



EUI Working Papers

MWP 2010/17
MAX WEBER PROGRAMME

ESTIMATING DYNAMIC CONTRACTS:
RISK SHARING IN VILLAGE ECONOMIES

Sarolta Laczó

EUROPEAN UNIVERSITY INSTITUTE, FLORENCE
MAX WEBER PROGRAMME

*Estimating Dynamic Contracts:
Risk Sharing in Village Economies*

SAROLTA LACZÓ

This text may be downloaded for personal research purposes only. Any additional reproduction for other purposes, whether in hard copy or electronically, requires the consent of the author(s), editor(s). If cited or quoted, reference should be made to the full name of the author(s), editor(s), the title, the working paper or other series, the year, and the publisher.

ISSN 1830-7728

© 2010 Sarolta Laczó

Printed in Italy
European University Institute
Badia Fiesolana
I – 50014 San Domenico di Fiesole (FI)
Italy
www.eui.eu
cadmus.eui.eu

Abstract

This paper studies the role of preference and income risk heterogeneity when risk sharing is partial due to limited commitment. I estimate the dynamic contract determining self-enforcing insurance transfers in a structural manner, and allow the coefficient of relative risk aversion and multiplicative income risk to differ across households. I find that the model explains the consumption allocation significantly better than the homogenous version and the benchmarks of perfect risk sharing and autarky, using household survey data from rural Pakistan. Enforcement constraints bind more often with heterogeneous households, implying less risk sharing. The paper then examines how social policies would interact with existing informal insurance arrangements. I simulate the effects of counterfactual transfers targeting the poor on consumption by both eligible and ineligible households. In the case of a one-time transfer, the heterogeneous (homogeneous) model predicts that consumption by ineligible households increases by one fifth (one fourth) of the transfer; if the transfer is permanent, the predicted share of ineligible households is one tenth (one fifth).

Keywords

Risk sharing, limited commitment, preference heterogeneity, incomplete markets, structural microeconometrics, rural Pakistan

1 Introduction

Households living in rural areas of low-income countries face a great amount of risk. Revenue from agricultural production is usually low and volatile, as a result of extreme weather conditions, such as erratic monsoon rains in South Asia. Further, outside job opportunities are often lacking. In addition, financial instruments to insure against consumption fluctuations are rarely available. In such an environment, households in a community have to rely on one another for insurance.

There exists ample empirical evidence that households in poor villages do not fully share the risks they face, although they do achieve a remarkable amount of insurance without formal contracts (Townsend, 1994; Grimard, 1997; Dubois, 2000; Dercon and Krishnan, 2003a,b; and others).¹ The literature has focused on two imperfections to explain the observed partial insurance, namely private information (Wang, 1995; Ligon, 1998; Ales and Maziero, 2009) and lack of commitment. The latter is arguably the preferable assumption for small communities in developing countries, since households can observe shocks faced by their neighbors (bad harvest, or illness), but no authority exists to enforce informal risk-sharing contracts. The model of risk sharing with limited commitment has been developed by Thomas and Worrall (1988), Coate and Ravallion (1993), Kocherlakota (1996), and Ligon, Thomas, and Worrall (2002),² and its implications are supported by mounting empirical evidence (Fafchamps, 1999; Attanasio and Ríos-Rull, 2000; Foster and Rosenzweig, 2001; Ligon, Thomas, and Worrall, 2002; and others). In this paper, insurance transfers are required to be voluntary, or, self-enforcing.

This paper is most related to the work of Ligon, Thomas, and Worrall (2002), LTW hereafter. To my knowledge, that paper is the only one estimating dynamic contracts in a structural manner in general, and the model of risk sharing with limited commitment in particular. They do not, however, perform any statistical tests on parameters or model selection. The first contribution of this paper is to provide a test of perfect risk sharing, where the alternative is a well-specified model of partial insurance. I apply the likelihood ratio-based tests introduced by Vuong (1989).

Second, I allow for preference heterogeneity across households. This is an important extension, because efficient risk sharing has two main implications. First, incomes should be pooled. Second, less risk-averse households should bear more uninsurable risk (Borch, 1962; Wilson, 1968). Assuming that risk preferences are homogeneous, we exclude an additional motive for risk sharing. Several papers have introduced preference heterogeneity in the case of perfect risk sharing. Dubois (2000) specifies an isoelastic utility function, and allows the coefficient of relative risk aversion to depend on observables, as well as for multiplicative preference shocks. Mazzocco and Saini (2009) construct nonparametric tests of perfect risk sharing allowing for preference heterogeneity. The present paper tests perfect risk sharing against a well-specified alternative, but considers only parametric models. In particular, households' relative risk aversion coefficients may depend on wealth, which I instrument by land owned and household characteristics, because of endogeneity concerns.

I also investigate whether differences in income risk faced by households matter. Differences in income processes are rarely taken into account in the literature on risk sharing. Schulhofer-Wohl (2007) uses an experimental measure of risk aversion, and finds evidence that occupational choice is affected by risk preferences in the United States. He argues that this should be taken into account when evaluating how well people are able to mitigate the adverse effects of risk they face. I allow groups of households to face different multiplicative income risks, in addition to differences in

¹See also the seminal papers by Cochrane (1991) and Mace (1991) for tests of perfect risk sharing in the United States.

²See also Alvarez and Jermann (2000) for a decentralization of the constrained-efficient allocation, trading Arrow-Debreu securities with endogenous solvency constraints.

expected income, and the presence of idiosyncratic risk. I compare the model of risk sharing with limited commitment, with and without preference and income risk heterogeneity, to the benchmarks of perfect risk sharing and autarky. I statistically evaluate how well these models are able to explain the allocation of consumption, i.e. the sharing rule, taking income as exogenous.

Afterwards, this paper studies how social policies would affect the consumption allocation, taking into account the fact that social transfers from the government interact with existing informal insurance arrangements. Since the estimation is done in a structural manner, the effects of counterfactual social policies on the consumption allocation can be simulated. Recent evidence on the Progresa program in Mexico suggests that a conditional cash transfer program targeting the poor also increases consumption by ineligible households (Angelucci and De Giorgi, 2009). In particular, a transfer of 10 Mexican pesos to eligible households increased consumption by ineligible households by about 1 peso. The model of risk sharing with limited commitment implies this partial sharing of the transfer. In addition, using the structural estimation results of this paper, the policy effects on both eligible and ineligible households can be predicted ex ante.³ I simulate the effects of counterfactual transfers targeting the poor. I compare the policy effects when the transfer is assumed to increase consumption directly, and when it raises income. In this way, I quantify the mistake made in predicting policy effects, if informal insurance arrangements are ignored. This research may also provide guidance for the evaluation and design of redistributive policies or micro-insurance programs, by showing how to take into account existing informal arrangements to share risk.

The paper is also related to the literature on explaining consumption inequality given income inequality. Krueger and Perri (2006) study whether limited risk sharing due to enforcement constraints can account for the fact that within-group cross-sectional consumption inequality increased less than income inequality in the United States over the period 1980-2003. Their model provides the desirable qualitative predictions. However, it implies too much risk sharing when calibrated to US data.⁴ On the other hand, Blundell, Pistaferri, and Preston (2008) document that income shocks have become less persistent, and thereby easier to insure against. The present paper only allows for transitory shocks, but builds and estimates a structural model of how consumption is allocated, given income.

This paper first details the theoretical models of risk sharing, building on Kocherlakota (1996), LTW, Kehoe and Perri (2002), and others. Households are infinitely lived and risk averse, and are endowed with some exogenous, random income. Both time and uncertainty are discrete. The distribution from which incomes are drawn is common knowledge ex ante, as well as income realizations ex post. In other words, there are no information problems. The only way households in a village may mitigate the adverse effects of the risk they face is to insure one another. Perfect risk sharing means that all idiosyncratic risks are insured, and aggregate risk is borne efficiently. However, formal insurance contracts are often not available in rural villages in developing countries, and households may choose to revert to autarky.

