about both measurement error and potential endogeneity of the income aggregate used on the right-hand side, in particular labor and transfer income. Despite the latter, we focus on the full income aggregate for two reasons: (i) as is well known from Townsend (1994), consumption is relatively well insured with respect to variation in crop income. Focussing mainly on this income source would therefore make it difficult to distinguish the limited commitment model (in either form) from full insurance. (ii) As noted in Ravallion and Chaudhuri (1997), there is a concern that changes in crop income and consumption vary systematically, biasing the coefficients in a regression of consumption on this income source. The authors suggest instead to use the full income aggregate (like we do) and instrument using all non-crop income sources, such as labor, trade and livestock income. In general, one should note that our estimates of the effect of income on consumption changes are somewhat larger than Townsend's (1994), which range from 0.08 to 0.14 depending on village and specification and of the same order of magnitude as Ravallion and Chaudhuri (1997) whose estimates vary between 0.11 and 0.34. Importantly, our conclusion that insurance is high but imperfect and that the degree of insurance is not significantly lower for households with income growth is robust to excluding labor and transfer income from the income aggregate.

[14]In the absence of information on actual transfers, these two stylised facts - consumption growth only weakly associated with income growth, and significantly less volatile - could, in principle, also be achieved by self-insurance through household saving and borrowing. As many previous studies, we interpret the evidence as risk sharing for two reasons: first, the degree of consumption smoothing is too strong, relative to that typically achieved by household self-insurance (see also Lim (1992)). Second, the limited commitment friction that we analyse in this paper, when applied to lending relationships, typically restricts borrowing to very small amounts. Households would thus largely have to rely on the accumulation and decumulation of individual asset holdings, which, however, are low in poor agricultural villages.

Figure 3: Consumption growth as a function of income growth in the ICRISAT dataset

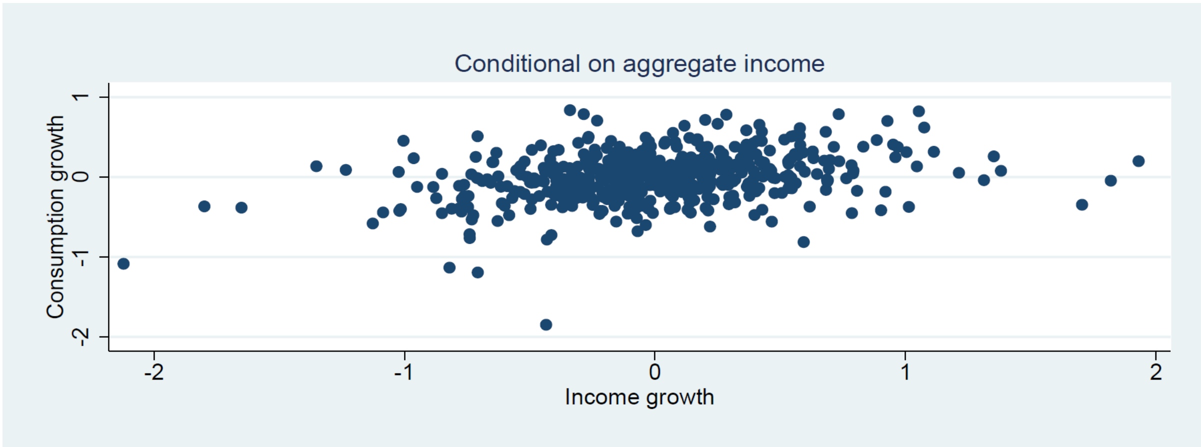

Notes: The figure shows a scatter plot of consumption and income growth in Aurepalle, Kanzara and Shirapur. The upper panel shows the raw data, the lower panel the differences in residuals from a regression of both log-income and log-consumption on time dummies.

zero income growth. Specifically, the point estimate of their difference (scaled by the variance of income growth to lie between 0 and 1), which is the moment we use in the estimation of the theoretical models in Section 5, is always negative, and for Aurepalle and Kanzara we cannot statistically reject the hypothesis of symmetry (corresponding to a 0 difference) at usual levels of confidence. The corresponding ratios of the variances for the two subsamples are 0.75, 0.64 and 0.60 in the three villages, which should be compared to a ratio larger than 50 for the standard theoretical model in Table 1.

In panel B of Table 3, we regress consumption growth on income growth and the interaction of income growth and a dummy for whether the household saw positive income growth.[15] Similar to the previous result, the data features an association of consumption and income growth whose point estimate is smaller for households with rising income, and again, the difference is not

[15]In Panel B of Table 3, 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 3. 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.

statistically different from 0 in two of the three villages, in this instance Kanzara and Shirapur. The difference in regression coefficients in Aurepalle, in contrast, is more strongly negative and statistically different from 0. This is partly due to a small number of observations with very large income observations (as seen in Figure 3).

Both sets of results point in the same direction: there is little difference in the amount of insurance obtained for positive and negative income growth, and where differences are significant they are not consistently so across the two moments capturing the asymmetry. In any case, in stark contrast to the implications of the dynamic limited commitment model presented in the previous section, the estimated magnitudes imply that, if anything, those with income losses have less insurance.

4 Consumption Risk Sharing in endogenous groups

This section proposes a quantitative model of dynamic risk sharing in endogenous groups. The key friction that gives the model predictive power for group sizes is the ability of households to deviate jointly as 'sub-coalitions' from any risk-sharing scheme, as in Genicot and Ray (2003).

4.1 Model setup

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 form a smaller insurance group in the future.

As proposed in Genicot and Ray (2003) and analysed in Bold (2009), the standard dynamic limited commitment model is made renegotiation-proof by abandoning the subgame perfect Nash equilibrium solution concept in favor of coalition-proofness, which allows subgroups 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 transfers and thus the degree of insurance.[16] Second, because for a given group of size n, the coalition-proof equilibrium risk sharing contract may fail to exist and the maximum sustainable size of an insurance group is known to be bounded in the alternative model (Genicot and Ray, 2003), it might also help to reduce counterfactual asymmetries, which we showed arise everywhere except in very small groups.[17] 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). This effectively implies that the risk sharing group should be as large as possible, which is naturally interpreted as the village in a developing country context or smaller exogenously bounded groups within it such as extended families.

We will now endeavor to derive a more realistic version of the standard dynamic limited commitment model in village economies – by allowing subgroups 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 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

[16] 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 letting agents return to an insurance mechanism with a certain probability (Broer, 2013).

[17] The intuition for the boundedness of groups in the model with coalitional deviations is as follows: in contrast to the standard model, both the outside option and the value from participating in the risk sharing arrangement increase in n. Now suppose that the largest stable group size was infinite. Then there would be a large group of size n that would be able to capture almost all the benefits from risk sharing, since the marginal benefit of adding members goes to zero as n goes to infinity. Because of this, increasing the group further to form stable groups of size greater n would not increase the benefit from risk sharing and the participating constraints would therefore bind tightly, implying transfers close to zero. The group of size n, however, would then destabilize the larger groups.

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).

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 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 swells to $2^n - 2$. Moreover, with individual deviations the outside option is completely determined by current income for any member of the insurance scheme. In the case of coalitional deviations, however, a deviating subgroup 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

$$(7) \qquad U_s^n(U_s^1, U_s^2, ..., U_s^{n-1}) = \max_{((U_r^i)_{r=1}^S)_{i=1}^{n-1}, (c_s^i)_{i=1}^n} u(c_s^n) + \delta \sum_{r=1}^S \pi_{sr} U_r^n \left(U_r^1, ..., U_r^{n-1}\right)$$

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

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

Here, the subscript on V_r indicates that the expectation of life-time utility is to be taken over the probability distribution that ensues after income realisation r.[18][19] The usual aggregate resource and promise-keeping constraints must also hold.

The set-up illustrates all the potential coalitional deviations from a group. To find the constrained-optimal contract for a group is not trivial, however, since the constraints in their form above cannot be used to define a constrained maximisation 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 cast this as a constrained dynamic social planner programme

[18] Note that we follow Genicot and Ray (2003) in not allowing deviating agents to share their income in the period of deviation. Since in our model simulations, deviations usually occur with households with similar (high) income realisations this is without loss of generality.

[19] For ease of interpretation of the stability constraints, consider an example with $n = 3$. In this case the set $\mathbf{J}$ of potential sub-coalitions has $m = 2^n - 2 = 6$ members. These are: $J_1 = \{1\}$, $J_2 = \{2\}$, $J_3 = \{3\}$, $J_4 = \{12\}$, $J_5 = \{13\}$, $J_6 = \{23\}$. Let's consider $J_6 = \{23\}$. The first member of this coalition is $j_{6,1} = 2$ and the second member is $j_{6,2} = 3$. Using the above notation, the enforcement constraints for this coalition state that there must be no $\mathbf{V}^{J_6} = \{V^{1,J_6}, V^{2,J_6}\} \in \mathbb{V}^{J_6}$ such that $U_r^2 < u(y_r^2) + \delta V_r^1(J_6)$ and $U_r^2 < u(y_r^2) + \delta V_r^2(J_6)$, where $V_r^1(J_6)$ is the continuation utility V^{1,J_6} if the previous state was r.

by restricting the set of possible deviations a subgroup can make. 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 subgroups of size m that lie on the constrained-optimal Pareto frontier of the deviating subgroup; (iii) to deter a threat it suffices to make a subgroup 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 subgroup, 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.

4.2 Approximate solution

As the previous discussion shows, the requirement of coalition-proofness is appealing because it makes risk sharing allocations immune to any collective deviation of any subgroup. The price of this conceptual appeal, however, is complexity, as even with the simplification in Bold (2009), this remains a high-dimensional problem.

