Risk Sharing in Village Economies Revisited

Tessa Bold, IIES, Stockholm University and CEPR

Tobias Broer, IIES, Stockholm University and CEPR*

December 2018

Abstract

We propose to replace the common notion of 'village' risk sharing by 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: first, the largest renegotiation-proof groups tend be substantially smaller than typical villages. Second, with strong insurance in small groups, individual consumption responds symmetrically to income rises and falls, while alternative models predict strong counterfactual asymmetry.

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, Swapnil Singh, Christian Stoltenberg, Kjetil Storesletten and Nicholas Barton, as well as the editor and five referees, 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, a symmetric reaction of consumption to positive and negative income shocks. Importantly, this is not trivial, because other popular models of risk sharing predict counterfactual asymmetries in consumption-income comovement, and are silent about group size. We therefore think that our results argue in favour of a model with endogenous insurance groups in the context of poor agricultural communities. 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 the resulting requirement of 'coalition-proofness' is particularly appealing in the context of village economies, where it seems difficult to prevent those 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 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 limited commitment 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.[1]

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 whose size is in line with the single-digit groups documented in other datasets, and with the evidence in the ICRISAT data (which, however, lacks explicit information on groups).[2] And second, it captures the approximate 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 that in the standard version of the model, without coalitional deviations, the asymmetry implied

[1]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.

[2]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).

by the limited commitment constraint increases quickly with the exogenous group size. And at the typical village sizes considered in the literature, it is infact substantially larger than in the ICRISAT data. In contrast, the model with coalitional deviations predicts negligible asymmetry. This is substantially closer to the ICRISAT data, where the asymmetry is negative but for the most part small and insignificant.

To show that these results do not depend on the particular environment we consider, we show that they are robust to alternative assumptions about outside options, about the income process, etc. We also show that, when average insurance is strong, as observed in our data, the counterfactual asymmetry in the standard model is little changed by including a stylised form of heterogeneity in preferences. Moreover, measurement error in income and consumption would have to be unrealistically large to attenuate this asymmetry sufficiently as to bring its predictions in line with the data from the ICRISAT villages. Finally, we show that the standard computational approximation of the equilibrium that we use (Ligon, Thomas and Worrall, 2002) produces results very similar to those from an exact computation. The largest groups for which we can compute the risk-sharing contract exactly is, with 4 households, close to the sizes we estimate using the approximation. In addition, we show that the largest renegotiation-proof groups are essentially identical as in our benchmark results (where insurance is strong but not perfect) when computing them exactly under the maintained assumption of full insurance (where we can consider any group size).

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 because we would expect the effect of policy interventions such as poverty reduction or income insurance to change substantially once one allows for coalitional deviations from risk sharing. In more standard environments, where limited commitment arises from the possibility of individual deviations only, it has been shown that policies to reduce income risk may be counterproductive as they reduce the punishment of exclusion from insurance, thus making the outside option to the contract more attractive and reducing the transfers households are willing to make to others. This can potentially crowd out private consumption insurance (Attanasio and Rios-Rull, 2000; Krueger and Perri, 2011).[3] With endogenous group sizes, in contrast, crowding out works through a completely new channel by affecting the sustainable group size.

Our focus in this study is on the limited commitment friction. This is, first, because formal institutions for contract enforcement are typically absent in poor agricultural villages, and, for example, information frictions are often viewed as less important. 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, it is the limited commitment friction that allows us to analyze endogenous group size, since without it group sizes are usually not endogeneously determined.

Importantly, our estimate of small single-digit group sizes depends on the observed income and consumption characteristics in the ICRISAT data. Other contexts may imply much larger insurance groups. And even for the ICRISAT data, it does not contradict the finding of insurance within (larger) 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 smaller endogenous groups that feature strong risk sharing and, for reasons we do not consider here, may not typically cross caste or kinship barriers. We think that the analysis of additional frictions, alternative environments and datasets, as well as of additional barriers to the formation of insurance groups should be fruitful areas for future research.[4]

The next section introduces the dynamic limited commitment model with coalitional deviations, and describes in detail our quantitative approximation. Section 3 describes the ICRISAT data and estimates the strength of insurance and its asymmetry in these villages. Section 4

[3]Broer (2013) shows, however, that this is less likely in the context of developed countries, where individuals have more assets they can pledge under the contract, and income shocks are typically more persistent.

[4]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.

presents the estimation results for the CD model, and compares them to those for the standard limited commitment model with individual deviations, and a simple model of self insurance. Section 5 provides additional evidence for small insurance groups in the ICRISAT data beyond that used in the estimation. It also compares the predictions of the coalitional deviations model to those of the standard limited commitment model when group sizes are small in both models. An online appendix contains numerous robustness checks and additional analysis.

2 Consumption insurance with limited commitment

This section describes a dynamic model of co-insurance under limited commitment with endogenous group formation. We also present two alternatives, a model with exogenous group formation and, in Section 2.6, a model where agents self-insure via savings.

The setting for mutual insurance is a standard economy where risk-averse households face idiosyncratic income risk but cannot commit to making the transfers implied by risk sharing. Consumption insurance is thus restricted by (ex-post) participation constraints: the utility value of continued participation in an insurance scheme must not be less than those of households' outside option in any state of the world.

2.1 A limited commitment village economy

We consider a village community with N households. In each period $t = 1, 2, ...\infty$, household i receives an endowment of the only consumption good $y^i(s_t)$, where s_t is the state of nature in period t drawn from a finite set $\mathcal{S} = \{1, ..., S\}$. 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 realizations and those of other villagers, but are not able to write binding contracts.

An insurance contract for a group of $n \leq N$ households in this environment 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, h_t, consisting of the previous states and transfers. Since the environment does not allow for legally binding agreements, an equilibrium insurance contract must be self-enforcing, which requires that in any state of nature $r \in \mathcal{S}$, the expected discounted utility implied by the contract for any household i must not be smaller than that of an outside option V_r^i

$$U_r^i \geq V_r^i, \forall i, r \in \mathcal{S} \tag{1}$$

Insurance transfers from households with high income realizations to those with low income can be sustained in such a context whenever reneging on the contract is costly, implying that the instantaneous benefit from doing so is traded off against a lower continuation value under the outside option V_r^i, for example because it implies the loss of future insurance possibilities.

To find the constrained-optimal insurance contract for a general outside option V_r^i, 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

$$\delta \gamma^i \pi_{sr} \phi_r^i \quad : \quad U_r^i \geq V_r^i \quad \forall i, r \in \mathcal{S} \tag{4}$$

and an aggregate resource constraint in each state and period.

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

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

The first-order and envelope conditions associated with this problem imply the following optimality condition that links the evolution of household i's consumption between period t and state r in period $t+1$ to that of a reference household n.

$$\gamma_r^i = \frac{u'(c_r^n)}{u'(c_r^i)} = \frac{1+\phi_r^i}{1+\phi_r^n}\gamma^i = \frac{1+\phi_r^i}{1+\phi_r^n}\frac{u'(c^n)}{u'(c^i)} \quad \forall r \in \mathcal{S}, \quad \forall i \neq n. \tag{6}$$

Equation (6) captures the essence of consumption insurance with limited commitment to contracts: relative marginal utility is constant, and insurance thus perfect, unless participation constraints bind, implying a strictly positive Lagrange multiplier $\phi_r^k, > 0$ for $k = i, n$. The model thus, in general, implies partial insurance with both perfect insurance and autarky as two limiting cases.

Inherent in the model is an asymmetry for the consumption process that is most easily illustrated with log-preferences, where the relative consumption of any two households $\frac{c_r^i}{c_r^j}$ 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 $Y_r = \sum_{i=1}^n c_r^i$ we can express household j's consumption as a function of village income Y_r and Lagrange multipliers: $c_r^j = Y_r \frac{\gamma^j(1+\phi_r^j)}{\sum_i \gamma^i(1+\phi_r^i)}$. Taking log-differences of both sides yields

$$(7) \qquad d\log(c_t^j) \;=\; d\log(Y_t) + \log(1+\phi_t^j) - \log\left(1 + \frac{\sum_{i=1}^n \gamma^i \phi_t^i}{\sum_{i=1}^n \gamma^i}\right).$$

where $d\log$ denotes the log difference and we suppress the dependence on state r in period t. 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 (7) illustrates the asymmetry inherent in risk sharing in a limited commitment environment: the consumption share of household i increases only when its participation constraint binds. Moreover, for a given vector of outside options of other villagers, its consumption share is increasing in her outside option V_r^i. Whenever the participation constraint is slack, the household shares the same decline in marginal utility with other unconstrained households, where the magnitude of the decline is independent of its outside option.

2.2 Implications of limited commitment for the joint distribution of consumption and income

Equation (7) does not typically lend itself to a direct test of the model because survey data are silent about membership in insurance groups within a village. Without information on group income, and about which households in the group are constrained, any test has to, essentially, rely on data about individual consumption, and its joint distribution with individual, and perhaps village, income. Typically, researchers study the joint distribution of the growth rates of individual consumption and income, whose average comovement has extensively been used as an intuitive measure of risk sharing. The fact that the value of the outside option V_r^i is, for many specifications, increasing in current income (as long as incomes are not too negatively correlated) suggests that this joint distribution should also capture the asymmetry suggested by equation (7). Specifically, consumption should be expected to respond more strongly to income increases, which are likely to tighten participations constraints, than to income decreases, which relax them. This has led previous studies (Krueger and Perri, 2005; Broer, 2013) to look at non-linearities in the conditional mean and variance functions of consumption growth around zero income growth.

To illustrate the implications of limited commitment for the joint distribution of consumption and income, we solve a simplified version of the limited commitment model presented in Ligon, Thomas and Worrall (2002), where the insurance group coincides with the village and the outside option V_r^i is specified as eternal individual autarky, such that $c_s^i = y_s^i$ for a deviating household in all states and periods after a deviation. Figure 1 illustrates the joint distribution of consumption and income by depicting a scatter plot of income and consumption growth (or their log-differences) based on a long simulation of the model when using the parameter estimates corresponding to the village of Aurepalle in the ICRISAT data set.[5] The figure illustrates how the asymmetry suggested by equation (7) translates to non-linearities in the joint distribution of income and consumption in this standard limited commitment economy. Consistent with the intuition suggested by equation (7), there is a pronounced kink in both the conditional mean and variance functions around 0 income growth.

Importantly, it can be seen from equation (7) that the asymmetry implied by limited commitment insurance and illustrated in Figure 1 depends on the degree of insurance, and the size of the insurance group. With full (or no) insurance, households are never (or always) constrained and consumption growth thus equals aggregate (individual) income growth, and is therefore

[5]Specifically, we use preferences with constant relative risk aversion equal to 0.95 and a discount factor of 0.85, equal to their estimates when targeting the observed log-changes in consumption and incomes, and an income process estimated on data for Aurepalle after partialling out time fixed effects and controlling for household fixed effects, see Section 3 for details. We use a simplified version of the model that abstracts from heterogeneity in income processes and direct utility penalties after a deviation.

Figure 1: Consumption and income growth in general equilibrium

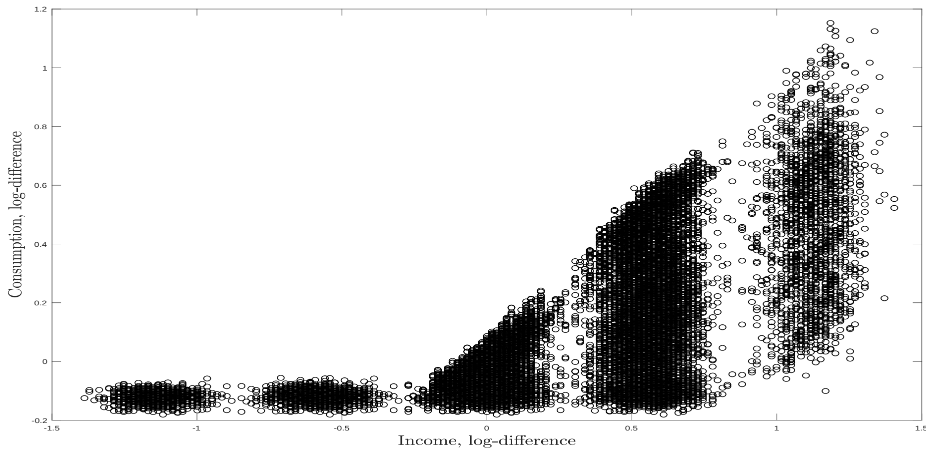

Notes: The figure shows a scatter plot of consumption and income growth (or their log-differences) for a simplified version of the model presented in Ligon, Thomas and Worrall (2002), and their parameter estimates corresponding to the village of Aurepalle in the ICRISAT data set. Specifically, we use preferences with constant relative risk aversion equal to 0.95 and a discount factor of 0.85, equal to their estimates when targeting the observed log-changes in consumption and incomes. We use a simplified version of the model that abstracts from heterogeneity in income processes and direct utility penalties after a deviation. The figure plots residuals from a regression that controls for movements in aggregate resources.

equally symmetric as incomes. Similarly, when two households partially insure each other, the budget constraint implies that consumption and income shares are complements. So consumption of the constrained and unconstrained responds symmetrically. With partial risk-sharing in a large insurance community, in contrast, the heterogeneous increases in consumption shares of constrained households are spread across many unconstrained households whose consumption shares decline by the same amount that is independent of their individual incomes and outside options, implying a more pronounced asymmetry.

2.3 Consumption-risk sharing in endogenous groups

We now present a model, first introduced by Genicot and Ray (2003), of coalition-proof dynamic risk sharing. The environment is the same as in Section 2.1: risk-averse and patient agents face a volatile income stream, which they can smooth by entering into a mutual insurance arrangement, that must be self-enforcing. Just as in the Ligon, Thomas and Worrall model estimated in Section 2.2, reneging on an insurance contract results in being excluded from the existing arrangement. The crucial difference is what happens after exclusion: rather than being barred from all smoothing possibilities, agents who have deviated can continue insuring each other. We call this model the 'CD' model (for 'coalitional deviations').

We begin by describing how stability of an insurance group of size n to such deviations can be assessed recursively. To do so, we need to define sets of stable expected payoff vectors $\mathcal{W}_s^*(m)$ for $m = 1, ..., n-1$ individuals and each state $s \in \mathcal{S}$, which contain the outside options for any subset of group members. For $m = 1$, there is just one stable payoff, the expected discounted utility of consuming volatile income forever:

$$W_s^*(1) = E_s \sum_{t=0}^{\infty} \delta^t u(y_t). \tag{8}$$

where E_s indicates that the expectation is taken over the probability distribution induced by an initial state s. Hence, in the case of only individual deviations, the set of stable payoff vectors consists of $\mathcal{W}_s^*(1) = \{W_s^*(1)\}$.

In the 'standard' limited commitment model with individual deviations (henceforth the 'ID' model), there would be only one stable payoff vector for any number of $m < n$ agents (i.e. any deviation, joint or not, implies individual autarky), consisting of the $m \times 1$ vector with each entry $i = 1, ..., m$ equal to $W_s^*(1)$. Building on the literature on equilibrium refinement in repeated games (Farrell and Maskin, 1989; Bernheim, Peleg and Whinston, 1987; Bernheim and Ray, 1989), Genicot and Ray (2003) argue that such a punishment may be too harsh to implement in groups of more than two households. Particularly in the context of risk-sharing groups within village economies, it seems difficult to prevent agents who deviate from the insurance arrangement in the current period to renegotiate their punishment of exclusion from any insurance possibilities and insure each other in the future. In other words, it seems reasonable to replace the outside option (or punishment path) of eternal autarky with continued insurance, albeit in a smaller group. For $m > 2$, the sets of stable expected payoff vectors therefore consist of all stable divisions of insurance surpluses for m players.

An insurance contract for a group of size n again consists of a vector of net transfers $\tau(s_t, h_t)$ for each state and history of the game. To find the constrained-efficient contract in a group of size n, the planner again maximizes the utility of the $n'th$ agent subject to delivering promised utilities to the other $n-1$ agents (see equation (2)) subject to a set of enforcement (or stability) constraints that embody the coalitional threats. Specifically, the insurance contract is stable if the following enforcement condition holds (see Genicot and Ray (2003), p.102):

There is no history of states and transfers h_t up to the current period, and no state r in the following period, such that for some subgroup of individuals (of size $m < n$) and some stable expected payoff vector $\mathbf{W}_r \in \mathcal{W}_r^*(m)$

$$u(y_r^i) + \delta W_r^i > U_r^i \quad \forall i \text{ in the subgroup} \tag{9}$$

where W_r^i is the i'th element of $\mathbf{W}_r$ and U_r^i is the continuation value of agent i in state r in the existing risk-sharing arrangement.

The coalition-proof outside options differ from the standard ID model in two main dimensions: first, by making deviations – to a new insurance group rather than individual autarky – more attractive, it restricts risk-sharing more strongly. Second, because the coalition-proof equilibrium risk-sharing contract may fail to exist for a given group size and the maximal sustainable size of an insurance group is known to be bounded (Genicot and Ray, 2003), the model endogenizes the size of insurance groups (or more precisely, the size of the largest stable group in a community of size N).[6] This contrasts with the ID model that has no endogenous mechanism to bound group size. Since the marginal benefit of adding additional members to the risk-sharing group is always positive (Murgai, Winters, Sadoulet and de Janvry, 2002), this has led researchers to consider as the relevant insurance group in a developing country context the

[6]The intuition for the boundedness of groups in the model with coalitional deviations is complex, but relies on the fact that the insurance benefit of increasing group size goes to zero as groups become large. If sustainable group size was infinite, one could choose a finite group such that it contains with probability 1 in every state a large enough coalition of n high income individuals that exhausts all diversification benefits. This coalition would have no incentives to make transfers, so autarky would be the only equilibrium in the larger group. See Genicot and Ray (2003), p. 94.

village or smaller exogenously bounded groups within it such as extended families. In contrast, the CD model allows us to derive sharp predictions for the equilibrium group size.