The model of risk sharing with limited commitment characterizes the case where enforcement constraints may bind. This is arguably a good description of informal insurance arrangements among households in rural communities (Attanasio and Ríos-Rull, 2000; LTW), or members of a family (Mazzocco, 2007).⁵ The constrained-efficient solution can be found by maximizing a (utilitarian)

³See Todd and Wolpin (2006, 2008) on ex-ante program evaluation.

⁴Note that, in their quantitative analysis, Krueger and Perri (2006) consider an economy with production, while I consider an endowment economy, and focus on the sharing rule.

⁵The model has also been used in a wide variety of other economic contexts, including risk sharing between an employee and an employer (Thomas and Worrall, 1988), and countries (Kehoe and Perri, 2002). Further, Schechter (2007) adopts the same model to examine the interaction between a farmer and a thief, and Dixit, Grossman, and

social welfare function, subject to resource constraints and enforcement constraints. The presence of future decision variables in today's enforcement constraints makes the problem non-recursive. Introducing the Pareto-weights of households in the social planner's objective, equal to the ratio of marginal utilities in equilibrium, as a state variable, the problem has a recursive structure (Marcet and Marimon, 2009). The solution of the model is fully characterized by a set of state-dependent intervals on the relative Pareto-weights (LTW). The first best requires keeping the ratio of marginal utilities constant. However, when an enforcement constraint binds (or, equivalently, last period's ratio is outside today's interval), we cannot do that, and consumption reacts to idiosyncratic income shocks.

In section 3, the empirical models are set up. In particular, I model how the consumption allocation today is explained by past consumption, incomes today and in the past, and household characteristics, according to the different models of risk sharing. Maximum likelihood estimators are derived assuming that multiplicative errors in the measurement of consumption are log-normally distributed. The estimation of the perfect risk sharing and autarky models is straightforward. To deal with measurement error in the limited commitment case, a simulated estimator is used. In addition, the solution of this last model can only be approximated. Therefore, I compare the model of risk sharing with limited commitment with a given precision to the benchmark models of perfect risk sharing and autarky.⁶ Vuong's (1989) model selection tests are appropriate in this context, see Fernández-Villaverde, Rubio-Ramírez, and Santos (2006). In general, the tests' great advantage is that we do not have to assume that any model is correctly specified, and both nested and non-nested models can be compared.

Afterwards, section 4 presents the data. I use an income-consumption panel from rural Pakistan, collected by the International Food Policy Research Institute (IFPRI). Section 5 contains the estimation results for the structural models, both with and without preference and income risk heterogeneity, as well as the statistical tests to compare the models. I find that limitations in the enforcement of risk sharing contracts, preference heterogeneity, and heterogeneity in income risk faced by households significantly improve the model's ability to explain the consumption allocation. The coefficient of relative risk aversion is found to increase with household wealth, instrumented by land owned and household characteristics. The model of risk sharing with limited commitment with heterogeneous households can explain 89.4 percent of the cross-sectional variation of consumption on average in the six villages in the sample. Therefore, the concern of LTW, namely, that the model cannot explain the allocation of consumption across households well, is addressed.⁷ In the case with heterogeneous households, enforcement constraints bind more often, with 45 percent probability versus 30 percent when preferences and income risk are homogeneous, implying less risk sharing.

Policy simulations are presented in section 6. I look at the effects of (i) a one-time transfer and (ii) a permanent transfer on consumption by eligible and ineligible households. The poor are defined as having below median average consumption. They receive a counterfactual transfer that is about 25 percent of the average income of the poor. If policy makers mistakenly assume that the transfer directly translates into consumption, only eligible households' consumption should increase. At the other extreme, if households share risk perfectly, who has received the transfer is not important, and any differences in the consumption response should come from differences in risk preferences. In the case of a one-time transfer, the model of risk sharing with limited commitment predicts that

Gul (2000) use a similar model to examine cooperation between opposing political parties.

⁶Robustness checks are performed comparing models with different levels of precision in approximating the solution.

⁷In the case of the model of risk sharing with limited commitment with homogeneous preferences, the correlation of consumption predicted by the model and actual consumption varies between 53 percent and 76 percent for two type of estimators in the three Indian villages LTW examine.

consumption by ineligible households increases by about one fifth of the transfer, when heterogeneity in both risk preferences and income risk is taken into account. If the transfer is expected to be received in all subsequent periods, the share of ineligible households is predicted to decrease to one tenth. When households are assumed to be homogenous, about one fourth of the one-time transfer and one fifth of the permanent transfer goes to ineligible households.

In section 7, some extensions to the empirical models are discussed. In particular, I consider a utility function that depends on time-varying household observables, unobservable individual effects, and preference shocks. In addition, I allow for measurement error in income, as well as in consumption. Concluding remarks are presented in section 8.

2 Models of risk sharing

Suppose that there are N infinitely-lived, risk-averse households in a community. They consume a private and perishable consumption good c . Each household i maximizes its expected lifetime utility,

$$E_0 \sum_{t=1}^{\infty} \delta^t u_i(c_{it}),$$

where E_0 is the expected value at time 0 calculated with respect to the probability measure describing the common beliefs, $\delta \in (0, 1)$ is the (common) discount factor, and c_{it} is the consumption of household i at time t . The instantaneous preferences of household i are described by the isoelastic (CRRA) utility function

$$u_i(c_{it}) = \frac{c_{it}^{1-\sigma_i} - 1}{1 - \sigma_i}, \quad (1)$$

where $\sigma_i > 0$ is the coefficient of relative risk aversion of household i .

Suppose that random income, denoted Y_i for household i , is independently and identically distributed (i.i.d.) over time for each household, independently across households, and follows some discrete distribution F_{Y_i} for household i . Let s_t denote the state of the world that describes the income realizations of all households in the community at time t . The distribution F_{Y_i} , $\forall i$, is common knowledge ex ante, and so are income realizations ex post at each time t . That is, there are no informational problems. Note also that income is exogenous. In other words, the effect of risk on choices among different income generating processes is ignored. In addition, individual savings are assumed to be absent.⁸

This section considers three models in turn. First, it considers the model of perfect risk sharing. Second, subsection 2.2 mentions the benchmark of autarky. Third, the model of risk sharing with limited commitment is detailed in subsection 2.3.

2.1 Perfect risk sharing

To find the Pareto-optimal allocations, we solve the social planner's problem. The (utilitarian) social planner maximizes a weighted sum of households' expected lifetime utilities,

$$\max_{\{c_{it}(s_t)\}} \sum_i \lambda_i \sum_{t=1}^{\infty} \sum_{s_t} \delta^t \pi(s_t) u_i(c_{it}(s_t)),$$

⁸Ligon, Thomas, and Worrall (2000) allow for individual savings in the model of risk sharing with limited commitment. In this case, randomization is needed to make the problem convex.

where λ_i is the (initial) Pareto-weight of household i in the social planner's objective, and $\pi(s_t)$ is the probability of state s_t occurring, subject to the resource constraint

$$\sum_i c_{it}(s_t) \leq \sum_i y_{it}(s_t), \forall s_t, \forall t,$$

where $y_{it}(s_t)$ is the income of household i at time t and state s_t .