As a result, it is not possible to compute the optimal coalition-proof allocation for the size of the villages under study here. 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, Thomas and Worrall (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 (7), 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.[20] 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 enters a dynamic risk sharing contract starting with equal weights, that adjust, just as in the original group, in response to constraints binding. 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

[20]We have also explored an alternative version in which the extent of risk sharing for the rest-of-village is constrained-efficient. The results are similar, but group sizes are, if anything, even smaller in this version of the model.

model. Since, as long as there is some risk sharing, income persistence and equal treatment together imply that 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 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, whichever 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 smallest 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 more complex than the standard one. In the latter, only aggregate income of the rest of the village matters, which determines the 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 a state variable. This puts an upper bound on the maximum village size N we can consider.

In summary, to find a quantitative solution of the dynamic limited commitment model with coalitional deviations, we

1. consider risk sharing between a household and a representative 'rest-of-the-village' household made up of the remainder,
2. allow individual households to deviate with those with the highest income realizations in a sustainable sub-coalition of size $n-1$ or smaller,
3. assume that, after one period of individual autarky following a deviation, any reneging sub-coalition can enter a dynamic risk sharing arrangement starting with an equal division of surplusses, or that required by any binding participation constraint, whichever larger.

4.3 Discussion

Throughout the analysis, we maintain a number of simplifying hypotheses. First, as in Genicot and Ray (2003), agents can form new subgroups only within an existing insurance group. While there may be many, unmodelled, reasons why (sub-)group-formation requires previous interaction, the motivation for this assumption is that group-formation without restriction typically causes problems even for the formulation of a solution concept, partly due to "cyclical blocking chains".[21] Also, we assume that agents experience different income realisations but are otherwise identical, that insurance takes place in groups, and that insurance transfers are only constrained by the group-level budget constraint and enforcement constraints. We thus abstract from heterogeneity in income processes (Ligon, Thomas and Worrall, 2002) or preferences (Laczo, 2014), as well as limits to information within groups (Kinnan, 2014; Ligon, 1998). This is partly because additional dimensions of heterogeneity and additional frictions would make the quantitative analysis of the model with coalitional deviations infeasible, but also because we

[21] See Genicot and Ray (2003), p. 97 for a discussion.

believe that the effect of coalitional deviations is best highlighted in the most standard version of the limited commitment model.[22] Importantly, as Section A.2 in the Online Appendix shows, the finding of strong asymmetry in the standard version of the model with individual deviations is not qualitatively affected by a stylised form of heterogeneity in preferences.[23] With coalitional deviations, however, apart from being computationally infeasible, additional dimensions of heterogeneity create a conceptual difficulty: in the presence of observed heterogeneity in incomes or preferences, some households are more attractive coalition partners than others, for example because they are less risk-averse or have lower income volatility. Moreover, preferences for attributes of coalition partners may also vary across heterogeneous households. As a first step, we therefore compare the standard model to the alternative with coalitional deviations under the assumption that households only differ in their history of ex post income realisations.

Our maintained assumptions also imply that we completely abstract from any network structure of the village or its subgroups. In fact, we view our work as complementary to studies analysing the formation of insurance networks with limited commitment (Bloch, Genicot and Ray, 2008; Ambrus, Mobius and Szeidl, 2014) where the focus on the structure of stable networks, however, requires a simplification of the analysis along dimensions that are central to our study.

5 Results

We now show how the quantitative model of risk sharing we proposed in the previous section predicts endogenous groups of between two and 5 households, substantially smaller than the ICRISAT villages and their samples analysed in Section 3, but also smaller than most alternative

[22]For example, Laczo (2014) finds evidence of preference heterogeneity when estimating the standard limited commitment model with individual deviations. And Mazzocco and Saini (2012) reject the joint hypothesis of full insurance and homogeneous risk preferences for caste groups in the ICRISAT villages, but cannot reject full insurance when allowing for heterogeneity in risk preferences.

[23]As the results show, if anything, the asymmetry becomes larger. For a version of the standard limited commitment model with a large number of households, Broer (2013) finds that the mean moments based on log-differences are essentially unaffected by heterogeneity in preferences as, for example, lower-than-average consumption volatility of some households offsets higher volatility of others.

exogenous groups such as typical extended families or castes within the village. Importantly, together with a more attractive outside option of coalitional deviations, this allows the model to predict both a correct degree of insurance and symmetric responses of consumption to positive and negative income shocks. Depending on preference parameters, the standard model with individual deviations, in contrast, can predict either a realistic degree of insurance or symmetric responses of consumption to income movements, but not both at the same time.[24]

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, Thomas and Worrall (2002)), or by making assumptions about the distribution of measurement error and derive the likelihood of different model specifications given the data conditional on the observed consumption shares in the first sample period, as in Laczo (2014).

The absence of information on group membership in the sample makes a conditional likelihood approach like that in Laczo (2014) infeasible, as the likelihood depends on the allocation of individuals to groups (that may comprise households not in the sample). In this paper, we therefore concentrate on a small number of moments that describe the key risk sharing properties of the data and the models, as well as the asymmetries in the joint distribution of household 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

[24]To derive these results we use the standard approximation, based on an individual household in an insurance relationship with a representative 'rest-of-the-village', adapted to the context of coalitional deviations. The Online Appendix, however, shows how our findings about the performance of the standard model relative to that with coalitional deviations are robust to the use of this approximation for the case of three-member insurance groups, where we can compute the exact solution and which is a rather typical group size in the version of the model with coalitional deviations. In fact, while Bold (2009) performs a similar analysis of the model without income persistence, we are, to our knowledge, the first to study the exact solution of the three agent-limited commitment model with an estimated, persistent income process and compare it to data from actual village economies.

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, the individual deviations model makes no prediction on group size - other than that groups should be as large as possible – and this has motivated researchers to examine risk sharing at the village level as a plausible exogenous risk sharing group. The villages in our data set, Aurepalle, Kanzara and Shirapur all contain several hundred households, but we only observe a sample of households in each village.[25] We therefore follow the standard practice in the literature (Ligon, Thomas and Worrall, 2002; Laczo, 2014) and estimate the model with village size equal to the number of households sampled by the ICRISAT. 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 assume that households can only form risk sharing arrangements with N other households where N equals the ICRISAT sample in each village. Hence, the largest group that can form has the same size as in the individual deviations model, but the largest stable group will typically be much smaller.[26] While the set of stable sizes may

[25]The fact that we observe a sample rather than the whole village may affect the estimated partial degree of insurance, especially when we condition on aggregate village resources. To address this, we also compare unconditional versions of the moments in the data and model.

[26]We make this assumption for practical purposes since estimation of the standard model and especially the recursive estimation of the alternative model with the full income distribution of the rest of the village are simply not computationally feasible for larger N. The fact that we limit the maximum group size to the sample has implications for the asymmetry the two models will generate, however, as seen in Section 2.3. Allowing group size to be larger than the sample would imply even more extreme values for the asymmetry in the standard model. It would only affect the results of the coalitional deviations model

contain several different group sizes, we select here the stable group size associated with the highest insurance benefit and thus concentrate on the largest group size n^{max}. 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 least the number of villagers in the data. Or more formally, we find the smallest k such that $N' = k \times n^{max} \geq N$, where N is again the sample size in the three ICRISAT villages. So N' is the number of individuals in our simulation of the coalitional deviations model. Since below we estimate n^{max} to be a small single-digit number, the resulting maximum difference in village size between the cases of individual and coalitional deviations is small.[27]

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

$$y_{it} = \alpha + \rho y_{it-1} + \epsilon_{it} \tag{8}$$

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

$$\begin{aligned} \rho &= \frac{Cov_y}{Var_y} \\ Var_\epsilon &= Var_y(1-\rho^2) \end{aligned} \tag{9}$$

Table 4 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

inasmuch as there are stable groups beyond the sample size, which is the maximum we consider.

[27] Similarly, because the largest stable group within the village (or rather the sample) tends to be small relative to the village, this is also very similar to an approach that holds village size constant and allocates households to stable groups in a way that maximises expected utility ex-ante.

Table 4: 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 (8).

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 three support points using Rouwenhorst's (1995) 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 subgroup, 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 (8), 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 rest of the village as an exogenous state variable and compute the transition probabilities on the basis of the individual income process in (8).