2.4 Approximating the model

To test whether the CD model is a good quantitative model for village risk sharing, we need to find the set of stable insurance group sizes and contracts in a typical village. Even the standard ID model, however, has so far been solved only for groups of up to three agents, which has led researchers to consider approximate solutions to the constrained-efficient contract. Relative to this, the CD model introduces an additional layer of complexity: in contrast to the individual deviations model, which gives each individual one threat point (conditional on the current state), a deviating subgroup can now continue to share risk starting from any division of surplusses in its set of stable expected payoff vectors. Since this set typically has infinitely many members, this leads to an infinite number of potential threats and, as shown in Bold (2009), possible strategies for the planner to deter them. In either model, an exact numerical solution for larger group sizes is out of reach.

To operationalize the CD model for quantitative analysis, we therefore adapt and extend the common approximation to the solution of the standard ID version. This approximation, 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 recasting the (simultaneous) $N-$ household insurance problem as a sequence of N two player problems in which an individual shares risk with an agent who represents the rest of the village of $N-1$ individuals, and thus has average preferences and receives an endowment equal to the average across the $N-1$ remaining villagers. The vector of outside options for both agents equals the consumption value of individual and average incomes respectively.

This approximation can be thought of as the constrained-efficient equilibrium of a simplified infinitely repeated insurance game where the planner abstracts from heterogeneity in the rest of the village when calculating the Lagrange multipliers ϕ_r^i in equation (4). Specifically, ϕ_r^i is derived in a simplified game where all other villagers $j \neq i$ are assigned a common multiplier, ϕ_r^{-i}, equivalent to pooling their income both inside the contract and during deviation. With the multipliers ϕ_r^i obtained in the simplified game in hand, the planner then solves the $N-1$ first-order conditions in equation (6) for consumption (thus never using the approximate multipliers ϕ_r^{-i}).

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. Our aim is to define an outside option of an individual i sharing risk in a group of size n that captures the idea of coalition-proofness and the dynamic nature of the contract. At the same time, we are looking for an approximation that does not require us to (1) track the entire history of shocks and transfers of the group, and (2) consider the entire stable set $\mathcal{W}_s^*(m)$ following state s as potential deviations for a group of size m and all possible strategies for the planner to deter these deviations.

To this end, we consider insurance contracts in groups of n villagers that increase in size $n = 2, 3, ...$ As in the ID model, the shadow value ϕ_r^i of a given outside option of individual i in state r is found by solving n sequential two-agent games where individual $i = 1, ..., n$ interacts with a rest of the group that pools income on and off the equilibrium path.

For given n the outside option of the rest of the group is unchanged relative to the standard approximation (so the planner abstracts from heterogeneity within the rest of the group when determining the shadow value of its outside option, ϕ_r^{-i}). For $n = 2$, the individual's outside option is autarky, as in the ID model. But for $n > 2$, an individual i who deviates now has the option to renegotiate the punishment of no-insurance with $k < n-1$ others who deviated in the same period. This renegotiation involves continued insurance in the largest sub-coalition of size $m \leq n-1$ that is stable to ulterior deviations. As in Genicot and Ray (2003), deviation is equivalent to consuming one's own income in the period of deviation. Starting from the second period after deviation, those who deviated can then enter a new constrained-efficient insurance arrangement. We assume that in this new arrangement utility is shared in the most equal fashion, as in Park, Mailath, Krueger and Cole (2018). This corresponds to equal initial Lagrange multipliers for individual i and households in the rest of the group, subject to any binding participation constraints (that may arise immediately).

For any given vector of current incomes in the group of n individuals, the relevant outside option for individual i corresponds to the sub-coalition that promises her the highest value subject to the constraint that it gives its other members at least the same value as continued insurance in the group of n. In other words individual i's relevant outside option depends not only on her own current income (as in the ID model) but also on the income distribution in the rest of the group, as it determines the set of profitable subcoalitions.

Having determined the relevant outside option of individual i, consumption and continuation values under the contract are (if possible) adjusted so that the members of the deviating subgroup are just indifferent between deviating and remaining in the n-household insurance scheme. The resulting shadow value ϕ_r^i of the individual's outside option following this state and history is recorded. If no consumption allocations and continuation values exist that deter a deviation at this given income vector, the group of size n is deemed unstable. Having considered all groups up to the size of the village N, we choose the largest such group that is stable (in each current state and for each history of past states and transfers) and compute the consumption allocation in the same fashion as in the ID model, by solving equation (6) using the sequentially determined Lagrange multipliers ϕ_r^i, for $i = 1, ..., n$.

Relative to the exact solution of the CD model, the adapted approximation reduces the potentially infinite set of deviations a coalition of size m can threaten to a single (the most equal treatment) allocation for which the planner satisfies the enforcement constraints. There remains one dimension, however, along which the CD approximation remains more complex than the standard one. In the latter, only the average income in the rest of the village matters. With coalitional deviations, however, the number of individuals with each income realization is needed to determine which coalitions can profitably deviate, implying that, rather than using a coarse discretization of rest-of-village income, the number of income states to consider for the remainder increases with group/village size. In what follows, we will use this adapted approximation method to solve for the set of stable sizes in the CD model and to compare its performance to the ID model with exogenous groups, and a simple self-insurance model.[7]

In the appendix, we examine the plausibility of the proposed model and its approximation in two ways. First, we examine the quality of the approximation in the coalitional deviations case for small group sizes by comparing it to the exact solution. Second, we consider the possibility that the model may be too complex for economic agents to solve. We therefore ask if there are simpler rules of thumb for calculating the constrained-efficient insurance contract that capture the essence of the model, namely dynamic risk-sharing and coalition-proofness, and deliver a similar fit to the data, but where solutions can be obtained in closed-form (see Deaton (1992) and Winter, Schlafmann and Rodepeter (2012) for a related discussion of rules of thumb for the optimal savings problem).

2.5 Discussion

It is important to stress that the model's focus is on group sustainability. The model is silent, however, on how sustainable groups are formed in the first place. While a formal treatment of this process is beyond the scope of this paper, we show in Bold and Broer (2018) in the context of an equal-sharing insurance contract, conditions under which the set of stable sizes identified here coincides with the absorbing states of an equilibrium process of coalition formation (EPCF), as defined in Ray and Vohra (2015). These conditions are, on the one hand, that the set of stable sizes is connected, and, on the other hand, that only subcoalitions can block an existing coalition in the EPCF. The two processes, i.e. examining stability of an insurance group and considering a dynamic process of coalition formation, then lead to equivalent outcomes because for any stable group size/absorbing state, the enforcement constraints examined by the two processes are the same, while for an unstable group size/transient state, the enforcement constraints examined

[7]For comparability, we use our adapted method also for the model with exogenous group size even though a simpler algorithm is available for the latter.

by the former are a subset of those examined by the latter. By implication, we show that any process of group coalition formation that starts with the grand coalition (i.e. the village) and allows only for internal blocking cannot come to rest at a group size that is larger than the largest stable group.

As a more concrete example of how small risk-sharing groups can emerge in an environment of explicit group formation, Ambrus, Mobius and Szeidl (2014) model a one-period risk-sharing game in a network, where agents are linked and links have an exogenously given value. The authors show that in this environment a simple model in which risk-sharing transfers between any two agents are sustained by the threat of the loss of the link between i and j should one of them cheat, is equivalent to a model in which transfers have to be coalition-proof, in the sense that any group of agents that deviate will be excluded from the larger network but can continue to share risk among each other. Even more importantly, the authors show that constrained-efficient risk-sharing arrangements in this context lead to risk-sharing islands with clearly delineated groups of agents who share risk fully among each other, but low levels of transfers across the risk-sharing islands (even when all links have the same value). We think these results are interesting, as they show (although in a more restricted context than in our model), how coalition-proof risk-sharing in groups can emerge as the outcome of a decentralized model of bilateral transfers in a network.

We also follow Genicot and Ray (2003) in assuming that households consume their income in the period of deviation. In other words, members of deviating coalitions can only start insuring each other starting in the period after deviations occur. This assumption follows the literature on infinitely repeated games where a deviation that triggers a punishment path is explicitly not part of the equilibrium strategy but defined as the action that maximizes a deviating player's one-period payoff taking the action of his opponent as given (Abreu, 1988). The introduction of renegotiation-proofness rules out equilibrium punishment paths that require players to play a Pareto-dominated equilibrium (thus allowing them to form risk-sharing coalitions, in our context), but does not change the assumptions on the out-of-equilibrium deviation from that literature (see for example Farrell and Maskin (1989)). While alternative assumptions are possible, we follow Genicot and Ray (2003) in our benchmark model, but show that our quantitative results are robust to allowing deviating coalitions to share risk already in the period of deviation.

It is also important to note that throughout the analysis, we maintain a number of simplifying assumptions. First, as in Genicot and Ray (2003), agents can form new subgroups only within an existing insurance group. We think this is reasonable as there may be many, unmodelled, reasons why (sub-)group-formation requires previous interaction. More importantly, group-formation without restriction typically causes problems for the formulation of a recursive solution concept and can lead to 'cyclical blocking chains'.[8] We therefore must for the moment rule out deviations with outsiders.

We also abstract from additional consumption-smoothing opportunities, such as saving or temporary migration. The opportunity to save improves the outside option relative to financial autarky, particularly for the rich. This would reduce insurance in the ID model. In the CD model, the effect would be more complicated: for example, more attractive individual deviations at given incomes would reduce co-insurance in groups of two, which could make groups of 3 more sustainable. Ultimately, it would be inconsistent to include individual consumption smoothing opportunities without also allowing groups to accumulate savings.

Migration of household members who expect to earn higher income in a nearby city, for example, adds additional insurance opportunities against common and idiosyncratic shocks and thus makes the risk-sharing scheme more attractive. It also makes the outside option of individual autarky more attractive, as households can use migration to insure against negative shocks, however. Morten (2016) finds the net effect on risk sharing to be negative in a standard ID model with endogenous migration. Because the degree of insurance in the small groups we estimate in our CD model is higher, and participation constraints at low income (that migration tightens most strongly) thus bind less often, we might expect the negative effect to be less pronounced. Other related issues are how migration affects the sustainable group size, and whether

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

permanent migration (which might be one reason for the attrition we observe in our data) would arise in equilibrium, and thus change group size over time. We leave an in-depth consideration of these issues to future research.

We also 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 believe that the effect of coalitional deviations is best highlighted in the most standard version of the limited commitment model.[9] Importantly, as Section A.8 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.[10]

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.

[9]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.

[10]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.

2.6 Self-insurance

In the quantitative analysis, we will compare the CD model to two other models of consumption smoothing: first, the standard ID model, and second a simple self-insurance (henceforth SI) model where, instead of engaging in mutual insurance, households build a buffer against income shocks by accumulating savings $b_t^i \geq 0$ with a village lender, or banker, remunerated at an exogenous interest rate R. Their period budget constraint in period t is thus

$$c_t^i = y_t^i + Rb_{t-1}^i - b_t^i \tag{10}$$

We think of this model more as a useful, standard comparison, rather than one that captures the particular institutions in the ICRISAT villages that our quantitative analysis focuses on.

3 The data

This section introduces the village economies that have been used most widely to study models of risk sharing: the ICRISAT panel. We describe the data and show scatter plots and key moments, motivated from the theory presented in Section 2.1 and 2.3 that allow us to determine what is a good quantitative model for village risk sharing. Importantly, since the ICRISAT data do not identify risk sharing units within the village, we, like many previous studies, focus on the joint distribution of individual consumption and income growth to evaluate risk sharing in the data. In contrast to previous work that concentrated on measures of the degree of insurance, we also study moments that capture the asymmetry suggested by the limited commitment mechanism in Section 2.1, namely the difference in comovement between consumption and income for those with income gains versus income losses. We confirm the strong degree of risk sharing found in previous studies of rural village economies. We also show how asymmetries in income and consumption comovement are negative, but small and insignificant for the most part, apart from two cases where one of the two measures we consider indicates asymmetry in the opposite direction.[11]

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. 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.[12] 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 other authors, we restrict our analysis to the years 1976-1981 and construct a fully balanced panel.[13]

The ICRISAT villages are poor with the average dweller living well below the $1 dollar a day poverty line (Table 18 in Section A.11 of the Online Appendix). On average, daily nondurable consumption per adult equivalent is 0.83 (Aurepalle), 1.10 (Kanzara) and 1.18 (Shirapur) 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 1, we report a first summary measure for the relative smoothness of consumption and

[11]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).

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

[13]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 1: 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.

Table 2: 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)***
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. 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. In both panels, consumption and income are demeaned period-by-period before the regression. Standard errors in parentheses are clustered at the household level.

income, namely 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 0.30 in Aurepalle to 0.56 in Kanzara.

In Table 2, we report the coefficient estimates for the following regression

$$d\tilde{c}_t^j = \alpha + \beta_{dcdy} d\tilde{y}_t^j + \varepsilon_t^j \tag{11}$$

where $d\tilde{c}_t^j$ and $d\tilde{y}_t^j$ denote the growth rate of adult-equivalent consumption and income, both demeaned with respect to the time dimension equivalent to a time fixed effect specification. The coefficient β_{dcdy} is a second measure of the degree of insurance, extensively studied in previous work (Townsend, 1994; Laczo, 2014), measuring which share of individual income movements passes through to consumption on average. The coefficients on the growth of adult-equivalent income imply that a 1% change in income leads to roughly a 0.2% change in consumption.[14] 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.[15]

As many previous studies, we take these two stylized facts – consumption growth only

[14] 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 2 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.

[15] Estimates without conditioning on village income are very similar (see Table 13 of Section A.6.)

Figure 2: Consumption growth and income growth in the ICRISAT dataset

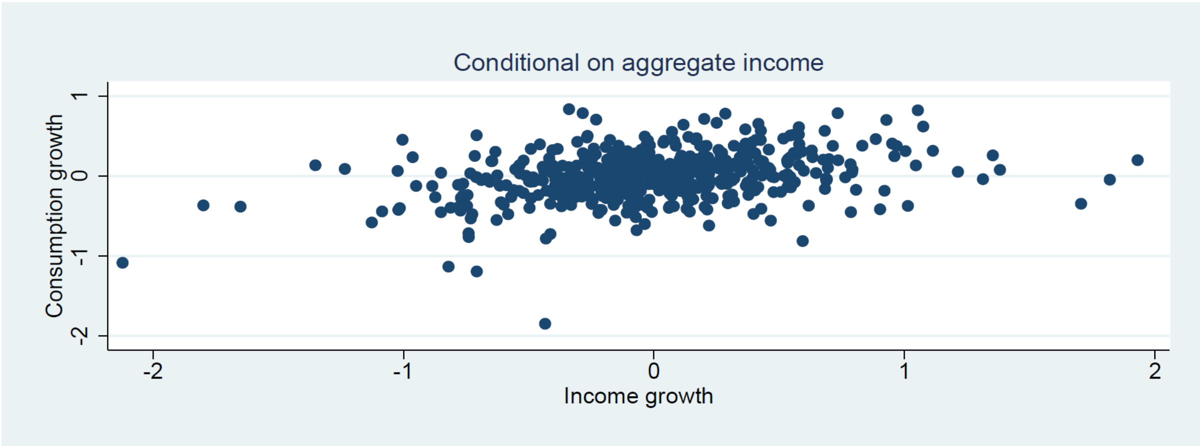

Notes: The figure shows a scatter plot of consumption and income growth for households in Aurepalle, Kanzara and Shirapur, where both measures are time-demeaned.

weakly associated with income growth, and significantly less volatile – as evidence of cross-sectional insurance. That income exceeds consumption suggests that there may be additional intertemporal mechanisms, and Lim and Townsend (1998) show in their comprehensive study of financial instruments in the ICRISAT data that accumulating crop inventory and currency, but not livestock and other real capital assets, contribute to consumption smoothing.[16] However, the authors also note that the degree of smoothing is too strong to be attributed entirely to saving and borrowing, and conclude that the ICRISAT villages "... appear to be economies in which there is nontrivial social interaction along insurance lines."

In Figure 2, we plot the joint distribution of the residual consumption and income growth (after time demeaning). The data shows little sign of the asymmetry suggested by the limited commitment mechanism in equation (7) and illustrated in the context of village-level risk sharing in the ID model in Figure 1: neither the variance of consumption nor the response of consumption to income look dramatically different as households move from negative to positive income growth.

[16] We therefore compare our quantitative results also to those of a simple self-insurance model, presented in Section 2.6.

We also test the impression of symmetry more formally using two summary moments. First, the second row of Table 1 reports the relative volatility of consumption and income growth for households with increasing vs non-increasing income. Consumption growth of households that experience positive income growth is, if anything, less volatile than that of those who experience negative or 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 4, 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.

In panel B of Table 2, we report a second measure of the asymmetry, based on the following regression that aims to capture non-linearities in the conditional mean function

$$d\tilde{c}_t^j = \alpha + \beta_{dcdy} d\tilde{y}_t^j + \beta_{dcdy|dy>0} d\tilde{y}_t^j * \tau_{dy>0} + \varepsilon_t^j \tag{12}$$

where $\tau_{dy>0}$ is a dummy variable that takes value 1 when income growth is positive, and 0 otherwise, and where we demean both consumption and income period-by-period before the regression.[17] The coefficient $\beta_{dcdy|dy>0}$ provides a second measure of the asymmetry of the joint distribution, by capturing the non-linearity in the conditional mean function of consumption growth around 0 income growth. 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 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 2).