The well-known result that

$$\frac{u'_k(c_{kt}(s_t))}{u'_i(c_{it}(s_t))} = \frac{\lambda_i}{\lambda_k}, \forall s_t, \forall t, \quad (2)$$

that is, the ratio of marginal utilities for any two households i and k is constant over time and across states of the world, follows from the first order conditions of the social planner's problem (Borch, 1962; Wilson, 1968). Equation (2) implies that all idiosyncratic risks are insured away, and households share aggregate risk efficiently. In particular, less risk-averse households bear more uninsurable risk. With the utility function (1), condition (2), for any s_t and t , is

$$\frac{c_{kt}^{-\sigma_k}}{c_{it}^{-\sigma_i}} = \frac{\lambda_i}{\lambda_k}. \quad (3)$$

2.2 Autarky

When households are in autarky, the problem is trivial, since individual savings have been assumed absent. The model predicts that

$$c_{it}(s_t) = y_{it}(s_t), \forall s_t, \forall t, \forall i. \quad (4)$$

Let $U_i^{aut}(s_t)$ denote the expected lifetime utility, or, the value function, of household i in autarky at state s_t and time t . It can be computed by iterating the Bellman equation

$$U_i^{aut}(s_t) = u_i(y_{it}(s_t)) + \delta \sum_{s_{t+1}} \pi(s_{t+1}) U_i^{aut}(s_{t+1}). \quad (5)$$

In the present i.i.d. case, the value function can also be written as

$$U_i^{aut}(s_t) = u_i(y_{it}(s_t)) + \frac{\delta}{1-\delta} \sum_s \pi(s) u_i(y_i(s)).$$

2.3 Risk sharing with limited commitment

To find the constrained-efficient consumption allocations, I follow Kehoe and Perri (2002) (for the case of an endowment economy), and solve the following problem. The social planner maximizes a weighted sum of households' expected lifetime utilities,

$$\max_{\{c_{it}(s^t)\}} \sum_i \lambda_i \sum_{t=1}^{\infty} \sum_{s^t} \delta^t \pi(s^t) u_i(c_{it}(s^t)),$$

where $\pi(s^t)$ is the probability of history $s^t = (s_1, s_2, \dots, s_t)$ occurring, and $c_{it}(s^t)$ denotes the consumption of household i when history s^t has occurred; subject to the resource constraints

$$\sum_i c_{it}(s^t) \leq \sum_i y_{it}(s_t), \forall s^t, \forall t, \quad (6)$$

and the enforcement constraints,

$$\sum_{r=t}^{\infty} \sum_{s^r} \delta^{r-t} \pi(s^r | s^t) u_i(c_{ir}(s^r)) \geq U_i^{aut}(s_t), \forall s^t, \forall t, \forall i, \quad (7)$$

where $\pi(s^r | s^t)$ is the probability of history s^r occurring given that history s^t has occurred up to time t . The right hand side has been defined in equation (5). The problem is not recursive, because future decision variables enter into today's enforcement constraints. Therefore, even if income is i.i.d., consumption may depend on the whole history of income realizations.

The enforcement constraint (7) assumes that, if a household deviates, other households in the community do not enter into any risk sharing arrangement with it in the future. Note that reversion to autarky is the most severe subgame perfect punishment in this environment. In other words, it is an optimal penal code in the sense of [Abreu \(1988\)](#). We might also call reversion to autarky a trigger strategy, or the breakdown of trust. Future research should examine whether alternative specifications of the outside option would improve the model's fit to data. Alternatives include allowing for hidden storage, community punishment for renegeing, and limiting the time length of exclusion from insurance arrangements.

Denoting the multiplier on the enforcement constraint of household i (7) by $\delta^t \pi(s^t) \mu_i(s^t)$, and the multiplier on the resource constraint (6) by $\delta^t \pi(s^t) \rho(s^t)$, when history s^t has occurred, the Lagrangian is

$$\begin{aligned} & \sum_{t=1}^{\infty} \sum_{s^t} \delta^t \pi(s^t) \left[\sum_i \lambda_i u_i(c_{it}(s^t)) \right. \\ & + \mu_i(s^t) \left(\sum_{r=t}^{\infty} \sum_{s^r} \delta^{r-t} \pi(s^r | s^t) u_i(c_{ir}(s^r)) - U_i^{aut}(s_t) \right) \\ & \left. + \rho(s^t) \left(\sum_i y_{it}(s_t) - c_{it}(s^t) \right) \right]. \end{aligned}$$

Using the ideas of [Marcet and Marimon \(2009\)](#), the Lagrangian can also be written in the form

$$\begin{aligned} & \sum_{t=1}^{\infty} \sum_{s^t} \delta^t \pi(s^t) \left[\sum_i M_i(s^{t-1}) u_i(c_{it}(s^t)) \right. \\ & \left. + \mu_i(s^t) (u_i(c_{it}(s^t)) - U_i^{aut}(s_t)) + \rho(s^t) \left(\sum_i y_{it}(s_t) - c_{it}(s^t) \right) \right], \end{aligned}$$

where $M_i(s^t) = M_i(s^{t-1}) + \mu_i(s^t)$ with $M_i(s^0) = \lambda_i$ (see also [Kehoe and Perri, 2002](#)). In words, $M_i(s^t)$ is the initial weight of household i plus the sum of the Lagrange multipliers on its enforcement constraints along the history s^t .

The first order condition with respect to $c_{it}(s^t)$ can be written as

$$M_i(s^t) u'_i(c_{it}(s^t)) - \rho(s^t) = 0. \quad (8)$$

There are also standard first order conditions relating to the resource and enforcement constraints, with complementarity slackness conditions. To illustrate this, let us consider two households sharing

risk, households i and k . Combining the first order conditions (8) for these two households for history s^t at time t , we have

$$\frac{u'_k(c_{kt}(s^t))}{u'_i(c_{it}(s^t))} = \frac{M_i(s^t)}{M_k(s^t)} = \frac{\lambda_i + \mu_i(s^1) + \mu_i(s^2) + \dots + \mu_i(s^t)}{\lambda_k + \mu_k(s^1) + \mu_k(s^2) + \dots + \mu_k(s^t)} \equiv x_i(s^t), \quad (9)$$

where $x_i(s^t)$ can be thought of as the relative Pareto-weight assigned to household i when history s^t has occurred, normalizing the weight of household k to 1 at each time t .

The vector of relative weights $x(s^t)$, with elements $x_i(s^t)$ defined in (9), can be used as a state variable in order to rewrite the problem in a recursive form (Marcet and Marimon, 2009). The current income state s_t does not tell us everything we need to know about the past. Only (s_t, x_{t-1}) does this, where x_{t-1} is the vector of relative weights, equal to the ratio of marginal utilities, inherited from the previous period. In other words, x_{t-1} is a sufficient statistic for everything that happened in the past. Denote by x_{it} the new relative weight of household i to be found at time t . The solution consists of policy functions for the consumption allocation and the new relative weight, with support over the extended state space (s_t, x_{t-1}) . That is, $c_{it}(s_t, x_{t-1})$, $\forall i$, and $x_t(s_t, x_{t-1})$ are to be determined. At last, the value functions can be written recursively as

$$V_i(s_t, x_{t-1}) = u_i(c_{it}(s_t, x_{t-1})) + \delta \sum_{s_{t+1}} \pi(s_{t+1}) V_i(s_{t+1}, x_t(s_t, x_{t-1})). \quad (10)$$

The solution is fully characterized by a set of state-dependent intervals on the relative weight x_i , that give the possible relative weights in each income state for household i (LTW). Denote the interval for household i for state s by $[\underline{x}_i^s, \bar{x}_i^s]$. Suppose that last period the ratio of marginal utilities was $x_{i,t-1}$, and today the income state is s . The relative weight of household i today is determined by the following updating rule (LTW):

$$x_{it} = \begin{cases} \bar{x}_i^s & \text{if } x_{i,t-1} > \bar{x}_i^s \\ x_{i,t-1} & \text{if } x_{i,t-1} \in [\underline{x}_i^s, \bar{x}_i^s] \\ \underline{x}_i^s & \text{if } x_{i,t-1} < \underline{x}_i^s \end{cases} \quad (11)$$

Numerical dynamic programming allows us to solve for the optimal intervals, and thereby the consumption allocation, given the income processes, utility functions, and discount rates of the households. As the number of periods tends to infinity, the initial weights in the social planner's objective only matters if perfect risk sharing is self-enforcing (Kocherlakota, 1996).

Remember that F_{Y_i} (F_{Y_k}) denotes the distribution from which household i 's (k 's) income is drawn, and that σ_i (σ_k) parametrizes the utility function of household i (k). Then, we can solve numerically for x_{it} given s_t (or, y_{it} and y_{kt}) and $x_{i,t-1}$, and the optimal intervals that depend on σ_i , σ_k , δ , F_{Y_i} , and F_{Y_k} . Details are in the Appendix. Once we know x_{it} , $\forall i$, the first order conditions (9) and the resource constraint (6) give the consumption allocation predicted by the model.