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

$$u(c_t) = \frac{c_t^{1-\sigma} - 1}{1 - \sigma} \tag{10}$$

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 realisations and then simulate consumption on the basis of the models' policy functions for N and N' households, the village sizes in, respectively, the standard and coalitional deviations model, in $T = 6,200$ periods. After discarding the first 200 periods, we then calculate moments from this simulated sample that summarise 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 moments similar to those in Table 1. Specifically, as summary measures for the degree of insurance, we use the same moments as before: the regression coefficient of consumption growth on income growth, β_{dcdy}, as a measure of the average effect of income changes on consumption; and 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. To capture non-linearities in the mean and variance function of consumption growth conditional on income growth, we also look at the relative consumption volatility of households with increasing vs. non-increasing income, measured by the difference in their respective variances of consumption growth $\frac{Var_{dc|dy>0} - Var_{dc|dy\leq 0}}{Var_{dy}}$ normalised by the average variance of income growth; and the difference of regression coefficients of consumption growth on income growth for households with rising and non-rising income ($\beta_{dcdy|dy>0} - \beta_{dcdy|dy\leq 0}$).[28]

We calculate all moments after subtracting period-specific village-averages from the individual data (observed and simulated) to make our results comparable to the empirical practice of including time dummies in risk sharing regressions (see e.g. Deaton (1990), Cochrane (1991), Ravallion and Chaudhuri (1997) or Laczo (2014)), and robust to any correlation in individual incomes not captured by the assumption of independent individual incomes. The results with-

[28]The latter two moments are now measured as differences, in contrast to the ratios presented in Table 1. This is because, for moments close to those observed in the data, the ratios are a more revealing summary measure. Whenever insurance is (close to) perfect, however, use of the ratio implies division of two zeros and thus undefined asymmetry moments. Since the grid on which we solve the model during our estimation encompasses parameter combinations for which full insurance is the solution, we avoid this problem by choosing the difference measure here (scaled by the average variance of income growth in case of the relative variance to ease interpretation).

out this conditioning on village-level aggregates, contained in the Online appendix, are very similar.[29]

5.4 Model performance with common preference parameters

To understand the different predictions of the two model specifications, Table 5 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. In line with Section 2.2, the simple version of the standard model predicts a degree of risk sharing that is even higher than the strong co-insurance observed in two of the three villages, with a regression coefficient β_{dcdy} predicted to equal about one half of the coefficient estimated from the data in Aurepalle and two thirds in Shirapur, and a relative consumption volatility $\frac{Var_{dc}}{Var_{dy}}$ an order of magnitude smaller than in the data for those villages.[30] Because incomes are more persistent and income volatility is smaller in Kanzara, autarky is estimated to be more attractive for high-income villagers there, which reduces insurance, and results in a regression coefficient of consumption on income growth in Table 5 of about fifty percent larger than the 22 percent observed in the data. Still, the model predicts less than half the consumption volatility observed in the data.

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-household risk sharing coalitions.[31] As a result, the model predicts less insurance than the data, with a

[29]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. Note that conditioning may affect the estimates of β_{dcdy} in the presence of preference heterogeneity when income of less risk-averse households comoves more strongly with aggregate income, as assumed by Mazzocco and Saini (2012). Without conditioning, however, we find very similar results, see Section A.4 in the Online Appendix.

[30]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 between 1500 and 2100 values corresponding to all different combinations of $N = 34, 37, 31$ incomes in our solution, rather than just 5 values).

[31]Strictly speaking, we have only shown here that the pair is the only stable group size within the

Table 5: Risk sharing moments for common preferences

	Aurepalle			Kanzara			Shirapur		
	Data	**DLC with deviations by: individuals**	**DLC with deviations by: coalitions**	**Data**	**DLC with deviations by: individuals**	**DLC with deviations by: coalitions**	**Data**	**DLC with deviations by: individuals**	**DLC with deviations by: coalitions**
n		34.00	2.00		37.00	2.00		31.00	2.00
δ		0.90	0.90		0.90	0.90		0.90	0.90
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.05	0.56	0.56	0.26	0.60	0.33	0.06	0.54
β_{dcdy}	0.21	0.10	0.53	0.22	0.33	0.63	0.17	0.12	0.53
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.10	0.05	-0.25	0.45	0.08	-0.16	0.10	0.03
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.41	0.19	-0.00	-0.14	0.54	0.00	-0.05	0.21	-0.00

Notes: For each of the three ICRISAT villages, the table shows four key moments summarising the joint distribution of consumption and income growth in the survey ("Data", in the first column for each village), and in simulations of the two versions of the dynamic limited commitment model ("DLC"): the standard model with individual deviations (in the second column for each village), and the alternative model with coalitional deviations (in the third column for each village). The four moments are: the relative variance of consumption and income growth $\frac{Var_{dc}}{Var_{dy}}$, the regression coefficient in a bivariate regression of consumption growth on income growth β_{dcdy}, the difference in the variances of consumption growth for agents with increasing and for those with non-increasing income (scaled by the variance of income growth) $\frac{Var_{dc|dy>0}-Var_{dc|dy\leq 0}}{Var_{dy}}$, and the difference in regression coefficients for those two subsamples $\beta_{dcdy>0} - \beta_{dcdy\leq 0}$. For the simulated model solutions, the table also presents the size of the insurance groups n and the discount factor δ.

regression coefficient β_{dcdy} estimated to be two and a half to three times larger than in the data. Predicted consumption volatility is also larger than in the data, apart from Kanzara, where the model only slightly overpredicts it.

In line with the results in Section 2.3, with 2-household insurance coalitions, the alternative model predicts a symmetric distribution of consumption and income changes. There is no difference in the linear association of income and consumption growth between households with income increases and those with decreases, and the difference in consumption growth variances of these households are small both in absolute terms and – even more so – relative to the large consumption growth volatility predicted by the model. In the standard model, in contrast, there is strong asymmetry: the difference in regression coefficients - zero in the alternative model - is almost twice as large as the average coefficient. Also, despite much smoother predicted consumption overall, the difference in consumption growth variances in the standard model is two to five times that in the alternative model. The alternative model thus fits better this dimension of the data, which, as already noted in Section 3, is largely symmetric and – in the few instances of a significant asymmetry – has asymmetries in the other direction. Note that, in contrast to what the point estimates of the data moments indicate in several cases, neither 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 tends to bind and therefore make consumption sensitive to income growth in periods of rising income.

The results above are obtained for homogeneous preferences. For the most part, we abstract from heterogeneity in preferences in this paper, despite evidence that this could improve the ability of the standard model to explain consumption and income comovement in ICRISAT data (Laczo, 2014). This is both for conceptual and computational reasons. Specifically, without strong auxiliary assumptions, it is unclear how to select small groups from a village whose

sample (i.e. $N = 34, N = 37$ and $N = 31$). We cannot rule out that there are stable sizes larger than the sample but smaller than the actual ICRISAT villages. In that sense, even the alternative version of the limited commitment model with endogenous group formation relies on an exogenous bound below which stable group sizes are identified.

households differ, for example, in their level of risk aversion. Section A.2 in the Online Appendix, however, presents simulations of the standard model with a stylised form of preference heterogeneity that show how the asymmetry, if anything, increases in that case.

5.5 Estimating the model

The parameterisation of preferences in Section 5.4 is close to estimates found in previous studies but ultimately arbitrary. In this section, we therefore use a simulated method of moments approach to choose preference parameters that minimise the distance between the selected moments from the ICRISAT villages and from simulated samples generated by the standard dynamic limited commitment model and the alternative with coalitional deviations.

As noted in Laczo (2014), a necessary condition for identification in the standard model with individual deviations is that at least one household has as many binding constraints as there are parameters to estimate. But because insurance is relatively strong in the ICRISAT villages, for most of the parameter space and in particular in the range where the model's fit is best, at most one constraint is binding. It is therefore not possible to identify time and risk preferences separately in any of the two models.[32] Figure 4 in the Online Appendix illustrates this for the case of Aurepalle. It shows that, for every value of the discount factor δ there is a value of risk aversion σ that delivers the same goodness of fit, and the same corresponding moments, for both models. We therefore normalise the coefficient of relative risk aversion σ to 1 in the estimation (log-preferences), and choose the discount factor that minimises the distance between the moments in the model simulation and in the data.[33]

[32]Note that we compute the model for three individual income states, which is the minimum needed for identification. However, robustness checks with a higher number of income states did equally not yield identification, simply because the targeted degree of insurance is too high. In the model with coalitional deviations the necessary condition for identification is for at least one household to have as many binding constraints as parameters to be estimated in any of the sustainable subcoalitions.

[33]Note that, given the lack of identification, we could equally have chosen to normalize the discount factor and estimate the value of risk aversion. We choose our normalization for comparability with previous studies, such as Laczo, who also normalizes the mean of risk aversion to one. In Section 6 and A.5 of the Online Appendix, we consider more extreme parameterisations with alternative values of σ, implying allocations characterised by stronger or less insurance than in the data that are not generated by log-preferences for the values of the discount factor we consider.

For the estimation of the models, the criterion to be minimised 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.[34] The variances are obtained by bootstrapping the data 1,000 times.

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.

Table 6 presents the moments of interest when the utility function is logarithmic and the discount factor δ is estimated to target $\frac{Var_{dc}}{Var_{dy}}$ and β_{dcdy}, the two moments that summarise the extent of insurance in the whole sample. In order to understand the estimates, it is useful to recall the role of the discount factor: since deviation delivers higher mean consumption in earlier periods at the price of eternally higher consumption volatility, higher discount factors, like higher risk aversion, deter deviation and increase risk sharing. Importantly, the estimated value of the discount factor is conditional on the normalisation of relative risk aversion to 1 (log-preferences). A normalisation to higher risk aversion would thus deliver lower estimated discount factors. In other words, when interpreting the estimation results below, the focus should be on the relative values across models and estimation criteria, not the absolute level of the discount factor.

[34]For a further description and application of this procedure see for example Blundell, Pistaferri and Preston (2008). To minimise the criterion we conduct a grid search on a fine discrete grid of $\delta \in [0.5, 0.99]$. 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. The results are robust to different weighting matrices and different functional forms for the asymmetric moments.