[17] Demeaning is equivalent to including time dummies in a linear regression of consumption growth on income growth. 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.

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.

4 Results

This section discusses the quantitative implications of the CD model of limited commitment to co-insurance with coalitional deviations presented in Section 2.3 and compares them to the data moments from the ICRISAT villages. The section also contrasts the main features of the CD model with those from the two alternatives presented in Section 2, namely the standard ID model of a village-level insurance scheme with limited commitment in the form of individual deviations and the SI model of self-insurance. The results show how the CD model predicts maximal sustainable group sizes of between four and five households, substantially smaller than the ICRISAT villages and their samples analysed in Section 3. Importantly, this allows the model to predict both the correct degree of insurance and symmetric responses of consumption to positive and negative income shocks. This is not trivial, since neither the ID nor the SI model can predict a realistic degree of insurance at the same time as symmetric responses of consumption to income movements.

4.1 Quantitative model evaluation

The main aim of this section is to structurally estimate the CD model of Section 2.3 and compare its quantitative implications to data from the three ICRISAT villages, as well as to the estimated versions of the ID and SI comparison models.

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). Instead, informed by a non-parametric analysis of the models as in Figure 1, we evaluate model implications in the form of four moments of the joint distribution of individual consumption and income growth that have a close link to intuitive features of the models such as the degree of risk sharing and the degree of asymmetry in their implied reaction of consumption to positive and negative income shocks. We complement this approach, which uses only a limited amount of information contained in the joint distribution of consumption and income, with a non-parametric alternative in the form of scatter plots of consumption and income growth (as in Ligon, Thomas and Worrall (2002)).

Focusing on the joint distribution of income and consumption *growth* has the advantage of being robust to unmodelled constant sources of heterogeneity, and has intuitive appeal, as, with the low serial correlation that we find in the data, income changes are approximately equal to income shocks. A standard measure of insurance based on that distribution is the average slope of the conditional mean (the regression coefficient β_{dcdy}, as in for example Townsend 1994). In addition, we also include in our vector of moments the relative variance of consumption and income growth $\frac{Var_{dc}}{Var_{dy}}$, which, whenever the model is non-linear like ours, is an additional important summary measure for the degree of insurance.

To capture the asymmetry suggested by equation (7), Figure 1 suggests to focus on the non-linearity in the conditional mean, and the conditional variance, around 0 income growth. We therefore consider first the difference between the variance of consumption growth of households that experience positive income growth and that of those that do not, $\frac{Var_{dc|dy>0}-Var_{dc|dy\leq 0}}{Var_{dy}}$ (scaled by the variance of income growth); and second the difference in the regression coefficient of consumption growth on income growth for households with rising and non-rising income, $\beta_{dcdy|dy>0}$ and $\beta_{dcdy|dy\leq 0}$.[18]

4.2 Parameter choice

To derive quantitative predictions for CD and ID models, we need to determine the size of the insurance group, the income process and preferences. The SI model, in contrast, has no

[18]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.

insurance groups, but requires an exogenous interest rate as input.

Although all three villages in the ICRISAT dataset comprised several hundred households at the time of the survey, it is standard practice in the literature (Ligon, Thomas and Worrall, 2002; Laczo, 2014) to estimate the ID model with village size equal to the number of households sampled by the ICRISAT (presumably because the standard solution requires an estimated process for village income). This amounts to setting $N = 34$ in Aurepalle, $N = 37$ in Kanzara, and $N = 31$ in the case of Shirapur. We follow this practice for the CD model and assume that households can only form risk sharing arrangements with $N - 1$ other households where N equals the ICRISAT sample in each village. Hence, the largest potential group that can form has the same size as in the individual deviations model, but the largest stable group will typically be much smaller.[19] While there are often several stable groups of different size, we concentrate on the largest stable group size n^{max}, for three reasons. First, this is the stationary size that implies the highest insurance benefit. So if individuals could choose group sizes ex ante, this is what they would choose. Second, this focus is consistent with the practice of focusing on the maximum group size in the ID model (where it is set equal to the village / sample size). And finally, since we found in Section 2 that the asymmetry tends to rise with group size, by looking at the largest sustainable groups we raise the bar for the CD model to rationalize the approximate symmetry found in the data. Since the maximum sustainable group size is substantially smaller than the village, the latter 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.

[19] Although standard in the literature, this assumption is not ultimately satisfactory. We nevertheless follow it 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 not computationally feasible for substantially larger N. Importantly, however, allowing group size to be larger than the sample would presumably imply even more extreme values for the asymmetry in the standard model. It would only affect the results of the coalitional deviations model inasmuch as there are stable groups beyond the sample size, which is the maximum we consider.

So N' is the number of individuals in our simulation of the coalitional deviations model.[20] Since below we estimate n^{max} to be a small single-digit number, the resulting maximum difference in village size between the ID and CD models is small.[21]

We identify processes for individual incomes for each of the three villages. For this, we assume that log-incomes of all village members follow an AR(1) process with common persistence parameter ρ

$$\tilde{y}_{it} = \rho \tilde{y}_{it-1} + \epsilon_{it} \tag{13}$$

where $\tilde{y}$ is the residual from a regression of income y_i on household fixed effects, and 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 $\tilde{y}$ as

$$\begin{aligned} \rho &= \frac{Cov_y}{Var_y} \\ Var_\epsilon &= Var_y(1-\rho^2) \end{aligned} \tag{14}$$

We partial out household-specific fixed effects to make our results robust to errors in adjusting for household size, and to permanent differences in household incomes, which would otherwise be identified as persistent shocks around homogeneous mean income.[22] We also partial out aggregate movements in income by de-meaning the data.[23]

Table 3 presents the estimates for the AR(1) parameter ρ and the shock variance Var_ϵ for

[20] Consider the village of Aurepalle, which contains 34 households, as an example. Following the algorithm above, if we found the largest stable group size in this village to be 4, we would simulate the model for 9 groups with 4 members each, giving a total of 36 households in the model simulations.

[21] 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.

[22] A previous version of this paper did not include household fixed effects. This implied estimates of ρ between 0.6 and 0.8 in the three villages. The qualitative conclusions from this specification were, however, the same.

[23] Section A.6 in the Online Appendix reports the results when estimating the income process using unconditional data.

Table 3: Estimated income processes

	Aurepalle	**Kanzara**	**Shirapur**
ρ	0.28	0.00	-0.18
Var_{α_i}	0.29	0.27	0.36
Var_{ϵ}	0.15	0.064	0.11

Notes: The table presents the point estimates for the persistence parameter ρ and the shock variance Var_{ϵ} for the AR(1) process (13), as well as the variance of household fixed effects $Var(\alpha_i)$.

the three villages. Incomes have low positive persistence in Aurepalle ($\rho = 0.28$), are serially uncorrelated in Kanzara, and have small negative serial correlation in Shirapur ($\rho = -0.18$). Aurepalle has the most volatile income shocks of the three villages. Table 3 also reports the variance of fixed effects Var_{α_i}, which account for a sizeable fraction of the total variance of individual incomes. 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. The Online Appendix A.10 reports the transition matrices.

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

$$u(c_t) = \frac{c_t^{1-\sigma} - 1}{1 - \sigma} \tag{15}$$

In the absence of prior information about discount factor and risk aversion, we choose preference parameters to bring model implications as much in line with the data as possible.

For the SI model, since we estimate the discount factor, the particular value of the exogenous interest rate is not important for the results. Neither is, as it turns out, the borrowing limit. We thus choose an interest rate equal to 4 percent and set the borrowing limit to 0.[24] Estimating the discount factor δ freely in the case of the SI model implies that we do not constrain the model to deliver any particular level of wealth holdings in the stationary equilibrium (for which we have no reliable data). In other words, we give the SI model the 'best chance' to fit the data.

[24] The reason for this choice is pragmatic: the moments that we look at are simply not affected by the choice of borrowing limit.

4.3 Simulated method of moments

Given solutions to the ID and CD models, as well as individual policy functions in the SI model, we draw a vector of income realisations and then simulate individual consumption for N households in the ID and SI models and N' households in the CD model, in $T = 6,200$ periods (starting with savings equal to 0 in the case of the SI model). After discarding the first 200 periods, we then calculate our four moments from this simulated sample. 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 partialling out aggregate income movements 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 without this conditioning on village-level aggregates are very similar, and contained in the Online appendix.[25]

We 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 CD model, and the two comparison models. As noted in Laczo (2014), a necessary condition for identification in the standard model with individual deviations is that households have binding constraints in as many income states as there are parameters to estimate. But because insurance is relatively strong in the ICRISAT villages, usually at most one constraint is binding. It is therefore not possible to identify time and risk preferences separately in either the ID or the CD model.[26] Figure 7 in the Online Appendix illustrates this for the case of Aurepalle. It shows that, within the range of discount factors that are consistent with

[25]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). This is an additional reason why we also match the models to unconditional moment estimates in Section A.6 in the Online Appendix.

[26]We compute the model for three individual income states, which is the minimum needed for identification. However, increasing the number of income states does not yield identification either, simply because the targeted degree of insurance is too high.

the observed degree of insurance, 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 (log-preferences) in the estimation, and choose the discount factor that minimises the distance between the moments in the model simulation and in the data.[27]

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

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

where f is the vector of moments calculated from the ICRISAT data and $g(\delta)$ 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.[28] The variances are obtained by bootstrapping the data 1,000 times.

4.4 Model estimates

Table 4 presents the moments of interest when 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 in the limited commitment environment: 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).

[27]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.

[28]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 $\delta \in [0.5, 0.98]$. 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.

Table 4: Preferences estimated to target degree of risk sharing - 2 moments

	Aurepalle				**Kanzara**				**Shirapur**			
	Data	**CD**	**ID**	**SI**	**Data**	**CD**	**ID**	**SI**	**Data**	**CD**	**ID**	**SI**
n		4.00	34.00			4.00	37.00			5.00	31.00	
δ		0.94	0.83	0.92		0.94	0.88	0.93		0.92	0.82	0.92
s.e.		0.01	0.01	0.01		0.01	0.02	0.01		0.01	0.02	0.02
σ		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	0.24	0.23	0.18	0.56	0.24	0.15	0.14	0.33	0.18	0.13	0.12
β_{dcdy}	0.21	0.24	0.33	0.34	0.22	0.28	0.26	0.28	0.17	0.22	0.25	0.25
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.01	0.30	0.10	-0.25	0.01	0.19	0.07	-0.16	0.01	0.17	0.06
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.41	0.01	0.48	0.24	-0.14	0.03	0.39	0.23	-0.05	0.03	0.36	0.24
Goodness of fit		2.09	6.51	10.08		3.88	5.71	6.12		4.17	8.30	8.80

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 CD model (in the second column for each village), the ID model (third column) and the SI model (fourth column). For the simulated model solutions, the table also presents the size of the insurance groups n (in the case of the ID and CD models) 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 $\delta \in [0.5, 0.98]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

Table 5: Preferences estimated to target degrees of risk sharing and asymmetry

	Aurepalle				Kanzara				Shirapur			
	Data	**CD**	**ID**	**SI**	**Data**	**CD**	**ID**	**SI**	**Data**	**CD**	**ID**	**SI**
n		4.00	34.00			4.00	37.00			5.00	31.00	
δ		0.94	0.89	0.93		0.94	0.90	0.94		0.92	0.88	0.94
s.e.		0.01	0.01	0.01		0.01	0.01	0.01		0.01	0.01	0.01
σ		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	0.24	0.05	0.12	0.56	0.24	0.07	0.10	0.33	0.18	0.02	0.06
β_{dcdy}	0.21	0.24	0.12	0.26	0.22	0.28	0.17	0.23	0.17	0.22	0.07	0.15
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.01	0.07	0.07	-0.25	0.01	0.11	0.05	-0.16	0.01	0.03	0.04
$\beta_{dcdy>0}-\beta_{dcdy\leq 0}$	-0.41	0.01	0.20	0.21	-0.14	0.03	0.28	0.20	-0.05	0.03	0.14	0.17
Goodness of fit		13.61	52.76	39.01		6.52	16.80	12.62		8.96	27.78	22.97

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 CD model (in the second column for each village), the ID model (third column) and the SI model (fourth column). For the simulated model solutions, the table also presents the size of the insurance groups n (in the case of the ID and CD models) 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 $\delta \in [0.5, 0.98]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

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.

The CD model predicts maximum sustainable insurance groups that comprise 4 households in Aurepalle and Kanzara, and 5 households in Shirapur, substantially smaller than the ICRISAT sample. The larger maximum sustainable group size in Shirapur reflects the moderately negative income persistence there, which increases the probability that high-income households become transfer recipients in the future and thus reduces the incentive of sub-coalitions of high-income households to deviate from an insurance group. For estimated values of the discount factor δ between 0.92 and 0.94, the CD model predicts an average association of consumption and income growth, as measured by the regression coefficient β_{dcdy}, slightly stronger than in the data, and somewhat underpredicts the relative volatility of consumption growth (most strongly in the village of Kanzara). We interpret these estimates as a reasonably good fit of the degree of insurance. Importantly, the CD model achieves this with an approximately symmetric consumption-income growth distribution: the coefficients summarizing the asymmetry are both close to 0.

These estimates of the CD model contrast strongly with those of the two comparison models. Both the ID and SI models fit the moments associated with the degree of insurance reasonably well, but less so than the CD model. More importantly, however, both predict strong asymmetry in the conditional mean function (as measured by the difference in regression coefficients $\beta_{dcdy>0} - \beta_{dcdy\leq 0}$), and - to a lesser degree in the case of the SI model - in the conditional variance function (as indicated by their implied difference in variances $\frac{Var_{dc|dy>0} - Var_{dc|dy\leq 0}}{Var_{dy}}$). The asymmetry in the SI model is, perhaps, more surprising, since the simplest version of self-insurance, the permanent income hypothesis (PIH), would predict the change in consumption to equal the change in permanent income, and thus a symmetric reaction to income rises and falls. As explained in Krueger and Perri (2005), however, with borrowing constraints consumption responds more to income changes when asset buffers are smaller. Since negative income shocks

Figure 3: Consumption and income growth in general equilibrium

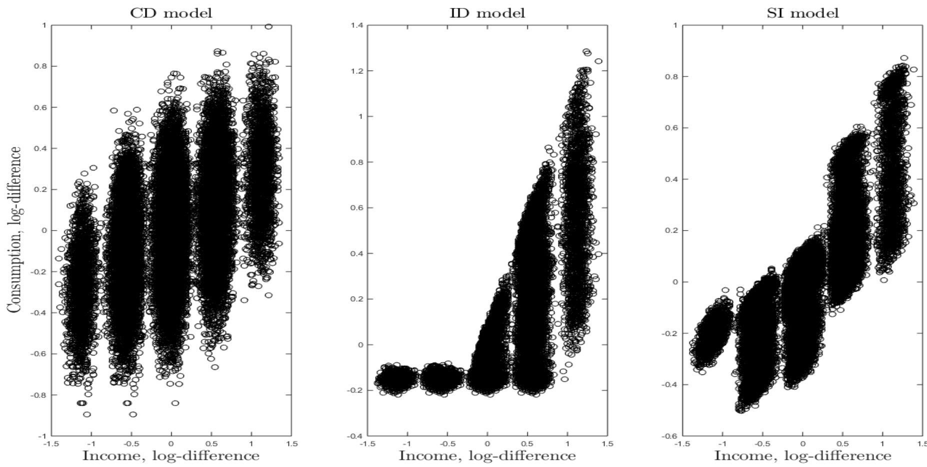

Notes: The figure shows scatter plots of consumption and income growth from a simulation of the CD, ID, and SI economies, for the income process and preferences estimated for Aurepalle in Tables 3 and 4, respectively. The figure plots residuals from a regression that controls for movements in aggregate resources.

reduce assets (as households dis-save) and make positive income growth more likely (as incomes are predicted to revert to their means), positive income shocks occur more often at low asset values, and are thus associated with larger consumption increases.

What do the estimated versions of the three models imply for the joint distributions of consumption and income? We examine this in Figure 3 with the help of scatter plots for the village of Aurepalle. The middle panel, depicting the ID model, is similar to Figure 1. Particularly, it features a similarly pronounced kink in both the conditional means and variances of consumption growth around zero income growth. The distribution generated by self-insurance, in the right panel, has an asymmetry that is somewhat smaller than that in the ID model. Particularly, the heteroscedasticity is less pronounced, in line with the small differences in variances $\frac{Var_{dc|dy>0} - Var_{dc|dy\leq 0}}{Var_{dy}}$ predicted by the SI model in Table 4. Finally, the left panel, depicting the CD model, shows a homoscedastic distribution around a linear conditional mean function. This is in line with the absence of asymmetry in its estimates in Table 4 and the joint distribution of income and consumption growth in the ICRISAT villages depicted in Figure 2 in Section 3.