3 Empirical models

Let us first specify the utility function (1). Assume that σ_i is a linear function of observables. In particular,

$$\sigma_i = 1 + z'_i \beta,$$

where β is a parameter vector to be estimated, and z_i represents a vector of time-invariant observable covariates of household i . Note that z_i does not contain an (additional) constant, as in Dubois (2000).

A normalization is needed, because the consumption risk borne by each household is determined by its risk tolerance relative to the average risk tolerance in the community. Further, if the coefficient on the constant were a free parameter, then, taking all households as risk neutral, any consumption allocation would be Pareto optimal.⁹ Remember that the above theoretical models assume perfect information, thus the preferences of each household are known to everybody in the community, but the econometrician only observes $z_i, \forall i$.

Assume that consumption is measured with a multiplicative measurement error that is log-normally distributed. Let c_{it}^* denote consumption observed by the econometrician, and let $\exp(\varepsilon_{it})$ be the multiplicative measurement error in household i 's consumption at time t . Then, we may write

$$c_{it}^* = \exp(\varepsilon_{it}) c_{it},$$

where ε_{it} is independently and identically distributed (i.i.d.) across households and time, and $\varepsilon_{it} \sim N(0, \gamma^2)$, where γ^2 is to be estimated.¹⁰ Note that true consumption c_{it} is observed by all households in the community. Measurement error in income is ignored for now, and is introduced in an extension in section 7.

I model the allocation of observed consumption, i.e. the sharing rule, $c_t^* \equiv (c_{1t}^*, \dots, c_{it}^*, \dots, c_{Nt}^*)$, for $t = 2, \dots, T$, determined by the history of income realizations, time-constant household characteristics, observed consumption at time 1, c_1^* , and parameters. In mathematical terms, we would like to know how the following conditional density could be specified based on the above models of risk sharing:

$$f(c_T^*, \dots, c_2^* \mid c_1^*, y_T, \dots, y_1, Z; \beta, \delta, \gamma^2, F_Y, \lambda), \quad (12)$$

where y_t , for $t = 1, \dots, T$, is the vector of income realizations for households at time t , $Z = [z_1, \dots, z_i, \dots, z_N]'$ is the matrix of household observables for all households, $\theta = (\beta, \delta, \gamma^2, F_Y)$ are the structural parameters to be estimated¹¹ with $F_Y = F_{Y_1, \dots, Y_i, \dots, Y_N}$ being the joint distribution of households' income processes, and the vector λ is a nuisance parameter. Each of the above theoretical models allows us to factorize the density (12). In particular, we may write

$$\prod_{t=2, \dots, T} f(c_t^* \mid c_{t-1}^*, y_t, Z; \beta, \delta, \gamma^2, F_Y, x_{t-1}), \quad (13)$$

where x_{t-1} is the state variable, which has elements $x_{i,t-1}$ and is not observed. I deal with this issue below.

For the limited commitment case, LTW have shown that the updating rule (11) holds for both the 2- and the N -household case. In the empirical part of their paper, they approximate the N -household economy by looking at each household i sharing risk with the 'rest of the community.' I follow their approach in this paper. This results in important gains in computation time, since the N -household case would require solving the model with $N - 1$ state variables, the number of relative Pareto-weights. I often call the rest of the community household k . Household k can also be thought of as the chief of the community, coordinating transfers. I evaluate the models in terms

⁹This is because marginal utility is constant for risk-neutral households. Thus any consumption allocation would keep the ratio of marginal utilities constant.

¹⁰I have allowed for measurement error in consumption to account for the error term in our estimating equations. In the consumption insurance literature, preference shocks are often used to introduce randomness, or, as in Cochrane (1991), consumption growth is measured with error. These alternative assumptions are not suitable in the case of risk sharing with limited commitment, as will be explained below.

¹¹Below θ often denotes a subset of these parameters, and is used as a short form for 'structural parameters to be estimated.'

of how well they explain the allocation of consumption in each community at each time t , but take changes in aggregate consumption as exogenous. Equivalently, I study how well the models explain each household's consumption relative to mean consumption in the community.

An additional issue is how to specify the preferences of household k . Let us subtract the community mean from each observable in the utility function for each household. This then means that the preferences of a typical household in the community are described by an isoelastic utility function, with coefficient of relative risk aversion equal to 1. That is, $u_k(c_{kt}) = \log c_{kt}$. Normalize also the Pareto-weight of household k to 1, that is, $\lambda_k = 1$. This is without loss of generality, since only relative Pareto-weights matter. Further, I assume that c_{kt} is well measured, since the variance of the measurement error in mean consumption in the community is only a fraction of the variance of the measurement error in each household's consumption. This assumption is only for notational simplicity. Think of explanatory variables in the utility function as deviations from their community mean hereafter, abusing notation.¹²

When preferences are heterogeneous, I use the logarithm of household wealth to capture heterogeneity in risk preferences, and instrument it because of endogeneity concerns. That is, before estimating the main equations derived from the structural models, I regress the logarithm of wealth on the logarithm of the value of land owned at time 1 and some household characteristics (household size, education of the head of the household, etc.), and include the predicted log wealth of household i as z_i . The details of this first-stage estimation are given below.

The next three subsections detail in turn how the model of perfect risk sharing (subsection 3.1), autarky (3.2), and risk sharing with limited commitment (3.3) are estimated. The estimations are done using (simulated) pseudo maximum likelihood estimators, and Vuong's (1989) tests are applied to statistically compare the models. Subsection 3.4 expands on model selection.

3.1 Perfect risk sharing

In the case of perfect risk sharing, the current consumption allocation should only depend on current and not past exogenous variables. It depends neither on the discount factor, nor on the distribution from which incomes are drawn. However, it depends on the time-constant unobservables, $x_{t-1} = \lambda$, $\forall t$. That is, the state variable is constant and equal to the initial relative Pareto-weights in the social planner's objective. Thus (13) can be written as

$$\prod_{t=2, \dots, T} f(c_t^* | y_t, Z; \beta, \gamma^2, \lambda). \quad (14)$$

Further, c_t^* only depends on today's income realizations through aggregate income.

Let us consider household i and the average household, k . Taking the logarithm of the first order condition with respect to (true) consumption for these two households, equation (3), noting that $\sigma_k = 1$ and $\lambda_k = 1$, we obtain

$$\sigma_i \log c_{it} - \log c_{kt} = \log \lambda_i.$$

Replacing for σ_i and rearranging give

$$\log \left(\frac{c_{it}}{c_{kt}} \right) = -z_i' \beta \log c_{it} + \log \lambda_i. \quad (15)$$

¹²Note also that, when preferences are homogeneous, meaning $\beta = 0$ here, the coefficient of relative risk aversion is normalized to 1, that is, $u_i(c_{it}) = \log c_{it}$, $\forall i$.

In terms of measured consumption c_{it}^* (15) reads

$$\log \left(\frac{c_{it}^*}{c_{kt}} \right) = -z'_i \beta \log c_{it}^* + \log \lambda_i + (1 + z'_i \beta) \varepsilon_{it}.$$

Now, let us take first differences to eliminate $\log \lambda_i$. Doing so and rearranging yields

$$\log \left(\frac{c_{it}^*}{c_{kt}} \right) = \log \left(\frac{c_{i,t-1}^*}{c_{k,t-1}} \right) - z'_i \beta \log \left(\frac{c_{it}^*}{c_{i,t-1}^*} \right) + (1 + z'_i \beta) (\varepsilon_{it} - \varepsilon_{i,t-1}). \quad (16)$$

Estimating (16), I implicitly assume that the ratio of marginal utilities observed at time $t-1$ contains all the information available on λ_i .