Table 6: Preferences estimated to target degree of risk sharing - 2 moments

	Aurepalle			**Kanzara**			**Shirapur**		
	Data	**DLC with deviations by: individuals**	**DLC with deviations by: coalitions**	**Data**	**DLC with deviations by: individuals**	**DLC with deviations by: coalitions**	**Data**	**DLC with deviations by: individuals**	**DLC with deviations by: coalitions**
n		34.00	4.00		37.00	4.00		31.00	5.00
δ		0.84	0.96		0.90	0.98		0.86	0.98
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.25	0.27	0.56	0.26	0.26	0.33	0.19	0.19
β_{dcdy}	0.21	0.32	0.29	0.22	0.33	0.25	0.17	0.27	0.19
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.41	0.04	-0.25	0.45	0.02	-0.16	0.31	0.01
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.41	0.52	0.07	-0.14	0.54	0.03	-0.05	0.44	0.02
Goodness of fit		4.97	2.06		5.37	3.06		6.50	3.34

Notes: For each of the three ICRISAT villages, the table shows four key moments summarising the joint distribution of consumption and income growth in the survey ("Data", in the first column for each village), and in simulations of the two versions of the dynamic limited commitment model ("DLC"): the standard model with individual deviations (in the second column for each village), and the alternative model with coalitional deviations (in the third column for each village). For the simulated model solutions, the table also presents the size of the insurance groups n and the discount factor δ, which is chosen to minimise the sum of differences between the first two moments ($\frac{Var_{dc}}{Var_{dy}}$ and β_{dcdy}) predicted by the models 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 $\delta \in [0.5, 0.99]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

Table 7: Preferences estimated to target degrees of risk sharing and asymmetry

	Aurepalle			Kanzara			Shirapur		
	Data	**DLC with deviations by: individuals**	**coalitions**	**Data**	**DLC with deviations by: individuals**	**coalitions**	**Data**	**DLC with deviations by: individuals**	**coalitions**
n		34.00	5.00		37.00	4.00		31.00	5.00
δ		0.90	0.98		0.94	0.98		0.92	0.98
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.05	0.20	0.56	0.04	0.26	0.33	0.01	0.19
β_{dcdy}	0.21	0.10	0.18	0.22	0.09	0.25	0.17	0.05	0.19
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.10	0.01	-0.25	0.08	0.02	-0.16	0.03	0.01
$\beta_{dcdy>0}-\beta_{dcdy\leq 0}$	-0.41	0.19	0.01	-0.14	0.17	0.03	-0.05	0.08	0.02
Goodness of fit		53.26	15.40		17.64	5.70		27.98	7.89

Notes: For each of the three ICRISAT villages, the table shows four key moments summarising the joint distribution of consumption and income growth in the survey ("Data", in the first column for each village), and in simulations of the two versions of the dynamic limited commitment model ("DLC"): the standard model with individual deviations (in the second column for each village), and the alternative model with coalitional deviations (in the third column for each village). For the simulated model solutions, the table also presents the size of the insurance groups n and the discount factor δ, which is chosen to minimise the sum of differences between all four moments predicted by the models 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 $\delta \in [0.5, 0.99]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

Allowing the discount factor 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 a lower discount factor in Aurepalle and Shirapur. In contrast, higher estimated discount factors in the alternative model with coalitional deviations increase the maximum number of households that share risk with each other to 4 (in Aurepalle and Kanzara) and 5 (in Shirapur). They also bring the degree of risk sharing close to that observed in the data – apart from the relative variance of consumption and income growth, which continues to be higher than that predicted by the model in Kanzara and less in Shirapur. We attribute the smaller maximum sustainable group size in Kanzara, one less than Shirapur despite identical preferences, to the less volatile and more persistent income process there, which reduces the benefits of insurance and makes deviation more attractive.

Note that the models rely on two very different mechanisms to arrive at their prediction of the observed degree of risk sharing. The standard model predicts moderate degrees of insurance taking place at the village level. The model with coalitional deviations predicts strong insurance among a smaller number of households, which translates into moderate degrees of insurance at the village level. Importantly, this strong insurance at the group level contains the increase in asymmetry that, according to Figure 2, would otherwise accompany the rise in group size from a pair in the common preference specification to 4 or 5 when preferences are estimated. In fact, at full insurance at the group level, the asymmetry would be zero independent of group size, as individual consumption simply follows (symmetric) group-income. The asymmetry in the standard model, in contrast, is strongly increased at the lower estimated discount factor that moves the degree of risk sharing closer to the data, and thus away from full insurance, in Aurepalle and Shirapur.[35]

[35]There is, potentially, an additional, more mechanical reason for lower asymmetry in the alternative model where, as explained in Section 5.1, we keep the number of village members 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-

Table 7 shows that the standard dynamic limited commitment model is unable to predict symmetry and realistic degrees of insurance. Specifically, when we include the two asymmetry moments in our estimation criterion, which results in higher estimated discount factors in Aurepalle and Shirapur, the model strongly overpredicts both the degrees of insurance and of asymmetry. In contrast, the coalitional deviations model continues to deliver symmetry and a degree of risk sharing close to the data (in fact closer to it in the case of Aurepalle where an increased estimate of the discount factor reduces the asymmetry at the price of underpredicting somewhat the relative variance of consumption growth). In line with these results, the goodness of fit of the individual deviations model, already substantially worse in Table 6, further deteriorates relative to the coalitional model when evaluated using all four moments of interest.[36]

In sum, we have shown that the standard specification of the dynamic limited model is able to predict realistic degrees of insurance or realistic degrees of asymmetry, 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 data. Note that neither the standard model nor the alternative, however, could simultaneously predict the relatively high volatility of consumption in the data together with a relatively low sensitivity of consumption to income growth. This may be indicative of measurement error in consumption, which increases the former but leaves the latter unaffected. Section A.3 of the Online Appendix uses the same simulated method of moments procedure as in this section to show how measurement error in consumption is estimated to be virtually the only source of variation in measured consumption according to

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 A.4 of the Online Appendix, where we find very similar results.

[36]Although it is tempting to compare the corresponding measures in the final rows of Tables 6 and 7 to a χ^2 distribution, this is valid only under the strong assumption of independent moment conditions. With, respectively, one and three degrees of freedom, this would imply p-values that are between 5 and 15 for the coalitional deviations model, at least one order of magnitude higher than for the standard version. One exception to this is the case of Aurepalle when targeting all four moments: the data is asymmetric in the 'other' direction, which the coalitional deviations model can also not capture. Ultimately, the aim of our exercise is, however, not to reject, or not, any of the two versions of the, still very stylized, limited commitment model, but to highlight their different implications for the structure of consumption risk sharing and their ability to capture key moments of the data.

the standard model. Trivially, this reduces the asymmetry in consumption movements to zero, at the price of predicting perfect insurance. The alternative model estimates more limited error in measured consumption, and slightly larger error in measured incomes, as well as reduced group sizes in Aurepalle and Shirapur, which allows it to match almost exactly the degree of insurance in the data while maintaining its prediction of symmetry.

6 Endogenous or exogenous insurance groups?

All the results so far point to the fact that the model with coalitional deviations performs better than the standard model because it is able to predict the degree of risk sharing observed in the ICRISAT data at the same time as symmetric consumption responses. Much of this superior performance results from its prediction of strong insurance among smaller groups of households, which is in line with a large literature in development economics that shows how risk sharing is often confined to smaller units such as clans, castes, extended families and networks.[37] This gives rise to a simple question: what do we gain from studying insurance groups with coalitional deviations, where small groups are an equilibrium outcome, relative to a version of the standard model with a smaller exogenous group size equal to that of such exogenous insurance units?

This section highlights two important features of the coalitional deviations model that allow it to predict the data better than the standard model even with smaller exogenous groups: first, the endogenous groups predicted by the coalitional deviations model are in fact smaller than typical exogenous ones, such as extended families or castes. This matters because, as we show below, the goodness of fit of the standard model deteriorates rapidly even at large single-digit group sizes. And second, even for identical small group sizes, as difficult as they may be to motivate in the standard model, the version with coalitional deviations predicts the

[37] 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).

observed degree of insurance better. This is because coalitional deviations imply more attractive outside options, which restrict insurance more and thus improve its fit with the data relative to the standard model, which is known to predict too strong insurance (Ábrahám and Cárceles-Poveda, 2009; Ligon, Thomas and Worrall, 2002).

In Table 8, we present estimates of the standard model with individual deviations but smaller exogenous insurance groups ranging from those estimated by the coalitional deviations model in Table 7 to a maximum of 10 households in all three villages[38]. The first thing to note is that, for 8 or 10 households in a group, the standard model predicts degrees of both insurance and asymmetry in consumption-income comovements that are similar to those of the village-level model. In fact, the goodness of fit deteriorates strongly as the number of household members rises, and is very similar to that of the village-level model for groups of 10 households. So the standard model with smaller groups whose size equals that of plausible exogenous insurance communities such as extended families does not fit the data well.

Interestingly, even for an exogenous group size equal to that estimated by the coalitional model in Table 7 (5 in Aurepalla and Shirapur, 4 in Kanzara), where both models predict roughly symmetric consumption-income comovements, the standard model fits the data less well as it strongly overpredicts the degree of insurance, despite lower estimated discount factors. This results from its more severe punishment of individual autarky that sustains too much risk sharing in equilibrium. For comparison, the sensitivity of consumption to income is between 30%-40% larger in the model with coalitional deviations in Table 7. That model achieves its better fit because, even for a given group size, its outside option to deviate with a sub-coalition of the group is always more attractive than individual autarky, implying less insurance in line with the data.[39]

[38]This range of sizes corresponds to other exogenous (and endogenous) sub-divisions related to extended families and networks documented in the literature. For example, Fitzsimons, Malde and Vera-Hernandez (2015) find that, in Malawian data, a household has on average 9.4 siblings of husband or wife (although not all of them live in the same village). And, looking at extended familiy networks in Mexico, Angelucci, de Giorgi and Rasul (2015) find that an average of 7.5 households belong to the same family within a village. Fafchamps and Lund (2003) find that households in the Philippines make and receive transfers from an average of 5 other households.