Table 5 shows that, when we include the two asymmetry moments in our estimation criterion, the estimates of the CD model are unchanged. At a higher estimated discount factor δ the ID model now substantially overpredicts the degree of insurance while retaining a counterfactual asymmetry. The estimates of the SI model are also of increased discount factors and insurance, but less so than in the ID model, implying a smaller reduction in the asymmetry. In line with these results, the goodness of fit of the two comparison models, already substantially worse in Table 4, further deteriorates relative to the CD model when evaluated using all four moments of interest.[29]

4.5 Inspecting the mechanism: the role of group size

The two limited commitment models rely on very different mechanisms to arrive at their prediction of the observed degree of risk sharing. The ID model predicts moderate degrees of insurance taking place at the village level. The CD model 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 reduces the increase in asymmetry when group size increases. This role of group size and the degree of insurance for the asymmetry is illustrated in Figure 4 which varies group size exogenously in the ID model. The figure presents the second and fourth moments in Tables 4 and 5, namely the regression coefficient β_{dcdy} (the dashed line) and the difference between the regression coefficients, $\beta_{dcdy|dy>0}$ and $\beta_{dcdy|dy\leq 0}$ (the solid line) as a function of group size (along the bottom axis) and for two values of δ implying moderate (top panel) and strong (bottom panel) insurance.

As expected, the degree of insurance (as indicated by the dashed lines) is lower when agents are more impatient. For a discount factor of 0.86 (the top panel), the model matches the observed insurance in Aurepalle when the insurance group comprises the entire village (of 34 households),

[29] It is tempting to compare the corresponding measures in the final rows of Tables 4 and 5 to a χ^2 distribution. This would indicate substantially higher p-values in the CD model than for the comparison models, but is valid only under the strong assumption of independent moment conditions. Ultimately, the aim of our exercise is, however, not to reject, or not, any of the, still very stylized, models, but to highlight their different implications for the structure of consumption risk sharing and their ability to capture key moments of the data.

Figure 4: Insurance and asymmetry in the standard model

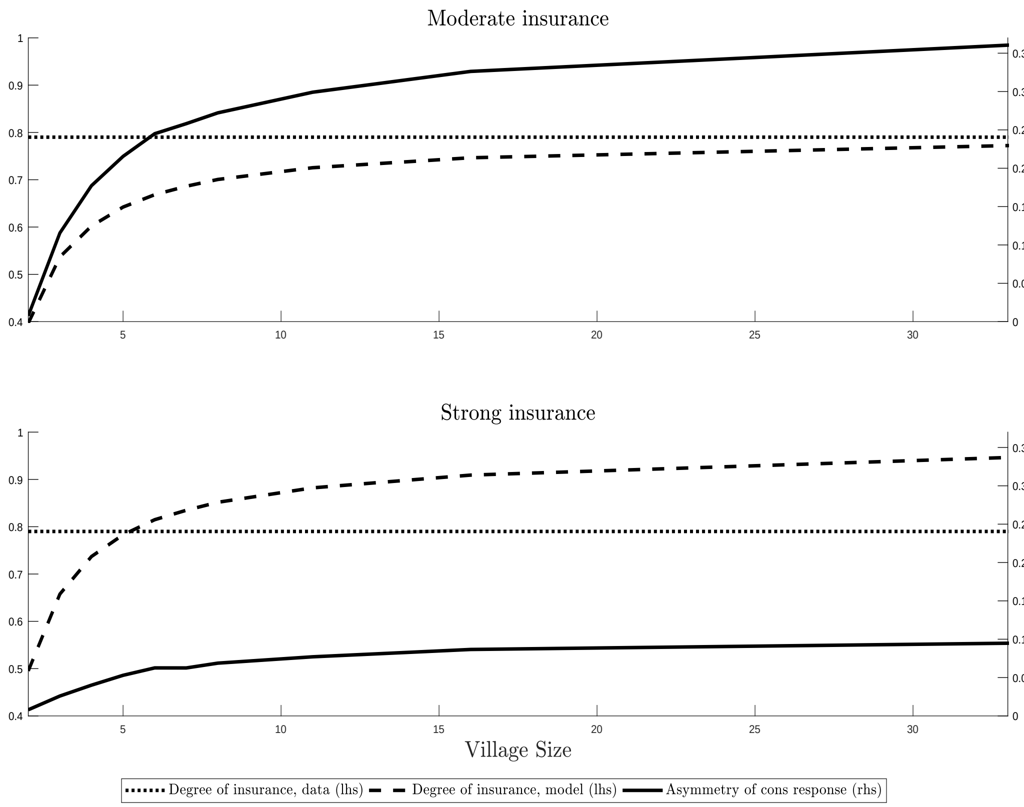

Notes: The figure presents results from a simulation of the standard ID limited commitment economy as in Ligon, Thomas and Worrall (2002), but using the whole cross-sectional distribution of incomes as a state variable to accomodate the rising group size, calculating transition probabilities through simulation of a panel of individual incomes. The figure presents two key moments that summarize the asymmetry and the degree of insurance as a function of village size (along the bottom axis) and for two values of δ implying moderate (top panel) and strong (bottom panel) insurance: first, the regression coefficient of consumption growth on income growth β_{dcdy} (the dashed line, indicated on the left axis); and second, the difference between the 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 solid line, indicated on the right axis).

and insurance within the village is only partial. When households are more patient ($\delta = 0.91$), in contrast, insurance is stronger and matches that in the data when insurance is essentially perfect within smaller groups that consist of only 5 households. Asymmetry is 0 when insurance groups are pairs, and increases monotonously with village size, but is substantially smaller when insurance is strong.[30]

Figure 4 therefore shows that the CD model achieves its superior fit largely because of its prediction of strong insurance in smaller group sizes rather than a superior performance than the ID model at given group size. This naturally gives rise to the question if the CD model of endogenous sustainable group size is preferable to a version of the ID model with exogenously small groups. We examine this question in Section 5.

4.6 Further analysis and sensitivity

A separate Online Appendix contains further analysis. It first corroborates the approximation underlying our benchmark estimates of the CD model. In Section A.2, we first compare the results from that approximation to the exact solution of the model for the small groups where the latter is feasible. In fact, for groups of three and four, the moments predicted by the approximation are extremely close to the exact ones. Thus, conditional on a small group size, we are confident that our approximation is accurate. This confidence rests on the maintained assumption, however, that the approximation indeed accurately identifies sustainable groups. While we cannot solve the exact model in its general version for larger group sizes, we can identify sustainable groups in the exact model whenever perfect insurance is sustainable. This

[30] There is, potentially, an additional, more mechanical reason for lower asymmetry in the alternative model where, as explained in Section 4.1, 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 partial out, is less than perfectly correlated with average incomes in the insurance group. Thus, with multiple risk sharing groups, conditioning on village variables leaves a group-component in household-level variables that makes the observed data more symmetric. Since the treatment of the simulated data follows directly from the standard conditioning we apply to the empirical data, this differential effect does not imply, in our view, any inconsistency. As a robustness exercise, however, we repeat the simulated method of moments estimation on the unconditional moments in Section A.6 of the Online Appendix.

is relevant in particular as insurance is indeed close to perfect in our benchmark estimation. In a second robustness check, we therefore report the maximum sustainable groups under a maintained norm of full insurance. The sizes we estimate in this alternative way are identical to the benchmark estimates for Aurepalle and Kanzara and one household smaller in Shirapur. Again, this underlines our confidence in the results.

Because the equilibrium transfer rules in the model we propose can be complex, we study in Section A.3 whether there is a simpler rule of thumb implementation of the constrained-efficient contract that keeps the main features of the full model but is easier to implement. Within the set of such simpler rules we estimate there to be full risk-sharing in small groups, in contrast to our benchmark estimates of strong but not perfect risk-sharing.

We then turn to examine the robustness of our benchmark results to changing some of the numerical inputs. Section A.4 illustrates how the model features are qualitatively unaffected when the income process has more than the three support points we consider in our benchmark analysis. Section A.5 shows robustness along two additional dimensions: first, with a different choice of outside option for the rest-of-village in the CD model, namely constrained-optimal rather than first-best risk sharing, the participation constraint for the rest of the village is relaxed, making the insurance mechanism more attractive, insurance stronger and insurance groups slightly larger (consisting of 6 households in the case of Aurepalle). Second, and again in line with intuition, when deviations are made more attractive by allowing instantaneous continued insurance in a sub-group, without a period of autarky, the opposite is true - the degree of insurance and the maximum group size are reduced (consisting of 3 households in the case of Aurepalle). Importantly, the approximate symmetry of consumption-income comovement in the CD model is unaffected by any of these changes in the specification.

Finally, we show how the results, and in particular the finding of a strong positive asymmetry in the standard ID limited commitment model, are not dependent on particular choices about the model environment or moments to target. Section A.6 shows that the results are unaffected when we do not partial out aggregate variation in the targeted data (and model) moments.

Second, Section A.7 reports results when we freely estimate the persistence parameter ρ in the CD and ID models to target the joint distribution of consumption and incomes in the data. Even with a serial correlation that is counterfactually negative relative to the moments we find in the data, the standard model is not able to simultaneously predict a symmetric distribution of consumption and income growth and a realistic observed degree of insurance. Section A.8 shows how the asymmetry actually increases when we allow a simple form of heterogeneity in risk-preferences in the ID model. Finally, Section A.9 estimates the three models with measurement error in incomes and consumption. Measurement error in consumption is estimated to be substantial in all three models, and brings the predicted volatility of consumption growth $Var_{dc|dy>0}/Var_{dc|dy\leq 0}$, which was counterfactually low in all benchmark estimates, in line with the data (thus essentially removing it from the estimation criteria). Measurement error in incomes attenuates the regression coefficient $\beta_{dcdy|dy}$ and 'blurs' the distinction between income increases and declines in measured data. This strongly improves the fit of the SI model. In the CD and ID models, in contrast, the estimates with income measurement error are not well identified and should therefore be treated with caution.

5 Risk sharing in endogenously small insurance groups

In this section, we provide evidence beyond the joint distribution of individual consumption and income, which we have focused on until now, that supports risk sharing in small, endogenous groups in our dataset. First, we review the empirical literature on mutual insurance networks and groups to show that there is ample evidence across a variety of contexts that insurance takes place in smaller subgroups within communities. Second, we show that the pattern of pairwise consumption correlations of households in the ICRISAT villages is inconsistent with risk sharing taking place at the village level, but strongly supports risk sharing in small groups. Third, we calculate the exact model solutions for single digit group sizes and show qualitative patterns that support the presence of coalitional deviation threats in the ICRISAT data.

5.1 Existing evidence for small risk-sharing groups

Our prediction of strong insurance among smaller groups of households is in line with a large literature in development economics that shows how risk sharing often takes place in smaller groups. Most relevant to our argument is the literature that maps relevant insurance networks by asking households to identify insurance partners they rely on in times of need (Fafchamps and Lund, 2003; De Weerdt, 2004). Fafchamps and Lund (2003) find that households in the Philippines make and receive transfers from an average of 5 other households. Our own calculations using data on social networks in South India (see Banerjee, Chandrasekhar, Duflo and Jackson (2013)) show that a household is connected to an average of just over 3 other households for the purposes of mutual help in times of need.

A group that has been found to be particularly important for risk-sharing is the immediate and extended family. For example, Dercon and Krishnan (2000) find some evidence of full risk sharing within nuclear households in Ethiopia, and Fitzsimons, Malde and Vera-Hernandez (2015) and Angelucci, de Giorgi and Rasul (2015) document the role of the extended family in risk-sharing in Malawi and Mexico respectively. The latter two studies also report the size of these extended family networks: In the Malawian data, a household has on average 9.4 siblings of husband or wife (although not all of them live in the same village), while an average of 7.5 households within a village belong to the same family in the Mexican sample.

Other studies find that insurance is typically stronger within (larger) exogenous groups such as clans, castes. For example, Grimard (1997) studies risk sharing among ethnic groups in Cote d'Ivoire, and most pertinently, Morduch (1991) and Mazzocco and Saini (2012) show for the ICRISAT panel that risk sharing is much stronger within castes than across.

Importantly, while these exogenous groups are often larger than the equilibrium group sizes we estimate here (though not much larger in the case of the extended family networks), this does not contradict our results: larger exogenous groups may contain several smaller endogenous groups that feature strong risk sharing and, for reasons we do not consider here, may not typically cross caste or kinship barriers. In fact, as we point out in Section 4, the degree of insurance observed in ICRISAT data is consistent with both large and small groups, but other moments of the joint distribution of individual consumption and income growth strongly point towards small groups.

5.2 Pairwise consumption correlations in the ICRISAT data support small risk-sharing groups

The original ICRISAT data does not allow for a mapping of insurance groups, since it does not include information on how many villagers a household cooperated with in situations of need, and/or who they received transfers from, or gave transfers to, in a given year. To provide additional evidence of small risk-sharing groups in this data set, we therefore follow an indirect approach to identify groups that builds on the intuition that consumption comovement should be stronger within than across insurance groups (see for example Ligon (2004)). This suggests that the distribution of bilateral consumption correlations within a village should be immediately informative about the size of insurance groups.

In the following we compare the distributions of bilateral consumption correlations implied by the ID model (where we can exogenously vary group size) to those in the data. For this, it is important to consider the possibility that households share risk with others in the village that are not in the data sample. Whenever both the sample and risk-sharing groups are small relative to the village population this increases the number of observed 'zero' correlations, corresponding to unconnected households, relative to the case where the village coincides with the smaller sample. Moreover, when we are interested in comparing the distribution of bilateral consumption correlations in the model to that in the data, we evidently need to modify our approach of comparing average population-moments implied by (a long simulation of) the model to average moments found in our data samples (accounting for sampling variation identified by bootstrapping the latter). Rather, in the following we compare the data distributions to distributions estimated both from a long census of villagers, and from short sample-panels whose cross-sectional and time-dimensions correspond to those in the ICRISAT data.

Figure 5: Model-implied bilateral consumption correlations within a village

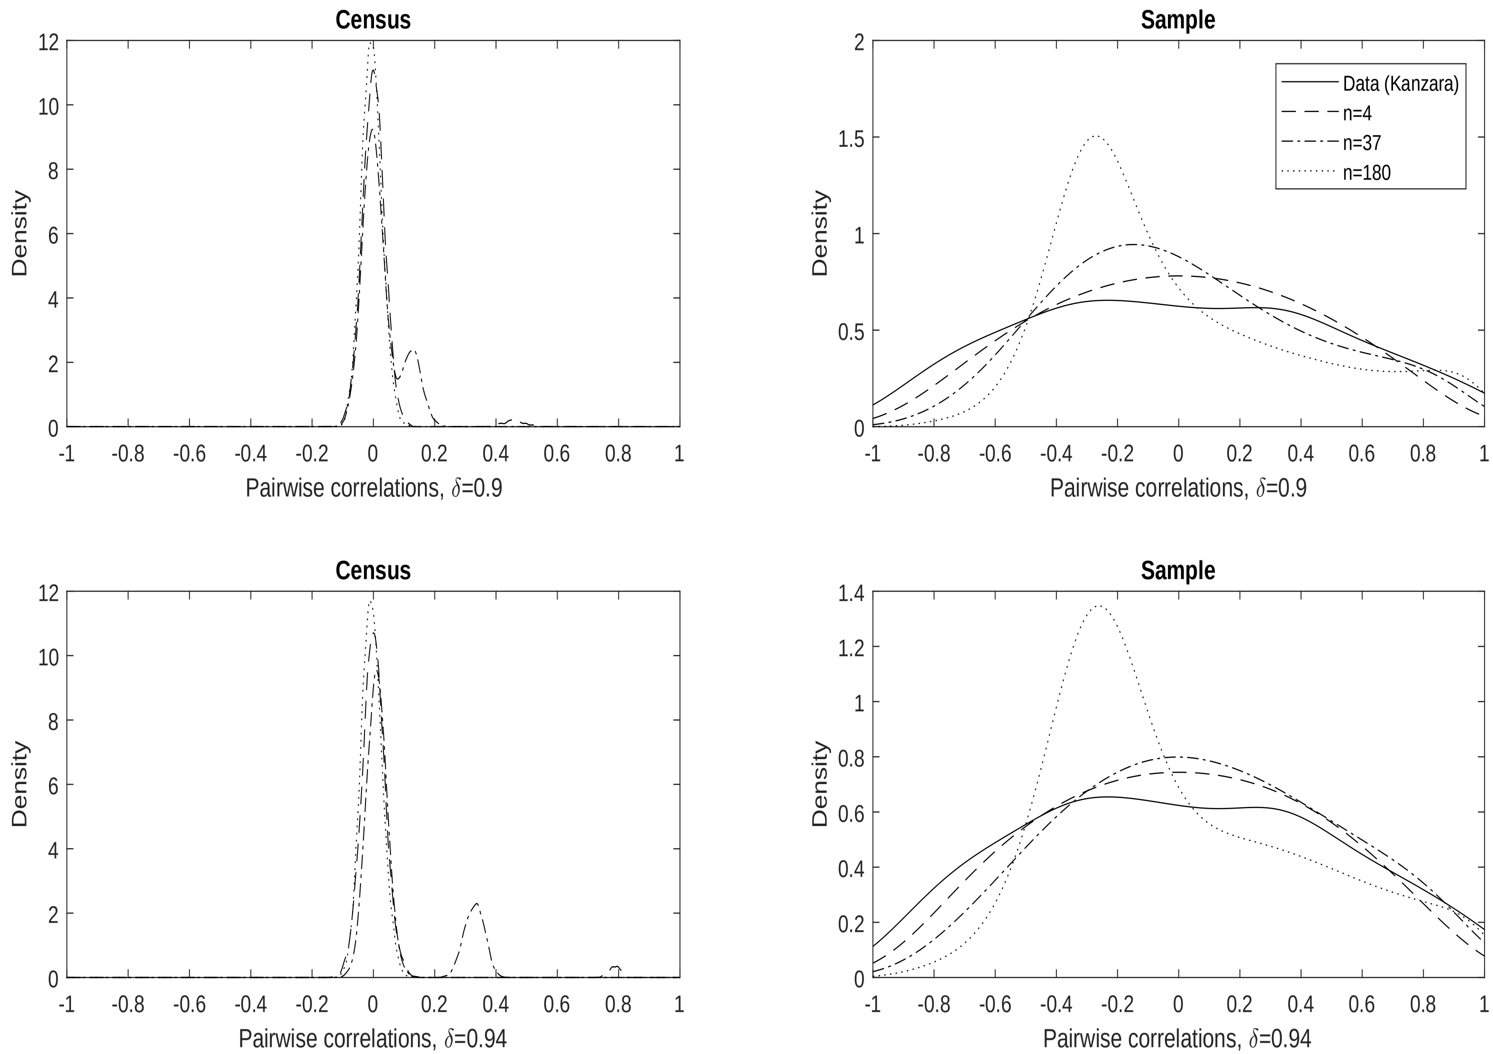

Notes: The figure shows kernel estimates of the density of pairwise consumption correlations in a village whose size and income process correspond to Kanzara, and where households are divided into insurance groups of size 4 (dashed line), 37 (dashed-dotted line) and 180 (dotted line). For the panels in the left-hand column of the figure ("Census") the data are generated by the ID model from a large number of simulated panels with $n = 180$ and a large number of time periods, controlling for time fixed effects. For the right-hand column ("Sample") the data are generated by a large number of 6-year samples of size $n = 37$, controlling for time fixed effects. The figure also plots a kernel estimate of the density of pairwise consumption correlations (after partialling out time fixed-effects) in the ICRISAT data for Kanzara (solid line).