Let $\psi^2(\theta) \equiv 2(1 + z'_i \beta)^2 \gamma^2$, and

$$d_{it}^{prs}(\theta) \equiv \left[\log \left(\frac{c_{it}^*}{c_{kt}} \right) - \log \left(\frac{c_{i,t-1}^*}{c_{k,t-1}} \right) + z'_i \beta \log \left(\frac{c_{it}^*}{c_{i,t-1}^*} \right) \right] / \psi(\theta).$$

Then, we may write the likelihood of observation it as

$$L_{it}^{prs}(\theta) = \phi(d_{it}^{prs}(\theta)),$$

where ϕ is the density of the standard normal distribution. Finally, the (pseudo) maximum likelihood estimator (MLE) maximizes

$$\ell^{prs}(\theta) = \sum_{i=1}^N \sum_{t=2}^T \log \phi(d_{it}^{prs}(\theta)), \quad (17)$$

with respect to θ , that is, β and the variance γ^2 . A preliminary step predicts household wealth, to be included as z_i , using land owned and household characteristics. The model is also estimated without preference heterogeneity for comparison. This means setting $\beta = 0$. Thus the only parameter that remains to be estimated is γ^2 .

I do not assume that the model is correctly specified, therefore I compute the variance-covariance matrix of the estimated parameters without assuming that the information matrix equality holds. I also take into account serial correlation. In particular, the variance-covariance matrix is estimated by $\hat{A}^{-1} \hat{B} \hat{A}^{-1}$, where \hat{A} is the estimated Hessian, that is,

$$\hat{A} = \sum_{i=1}^N \sum_{t=2}^T -\nabla_{\theta}^2 \ell_{it}(\hat{\theta}), \quad \text{and} \quad \hat{B} = \sum_{i=1}^N \sum_{t=2}^T \hat{s}_{it} \hat{s}'_{it} + \sum_{i=1}^N \sum_{t=2}^T \sum_{r \neq t} \hat{s}_{ir} \hat{s}'_{it},$$

where $\hat{s}_{it} = \nabla_{\theta} \ell_{it}(\hat{\theta})'$ is the score evaluated at the estimated parameters, and where the second term in the expression for \hat{B} accounts for serial correlation (Wooldridge, 2002). Both the first and second derivatives of the log-likelihood function can be computed analytically here.

3.2 Autarky

In the autarky case, there is no state variable. Further, preferences do not matter. Therefore, we may simply write the likelihood of the consumption allocation for $t = 2, \dots, T$ as

$$\prod_{t=2, \dots, T} f(c_t^* | y_t; \gamma^2). \quad (18)$$

Taking the logarithm of (4) and introducing measured consumption give

$$\log c_{it}^* = \log y_{it} + \varepsilon_{it}.$$

The consumption of household i relative to mean consumption is

$$\log \left(\frac{c_{it}^*}{c_{kt}} \right) = \log \left(\frac{y_{it}}{y_{kt}} \right) + \varepsilon_{it}, \quad (19)$$

where I have just added and subtracted $\log c_{kt} = \log y_{kt}$ to have the same dependent variable in the equation to be estimated as above. In terms of the allocation of consumption within a community, the autarky model says that the consumption share of household i should be the same as its income share.

Let

$$d_{it}^{aut}(\theta) \equiv (\log c_{it}^* - \log y_{it}) / \gamma.$$

The likelihood of observation it is

$$L_{it}^{aut}(\theta) = \phi(d_{it}^{aut}(\theta)),$$

and the log-likelihood function to be maximized is

$$\ell^{aut}(\theta) = \sum_{i=1}^N \sum_{t=2}^T \log \phi(d_{it}^{aut}(\theta)). \quad (20)$$

The only parameter to be estimated is γ^2 . I allow for misspecification and serial correlation when computing the variance of the estimated parameter.

3.3 Risk sharing with limited commitment

Remember that in the limited commitment case, the (true) ratio of marginal utilities from the last period x_{t-1} is the state variable in the recursive version of the model. It is a sufficient statistic for everything that happened in the past, including the initial condition. In other words, instead of conditioning on the whole history of income state realizations s^t , and the initial Pareto-weights in the social planner's objective λ , it is sufficient to condition on the current income state s_t and x_{t-1} . However, unlike in the perfect risk sharing case, the consumption allocation may also depend on the discount factor δ and the distribution from which incomes are drawn F_Y . Remember also that x_{t-1} is not observed in (13). Therefore, we have to condition on x_{t-1}^* with elements $x_{i,t-1}^* = (c_{i,t-1}^*)^{1+z_i'\beta} / c_{k,t-1}$, the observable ratio of marginal utilities at time $t-1$, instead of x_{t-1} with elements $x_{i,t-1} = (c_{i,t-1})^{1+z_i'\beta} / c_{k,t-1}$. With measurement error, all past values of consumption could be informative of $x_{i,t-1}$. For tractability, I only deal with the density (13).

Let us consider once again household i sharing risk with household k . According to the theoretical model of section 2.3, the first order condition can be written as

$$\log \left(\frac{c_{it}}{c_{kt}} \right) = -z_i'\beta \log c_{it} + \log x_{it}, \quad (21)$$

replacing x_{it} for λ_i in (15). According to the model of risk sharing with limited commitment, given x_{t-1} , preferences, current income realizations, and the distribution from which incomes are drawn,

we can solve numerically for the relative Pareto-weight of household i at time t . In mathematical terms, we can compute $x_{it}(s_t, x_{t-1})$, $\forall i$ (see equations (10) and (11)).

Let $g(\cdot)$ denote the function relating x_{it} to observables, parameters, and $x_{i,t-1}$. That is, $x_{it} = g(y_t, x_{i,t-1}, z_i; \theta)$, with $\theta = (\beta, \delta, F_Y)$. Note that, in general, $g(\cdot)$ cannot be expressed analytically, but its value can be computed given any set of argument values. Replacing for x_{it} in (21) gives

$$\log\left(\frac{c_{it}}{c_{kt}}\right) = -z'_i\beta \log c_{it} + \log g(y_t, x_{i,t-1}, z_i; \theta). \quad (22)$$

Remember that, instead of c_{it} , the econometrician observes $c_{it}^* = \exp(\varepsilon_{it})c_{it}$. Then, in terms of observable consumption (22) is

$$\log\left(\frac{c_{it}^*}{c_{kt}}\right) = -z'_i\beta \log c_{it}^* + \log g(y_t, x_{i,t-1}, z_i; \theta) + (1 + z'_i\beta)\varepsilon_{it}. \quad (23)$$

The first econometric issue is the identification of the discount factor δ . The second issue is that measurement error influences the updating of the state variable. That is, among the arguments of $g(\cdot)$ in equation (23), instead of $x_{i,t-1}$ only

$$x_{i,t-1}^* = (\exp(\varepsilon_{i,t-1}))^{1+z'_i\beta} x_{i,t-1} \quad (24)$$

is observed. I deal with these issues in the next two subsections.

3.3.1 Identifying the discount factor

In the perfect risk sharing case, the predicted consumption allocation was independent of δ , the discount factor. The question now is whether we can identify this parameter in the case of risk sharing with limited commitment, without observing the interest rate, and without considering intertemporal consumption decisions. Proposition 2 states that the answer is yes, if some but not perfect risk sharing occurs.

Before giving a formal proof, I illustrate the role of δ in determining the consumption allocation with an example. Let us consider two ex-ante identical households, whose preferences are described by the utility function $u(\cdot) = \log(\cdot)$. Both households face the random prospect $\tilde{Y} = (20, 1/2; 10, 1/2)$, and I assume that their incomes are perfectly negatively correlated. Thus, there are only two income states, $\{(y_{1t} = 20, y_{2t} = 10), (y_{1t} = 10, y_{2t} = 20)\}$, $\forall t$. Finally, assume that $\lambda = 1$. If agents stay in autarky, nothing is transferred in both states. If perfect risk sharing occurs, a transfer of 5 is made in both states, and both households consume 15. The model of risk sharing with limited commitment can predict any transfer between 0 and 5, and the exact amount depends on the households' patience, that is, on δ (see Figure 1). For example, if $\delta = 0.7$, the transfer is 1.0074; if $\delta = 0.75$, it is 2.5382; if $\delta = 0.8$, 4.0772 is transferred.