[39]Importantly, this implication of stronger insurance in the standard model even at given group size is

It might appear that the small group sizes estimated in the coalitional deviations model stand in contrast to a literature that has established the importance of larger non-village risk sharing units such as castes. This is not a contradiction, however, inasmuch as these papers do not test for full insurance in these risk sharing units, as is the case in Mobarak and Rosenzweig (2012).[40] As should be clear from the analysis in Section 5.5, moderate insurance in a larger exogenous group could equally be generated through relatively high insurance in several smaller groups, that are not clearly delineated by virtue of their endogeneity. That is to say, the coalitional deviations model does not require clearly discernable exogenous barriers to group size, since it delivers small groups simply through the threat of coalitional deviations. Taken together, we therefore think the prediction of small groups is plausible in the coalitional deviations model even if these may be absorbed in larger exogenous groups. Conversely, our result show that the limited commitment friction alone, in its generalised version with coalitional deviations but without other costs of deviation, is not able to explain insurance transfers within substantially larger groups for the income processes we identify in ICRISAT data. Importantly, small group sizes are not an intrinsic feature of the new model we propose per se. It is therefore possible that our model of risk sharing in endogenous groups, together with other frictions or in contexts with different income risk, is able to predict larger group sizes. Whether this implies degrees of insurance and of symmetry in consumption responses to income shocks in line with the particular data under study would then have to be investigated.

---

not an artifact of the particular group sizes and preferences estimated by fitting the model to the ICRISAT data. Comparing the degree of insurance in the standard model and the alternative at identical group sizes across a wide range of identical preference parameters δ and σ for which the coalitional deviations model has stable group sizes and limited risk sharing, we find that insurance is always stronger in the former. Moreover, we predict the degree of insurance to be 30% lower in the coalitional deviations model for groups of four and 45% lower for groups of 5 (see Figure 7 in Section A.5 of the Online Appendix), very similar to the relative degree of insurance comparing the fitted models.

[40] When allowing for preference heterogeneity, Chiappori, Samphantharak, Schulhofer-Wohl and Townsend (2014) and Mazzocco and Saini (2012) cannot reject the null of full insurance in villages and castes respectively. Our study is, however, quite different since it is implicitly based on a null hypothesis of limited commitment to contracts being a friction to risk sharing.

Table 8: Moments for given group size - ID model

	Aurepalle					Kanzara					Shirapur				
	Data	**DLC with deviations by: individuals**				**Data**	**DLC with deviations by: individuals**				**Data**	**DLC with deviations by: individuals**			
n		5.00	6.00	8.00	10.00		4.00	6.00	8.00	10.00		5.00	6.00	8.00	10.00
δ		0.97	0.92	0.92	0.90		0.96	0.94	0.94	0.93		0.93	0.93	0.92	0.91
σ		1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00		1.00	1.00	1.00	1.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.13	0.11	0.09	0.04	0.56	0.20	0.13	0.11	0.07	0.33	0.13	0.11	0.09	0.03
β_{dcdy}	0.21	0.11	0.12	0.10	0.09	0.22	0.19	0.16	0.14	0.13	0.17	0.14	0.11	0.11	0.08
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.00	0.03	0.03	0.08	-0.25	0.02	0.08	0.08	0.13	-0.16	0.01	0.01	0.03	0.05
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.41	-0.00	0.05	0.05	0.16	-0.14	0.01	0.11	0.13	0.23	-0.05	0.03	0.04	0.07	0.14
Goodness of fit		24.84	29.46	35.06	51.86		6.72	11.15	12.64	17.05		11.95	14.08	17.40	27.69

Notes: For each of the three ICRISAT villages, the table shows four key moments summarising the joint distribution of consumption and income growth in the survey ("Data", in the first column for each village), and in simulations of the standard dynamic limited commitment model ("DLC") with individual deviations for group sizes between that estimated by the alternative coalitional deviations model in table 7 and $n = 10$ households. For the simulated model solutions, the table also presents the size of the insurance groups n and the discount factor δ, which is chosen to minimise the sum of differences between all four moments predicted by the models 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 $\delta \in [0.5, 0.99]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

7 Conclusion

In this paper, we have argued to replace the 'village', or in fact any other exogenous risk sharing group in poor agricultural communities, with a concept of endogenous groups of mutual insurance. For this, we have proposed a quantitative model of dynamic risk sharing with limited commitment whose predictive power for group sizes arises from the ability of households to deviate from any risk-sharing scheme jointly as 'sub-coalitions', as in Genicot and Ray (2003). We showed how, for realistic income risk and preferences, this renegotiation-proof coalitional deviations model predicts insurance groups of up to five households, smaller than the village, and smaller also than typical exogenous groups such as extended families or castes within the village. Importantly, it is precisely this prediction of small insurance groups, together with a more attractive outside option of continued co-insurance, that enables the model to predict a realistic degree of insurance at the same time as symmetric responses of consumption to income shocks. In a standard limited commitment model, in contrast, the fact that participation constraints are more likely to bind in periods of high income implies an asymmetry in the consumption response to income shocks that is counterfactually strong both in village-sized groups and smaller ones, such as extended families, unless preference parameters are such that insurance is, essentially, predicted to be perfect.

We think that our results raise several interesting questions for future research. First of all, although motivated by the quest for better policies, this paper has not analysed how the model with endogenous group sizes responds to policy interventions, such as public income insurance policies. Our results have shown that for the less volatile (but more persistent) income process estimated for Kanzara, sustainable insurance groups are smaller. This suggests that interventions that change income processes may also change the size of insurance groups, an effect that previous contributions based on the standard model, such as Attanasio and Rios-Rull (2000), were not able to analyse. More generally, it would be interesting to perform a comprehensive comparative statics exercise that studies how group size, as well as the degree of insurance and symmetry of the resulting joint consumption-income process, respond to changes in the environment such as a change in income risk, access to formal financial markets, and others. Finally, it would be interesting to analyse formally the dynamics governing the formation and break-up of risk sharing groups when there are (anticipated) changes in the environment and groups are known to (potentially) be temporary, as in Bloch, Genicot and Ray (2007). We think that, beyond the application to village risk sharing used in the present study, these issues should also be important for analysing the stability of, for example, nation-states made up of different regions, or groups of countries that share risk in international organisations or regional unions.

References

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

Ambrus, Attila, Markus Mobius, and Adam Szeidl, "Consumption Risk-Sharing in Social Networks," *American Economic Review*, 2014, *104* (1), 149–82.

Angelucci, Manuela, Giacomo de Giorgi, and Imran Rasul, "Resource Pooling within Family Networks: Insurance and Investment," 2015. Mimeo.

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.

Bloch, Francis, Garance Genicot, and Debraj Ray, "Reciprocity in Groups and the Limits to Social Capital," *American Economic Review*, 2007, *97* (2), 65–69.

———, ———, and ———, "Informal insurance in social networks," *Journal of Economic Theory*, November 2008, *143* (1), 36–58.

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

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

——— and Stefan Dercon, "Insurance Companies of the Poor," *CEPR Discussion Papers*, 2014, *10278.*

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.

———, Marek Kapička, and Paul Klein, "Consumption Risk Sharing with Private Information and Limited Enforcement," CERGE-EI Working Papers wp531, The Center for Economic Research and Graduate Education - Economics Institute, Prague February 2015.

Chiappori, Pierre-André, Krislert Samphantharak, Sam Schulhofer-Wohl, and Robert M Townsend, "Heterogeneity and risk sharing in village economies," *Quantitative economics*, 2014, *5* (1), 1–27.

Cochrane, John H., "A Simple Test of Consumption Insurance," *Journal of Political Economy*, February 1991, *99*, 957–976.

Deaton, Angus, "On Risk, Insurance and Intra-Village Consumption Smoothing," 1990. Mimeo.

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, Pierre, "Heterogeneity of Preferences, Limited Commitment and Coalitions: Empirical Evidence on the Limits to Risk Sharing in Rural Pakistan," CEPR Discussion Papers 6004, C.E.P.R. Discussion Papers December 2006.

——, **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.

Fitzsimons, Emla, Bansi Malde, and Marcos Vera-Hernandez, "Group Size and the Efficiency of Informal Risk Sharing," mimeo, University College London 2015.

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," 2014.

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

Krueger, Dirk and Fabrizio Perri, "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., "Risk-Sharing and Information in Village Economies," *Review of Economic Studies*, 1998, *65*, 847–864.

____ **, 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.

Lim, Yongjae, "Disentangling permanent income from risk sharing: a general equilibrium perspective on credit markets in rural South India." PhD dissertation, University of Chicago, Department of Economics 1992.

Mazzocco, Maurizio and Shiv Saini, "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," *American Economic Review*, February 2012, *102* (1), 428–68.

Mobarak, A. Mushfiq and Mark Rosenzweig, "Selling Formal Insurance to the Informally Insured," Technical Report, Yale University 2012.

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.

A FOR ONLINE PUBLICATION: Appendix

In this online appendix, we discuss the results and their robustness along a number of dimensions. First, we illustrate the identification problem in the limited commitment model at high levels of insurance. Second, we estimate the standard limited commitment model with preference heterogeneity. Third, we estimate both models on the basis of unconditional moments, rather than the residuals from a regression on time dummies. Fourth, we introduce measurement error. Fifth, we compare the degree of insurance in the standard model and the alternative for a given group size and a wide range of preference specifications. Sixth, we compare the exact and approximated solution for small group sizes. Finally, we present additional descriptive statistics from the ICRISAT data.

A.1 Identification of preference parameters

Figure 4 shows how neither version of the model separately identifies the two preference parameters, risk aversion σ and discount factor δ. Rather when σ declines - at a higher level in the case of coalitional deviations where average insurance is lower - as δ rises, neither the moments nor the goodness of fit changes. We thus normalise σ to 1 (log-preferenes) in the main text.