Figure 5 documents how the size of insurance groups affects the distribution of within-village consumption correlations in the ID model. The panels in the left-hand column of Figure 5 show, for low (top row) and high insurance (bottom row), a kernel estimate of the distribution of bilateral consumption correlations in model-generated census data of a village population of 180 households across many time periods, where village size and income process correspond to Kanzara, and where households are divided into equal-sized insurance groups of different sizes (we refer to this henceforth as the population bilateral consumption correlations).[31] We choose groups of size 4 (dashed lines) and 37 (dashed-dotted lines), corresponding to our estimated group size in the CD model and the sample size, respectively and solve the model for $\delta = 0.9$ and $\delta = 0.94$, the discount factor estimated by the CD and the ID model in Table 5. We compare these to a village-level group equal to 180 households (the number of households in Kanzara, solid lines).[32] As in previous results, we control for time fixed effects.

For groups smaller than the village, the distribution in Figure 5 is bimodal, as bilateral correlations between villagers in different groups cluster around zero, but are positive for those within the same group. The correlation for those sharing membership of a risk-sharing group is increasing in the discount factor (as a higher discount factor implies more insurance), and decreasing in group size (as group-income is more volatile in smaller groups whose consumption thus comoves more strongly relative to idiosyncratic consumption and income movements). For the village-level model, in contrast, time fixed effects capture the entire variation in group-level resources, leaving only idiosyncratic consumption movements that imply a zero correlation in the village.

Figure 5 clearly shows that the population patterns of bilateral consumption correlations differ strongly depending on whether risk sharing takes place at or below the village level. To

[31] We focus on Kanzara in this analysis, because it is the village with the highest ratio of sampled households (37) to households living in the village at the time of the survey (180), making it easier to identify risk-sharing groups within the village (see Binswanger and Jodha (1978)).

[32] Using the whole income distribution as a state variable becomes infeasible for the village-level model with $n = 180$. We thus use a procedure similar to that in Laczo (2014), based on a discretized support of aggregate village income, rather than the full support of all possible aggregate income realizations. This, however, makes the degree of insurance at given preference parameters difficult to compare to those at smaller group sizes.

Table 6: Testing for equality of the data and model generated distribution of pairwise consumption correlations

Kanzara			
	CD	**ID**	**ID**
n	4.00	37.00	180
δ	0.94	0.90	0.90
Goodness of fit	0.0161	0.0234	0.0351
p-value	0.994	0.844	0.001

Notes: For the village of Kanzara, the table reports output from a Kolmogorov-Smirnov test that tests for equality of the distribution of bivariate consumption correlations in the data and the distribution generated by a model of risk sharing in groups of $n = 4$ and $\delta = 0.94$ (the estimated group size and discount factor in the CD model), $n = 37$ and $\delta = 0.9$ (the sample size in Kanzara and estimated discount factor in the ID model), and $n = 180$ and $\delta = 0.9$ (the number of households living in Kanzara). The statistics are based on 10,000 samples consisting of 37 households in 6 time periods drawn from the simulated model data. Over these samples, the minimum KS test statistic and the associated p-value are recorded. The table reports the average of these statistics over 1,000 repetitions of this exercise.

properly compare the model-generated pairwise correlations to the data for Kanzara, however, we need to take account of the small sample size of $n = 37$ and the short time dimension of only 6 years in the ICRISAT panel. The right-hand panel of Figure 5 therefore shows the density of bilateral consumption correlations averaging over 10,000 6-year panels of 37 households drawn from the population of 180 simulated households (we refer to this hence-forth as the sample bilateral correlations). Compared to the census, the model-generated densities based on samples are now much flatter. And the small-group density ($n = 4$) in particular traces out the shape of the density estimated on the data for Kanzara (the solid line) quite closely. Importantly, both the data generated by sampling from the model simulation and the empirical data contain many pairs of households whose consumption is strongly negatively and strongly positively correlated. The village-level model, in contrast, does not capture the dispersion, as its mass remains too concentrated around the mean.

The reason for the superior fit of the sample bilateral correlations is two-fold: (i) because the correlation between any two households is calculated on the basis of only six time periods, there is a much wider range of possible correlations, both for households who share membership of a group (varying around a positive average correlation), but especially for those who do not (varying around a zero average correlation). (ii) When risk-sharing groups are small relative to the sample (and village) size, the sample is mostly made up of households whose consumption is not connected through membership in a group. This makes the model-generated densities more similar to the relatively symmetric and dispersed density for Kanzara.

The graphical result is confirmed by a Kolmogorov-Smirnov test, which tests for equality of any two densities. Specifically, we repeat the above sampling exercise 1000 times, each time recording the KS p-value and test statistic (a measure of the average deviation between the density of bilateral correlations observed in Kanzara and those generated by the model for groups of size $n = 4$, $n = 37$ and $n = 180$) that produce the best fit between data and model over the 10,000 (37-households sampled from 180, six time periods) panels. We report the average of these statistics across the 1,000 repetitions in Table 6.[33] The deviation between model and data is smallest for the small group model with $n = 4$: the test of equality of the distributions has a p-value of 0.99. Risk sharing at the village level, on the other hand, is soundly rejected. Moreover, the superior fit of the small group model is extremely consistent across the 10,000 $\times$ 1,000 panels and repetitions: groups of $n = 4$ give a better fit to the data than groups of $n = 37$ in 93% of the cases and a better fit than groups of $n = 180$ in 99% of the cases.

5.3 Are risk-sharing groups endogenously small?

If small insurance groups exist in the ICRISAT villages, this begs the question whether they are endogenously small. Put another way, 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 group size equal to that of an exogenous insurance unit? This is clearly an important question, particularly because the policy implications of the two models may be quite different: in the CD model group size changes endogenously in response to policy

[33]When insurance groups are smaller than the village, the ability of the model to explain the observed correlation pattern crucially depends on picking the same mix of connected and unconnected households as in the ICRISAT sample. Since the latter is unobserved, we focus for each group size on the simulated panel that gives the best fit between model and data, assuming that this is the one that comes closest in terms of getting the mix of members and non-members right.

changes, while in the ID model, group size does not respond.

Conceptually, we would argue that there is much to like about the coalitional deviations model. Or conversely, the equilibrium group sizes we estimate are hard to motivate in a model with exogenously small groups unless one wants to argue that the number of actual risk-sharing partners is always equal to the number of potential risk-sharing partners. In contrast, 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. By implication, moderate insurance in larger exogenous groups, as observed for example in Mobarak and Rosenzweig (2012), is consistent in the CD (but not the ID) model with relatively high insurance in several smaller groups, that are not clearly delineated by virtue of their endogeneity.

So far, the main mechanism we used to distinguish between the CD and ID model is the different implication for group size. Given the overwhelming evidence for small insurance group sizes documented in Section 4.4 and Section 5.2, we now solve the ID and CD model for the same small group size and compare the implied joint distributions of individual income and consumption. Specifically, we solve an exact version of the CD and ID model for groups of size 3 and 4, which conditions on the full history of shocks and transfers in the group (rather than reducing the N-agent contract to a sequence of two-player contracts between an individual and rest-of-village) and, in the case of the CD model, allows deviating subgroups to continue with any division of surplus that is stable, not just the most equal one.

For the exact solution of the CD model, it is easy to see from equation (9) that the planner can simultaneously deter all potential deviations by a subgroup by offering its members consumption and continuation payoffs that make all members of the subgroup as well off as one period of autarky followed by some allocation on their constrained-efficient Pareto frontier. This follows, since by definition of Pareto efficiency, there are no further deviations from such an allocation that are profitable for all members of the deviating sub-coalition. As shown in Bold (2009), the optimal deterrence of coalitional deviations therefore requires the planner to optimally choose the allocation on the constrained Pareto frontier for which to make the subgroup indifferent between staying or leaving. This introduces a trade-off in the equilibrium allocations of constrained agents, which is not present in the ID model, where consumption and continuation payoffs of constrained agents are always optimally found by setting them equal to autarky.

We focus here on the exact solution of the model, (i) because when we solve the model only for small groups, we do not have to resort to the approximation, (ii) since the approximation treats the rest of the village in the same way in either model, this blurs some of the differences between the two models that may help us to distinguish them for a given group size. Table 7 and Figure 8 report the results for the income process estimated for Aurepalle.[34]

For high discount factors, the two models behave similarly at small group sizes and both versions of the exact limited commitment model provide a good fit to the data (see Table 7). Both predict groups of size four with similar discount factors of .92 in the CD model and of .9 in the ID model.[35] In the CD case, the estimated moments and preferences are very similar to Section 4.4, while the ID model is much improved, estimating almost the same preferences, but now, because of its smaller group size, achieving symmetry together with the moderate degree of insurance seen in the data. Though the goodness of fit is slightly better in the CD model, the difference between the two models at given group size is negligible when fitting to the ICRISAT data.

The same is not true at lower values of δ, however, where there are important differences in the simulated moments. In Figure 6, we plot the four moments across a range of discount factors in a group of size 3 (groups of size four do not attain stability for discount factors below .9). For each discount factor, the predicted degree of insurance is on average lower in the model with coalitional deviations than in the standard alternative, and the differences are stark for

[34] Just as in Section 4.4, we estimate the models for $n = 3$ on an income grid with three support points. However, for groups of size $n = 4$, this becomes infeasible, and we therefore estimate the model on an income grid with only two support points. In the CD case, the recursive nature of the algorithm also requires estimation of the model for $n = 2$ and $n = 3$ in order to calculate the deviation payoffs in a group of size $n = 4$ and for consistency, we therefore solve the model for these group sizes, when they are used as an input into the solution for $n = 4$, also on the smaller income grid.

[35] Note that the choice is here only over groups of size 2, 3 and 4 in the case of the ID model and over the largest stable group in the set of 2, 3 and 4 for each discount factor in the case of the CD model. The worse goodness of fit is explained by estimating the model on a coarser grid for both preferences and income compared to Section 4.4

Figure 6: Simulated moments from exact solution of CD and ID model for groups of size three.

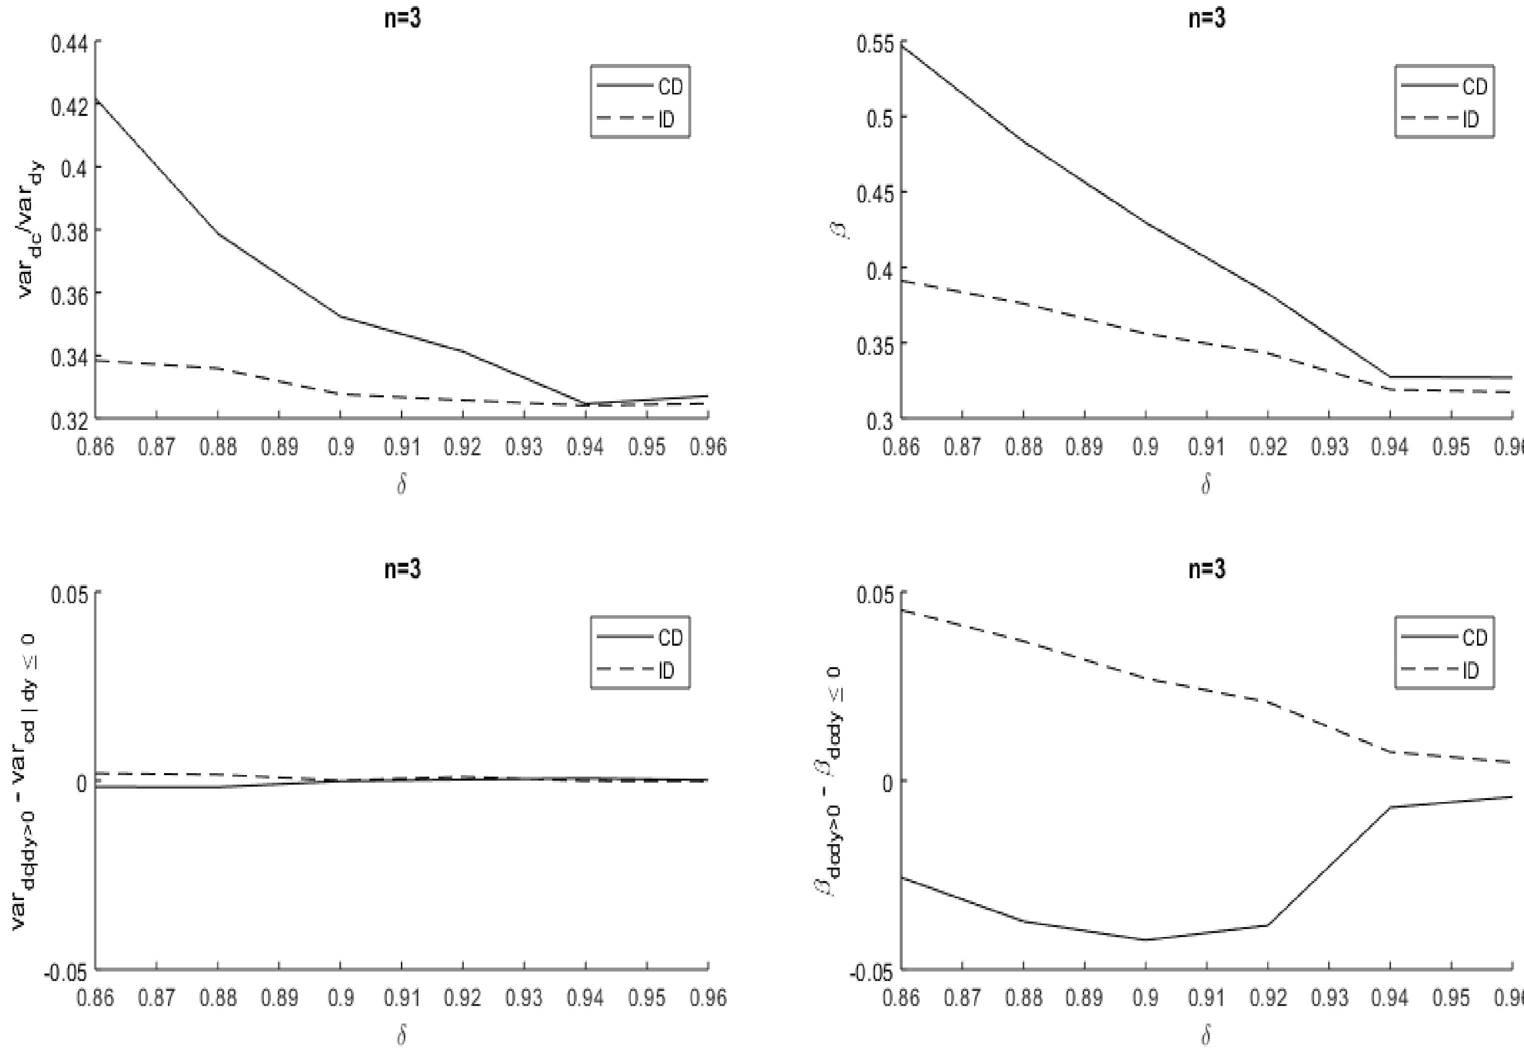

The table shows four key moments summarising the joint distribution of consumption and income growth in simulations of the CD model vs. the ID model with groups of size 3 and $\delta \in [0.86, 0.96]$. The left upper panel plots the ratio of the variance of individual consumption changes to the variance of individual income changes. The right upper panel shows the slope coefficient in a regression of consumption growth on income growth. The left bottom panel reports the asymmetry in the variance ratio for income winners and losers on the left axis. The difference between income winners and losers of the slope coefficients in a regression of consumption growth on income growth is reported in the bottom right panel. In all panels, the solid line reports the moment in the CD model and the dashed line in the ID model.

low discount factors – a consequence of the more attractive outside option in the CD model. This suggests that the CD model may be better able to explain data patterns in the context of low degrees of risk sharing where the ID model does a poor job (for example in economies with capital, see e.g. Ábrahám and Cárceles-Poveda (2009)).