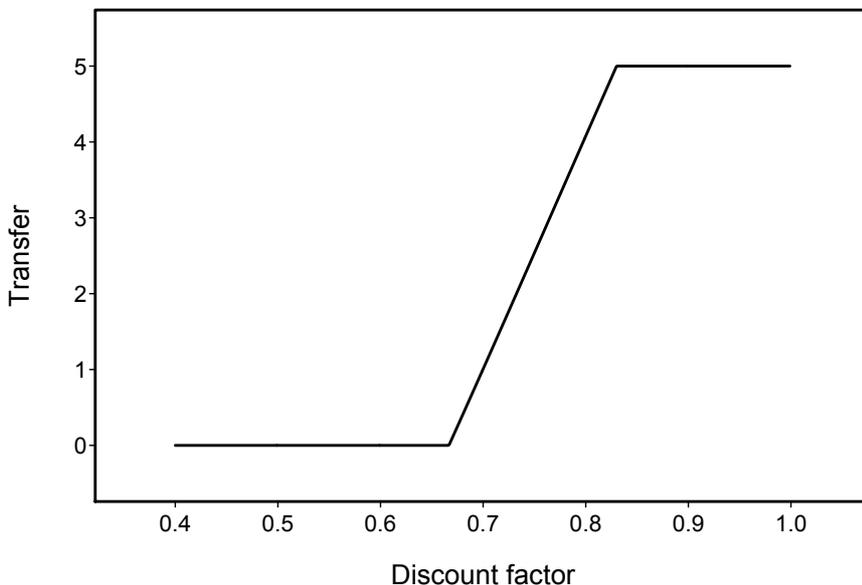
Denote by $\bar{\delta}$ the discount factor such that, $\forall \delta \geq \bar{\delta}$, perfect risk sharing is self-enforcing, and by $\underline{\delta}$ the discount factor such that, $\forall \delta \leq \underline{\delta}$, all households stay in autarky.¹³

Proposition 1. *The parameter δ is identified if $\delta \in (\underline{\delta}, \bar{\delta})$, that is, if some informal insurance is achieved and at least one enforcement constraint binds.*

Proof. Let us prove this for the case with homogeneous risk preferences. The argument for the heterogeneous case is similar. Compared to the perfect risk sharing case, additional information can

¹³LTW have shown that $\bar{\delta}$ and $\underline{\delta}$ exist.

Figure 1: The transfer from the household earning 20 today, to its risk sharing partner receiving 10, as a function of δ



only come from binding enforcement constraints. Suppose that at time t household i 's enforcement constraint is binding. That is, observed consumption is such that a positive transfer is made, but the ratio of marginal utilities is not the same as at time $t - 1$.¹⁴ Let us rewrite (26) with equality and with $z'_i\beta = 1 - \sigma = 0$. Simple algebra then gives

$$\log y_{it}(s_t) - \log c_{it}(s_t, x_{t-1}) = \delta \sum_{s_{t+1}} \pi(s_{t+1}) [V_i(s_{t+1}, x_t(s_t, x_{t-1})) - V_i^{aut}(s_{t+1})],$$

where the left hand side is the utility cost of the transfer household i makes today, and the right hand side is the welfare gain of sharing risk according to the informal insurance contract, rather than staying in autarky in the future. If the right hand side is strictly monotonic and continuous in δ , and only this constraint ever binds, we could perfectly match household i 's consumption at time t from the data, with an appropriately chosen unique δ .

The expected future gain of insurance is strictly increasing in δ , for $\delta \in (\underline{\delta}, \bar{\delta})$, since a higher δ relaxes all enforcement constraints. Note that as δ approaches 1, perfect risk sharing (the first best) is self-enforcing by the well-known folk theorem. On the other extreme, when it is close to 0, no voluntary transfers are made. In between, the higher δ is, the closer transfers get to their first-best level. That is, when δ is higher, more informal insurance is achieved, and consumption is smoother across income states. In other words, a higher δ means a better enforcement technology.

It is easy to see that $V_i^{aut}(s_{t+1})$ is continuous in δ . As for $V_i(s_{t+1}, x_t)$, LTW have shown that the limits of the optimal state-dependent intervals, which fully characterize the solution of the model,

¹⁴Considering the example above, a transfer strictly greater than 0, but strictly less than 5 is observed.

are continuous in δ . Since $V_i(s_{t+1}, x_t)$ is a continuous function of these limits, it is itself continuous in δ . It follows that one binding enforcement constraint identifies δ . \square

3.3.2 Measurement error

Let $\varepsilon_{i,t-1}^j$ denote a realization of measurement error in household i 's consumption at time $t-1$, drawn from the distribution $N(0, \gamma^2)$. Knowing $x_{i,t-1}^*$ and $\varepsilon_{i,t-1}^j$, we can easily compute $x_{i,t-1}$ (see equation (24)). Then, $g(y_t, x_{i,t-1}, z_i; \theta) = g(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta)$. To deal with the fact that measurement error enters the updating of the state variable, I first write the likelihood of each it observation conditional on $\varepsilon_{i,t-1}^j$. Then, averaging the conditional likelihood over J draws, I integrate $\varepsilon_{i,t-1}$ out. That is, in the case of risk sharing with limited commitment with measurement error, I use a simulated (pseudo) maximum likelihood estimator (SMLE).

Conditional on $\varepsilon_{i,t-1}^j$, (23) becomes

$$\begin{aligned} \log\left(\frac{c_{it}^*}{c_{kt}^*}\right) &= \log\left(\frac{\hat{c}_{it}(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta)}{\hat{c}_{kt}(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta)}\right) - z'_i \beta \log\left(\frac{c_{it}^*}{\hat{c}_{it}(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta)}\right) \\ &\quad + (1 + z'_i \beta) \varepsilon_{it}^j. \end{aligned} \quad (25)$$

Note that, when perfect risk sharing is self-enforcing, (25) is equivalent to (16). This is because

$$\begin{aligned} &\log\left(\frac{\hat{c}_{it}(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta)}{\hat{c}_{kt}(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta)}\right) + z'_i \beta \log \hat{c}_{it}(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta) \\ &= \log\left(\frac{c_{i,t-1}^*}{c_{k,t-1}^*}\right) + z'_i \beta \log c_{i,t-1}^* - (1 + z'_i \beta) \varepsilon_{it-1} \end{aligned}$$

in that case. Similarly, we get back the estimating equation of autarky, equation (19), if $\delta \leq \underline{\delta}$.

Using (5) and (10), and replacing for the utility function using (1) with $\sigma_i = 1 + z'_i \beta$, the enforcement constraint of household i at time t , that the predicted consumption allocation has to satisfy, can be written in a recursive form as

$$\begin{aligned} &\frac{c_{it}(s_t, x_{t-1})^{-z'_i \beta} - 1}{-z'_i \beta} + \delta \sum_{s_{t+1}} \pi(s_{t+1}) V_i(s_{t+1}, x_t(s_t, x_{t-1})) \geq \\ &\geq \frac{y_i(s_t)^{-z'_i \beta} - 1}{-z'_i \beta} + \delta \sum_{s_{t+1}} \pi(s_{t+1}) V_i^{aut}(s_{t+1}). \end{aligned} \quad (26)$$

This inequality is to be used in the numerical solution of the model, with

$$c_{it}(s_t, x_{t-1}) = \hat{c}_{it}(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta) \quad \text{and} \quad x_{it}(s_t, x_{t-1}) = g(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta).$$

Let us now consider alternative assumptions to generate residuals in the estimating equation (25). Preference shocks are not suitable, because today's shock would drop out of the equation. Assuming that consumption growth is measured with error, as in Cochrane (1991) for example, we could not take measurement error properly into account in the limited commitment case, since we would have to draw $\varepsilon_{i,t-1}^j$ from a random walk.