A.2 Preference heterogeneity in the standard model

First, note that, trivially, preference heterogeneity is in principle able to reconcile a realistic degree of insurance and symmetry in consumption. This is because both autarkic allocations (where households consume their income) and full insurance imply symmetry. The right mix of approximately risk-neutral and highly risk-averse households, experiencing perfect and zero income-consumption comovement respectively, can thus always deliver the right average co-

Figure 4: Estimated parameters and moments in Aurepalle

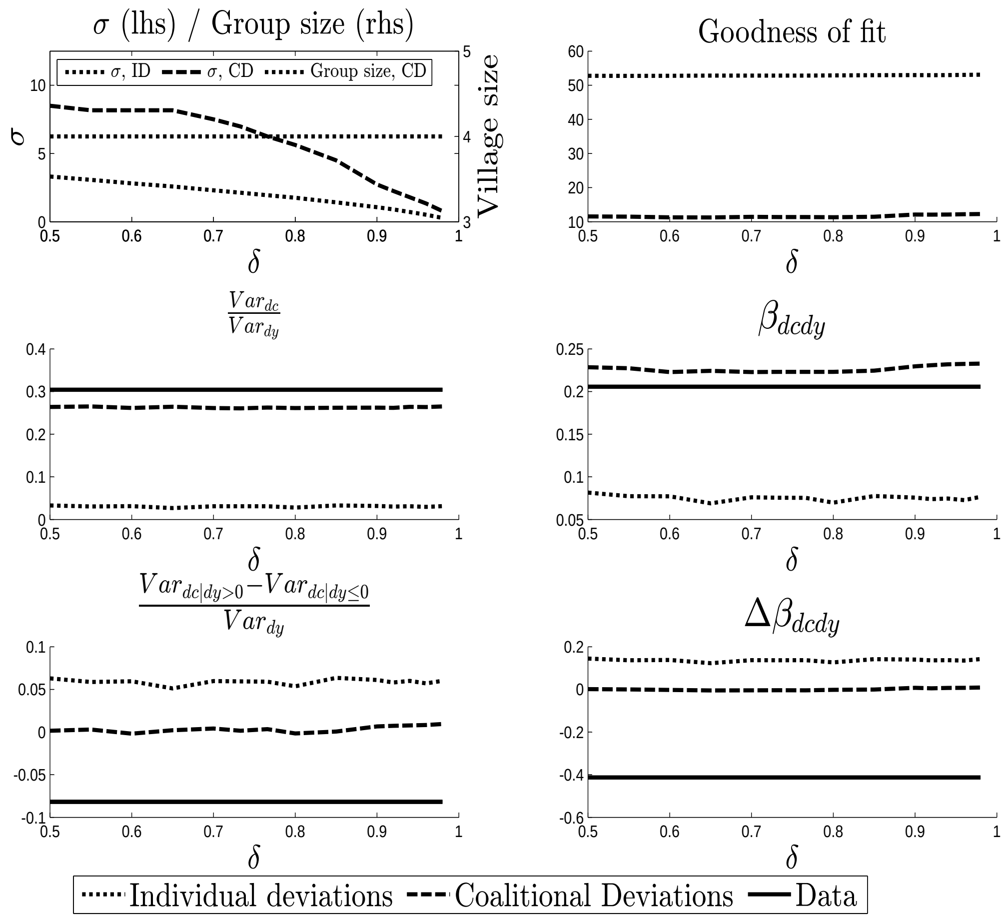

Notes: The figure shows the estimated risk aversion σ and (in the case of coalitional deviations) village size (panel 1), the goodness of fit (panel 2) and the 4 moments of interested for different values of the discount factor δ.

movement and symmetry in our moments-based approach. For moderate degrees of preference heterogeneity, in contrast, the asymmetry predicted by the standard model is likely to be stronger than in the benchmark, however. This is because insurance is predicted to be powerful in the standard model. Introducing dispersion in, for example, risk aversion around its estimated mean moves some households even closer to perfect insurance, which does not change their consumption moments much. It moves others further away from full insurance, however, which typically implies an increase in asymmetry (as they move up the right hand side of the inverse U-shaped relation between asymmetry and the degree of risk sharing). This effect is thus likely to increase the asymmetry on average in the village, and to decrease the average degree of insurance. It is counteracted by the fact that differences in risk aversion imply a higher speed of consumption decline for less risk-averse unconstrained agents, whose consumption is optimally more volatile. The variance of consumption growth of the unconstrained thus increases. As it turns out, at the strong insurance predicted by the standard model, the first effect dominates, and the asymmetry increases when we risk-preferences are heterogeneous.

Table 9 illustrates this point, by showing the four moments of interest when assuming the same time discount factor as in table 5, but an equal number of villagers with risk aversion coefficients equal to 0.5 and 1.5. In all three villages, the asymmetry increases in line with the decrease in average insurance.

A.3 Measurement Error in Consumption and Incomes

As Section 5.5 in the main paper shows, neither the standard model nor the alternative could simultaneously predict the relatively high volatility of consumption in the data together with a relatively low sensitivity of consumption to income growth. This may be indicative of measurement error in consumption, which increases the former but leaves the latter unaffected. In this section we thus generalise our model by assuming that measured consumption, as well as

Table 9: Risk sharing moments for heterogeneous preferences

	Aurepalle		**Kanzara**		**Shirapur**	
	Data	**DLC model, individual deviations**	**Data**	**DLC model, individual deviations**	**Data**	**DLC model, individual deviations**
n		34.00		36.00		30.00
δ		0.90		0.90		0.90
σ		1.50		1.50		1.50
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.10	0.56	0.18	0.33	0.11
β_{dcdy}	0.21	0.15	0.22	0.24	0.17	0.16
$\frac{Var_{dc\vert dy>0}-Var_{dc\vert dy\leq 0}}{Var_{dy}}$	-0.08	0.18	-0.25	0.32	-0.16	0.18
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.41	0.26	-0.14	0.40	-0.05	0.27

Notes: For each of the three ICRISAT villages, the table shows the four key moments summarising the joint distribution of consumption and income growth in the survey ("Data", in the first column for each village), and in simulations of the standard dynamic limited commitment model with individual deviations (in the second column for each village) and a simple form of preference heterogeneity, where the village population comprises two groups of equal size whose risk aversion σ equals 0.5 and 1.5 respectively.

income, include measurement error given by

$$\widehat{c_{it}} = c_{it} + \xi_{it} \tag{11}$$

$$\widehat{y_{it}} = y_{it} + \nu_{it} \tag{12}$$

where $\widehat{x_{it}}$ and x_{it} are the measured and true levels of variable x, and ξ_{it} and ν_{it} denote measurement error in consumption and income that is identically and independently distributed across individuals and time with variance Var_{ξ} and Var_{ν}.

Note that, for a given measured variance and autocovariance of income in the data $Var_{\widehat{y}}$ and $Cov_{\widehat{y}}$, the persistence parameter ρ and variance of 'true' income shocks Var_{ϵ}, which are both an input to the model, are now a function of Var_{ν}:

$$\rho = \frac{Cov_{\widehat{y}}}{Var_{\widehat{y}} - Var_{\nu}} \tag{13}$$

$$Var_{\epsilon} = (Var_{\widehat{y}} - Var_{\nu}) * (1 - \rho^2). \tag{14}$$

This necessity of solving the model afresh for all parameters, villages, and both models when Var_ν takes a new value constrains us to a small number of values for Var_ν. We therefore constrain measurement error by allowing the variance of measured income growth to be at most three times that of true income y_{it}.

Also note that, for a given measured variance and autocovariance of incomes $Var_{\widehat{y}}$ and $Cov_{\widehat{y}}$, an increase in measurement error increases the persistence of income shocks ρ. We would thus expect measurement error in incomes to have two (potentially opposing) effects: first, for a given degree of income risk sharing, it attenuates the measured regression coefficient β_{dcdy}. Second, by increasing income persistence and thus the dispersion of autarky values, we would expect it to reduce the degree to which 'true' income risk is shared.

Table 10 presents the results. Estimates of measurement error in both consumption and income are substantial in almost all versions of both models, and strongly improves the fit of the model with individual deviations. Specifically, measurement error in consumption of about 90 percent in the ID model, and of between 20 and 50 percent in the alternative with coalitional deviations, allows both models to fit the relative variance of consumption and income $\frac{Var_{dc}}{Var_{dy}}$ almost perfectly.[41] The improved fit in the standard model comes at the cost of reducing its empirical content: since consumption measurement error leaves all other moments unchanged, it effectively removes the relative variance of consumption and income growth from the objective function. Relieved from the objective of generating volatile consumption paths, the standard model predicts, with a higher estimated discount factor δ and income measurement error of 30 to 50 percent, insurance that is perfect, or close to perfect, in Aurepalle and Shirapur. In Kanzara, in contrast, the estimates of the standard model are unchanged and income measurement error is estimated to be absent.

In the model with coalitional deviations, measurement error in incomes is on average larger, accounting for between 30 and 70 percent of the measured variance in income growth. As discussed before, this reduces the volatility of income shocks, but increases their persistence,

[41] We constrain consumption measurement error to lie on a grid that takes 31 values, hence the small difference between the data and model moments.