Second, the coalitional deviations model, in line with the data, is able to generate a negative asymmetry in the consumption response of income winners and income losers, $\beta_{dcdy>0} - \beta_{dcdy\leq 0}$. That is, consumption changes of income winners are *less* responsive to income changes than those of income losers. While the absolute size of the asymmetry is below that observed in the data for Aurepalle and Kanzara, it is of the same order of magnitude as in Shirapur for lower discount factors in groups of size 3.

How does the coalitional deviations model generate a negative asymmetry in the regression coefficient? The pattern arises in periods when $n-1$ constrained individuals face a single unconstrained, typically low income, individual. The budget constraint then implies that the former's consumption is almost exclusively determined by income, while a small change in income of a constrained individual would affect her consumption only via its effect on joint group income of the $n-1$ constrained agents. As a result, there is negative asymmetry in the degree of risk-sharing in these states, which occur frequently when discount factors are low and groups small.

In sum, we think that the additional evidence presented here, together with our core results in Section 4.4, strongly supports the interpretation that the aggregate patterns we observe in the ICRISAT data are generated by risk-sharing in small groups. We can also show that even for a given small group size, the two models have different implications for the joint distribution of consumption and income, both in terms of the degree of risk-sharing and its asymmetry, which become quantitatively important at discount factors lower than those predicted for the ICRISAT villages. These differences notwithstanding, we would argue that any test that attempts to distinguish between the two models at given small group size, i.e. without exploiting their different group size predictions, will struggle in an environment with moderate to high aggregate insurance: both models generate moderate aggregate insurance through almost perfect insurance in small groups, implying that the enforcement constraints which allow us to distinguish the two models seldom bind.

Table 7: Preferences estimated to target all 4 moments - exact solution for groups of size 3 and 4

Aurepalle			
	Data	**CD**	**ID**
n		4.00	4.00
δ		0.92	0.90
s.e.		0.01	0.01
σ		1.00	1.00
$\frac{Var_{dc}}{Var_{dy}}$	0.22	0.23	0.22
β_{dcdy}	0.21	0.27	0.23
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	-0.000	0.001
$\beta_{dcdy>0}-\beta_{dcdy\leq 0}$	-0.41	-0.002	-0.016
Goodness of fit		17.58	17.81

Notes: For the village of Aurepalle, 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 CD model (second column) and the ID model (column 3). 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 $\delta \in [0.86, 0.96]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

6 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). Our estimation of the model showed that, 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 strong insurance in small insurance groups that enables the model to predict a realistic degree of insurance at the same time as symmetric responses of consumption to income shocks. Moreover, this is not a trivial feature of the model we propose: both in the standard limited commitment model and in a simple model of bufferstock savings reacts consumption substantially more to income rises than to income declines when the measured degree of insurance is in line with the data. These models can thus predict either the observed degree of insurance or approximate symmetry, but not both, unless measurement error is so large as to dominate both the variation in consumption and income.

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. Our results suggest, however, that interventions that change income processes may not only affect insurance in given group sizes through incentives to deviate (as in Attanasio and Rios-Rull (2000), Krueger and Perri (2011), and Broer (2011)) but also change the size of insurance groups, an effect that the standard ID model cannot capture. 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.

Abreu, Dilip, "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, 1988, *56*, 383–396.

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.

Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson, "The Diffusion of Microfinance," *Science*, 2013, *341* (6144).

Bernheim, B. D. and D. Ray, "Collective Dynamic Consistency in Repeated Games," *Games and Economic Behavior*, 1989, *1*, 295–326.

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

Binswanger, Hans P. and N.S. Jodha, "Manual of Instructions for Economic Investigators in ICRISAT's Village Level Studies," Technical Report, Village Level Studies Series, Economics Program, ICRISAT 1978.

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 Tobias Broer, "Risk-Sharing in Village Economies Revisited," CEPR Discussion Papers 11143, C.E.P.R. Discussion Papers March 2016.

——— and ———, "A Note on Group Formation," December 2018. .

Broer, Tobias, "Crowding out and crowding in: When does redistribution improve risk-sharing in limited commitment economies?," *Journal of Economic Theory*, May 2011, *146* (3), 957–975.

———, "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.

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

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

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.

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

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.

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

——— **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.

Ligon, Ethan, *Insurance Against Poverty*, ed., Vol. of , : Oxford University Press, .

Lim, Youngjae and Robert M. Townsend, "General Equilibiurm Models of Financial Systems: Theory and Measurement in Village Economies," *Review of Economic Dynamics*, 1998, *1*, 59–118.

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.

Morten, Melanie, "Temporary migration and endogenous risk sharing in village india," Technical Report, National Bureau of Economic Research 2016.

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.

Park, Yena, George Mailath, Dirk Krueger, and Harold Cole, "Fragile Financial Coalitions Under Belief Coordination Frictions," mimeo 2018.

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

Ray, Debraj and Rajiv Vohra, "Coalition formation," in "Handbook of game theory with economic applications," Vol. 4, Elsevier, 2015, pp. 239–326.

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.

Winter, Joachim K, Kathrin Schlafmann, and Ralf Rodepeter, "Rules of Thumb in Life-cycle Saving Decisions," *The Economic Journal*, 2012, *122* (560), 479–501.

A FOR ONLINE PUBLICATION: Appendix

In this online appendix, we discuss the results and their robustness along a number of dimensions. In Section A.1, we illustrate the identification problem in the limited commitment model at high levels of insurance. In Section A.2, we evaluate the approximation of the CD and ID models that we use in our benchmark results. For this, we show, first, that for the small group sizes of three and four households where we can calculate it, the exact solution of the model implies key moments that are very similar to those in the approximation. Second, the sustainable group sizes in a version of the CD model that conditions on full insurance, where we can easily consider much larger group sizes, are also very close to those in the benchmark model. Section A.3 studies whether there is a simpler rule of thumb implementation of the constrained-efficient contract that keeps the main features of the full model but is easier to implement. Within the set of such simpler rules we estimate there to be full risk-sharing in small groups, in contrast to our benchmark estimates of strong but not perfect risk-sharing. Section A.4 shows evidence that our results are robust to the number of support points in the process of individual incomes, while Section A.5 considers alternative outside options. Section A.6 shows how estimating the models on the raw data, rather than demeaning period by period, yields results similar to our benchmark estimates. Section A.7 discusses the role of income persistence and presents results for the CD and ID models where we also estimate income persistence to target the consumption moments. Section A.8 considers a simple form of preference heterogeneity. Section A.9 generalizes our benchmark estimates to include measurement error in consumption and incomes. Sections A.10 and A.11 report, respectively, the transition matrices for income we use in our computation, and additional descriptive statics for the ICRISAT data.

A.1 Identification of preference parameters

Figure 7 shows how neither the CD nor the ID 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 structural estimation. Note that the endogeneity of group size in the CD model can potentially help identification: at a discount factor above 0.96 we choose a risk aversion parameter that is consistent with unchanged moments and goodness of fit at a given group size of 4. But as the figure shows, now groups of 5 become sustainable, which increases, for example, the degree of insurance, and decreases the goodness of fit. This may bound the range of discount factors for which the model is not identified. Within those bounds, however, the model is still not identified. This is why we choose to normalize $\sigma = 1$ in both the ID and CD model.

Figure 7: Estimated parameters and moments in Aurepalle

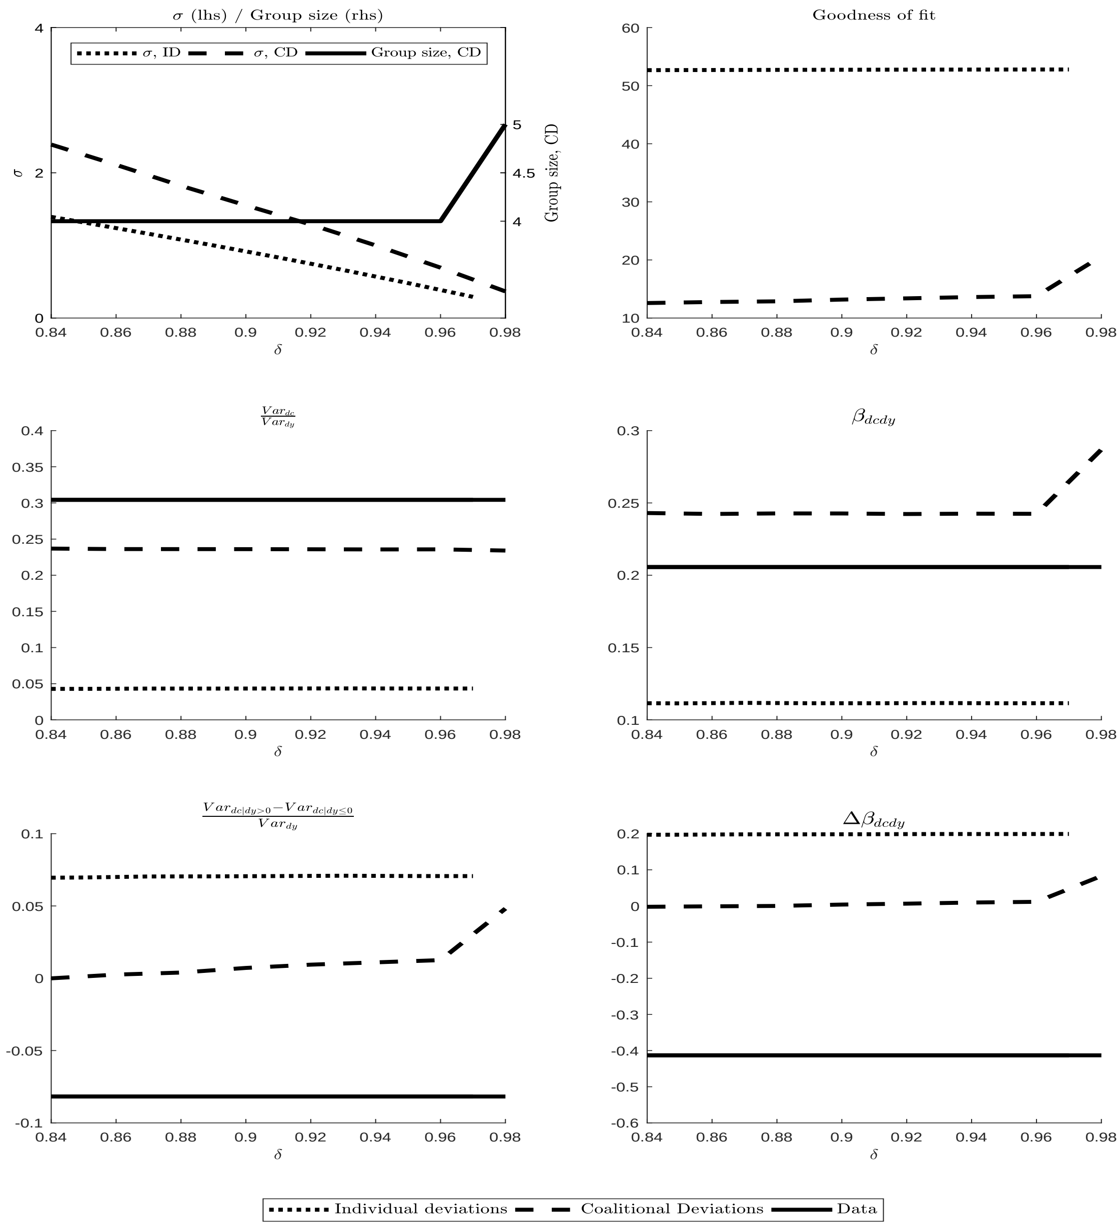

Notes: The figure shows the estimated risk aversion σ and (in the case of coalitional deviations) group size (panel 1), the goodness of fit (panel 2) and the 4 moments of interest for different values of the discount factor δ.

A.2 Comparing the exact model and approximation

The analysis in our paper adapts the standard approximation of the dynamic limited commitment model (Ligon, Thomas and Worrall, 2002) that derives the policies for the Lagrange multipliers, ϕ, in equation (6) from a ficticious setting where an individual interacts with a homogeneous rest of the village whose outside option depends only on aggregate income. The approximation then calculates consumption shares using the relative multipliers that result when several such individuals interact. While this approximation is exact for $n = 2$ and accurate for large groups of several hundred households (Ligon, Thomas and Worrall, 2002), how it evolves for $n > 2$ is unknown. This appendix compares the approximation to an exact solution along two dimensions: first, we compute the exact solution of both the CD and ID models for small groups of size 3 and 4, the largest sizes where this is feasible, and compare the implied features to those of the approximation. Second, we identify the maximum sustainable size of insurance groups exactly, under the assumption of full insurance.

A.2.1 Exact and approximate model solutions at small group sizes

While the approximation trivially coincides with the exact solution when insurance is perfect, we expect it to be imperfect whenever risk sharing is partial. In particular, the assumption of income-pooling in the rest of the village in the approximation introduces two sources of error relative to the exact solution: (i) the Lagrange multiplier on the incentive constraint of the rest of the village is calculated at its average income realisation. In an exact solution, in contrast, it is the average of the Lagrange multipliers across the income distribution in the rest of the village that determines consumption of any remaining unconstrained individual. Given the convexity of the Lagrange multiplier around zero, 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) in an exact solution, deviating individuals continue in autarky (in the ID model) or in a smaller stable group (in the CD model).

Figure 8: Simulated moments from exact solution of CD and ID model for groups of size three and four.

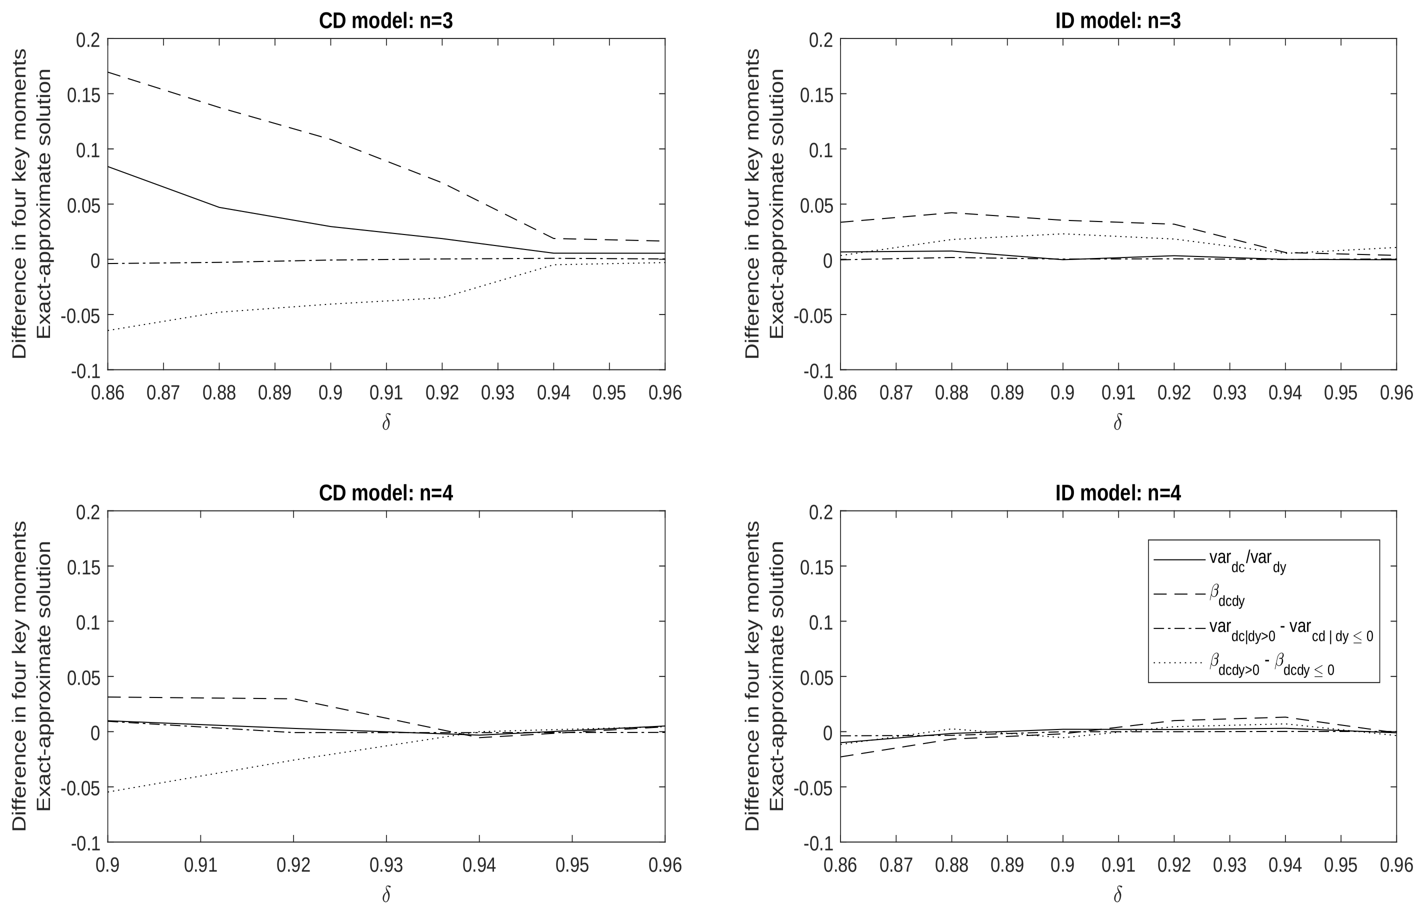

The table shows four key moments summarising the joint distribution of consumption and income growth in simulations of the CD model vs. the ID model with groups of size 3 and 4 and $\delta \in [0.86, 0.96]$. In each panel, the solid black line plots the difference between exact and approximate solution in the ratio of the variance of individual consumption changes to the variance of individual income changes. The dashed line plots the difference between exact and approximate solution in the slope coefficient in a regression of consumption growth on income growth. The dashed-dotted line plots the difference between exact and approximate solution in the asymmetry of the variance ratio for income winners and losers. The dotted line reports the difference between exact and approximate solution in the asymmetry of the slope coefficients in a regression of consumption growth on income growth.