3.3.3 The simulated pseudo maximum likelihood function

Let $\psi^2(\theta) \equiv (1 + z'_i\beta)^2 \gamma^2$, and

$$d_{it}^{lc} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right) \equiv \left[\log \left(\frac{c_{it}^*}{c_{kt}} \right) - \log \left(\frac{\hat{c}_{it} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right)}{\hat{c}_{kt} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right)} \right) \right. \\ \left. + z'_i\beta \log \left(\frac{c_{it}^*}{\hat{c}_{it} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right)} \right) - (1 + z'_i\beta) \varepsilon_{i,t-1}^j \right] / \psi(\theta). \quad (27)$$

Then, the likelihood of observation it given $\varepsilon_{i,t-1}^j$ is

$$L_{it}^{lc}(\theta) = \phi \left(d_{it}^{lc} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right) \right).$$

Making J draws for $\varepsilon_{i,t-1}^j$, the simulated likelihood of observation it is

$$L_{it}^{lc}(\theta) = \frac{1}{J} \sum_{j=1}^J \phi \left(d_{it}^{lc} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right) \right).$$

Finally, the simulated (pseudo) log-likelihood function to be maximized is

$$\ell^{lc}(\theta) = \sum_{i=1}^N \sum_{t=2}^T \log \left(\frac{1}{J} \sum_{j=1}^J \phi \left(d_{it}^{lc} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right) \right) \right). \quad (28)$$

Allowing for misspecification, the SMLE consistently estimates the pseudo-true values of the parameters and is asymptotically normal, if both the number of it observations, that I denote by M , and the number of simulations J tend to infinity, and $\sqrt{M}/J \rightarrow 0$ (see Gouriéroux and Monfort, 1997, for example). When computing the variance-covariance matrix, the information matrix equality is not assumed to hold, and possible serial correlation is taken into account, as in the perfect risk sharing case. The score and the Hessian are computed numerically.

3.3.4 Estimation

The estimation is done in three steps. A preliminary step (i) involves regressing the logarithm of household wealth on the logarithm of the value of land owned at time 1 and some household characteristics to obtain predicted wealth, as in the case of perfect risk sharing, and estimating the distribution from which income is drawn, F_{Y_i} , for each household i . Then, (ii) the inside optimization computes the consumption allocation predicted by the model, given the observable covariates and parameters. Finally, (iii) the log-likelihood (28) is maximized over the remaining structural parameters, $\theta = (\beta, \delta, \gamma^2)$. Now I turn to the details of each of these steps.

(i) Estimation of the income processes. The discrete distribution from which income is drawn, F_{Y_i} , $\forall i$, has to be estimated. This cannot be done in general, because the time dimension of the panel is not large enough (maximum 12 for each household). First of all, I estimate a household-specific mean. Mean income can be thought of as a proxy for the income generating capacity of the household. Then, I divide income observations by household mean income, because it is well known that CRRA agents care about multiplicative risk, and estimate quantiles. I do this under

two alternative assumptions: (i) I assume that all households face the same multiplicative risk - the homogeneous income risk case; and (ii) I create deciles based on the coefficient of variation of household income, and estimate quantiles on the multiplicative risk separately for each group - the heterogeneous income risk case. For the rest of the community, quantiles are computed over mean income in the community. I permit 7 income states for each household, and 4 for the rest of the community.¹⁵ The income states for household i input in the model are the estimated quantiles over the multiplicative risk times mean income of household i .

(ii) Inside optimization. We have to solve the model of risk sharing with limited commitment to find the predicted consumption allocation. The Bellman equation (10) is solved by iteration. A grid is defined over the continuous state variable x_i . At iteration h , we solve for the new consumption values in states where an enforcement constraint is binding using (26) with equality, while the ratio of marginal utilities stays constant in other states. The values from iteration $h - 1$ are kept for $V_i(s_{t+1}, x_t)$ in (26). At the first iteration, the values of perfect risk sharing are used. We continue iterating until the policy function converges, that is, the optimal state-dependent intervals on x_i do not change. The algorithm to solve for the constrained-efficient risk sharing contract, given observables and structural parameters, does not impose much additional difficulty relative to the case without preference heterogeneity and heterogeneity in Y_i , except for computation time. Computation time is proportional to the number of households. The appendix gives more details on the algorithm. This step leads to the predicted consumption values, $\hat{c}_{it} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right)$ and $\hat{c}_{kt} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right), \forall i$, to be replaced in equation (25).

(iii) Outside optimization. The likelihood maximization is done, after a grid search, using a standard optimization algorithm available in R.¹⁶ The preliminary grid search is necessary, since we can only identify δ on the interval $(\underline{\delta}, \bar{\delta})$ (see Proposition 1). It also provides a good initial guess for the parameter values.

The maximum-likelihood estimates of the structural parameters (β , δ , and γ^2) are obtained by iterating between the dynamic program that solves for the predicted consumption allocation and the likelihood maximization routine. For comparison, the model is also estimated without preference heterogeneity. In that case only δ and γ^2 are to be estimated.

3.3.5 Simulation and approximation errors

The estimation of the model of the risk sharing with limited commitment involves both simulation and approximation. I take the number of simulations $J = 100$, because I use close to 1500 it observations, thus $J \gg \sqrt{M}$. Computation time is only moderately increased when adding additional draws, because the optimal intervals, which fully characterize the solution of the model, do not have to be recomputed.¹⁷

The continuous state variable x_i has to be discretized, and I use a 30-point grid.¹⁸ Computation time is approximatively proportional to the number of income states, which I take as $7 \times 4 = 28$, and the square of the number of grid points on x_i . Increasing the number of grid points would be beneficial in better approximating the true solution of the model. We are limited by the cost in

¹⁵I later perform a robustness check with 8 income states for each household.

¹⁶See www.r-project.org.

¹⁷Below I check that the results are robust with respect to changing J . In particular, I repeat the estimation of the model of risk sharing with limited commitment with heterogeneous households setting $J = 200$.

¹⁸A robustness check with 60 grid points verifies that the parameter estimates are not sensitive to changing the number of grid points.

terms of computation time. Additional approximation error may come from the fact that I limit the number of iterations when solving the model.

Using Vuong’s (1989) tests below, I allow for approximation errors, as a form of misspecification. We may say that what is compared is the limited commitment model with 30 grid points on x_i (instead of a continuous x_i). Fernández-Villaverde, Rubio-Ramírez, and Santos (2006) have recently proposed a way to test whether approximation errors are important when estimating computed dynamic models. They propose comparing models with different precisions, e.g. in terms of the number of grid points, using Vuong’s test. They provide Monte Carlo evidence that the method indeed works, assuming that the model estimated is the correct one. In this paper, I do not assume that any of the models are correctly specified, and it is possible that, in the case of risk sharing with limited commitment, a higher number of grid points would lead to a decrease in the model’s fit to data.

In a recent paper, Akerberg, Geweke, and Hahn (2009) argue that, in terms of the asymptotic properties of the maximum likelihood estimator, approximation error in computed dynamic models has similar effects as a limited number of simulations. In particular, they show that the difference between where a parameter estimate converges and the true value is of the same order of magnitude as approximation errors in the policy function. They also show that as approximation error converges to zero at any rate, the estimator is consistent. These results provide ground for asymptotic inference about the parameters of computed dynamic models that are estimated by maximum likelihood. To my knowledge, the small sample properties of such estimators have not been studied, and exploring this issue is an important task for future research.

3.4 Model selection

To statistically compare the above models, I use model selection tests introduced by Vuong (1989). Vuong proposes likelihood ratio-based statistics to compare nested and non-nested models. These statistics allow us to test the null hypothesis that two competing models are equally close to the true data generating process, against the alternative that one model is closer. Neither model has to be correctly specified.

The tests are based on the difference between the log likelihood values of the two models being compared. Suppose that we want to compare model 1 and model 2 using M observations. Denote the log likelihood of observation m for model 1 (2) at the estimated parameter vector by ℓ_m^1 (ℓ_m^2). The likelihood ratio is defined as

$$LR = \sum_{m=1}^M (\ell_m^1 - \ell_m^2).$$

Denote the number of parameters to be estimated by q_1 (q_2) for model 1 (2).