Table 10: Preferences estimated to target all 4 moments and measurement error

	Aurepalle			Kanzara			Shirapur		
	Data	DLC with deviations by: individuals	DLC with deviations by: coalitions	Data	DLC with deviations by: individuals	DLC with deviations by: coalitions	Data	DLC with deviations by: individuals	DLC with deviations by: coalitions
n		34.00	2.00		37.00	2.00		31.00	4.00
δ		0.96	0.96		0.94	0.96		0.94	0.98
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.31	0.28	0.56	0.57	0.57	0.33	0.35	0.34
β_{dcdy}	0.21	0.00	0.21	0.22	0.09	0.22	0.17	0.03	0.16
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	-0.00	0.01	-0.25	0.09	0.01	-0.16	0.02	0.00
$\beta_{dcdy>0}-\beta_{dcdy\leq 0}$	-0.41	0.01	-0.00	-0.14	0.17	-0.01	-0.05	0.06	0.01
$\frac{Var_{d\epsilon}}{Var_{dc}}$		0.94	0.19		0.88	0.62		0.91	0.47
$\frac{Var_{d\nu}}{Var_{dy}}$		0.47	0.67		0.00	0.67		0.27	0.27
Goodness of fit		21.67	11.12		8.89	2.02		11.19	4.06

Notes: For each of the three ICRISAT villages, the table shows four key moments summarising the joint distribution of consumption and income growth in the survey ("Data", in the first column for each village), and in simulations of the two versions of the dynamic limited commitment model ("DLC"): the standard model with individual deviations (in the second column for each village), and the alternative model with coalitional deviations (in the third column for each village). For the simulated model solutions, the table also presents the size of the insurance groups n and the discount factor δ, and measurement error in consumption and income $\frac{Var_{d\epsilon}}{Var_{dc}}$ and $\frac{Var_{d\nu}}{Var_{dy}}$. δ, $\frac{Var_{d\epsilon}}{Var_{dc}}$ and $\frac{Var_{d\nu}}{Var_{dy}}$ are chosen to minimise the sum of differences between all four moments predicted by the models and those observed in the data, weighted by the inverse variance of the data moments calculated using a bootstrap procedure. The estimated parameters are those that minimise this criterion on a grid of $\delta \in [0.5, 0.99]$. The goodness of fit reported is the value of the criterion function at the chosen parameters.

reducing the benefit of insurance and the maximum sustainable group size. This can be seen in Shirapur, where the estimated discount factor is unchanged but the size of insurance groups is reduced from 5 to 4, and attenuation bias brings the regression coefficient β_{dcdy} closer to the data. In Aurepalle and Kanzara, the estimate of the discount factor is reduced from 0.98 to 0.96. This adds to the effect of strong estimated income measurement error, and makes only pairs of households sustainable. Interestingly, in Kanzara, the resulting reduction in insurance, and the corresponding rise in the regression coefficient β_{dcdy}, is more than offset by the attenuation bias of measurement error.

In sum, measurement error in consumption and income allows both the standard model, and its alternative specification with coalitional deviations, to better match the data. In the standard model, however, this comes at the price of predicting (almost) perfect insurance (or in the case of Kanzara unchanged and strong insurance): measurement error is estimated to account for 90 percent of consumption variations which are almost insensitive to income movements. While that latter prediction is counterfactual, it allows the standard model to generate more symmetry than it could achieve without measurement error.[42] In the version of the model with coalitional deviations, in contrast, modest measurement error in consumption, and stronger error in incomes, allow the model to match the degree of insurance (almost) perfectly, without affecting the predicted symmetry of the joint consumption and income distribution.

A.4 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

[42]Note that once we introduce measurement error in incomes, there is another way in which the standard model could achieve a better fit: start from a prediction of autarky, which implies symmetry, and then – by use of measurement error in income – attenuate the sensitivity of consumption to income to realistic levels. Given the average degree of risk sharing in the ICRISAT data of 0.2, this would however require the variance of measured income growth to be five times that of true income, which we deem unsatisfactorily large.

regression on time dummies. As we discussed in Section 5.5, this procedure has somewhat different effects 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.

Figure 5: Consumption and income growth in general equilibrium

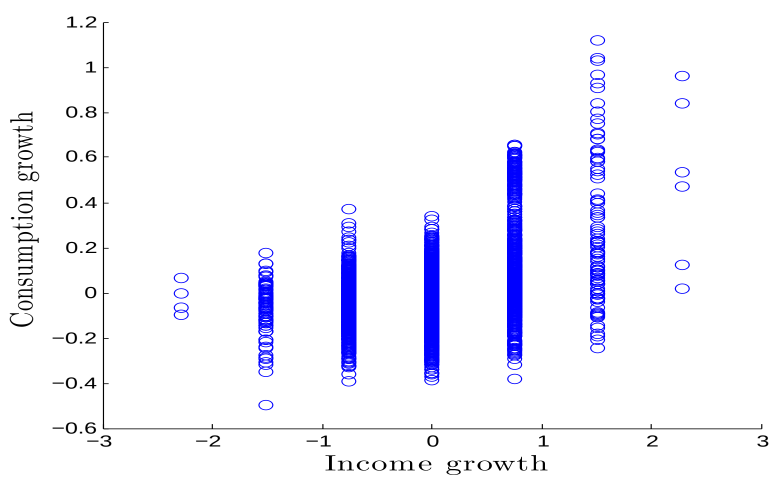

Notes: The figure shows a scatter plot of consumption and income growth from a simulation of the standard general equilibrium version of the limited commitment economy using the simulated data without conditioning on income growth.

Figure 5 and 6 correspond to figures 1 and 2 in the main paper, but use raw data instead of residuals from a first-stage regression on time dummies. The figures are very similar. The conditional variance of consumption for negative income growth is stronger in Figure 5, as consumption movements now include those proportional to aggregate village income.[43] As expected, the asymmetry in the conditional mean and variance, as well as the increase in the

[43] Also, the discrete nature of aggregate income is more visible in Figure 6 as raw aggregate income lies on a grid, but residuals do not.

Table 11: Preferences estimated to target all 4 moments (unconditional)

	Aurepalle			**Kanzara**			**Shirapur**		
	Data	**DLC with deviations by: individuals**	**coalitions**	**Data**	**DLC with deviations by: individuals**	**coalitions**	**Data**	**DLC with deviations by: individuals**	**coalitions**
n		34.00	5.00		37.00	6.00		31.00	5.00
δ		0.94	0.98		0.94	0.99		0.94	0.98
$\frac{Var_{dc}}{Var_{dy}}$	0.33	0.03	0.23	0.61	0.07	0.18	0.37	0.03	0.21
β_{dcdy}	0.22	0.03	0.20	0.18	0.12	0.18	0.15	0.04	0.21
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy\leq 0}}$	-0.01	0.00	0.02	-0.17	0.26	0.02	-0.14	0.00	0.00
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.39	0.01	0.02	-0.02	0.23	0.03	0.01	0.02	0.04
Goodness of fit		49.41	17.72		15.09	5.84		18.07	5.46

Notes: For each of the three ICRISAT villages, the table shows four key moments summarising the joint distribution of consumption and income growth in the survey ("Data", in the first column for each village), and in simulations of the two versions of the dynamic limited commitment model ("DLC"): the standard model with individual deviations (in the second column for each village), and the alternative model with coalitional deviations (in the third column for each village). For the simulated model solutions, the table also presents the size of the insurance groups n and the discount factor δ, which is chosen to minimise the sum of differences between all four moments predicted by the models 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 $\delta \in [0.5, 0.99]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

Figure 6: Key moments in village economies of different size

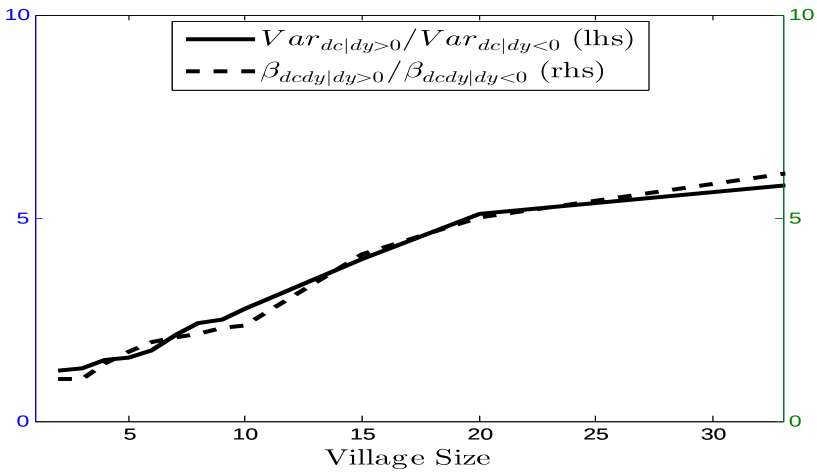

Notes: The figure plots $\frac{Var_{dc|dy>0}}{Var_{dc|dy\leq 0}}$ (the relative variance of consumption of households that experience positive income growth divided by that of those that do not, as a solid line) and $\frac{\beta_{dcdy|dy>0}}{\beta_{dcdy|dy\leq 0}}$ (the ratio of regression coefficients of consumption growth on income growth for households with rising and non-rising income, as a dashed line) from simulations of the model with increasing village size (reported along the bottom axis), based on the simulated data without conditioning on income growth..

moments as village size increases in Figure 2, are thus somewhat less pronounced but still strong.

In Table 11 we repeat our benchmark estimation based on 'raw' data, rather than residuals from a regression on time dummies. This makes little difference to the degree of insurance measured in the three villages, but reduces the negative asymmetry observed in all three villages. Since the model with coalitional deviations predicts symmetry already, 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. An exception to this is Kanzara, where the estimated group size is increased to 6, which increases the degree of insurance.

More interestingly, the performance of the standard model is equally changed very little: while the discount factor is estimated to be somewhat higher than in the benchmark estimation, risk sharing is predicted to be even stronger than with conditioning. In other words, even in the raw data, idiosyncratic movements in individual consumption and income are sufficiently important that the standard model cannot at the same time predict symmetry and a realistic degree of insurance.