This contrasts with an outside option of perfect risk sharing within the rest-of-the-village in the approximation. Without persistence in incomes (as is approximately true in our data), this makes the outside option strictly better in the approximation and hence its participation constraint more binding, resulting in less risk sharing (and thus potentially offsetting effect (i)). In the CD model, there is an additional source of error resulting from our assumption of equal treatment (subject to binding participation constraints). This prevents the planner from treating coalition members asymmetrically despite the fact that this might increase her ability to deter joint deviations and may be optimal given the history of the contract. Importantly, while it is a priori unclear which of the three effects dominates, approximation error should decrease the higher is the degree of insurance, and the less often participation constraints bind in equilibrium.

We now compare the exact and the approximate solution for the small group sizes, $n = 3$ and $n = 4$, where we can calculate the former. Note that, while Bold (2009) performs a similar analysis of the model in groups of three without income persistence, we are, to our knowledge, the first to study the exact solution of the three and four agent-limited commitment model with an estimated, persistent income process and compare it to data from actual village economies.

In Figure 8, we compare the approximate and the exact solution of ID and CD models. We solve the models with log utility and discount factor δ in intervals that cover the estimates for the CD model reported in Table 5, for the income process estimated in Aurepalle. The figure plots the difference between the exact and approximate versions of the four key moments used in the analysis, namely measures of the sensitivity of consumption with respect to income changes and its asymmetry in the sub-samples of income winners and losers.

Most importantly, the approximation errors in the CD model are indeed very small at the estimated values of δ in Table 5 ($\delta = .94$ for $n = 4$), giving us confidence in our main result. Moreover, at values of δ greater than the estimated values, they remain close to zero. And the same holds for the small-group version of the ID model. Again, this is because any difference between the exact and approximate solutions arise from states when participation constraints bind, and there are fewer of those at high discount factors.

This also implies that we would expect approximation errors to rise at lower discount factors. Figure 8 shows this to be the case for both the ID and CD models at $n = 3$, if much less so for $n = 4$, where insurance is higher for a given discount factor in the ID model and CD groups are not stable for discount factors below .9. Speci, the exact solution of the CD model predicts somewhat lower insurance than the approximation at low discount factors, suggesting that the use of an average participation constraint for the rest of village in the approximate solution (as discussed above) may be the dominant source of error. And again, the exact solution of the CD model features negative values of the regression asymmetry moment for lower discount factors (see our discussion in Section 5.3).

Overall, the increase in the approximation error at lower discount factors seems modest in relation to the level of the moments (see Figure 6), although, as one might expect, the approximate solution does understate somewhat the increasing difference in the predicted degree of insurance between the two models in groups of 3 households. Importantly, the rise in approximation errors is concentrated at group sizes and values of δ smaller than those we estimate, where there are substantial differences between the two models even at given group size. Our conclusion from this evidence is therefore that for the income processes and high degrees of insurance observed in ICRISAT data, and the resulting estimated group sizes, the approximation we use is accurate.

A.2.2 The set of stable group sizes with full insurance

The results in the previous subsection make us confident that our results accurately capture the features implied by limited commitment insurance in small groups. Since we cannot solve the models exactly for larger groups, however, these results are strictly speaking informative only under the maintained assumption that the approximation indeed accurately identifies stable groups. That is, our comparison of the exact and approximate solution at small group sizes justifies the use of the approximation only inasmuch as there are no larger group sizes that would be stable in the exact (but not the approximate) solution. In this section, we examine this assumption.

While we cannot solve the exact model in its general version for larger group sizes, we can identify sustainable groups in the exact model whenever perfect insurance is sustainable. This is relevant in particular as insurance is indeed close to perfect in our benchmark estimation. In this section, we therefore look at sustainable groups with perfect income-pooling. Specifically, we maintain equal sharing of consumption resources in groups whose size varies between $n = 2$ and the village size, and identify the maximum group size such that individuals would in no income state find it optimal to deviate and form any smaller group (that continues to share resources equally after one period of autarky).

To understand the results in Table 8, note that with full insurance, the predicted moments are insensitive to changes in preferences that are consistent with a given sustainable group size. In other words, even after normalizing the risk aversion parameter, the discount factor δ is only set identified between a lower and an upper value, δ_{min} and δ_{max} respectively, indicated in the table. It turns out that this range only comprises one value of δ on our grid, however, which is, as expected, higher than in our benchmark estimates, as required to make full insurance sustainable relative to the outside option of individual autarky. Importantly, the estimated maximum sustainable group sizes are close to the benchmark sizes. Indeed they are identical for Aurepalle and Kanzara where the estimated degree of insurance is slightly larger than in the benchmark estimation where it was close to, but not equal to, full insurance. At a smaller maximum sustainable group size of 4, the degree of insurance in Shirapur is lower than that in the benchmark results. Again, these results give us confidence that our benchmark estimation, based on our approximation, accurately captures the data generating process.

A.3 Implementing the constrained-optimal contract with a rule of thumb

The transfer function in the CD model is potentially complex. In this subsection, we examine whether there is a simpler rule of thumb for the constrained-efficient contract that fulfills the following properties (see Winter, Schlafmann and Rodepeter (2012) for a related discussion in the context of the optimal savings problem): (i) the rule of thumb for transfers is simple and closed-form, (ii) it retains the essence of the exact solution, namely history dependence (whereby households are rewarded for contributions later in the contract) and coalition-proofness.

In general, a closed-form solution of the limited commitment risk-sharing contract can only be obtained when a candidate transfer rule is determined independently of participation constraints. To find a rule of thumb, we therefore do not look for the constrained optimal transfer that satisfies the participation constraints exactly (as in Section 2), but rather examine whether a given fixed transfer rule is stable with respect to deviations. The simplest fixed transfer rule is presumably equal sharing of resources among households. Inspired by this, we concentrate on a class of simple transfer rules according to which household transfers equal a constant percentage x ($50 < x < 100$) of those implied by equal sharing every period, adjusted for a simple form of history dependence: when all group members have the same income, those who have received transfers in the previous period repay y% of the received transfer to those who made it.[36]

The set of stable groups and insurance contracts is then derived recursively, as in the full model: Suppose that for each $m < n$, we have identified the transfer rule $\sigma(x^*, y^*, m)$ that maximizes expected discounted life-time utility subject to the insurance contract being stable with respect to deviations by sub-groups (which themselves must be stable with respect to further deviations). For a group of size n, we then calculate for each insurance contract $\sigma(x, y, n)$ with $x \in (0, 100)$ and $y \in (0, 25)$ the implied discounted life-time utility. We then check in

[36]We restrict history dependence to these symmetric states because this is when the full model suggests that history dependence matters most: participation constraints typically do not bind in these states, such that consumption shares are a direct function of the history of states and transfers through the updated relative Pareto weights, γ_r^i for $i = 1, ...n-1$.

each state whether the utility the k individuals with a high income realization derive from the contract is higher than their expected payoff of deviating as a group with $m \leq k$ individuals that implements its preferred stable contract $\sigma(x^*, y^*, m)$ following deviation (since this is the most binding constraint).

If all participation constraints are satisfied for $\sigma(x, y, n)$, then this contract is part of the set of stable contracts for group n. Having made this calculation for each possible contract $\sigma(x, y, n)$, the constrained-optimal contract is found as the one among those that are stable, which maximizes expected discounted life time utility. If no contract satisfies the participation constraints, the group of size n is deemed unstable.

Implementing this rule of thumb for the CD model, we estimate group sizes of 4 in all villages, very similar to the full model. Rather than close-to perfect insurance in small groups, as in the full model, the rule of thumb predicts full risk-sharing in these groups ($x = 100$, $y = 0$). In other words, the rule-of-thumb estimates are identical to those in Table 8, where we looked at maximum stable group sizes under the maintained assumption of full insurance (and we therefore omit reporting them).[37] The simpler models match the data almost as well as the exact or more complex approximation of the contract. Hence, it is plausible that the observed data is generated by a rule-of-thumb implementation of the constrained-efficient risk-sharing contract that is robust to coalitional deviations.

A.4 A finer income grid

For our simulations, we use a discrete version of the estimated income processes in Table 3. Since most of our results focus on (moments of) the joint distribution of consumption and income growth, the way in which we discretize incomes might potentially affect our results. In particular, one might worry that the degree of asymmetry in the ID model, which was a primary reason for its inferior fit, may be a consequence of our particular choice of income process.

[37] Our quantitative analysis of rules of thumb only considers two income states which leads to small quantitative differences in the moments relative to those reported in Table 8.

Although we are constrained in the number of income states for our benchmark results (which use the full joint distribution of incomes in the insurance group as a state variable), we can compute the ID model for different income processes using a procedure similar to that in Laczo (2014), where only the aggregate income in the rest of the village is relevant. Table 9 reports the results for one such exercise, where we double the number of support points of individual income from 3 to 6 and approximate the income process for the rest of the village as a discretized AR(1) process with 5 support points.[38]

To understand the results, note that more dispersion in incomes (of either the individual or the rest of the village) reduces insurance. This is because higher maximum income typically implies a higher maximum autarky value, and therefore a more binding participation constraint for the highest income individuals. This explains the stronger insurance in this specification of the model (with a coarse approximation of rest-of-village income) relative to our benchmark: the former does not allow for the extreme income realizations that are present in our benchmark economy (with its full set of possible cross-sectional distributions).

When the process for individual income has three support points ($n_y = 3$), insurance is strong, and in fact close to perfect for the village of Shirapur (where the negative serial correlation in incomes implies less dispersed outside options). In line with the intuition in the previous paragraph, increasing the support points of the individual income process from three to six implies more extreme individual income realizations, and thus more extreme values of autarky, more binding participation constraints, and therefore less insurance. Importantly, however, the asymmetry of the standard model increases as the degree of insurance falls. Thus, whenever there is strong insurance with a coarse income support (as in our benchmark estimates), an increase in the income support lowers the degree of insurance but increases the asymmetry.

[38] As in Laczo (2014), we estimate a continuous AR(1) process on simulated data and use Rouwenhorst's (1995) method to dicretize it.

Table 8: Sustainable groups with full insurance

	Aurepalle		**Kanzara**		**Shirapur**	
	Data	**CD**	**Data**	**CD**	**Data**	**CD**
n		4.00		4.00		4.00
δ_{min}		0.98		0.98		0.97
δ_{max}		0.98		0.98		0.97
σ		1.00		1.00		1.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.23	0.56	0.23	0.33	0.23
β_{dcdy}	0.21	0.22	0.22	0.23	0.17	0.23
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.00	-0.25	0.00	-0.16	-0.00
$\beta_{dcdy>0}-\beta_{dcdy\leq 0}$	-0.41	0.00	-0.14	-0.00	-0.05	0.00
Goodness of fit		12.84		5.53		6.48

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 where groups fully insure consumption of its members, but are required to be sustainable in the long run with respect to deviations to smaller full-insurance groups. 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 $\delta \in [0.9, 0.99]$. The goodness of fit reported is the value of the criterion function at the chosen parameters.

Table 9: The ID model with different income processes

	Aurepalle			Kanzara			Shirapur		
	Data	**ID, n_y=3**	**ID, n_y=6**	**Data**	**ID, n_y=3**	**ID, n_y=6**	**Data**	**ID, n_y=3**	**ID, n_y=6**
n		34.00	34.00		37.00	37.00		31.00	31.00
β		0.89	0.89		0.90	0.90		0.88	0.88
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.04	0.05	0.56	0.06	0.08	0.33	0.00	0.01
β_{dcdy}	0.21	0.11	0.11	0.22	0.17	0.18	0.17	0.02	0.05
$\frac{Var_{dc\|dy>0}}{Var_{dc\|dy\leq 0}}$	-0.08	0.06	0.07	-0.25	0.09	0.11	-0.16	0.00	0.02
$\frac{\beta dcdy>0}{\beta dcdy\leq 0}$	-0.41	0.19	0.20	-0.14	0.24	0.28	-0.05	0.03	0.11

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 ID model with two discretizations of the individual income process (13) that both use Rouwenhorst's (1995) method but differ in the number of support points n_y, equal to three and six, respectively. The income process for the rest of the village is approximated as a discretized AR(1) process with 5 support points.

A.5 Alternative outside options

In this section, we briefly investigate how the predictions of the CD model change when we consider alternative specifications of the outside option in the case of the village of Aurepalle. We first consider the case where individuals can immediately join new coalitions, rather than having to go through one period of autarky. We then look at an alternative specification of the outside option of the 'rest-of-village' in our approximation of the CD model.

A.5.1 Immediate coalition formation

Following Genicot and Ray (2003), in our benchmark setting we assumed that individuals go through one period of autarky before they can renegotiate and share risk with others in a smaller coalition (see discussion in Section 2.5). In this section we consider an alternative specification where deviating households can share risk starting from the first period.

Table 10 shows how this more attractive outside option reduces the maximum group size to 3, and has a reducing effect on the degree of insurance that is only partly offset by a higher estimated discount factor of $\delta = 0.96$. The asymmetry continues to be zero and the goodness of fit is essentially the same as in Section 4.4, as the better fit of the variance moment offsets the worse fit of the regression moment.

A.5.2 Second-best risk sharing in the 'rest-of-village'

The approximate solution of the CD model follows Ligon, Thomas and Worrall (2002) by studying a two-agent contract between an individual and the rest of an insurance group, and by approximating the rest of the group as a single agent in the solution of the constrained-efficient allocation (although not in the simulation, where the full vector of incomes is taken into account). This simplification corresponds to assuming, for the solution of the contract, that the rest of the village is able to share risk perfectly. This implies an inconsistency, as the households in the rest of the village are better able to share risk among themselves than with the individual

Table 10: Preferences estimated to target all 4 moments - no period of autarky

	Aurepalle	
	Data	**CD**
n		3.00
δ		0.96
s.e.		0.48
σ		1.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.32
β_{dcdy}	0.21	0.31
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.00
$\beta_{dcdy>0}-\beta_{dcdy\leq 0}$	-0.41	-0.00
Goodness of fit		13.66

Notes: For the village of Aurepalle, 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 CD model without a period of autarky following a deviation (in the second column). 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 $\delta \in [0.5, 0.98]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

under consideration. In principle, one could try to find a fixed point by iteratively adjusting the outside option for the rest of the village to be consistent with a candidate solution to the two-agent problem until the two imply the same degree of risk-sharing. This is, unfortunately, computationally too demanding for our estimation approach. The recursive solution of the CD model, however, where we sequentially solve for the contract in groups of $n = 2, 3, 4$ etc., offers an alternative approximation of rest-of-group utility. Specifically, it allows us to use as the outside option of the rest of the group the average expected utility that its members expect to get from a risk-sharing contract with $n - 1$ households. Making the same assumption as in our benchmark approximation that deviating households share the surplus of the group as equally as possible after one period of autarky consumption then allows us to calculate an alternative approximate solution of contracts in the CD model where there are the same frictions within the rest of the village as between it and the individual under consideration. Importantly, we would expect this outside option to improve risk sharing between the individual and the rest the group, whose less attractive outside option of constrained (as opposed to perfect) risk sharing corresponds to a relaxation of its participation constraints relative to the benchmark model.

Table 11 presents the results for the village of Aurepalle. As expected, insurance is improved. This is mainly, however, due to an increase in the maximum sustainable group size to 6 households, at an unchanged estimate of the discount factor, equal to 0.94. This is because when the participation constraint of the rest of the group is relaxed, a group can provide better insurance to the individual, and is thus more likely to meet participation constraints. So larger groups become sustainable. We leave an in-depth study of this alternative model to future research.

Table 11: Preferences estimated to target all 4 moments - second-best outside option for the rest of the village

	Aurepalle	
	Data	**CD**
n		6.00
δ		0.94
s.e.		0.01
σ		1.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.15
β_{dcdy}	0.21	0.17
$\frac{Var_{dc\|dy>0} - Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.01
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.41	0.03
Goodness of fit		21.24

Notes: For the village of Aurepalle, 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 CD model where the outside option of the the rest of the village in the approximation does not assume perfect risk sharing as detailed in the main text (in the second 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 $\delta \in [0.5, 0.98]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

A.6 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 consumption and income data both from the ICRISAT panel and from model simulations on movements in aggregate resources, using residuals from a regression on time dummies (or by demeaning the sample period-by-period). As we discussed in Section 4.3, 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. Demeaning the sample period-by-period, 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. This section therefore presents estimates of the CD, ID and SI models using raw data.

Table 12: Estimated income processes

	Aurepalle	**Kanzara**	**Shirapur**
ρ	0.20	0.039	-0.17
Var_{α_i}	0.29	0.27	0.36
Var_{ϵ}	0.19	0.079	0.12

Notes: The table presents the point estimates for the persistence parameter ρ and the shock variance Var_{ϵ} for the AR(1) process (13) using raw data.

Table 12 reports estimates of the income process (13) when using raw income data, which mechanically increases the variance of shocks but has an ambiguous effect on the AR(1) parameter ρ. The estimated process is virtually identical to the benchmark in Shirapur. Raw incomes are slightly less (more) persistent in Aurepalle (Kanzara). The increase in the variance of shocks is most pronounced in Aurepalle.

In Table 13 we repeat our benchmark estimation based on 'raw' data, rather than demeaning the sample period-by-period. This makes little difference to the degree of insurance measured in the three villages, but reduces the negative asymmetry observed in all three villages. The estimated group sizes are unchanged by the slightly different income process, with the exception of Kanzara, where groups of 5 are now sustainable. Together with a slight increase in the estimated discount factor, this increases the predicted degree of insurance there. The degree of insurance in the other two villages, and the prediction of symmetry, are essentially identical to the benchmark estimation. The predictions of the two comparison models change only very little when using raw data for income process and target moments. Importantly, the prediction of too much insurance and asymmetry in the opposite direction relative to the data remain.

Table 13: Preferences estimated to target all 4 moments (raw data)

	Aurepalle				Kanzara				Shirapur			
	Data	**CD**	**ID**	**SI**	**Data**	**CD**	**ID**	**SI**	**Data**	**CD**	**ID**	**SI**
n		4.00	34.00			5.00	37.00			5.00	31.00	
δ		0.92	0.88	0.93		0.95	0.90	0.95		0.92	0.87	0.93
s.e.		0.01	0.01	0.01		0.01	0.01	0.01		0.01	0.01	0.01
σ		1.00	1.00	1.00		1.00	1.00	1.00		1.00	1.00	1.00
$\frac{Var_{dc}}{Var_{dy}}$	0.33	0.24	0.03	0.10	0.61	0.19	0.06	0.07	0.37	0.18	0.03	0.08
β_{dcdy}	0.22	0.25	0.08	0.23	0.18	0.22	0.14	0.17	0.15	0.21	0.10	0.19
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.01	0.01	0.04	0.06	-0.17	0.01	0.08	0.04	-0.14	0.01	0.05	0.05
$\beta_{dcdy>0}-\beta_{dcdy\leq 0}$	-0.39	0.01	0.15	0.20	-0.02	0.02	0.24	0.16	0.01	0.03	0.18	0.19
Goodness of fit		17.02	57.17	45.53		5.91	12.67	10.20		6.94	19.94	16.34

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 raw data of the survey ("Data", in the first column for each village), and in simulations using the income process estimated with unconditional data (see Table 12) in the CD model (in the second column for each village), the ID model (third column) and the SI model (fourth column). For the simulated model solutions, the table also presents the size of the insurance groups n (in the case of the ID and CD models) 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 $\delta \in [0.5, 0.98]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

A.7 The role of income persistence

The persistence of income shocks is a key parameter in both the CD and the ID model. This is because in both models, the value of the outside option that enters the participation constraint depends on income persistence. In the ID model this outside option is autarky, which becomes more attractive for high income households when incomes are more persistent, implying that future incomes are expected to also be high. At the same time, higher persistence reduces the probability of receiving transfers in the future. Together, this makes a deviation more attractive for high income individuals when persistence is high. In the CD model there is an additional effect, as the persistence of incomes changes the relative attractiveness of high vs. low income villagers as partners in a deviating coalition. Thus, when incomes are persistent, those with high income today are attractive partners as they are likely to also be income-rich tomorrow, and thus able to provide consumption resources for others with low income draws. With negative serial correlation, the same is true for villagers with low income (who, however, may not be willing to deviate as they are receivers of transfers today, and are expected to be income rich only 'every other' period).

In our benchmark calibration, we estimated the income process directly from the data. This resulted in persistence that was small in absolute magnitude for all villages. We argued why we chose this benchmark process for incomes, based on an estimation that allowed households to differ in their mean incomes, but assumed common persistence. Alternative approaches, such as an estimation for different income groups (Laczo, 2014), household-specific income processes (Ligon, Thomas and Worrall, 2002), or a homogeneous income process for all villagers (as in the working paper version of this paper, Bold and Broer (2016)), have also been used in previous work, and imply different estimates of income persistence. In this section we study whether the standard ID model is able to capture the observed income-consumption distribution better when we choose the value of the persistence parameter ρ that best matches the observed consumption-income comovement in data. Intuitively, one might expect negative serial correlation in incomes to reduce the asymmetry of the joint distribution of consumption and income in that model by flattening the relationship between income and autarky values. For each of the three villages, and given a common village-specific cross-sectional dispersion of incomes that we take from the data, we therefore estimate the value of the persistence parameter ρ freely on a grid between -0.6 and $+0.6$ to best match our four target moments.

Table 14 reports the results. The CD model predicts positive persistence and estimates a higher discount factor, which interact to leave the predicted moments approximately unchanged, apart from a fall in the maximum sustainable group size in Shirapur by 1 household that brings insurance there closer to the data. The fact that serial correlation does not much affect the moments predicted by the CD model is also evidenced by the large uncertainty surrounding the estimate of ρ. At an estimated value of ρ equal to the lower bound of -0.6 in all three villages, and a discount factor δ substantially lower than in our benchmark results in Table 5, the ID model now predicts a degree of insurance that is lower, and thus closer to the data, at the price of a somewhat increased asymmetry. The overall goodness of fit is only marginally improved, however. So even with a serial correlation that is counterfactually negative relative to the moments we find in the data, the ID model is not able to simultaneously predict a symmetric distribution of consumption and income growth and a realistic observed degree of insurance.

A.8 Preference heterogeneity in the ID model

The estimates in Table 5 suggested that the benchmark specification of the ID model, where all households were assumed to have identical preferences, was not able to explain the degree of risk-sharing observed in the data at the same time as approximately symmetric consumption-income comovement. Heterogeneity in preferences, in contrast, is, rather trivially, able to reconcile a realistic degree of insurance with symmetry in consumption in the ID model. 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 comovement and symmetry in our moments-based approach. With a less extreme degree

Table 14: Preferences and persistence estimated to target degrees of risk sharing and asymmetry

	Aurepalle			**Kanzara**			**Shirapur**		
	Data	**CD**	**ID**	**Data**	**CD**	**ID**	**Data**	**CD**	**ID**
n		4.00	34.00		4.00	37.00		4.00	31.00
δ		0.97	0.74		0.96	0.84		0.98	0.83
s.e.		0.10	0.11		0.05	0.08		0.04	0.12
σ		1.00	1.00		1.00	1.00		1.00	1.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.24	0.08	0.56	0.24	0.08	0.33	0.24	0.03
β_{dcdy}	0.21	0.25	0.20	0.22	0.25	0.20	0.17	0.23	0.10
$\frac{Var_{dc\|dy>0} - Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.01	0.10	-0.25	0.01	0.10	-0.16	0.00	0.04
$\beta_{dcdy>0} - \beta_{dcdy\leq 0}$	-0.41	0.00	0.24	-0.14	0.01	0.24	-0.05	0.00	0.15
ρ		0.60	-0.60		0.20	-0.60		0.60	-0.60
s.e.		1.89	0.44		3.59	0.57		5.36	0.73
Goodness of fit		12.82	48.44		5.74	14.79		6.49	26.77

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 CD model (in the second column for each village) and the ID model (third column). For the simulated model solutions, apart from the size of the insurance groups n, the table also presents the discount factor δ and the value of income persistence, which 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 $\delta \in [0.5, 0.98]$ and $\rho \in [-0.6, 0.6]$ and the goodness of fit reported is the value of the criterion function at the chosen parameters.

of preference heterogeneity, in contrast, the asymmetry predicted by the standard model may actually be stronger than in the benchmark. This is because when insurance is strong, as in the ICRISAT data, introducing dispersion in, for example, risk aversion around its estimated mean moves more risk-averse households even closer to perfect insurance, which may not change their consumption moments much. Insurance declines, however, for less risk-averse households, who care less about insurance and whose participation constraints thus bind more often, and whose consumption declines faster when they are unconstrained. This implies a decrease in insurance, and an increased asymmetry in their consumption process (as they move up the right hand side of the inverse U-shaped relation between asymmetry and the degree of risk sharing). At the strong insurance predicted by the standard model, this increase in the asymmetry for the less risk-averse may dominate the decline for more risk-averse households (and the increased dispersion of consumption growth for the unconstrained).

Table 15 compares our benchmark estimates to an alternative where villages consist of an equal number of villagers with risk aversion coefficients equal to 0.5 and 1.5 (implying a number of villagers divisible by two, which reduces those in Kanzara and Shirapur by 1). The fall in insurance implied by the presence of agents that are substantially less risk averse is counteracted by an increase in the estimated discount factor δ in all villages. The heterogeneity acts to increase the relative variance of consumption growth, while the increased discount factor decreases the regression coefficient β_{dcdy} and the corresponding asymmetry moments. Overall, the estimates are changed little by the inclusion of this stylised form of heterogeneity.

Table 15: Risk sharing moments for heterogeneous preferences

	Aurepalle			**Kanzara**			**Shirapur**		
	Data	**ID**	**ID Het**	**Data**	**ID**	**ID Het**	**Data**	**ID**	**ID Het**
n		34.00	34.00		37.00	36.00		31.00	30.00
δ		0.89	0.92		0.90	0.92		0.88	0.92
s.e.		0.01	0.01		0.01	0.02		0.01	0.01
σ		1.00	1.50		1.00	1.50		1.00	1.50
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.05	0.06	0.56	0.07	0.09	0.33	0.02	0.03
β_{dcdy}	0.21	0.12	0.11	0.22	0.17	0.16	0.17	0.07	0.07
$\frac{Var_{dc\|dy>0}-Var_{dc\|dy\leq 0}}{Var_{dy}}$	-0.08	0.07	0.09	-0.25	0.11	0.12	-0.16	0.03	0.04
$\beta_{dcdy>0}-\beta_{dcdy\leq 0}$	-0.41	0.20	0.19	-0.14	0.28	0.25	-0.05	0.14	0.13
Goodness of fit		52.76	50.64		16.80	16.10		27.78	26.76

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 ID model (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.

A.9 Measurement error in consumption and incomes

We now examine the extent to which the introduction of measurement error may change our results. As argued in section 4.4, measurement error in consumption may help the models to reconcile a sizeable volatility of measured consumption with a modest comovement of consumption and income. Measurement error in incomes, on the other hand, has three main effects on the moments we focus on: first, it attenuates the slope coefficient in a regression of consumption growth on income growth. Second, when more of the measured variance of incomes comes from random error, the true income process becomes less variable and its autocorrelation increases in absolute magnitude. This is because in order to predict a given measured autocovariance, which is unaffected by classical measurement error, less volatile income shocks have to be more persistent. Finally, measurement error "blurs" the asymmetry moments that condition on the sign of income changes. Whenever the latter is mainly determined by measurement error, we would thus expect consumption moments to be identical for households with rising and falling incomes.

In this section we generalise our model by assuming that measured consumption, as well as income, include measurement error given by

$$\widehat{c_{it}} = c_{it} + \xi_{it} \tag{16}$$

$$\widehat{y_{it}} = y_{it} + \nu_{it} \tag{17}$$

where $\widehat{x_{it}}$ and x_{it} are the measured and true levels of the logarithm 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:

$$\rho = \frac{Cov_{\widehat{y}}}{Var_{\widehat{y}} - Var_\nu} \tag{18}$$

$$Var_\epsilon = (Var_{\widehat{y}} - Var_\nu) * (1 - \rho^2). \tag{19}$$

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, as mentioned above.

Table 16 presents the results. Estimates of measurement error in both consumption and income are substantial in all models. Specifically, measurement error in consumption, which accounts for between 37 and 94 percent of the variance of measured consumption growth, allows all models to fit the relative variance of consumption and income growth $\frac{Var_{dc}}{Var_{dy}}$ almost perfectly.[39] The improved fit comes at the cost of reducing the empirical content of the models: since consumption measurement error leaves all other moments unchanged, it effectively removes the relative variance of consumption and income growth from the objective function.

Measurement error in incomes is also estimated to be substantial and equals the maximum value of our grid, where it accounts for two thirds of the variance in measured income growth, in 7 of the 9 cases we consider. This results in a strongly improved fit of the comparison models. In the SI model, in particular, it reduces the asymmetry and attenuates the regression coefficient β_{dcdy}, bringing it in line with values observed in the data. Measurement error also reduces the asymmetry in the ID model, which, however, continues to predict a counterfactually low coefficient β_{dcdy}.

Generally, the estimates in Table 16 imply that the joint distribution of consumption and

[39]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 16: Preferences estimated to target all 4 moments and measurement error

	Aurepalle				Kanzara				Shirapur			
	Data	**CD**	**ID**	**SI**	**Data**	**CD**	**ID**	**SI**	**Data**	**CD**	**ID**	**SI**
n		2.00	34.00			2.00	37.00			3.00	31.00	
δ		0.93	0.95	0.91		0.93	0.54	0.88		0.88	0.85	0.82
s.e.		0.41	0.27	0.05		1.59	N.I.	0.19		0.48	0.83	0.12
σ		1.00	1.00	1.00		1.00	1.00	1.00		1.00	1.00	1.00
ϕ				0.00				0.00				0.00
$\frac{Var_{dc}}{Var_{dy}}$	0.30	0.30	0.31	0.30	0.56	0.54	0.53	0.56	0.33	0.33	0.34	0.32
β_{dcdy}	0.21	0.21	0.03	0.22	0.22	0.22	0.14	0.23	0.17	0.17	0.09	0.16
ΔV_{dc}	-0.08	0.00	0.01	0.04	-0.25	-0.00	0.09	0.03	-0.16	0.00	0.03	0.02
$\Delta\beta_{dc}$	-0.41	-0.00	0.04	0.06	-0.14	-0.00	0.11	0.05	-0.05	0.01	0.04	0.03
$\frac{Var_{d\epsilon}}{Var_{dc}}$		0.37	0.94	0.45		0.64	0.80	0.69		0.59	0.84	0.72
s.e.		0.26	0.39	0.30		1.28	0.09	0.08		0.12	0.09	0.20
$\frac{Var_{d\nu}}{Var_{dy}}$		0.64	0.65	0.65		0.65	0.60	0.65		0.65	0.64	0.64
s.e.		1.56	7.20	0.31		7.04	0.68	0.47		0.32	0.02	0.02
Goodness of fit		10.60	21.52	15.70		2.06	6.06	3.03		4.05	7.66	5.10

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 CD model (in the second column for each village), the ID model (third column) and the SI model (fourth column). 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 $\delta \in [0.5, 0.98]$. The goodness of fit reported is the value of the criterion function at the chosen parameters. The discount factors associated with the ID model in Kanzara are not identified (N.I.), as the model predicts autarky there.

income is dominated by measurement error, particularly in the case of the ID model. Apart from a counterfactually low coefficient β_{dcdy} this implies a counterfactually low fit of the consumption data in a Townsend (1994)-type regression of individual consumption growth on individual income growth and time fixed effects. Specifically, the R^2 of this regression in the ICRISAT data is 0.40, 0.31 and 0.20 for, respectively, Aurepalle, Kanzara and Shirapur. In contrast, that predicted by the ID model is 0.063, 0.067 and 0.081.

There is an additional reason why the results in Table 16, particularly for the CD and ID models, should be treated with caution, namely the extreme uncertainty surrounding several of the point estimates. In the ID model, for example, the discount factor δ is not identified in the case of Kanzara (where the model predicts autarky plus strong attenuation, and predicted moments are thus unchanged from small changes in δ). In the CD model, while consumption measurement error is clearly identified by the relative volatility of consumption and incomes, income measurement error is often not identified, as illustrated by the extreme standard errors surrounding some of its estimates. This is because the relative volatility of consumption and incomes is fitted perfectly by measurement error in consumption and the asymmetry moments are small and only little affected by income measurement error and changes in the discount factor. The regression coefficient β_{dcdy} can then often be matched in two ways: with a high discount factor, implying strong insurance in larger groups of size 4 and 5, and zero measurement error (as seen in Section 4). Or with substantial measurement error, which reduces the true volatility of incomes, making membership in large insurance groups less attractive, and thus raising the 'fundamental' regression coefficient by reducing group size. A further increase in measurement error can then often attenuate the measured β_{dcdy} to match the data. In fact, in our solution maximum group sizes decline monotonically with income measurement error, to 2 or 3 households at the high values estimated in Table 16. Our approach, which uses a grid of measurement error and discount factors does not capture this lack of identification exactly. But the goodness of fit in Table 16 is very similar to that of the CD model with measurement error only in consumption, equal to 11.5, 3.0 and 5.2, with maximum stable group sizes of 5, 5, and 6, in Aurepalle, Kanzara and Shirapur respectively. These group sizes can be viewed as an upper bound on the largest

Table 17: Transition probabilities for household incomes

	Aurepalle			Kanzara			Shirapur		
	y_1	y_2	y_3	y_1	y_2	y_3	y_1	y_2	y_3
y_1	0.4085	0.4613	0.1302	0.2487	0.5000	0.2513	0.1686	0.4840	0.3474
y_2	0.2306	0.5387	0.2306	0.2500	0.5000	0.2500	0.2420	0.5160	0.2420
y_3	0.1302	0.4613	0.4085	0.2513	0.5000	0.2487	0.3474	0.4840	0.1686

sustainable sizes, which would be reduced by measurement error in incomes, for which the four moments we focus on are, however, a poor guide in the case of the CD model.[40]

A.10 Transition matrices

Table A.10 reports the probabilities of moving from income state y_i to y_j for Aurepalle, Kanzara, and Shirapur.

A.11 Descriptive statistics

Table 18 presents descriptive statistics from the ICRISAT data.

Table 18: 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).

[40] It is tempting to compare these numbers to those in Table 16 using a Chi-Sqare distribution. But since the estimated income measurement error equals the bound of our grid, this would be invalid.