A.5 Comparing endogenous and exogenous group size

In Figure 7, we plot the ratio of the β coefficients in the coalitional deviations model and the individual deviations model across a wide range of preference parameters for group sizes for which the coalitional deviations model is stable. Since we are interested in cases that deliver realistic degrees of partial insurance, we 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 perfect, i.e. where there is idiosyncratic variation in consumption even after controlling for group resources.

The figure shows how the standard model always predicts stronger insurance, and increasingly so at larger group sizes. In fact, the degree of insurance is 30% lower in the coalitional deviations model for groups of four and 45% lower for groups of 5, very similar to the relative degree of insurance comparing the fitted models in the main paper. Relating the ratio of β coefficients to group size in a linear regression, and plotting the regression line in Figure 7, we find that for every additional member of an insurance group, the ratio of the two coefficients increases by 15%.

A.6 Comparing the exact model and approximation

Finally, we examine how solving the model with the standard rest-of-the village approximation rather than solving it exactly impacts the result and in particular to what extent it differentially affects the standard model versus the alternative.

First, the approximation is exact for two households and gets worse as group size increases. Second, the approximation and the exact solution are identical for full insurance as differences between the two arise only when enforcement constraints bind and, by implication, the approx-

Figure 7: The degree of insurance is smaller in the coalitional deviations model even for given group size

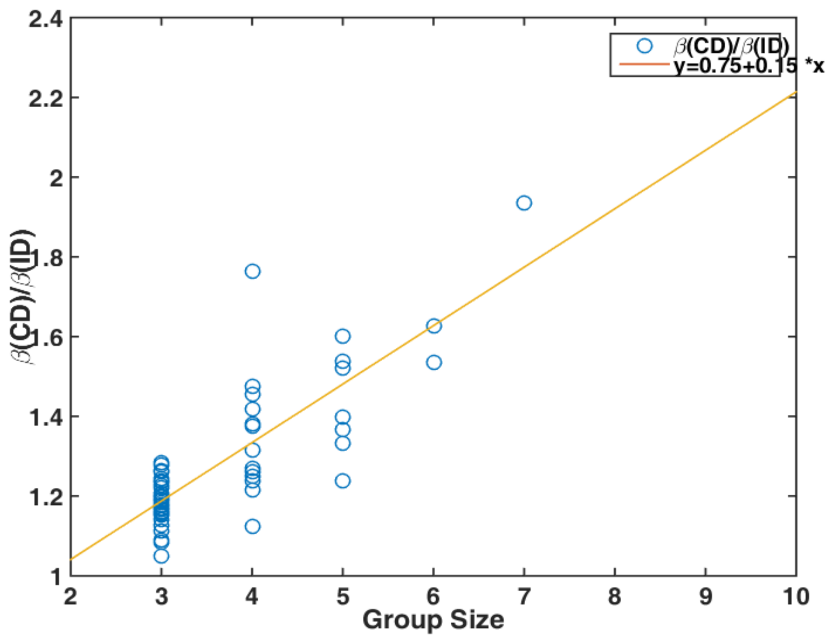

Notes: The figure shows the ratio of the degree of insurance in the two models for all parameter combinations where the coalitional deviations model is stable, the degree of risk sharing lies between 0.1 and 0.35 and risk sharing is not perfect at the group level. Each point corresponds to a particular combination of patience, risk aversion and group size. The line is the prediction from a linear model regressing the ratio on group size and a constant.

Table 12: Comparing the exact and approximate model solutions for $n = 3$

	δ=0.6	δ=0.7	δ=0.8	δ=0.9	δ=0.95
DLC with individual deviations:					
$\beta_{dcdy}/\widehat{\beta_{dcdy}}$	1.11	1.07	1.04	1.06	1.01
$\beta_{dcdy>0} - \beta_{dcdy\leq 0} - (\widehat{\beta_{dcdy>0} - \beta}_{dcdy\leq 0})$	-0.07	-0.04	-0.02	0.00	0.00
DLC with coalitional deviations:					
$\beta_{dcdy}/\widehat{\beta_{dcdy}}$					1.04
$\beta_{dcdy>0} - \beta_{dcdy\leq 0} - (\widehat{\beta_{dcdy>0} - \beta}_{dcdy\leq 0})$					-0.00

The table shows two summary measures for the difference between the exact solution of the dynamic limited commitment model with $n = 3$ agents, and its standard approximation: first, the ratio of the two resulting coefficients in a bivariate regression of consumption growth on income growth $\beta_{dcdy}/\widehat{\beta_{dcdy}}$; and second, the difference in the resulting relative regression coefficients for households with increasing and non-increasing incomes $\beta_{dcdy>0} - \beta_{dcdy\leq 0} - (\widehat{\beta_{dcdy>0} - \beta}_{dcdy\leq 0})$. For the coalitional deviations model (bottom rows), we only show the results for preference parameters where the $n = 3$ group is stable.

imation error becomes larger the smaller the degree of insurance.

Together this implies that the model with coalitional deviations is closer to the exact solution than the standard model: (i) because the model predicts smaller group sizes, and (ii) within these smaller groups, the degree of insurance tends to be quite high. There is an important caveat to this argument, however, namely that it holds only if the set of group sizes stable with respect to coalitional deviations is not significantly affected by approximating the solution. While we cannot test this, the fact that equilibrium group sizes tend to be small has also been observed in other – stationary – contexts (Bold and Dercon, 2014; Fitzsimons, Malde and Vera-Hernandez, 2015).

When the approximation departs from the exact solution, this is because of two sources of error, which potentially counteract each other: (i) the approximation assigns everyone in the rest of the village equal income during the period of deviation and therefore calculates the Lagrange multiplier on the incentive constraint at the average income realisation. The true measure of how constrained the risk sharing group is, is however the average of the Lagrange multipliers across the income distribution in the rest of the village. Given the convexity of the Lagrange multiplier, the multiplier at the average income realisation will in general be smaller than the average of the multipliers across the idiosyncratic income realisations. Hence, for given preferences the approximation will tend to predict more risk sharing than the exact solution and this difference increases with group size. (ii) On the other hand, following deviation the $n-1$ individuals who break away from the group continue in autarky or in the next smaller stable group in the exact solution, while they continue to share risk perfectly in the approximation. Without persistence in incomes, this would make the outside option strictly better in the approximation and hence the degree of risk sharing smaller. With persistence, however, autarky with high income or continuing in a group with others with high income may dominate equal sharing in a larger group.

How well the approximation tracks the exact solution will therefore be a function of the estimated group size and of the preferences of the group's members. When patience is low, the instantaneous effect of the approximation error that leads to an overestimation of the degree of risk sharing will dominate. When patience is high, the long-term effect – which is ambiguous with persistence – will dominate. Of course, when patience is high, any solution will be closer to first-best which in itself serves to reduce the approximation error.

Since the standard dynamic limited commitment model and the alternative make different predictions about group size, preferences and the degree of risk sharing within a given group, we now examine if they are differentially affected by approximation error, and – more importantly – discuss whether this could bias our qualitative conclusions about which model is superior. To this end, we compare the exact and the approximate solution for small group sizes $n = 3$. Note that, while Bold (2009) performs a similar analysis of the model without income persistence, we are, to our knowledge, the first to study the exact solution of the three agent-limited commitment model with an estimated, persistent income process and compare it to data from actual village economies.

In Table 12, we calculate the approximate and the exact solution of the standard dynamic limited commitment model and its alternative with coalitional deviations following Bold (2009). We solve the models with log utility and discount factor $\delta = \{0.6, 0.7, 0.8, 0.9, 0.95\}$, an interval which contains both full and intermediate degrees of insurance. The table compares simulated moments based on the exact and the approximate solution, namely the sensitivity of consumption with respect to income changes and its asymmetry in the sub-samples of income winners and losers.

Four things emerge from this exercise. (1) For small group sizes, the approximation and the exact solution lie extremely close together. (2) The difference between approximation and exact solution converges towards 0, as the degree of risk sharing increases towards full insurance. (3) As expected, for low values of the discount factor the instantaneous effect of the approximation error dominates and the approximation therefore predicts more risk sharing than the exact solution (indicated by ratios > 1 in the table). (4) The asymmetry in the approximation is larger than in the exact solution for low values of the discount factor δ, but the difference between the two is always small relative to the level (falling from 0.07 - equivalent to roughly 10 percent of the difference in regression coefficients in the exact solution $\beta_{dcdy>0} - \beta_{dcdy\leq 0}$ - to 0). The stronger asymmetry in the approximation arises because stronger insurance in the approximation implies a rise along the left-hand side of the inverse U-shaped relationship that links the degree of insurance to asymmetry (which equals 0 for both autarky and full insurance but is positive between the two).

Taken together, this implies that the approximation works well for small groups and when the degree of risk sharing is high. Since the average group size is low in the model with coalitional deviations and larger groups tend to exist only when risk sharing is high, this implies that the alternative model is hardly affected by approximation error and the estimated preferences should therefore be similar to the exact solution.[44]

Group sizes in the standard model are set equal to village-level sample sizes in the ICRISAT panel. In line with the data, the model predicts moderate to large degrees of risk sharing. This implies that the model's solution lies in a region where the approximation error will tend to overpredict the degree of risk sharing. As seen in Table 6, this is not an issue per se, however, since the model can match the average degree of risk sharing in the data without a problem – the only caveat being that the estimated discount factor is probably slightly lower than in the

[44]The implicit assumption here is that the set of stable sizes itself is not affected by the approximation and the fact that we assign equal sharing to deviating groups.

Table 13: 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, 8 Indian rupees were worth about 1 US dollar, which is about 4.60 dollars in 2016 (see Laczo (2014) for calculations).

exact solution.

A.7 Descriptive statistics

Table 13 presents descriptive statistics from the ICRISAT data.