If the two models are non-nested, then, under the null hypothesis that the two models are equally close to the true data generating process,

$$\frac{LR}{\sqrt{M}\hat{\omega}} \Rightarrow N(0, 1),$$

where $\hat{\omega}$ is the estimated standard deviation of the likelihood ratio, that is,

$$\hat{\omega}^2 = \frac{1}{M} \sum_{m=1}^M (\ell_m^1 - \ell_m^2)^2 - \left(\frac{1}{M} \sum_{m=1}^M (\ell_m^1 - \ell_m^2) \right)^2,$$

and where \Rightarrow means convergence in distribution. If the two models are nested, and we want to allow for the possibility that the unconstrained model is not correctly specified, then, under the null,

$$2LR \Rightarrow M_{q_1+q_2}(\cdot; \hat{\kappa}),$$

where $M_{q_1+q_2}(\cdot; \hat{\kappa})$ is the cumulative distribution function of a weighted sum of $q_1 + q_2$ χ^2 distributions with degrees of freedom equal to 1 (Vuong, 1989).¹⁹ The p -values of the weighted χ^2 distribution are simulated. I do 100,000 replications.

I compare seven models: risk sharing with limited commitment with heterogeneous preferences and heterogeneous income risk (LC^{u_i, r_i}), with heterogeneous preferences but homogeneous income risk ($LC^{u_i, r}$), with homogeneous preferences but heterogeneity in income risk (LC^{u, r_i}), the homogeneous case ($LC^{u, r}$), and the benchmark models of perfect risk sharing with heterogeneous preferences (PRS^{u_i}), perfect risk sharing with homogeneous preferences (PRS^u), and autarky (AUT). LC^{u_i, r_i} nests LC^{u, r_i} and the benchmark models. $LC^{u_i, r}$ nests $LC^{u, r}$ as well as the benchmark models. LC^{u_i, r_i} and $LC^{u, r}$ nest PRS^u and AUT. PRS^{u_i} nests PRS^u . The remaining combinations are non-nested.

4 Data

The data come from an income-consumption survey conducted by the International Food Policy Research Institute (IFPRI) in rural Pakistan between July 1986 and September 1989. Almost 1000 households were interviewed over 12 rounds in 46 villages in 4 districts of Pakistan. The districts were not chosen randomly: 3 are the least-developed districts in their respective provinces (Attock in Punjab, Badin in Sind, and Dir in North-West Frontier Province), while the 4th is a more prosperous district (Faisalabad in Punjab). Then, in each district, two markets were chosen, and villages were randomly selected from a stratified sample based on distance from these markets. Finally, households were chosen randomly within each village. Attrition is mainly due to administrative problems, and not households' self-selection, and I assume that attrition is random. In each household, both the male and female head were interviewed. In addition, village questionnaires were administered, which give information about prices, for example. For further details see Alderman and Garcia (1993).

Due to computation time constraints, I only present results for Faisalabad district in Punjab. Pakistani Punjab has well developed factor and product markets, compared to the poor semi-arid areas on which much work on risk sharing in developing countries is based (Kurosaki and Fafchamps, 2002). Therefore, finding limits to risk sharing in Punjab makes it likely that similar constraints to insurance exist elsewhere.

This data set is attractive for the purposes of the present paper for several reasons. First of all, both the cross-sectional and time dimensions are relatively big compared to other similar data sets. Second, consumption data were collected from the female head of the household, while income data from the male head. Thus the assumption that measurement error in consumption is independent of the income measure is more compelling than usual. Further, both small and larger communities (villages and districts) could be examined in order to shed light on whether enforcement problems are more important in bigger communities. This is, however, left for future work.

¹⁹The weights $\hat{\kappa}$ can be computed by finding the real, nonzero eigenvalues of the matrix

$$\begin{bmatrix} -\hat{B}^1(\hat{A}^1)^{-1} & -\hat{B}^{1,2}(\hat{A}^2)^{-1} \\ \hat{B}^{2,1}(\hat{A}^1)^{-1} & \hat{B}^2(\hat{A}^2)^{-1} \end{bmatrix},$$

where $\hat{A}^1 = \sum_{i=1}^N \sum_{t=2}^T -\nabla_{\theta}^2 \ell_{it}^1$, $\hat{B}^1 = \sum_{i=1}^N \sum_{t=2}^T \hat{s}_{it}^1 \hat{s}_{it}^{1'}$, similarly for model 2, and $\hat{B}^{1,2} = \hat{B}^{2,1'} = \sum_{i=1}^N \sum_{t=2}^T \hat{s}_{it}^1 \hat{s}_{it}^{2'}$.

4.1 Variables used

For the purposes of this paper, we need measures of consumption, income, wealth, and some household characteristics. To measure consumption, I use nondurable consumption per adult equivalent. Nondurable consumption is constructed as the sum of food consumption and expenditure on clothing, hygiene items, tobacco, and cinema. To compute the adult equivalent household size I use the same age-gender weights as Townsend (1994).²⁰ Consumption is weekly and is expressed in 1986 Pakistani rupees per adult equivalent. In 1986 about 16 rupees were worth one US dollar, that is about two 2009 US dollars.

Income is constructed as the sum of net income from crop production, net income from poultry and livestock, net income from craft work, net income from produce from orchards, income from assets (hiring out bullocks, tractors, threshers, land, income from mills owned, and selling water), wage income minus wages paid for hired labor (that is not used in agricultural production), transfers from the state, and transfers from abroad, that is, all transfers from outside the community (transfers from within Pakistan from friends, relatives, or religious organizations, typically the local mosque, are thus excluded). Medical expenditures and educational investment are subtracted (as in Gourinchas and Parker, 2002). Income is then divided by the adult equivalent size of the household. Income is expressed in 1986 Pakistani rupees per week per adult equivalent, as is consumption.

I allow the coefficient of relative risk aversion to depend on wealth, measured as the logarithm of thousands of 1986 Pakistani rupees per adult equivalent. In mathematical terms,

$$z'_i\beta = \beta\text{wealth}_i.$$

Household wealth is the sum of the value of land, houses, other assets (such as TV sets, watches), tools, and livestock owned. These are measured at different dates, unfortunately mostly not at time 1, so I take a time-average, lacking a better alternative. I value all items using prices available in the survey, except for land. I value land starting from the median rent per acres for different types of land (rainfed, or canal irrigated, for example). Then I use data from Renkow (1993) that rents are 2 percent and 2.6 percent of land prices for rainfed and better-quality land respectively in the 1986-1989 period in Pakistani Punjab.

To deal with endogeneity concerns, I instrument wealth with land and exogenous household characteristics. **land** is the logarithm of the value of land owned per adult equivalent at time 1 plus 1. Let $I(\text{village}_l)$ equal 1 if household i lives in village l , and 0 otherwise. I include the predicted values from the linear regression

$$\begin{aligned} \text{wealth}_i = & \alpha_0 + \alpha_1\text{land}_i + \alpha_2\text{hhsz}_i + \alpha_3\text{age}_i \\ & + \alpha_4\text{gender}_i + \alpha_5\text{educ}_i + \sum_l \alpha_{6,l} I(\text{village}_l) + \text{residual}, \end{aligned} \quad (29)$$

as z_i in the estimating equations explaining the consumption allocation to capture differences in risk aversion across households.

Hhsz is the average adult-equivalent size of the household. **Age** is the age of the head of the household at time 1. Household heads are almost exclusively male in the sample, and I construct a measure of the gender composition of households. In particular, **gender** equals the time-average of the proportion of women among adults. To measure education, a categorical variable is used based

²⁰These weights are: 1 for adult males, 0.9 for adult females, 0.94 and 0.83 for males and females aged 13-18 respectively, 0.67 for children aged 7-12, 0.52 for children aged 4-6, 0.32 for children aged 1-3, and 0.05 for infants below 1 year of age.

Table 1: Descriptive statistics, Faisalabad district

Variable	Mean	Sd	Min	Max	Observations
Nondurable consumption ^a	781.42	836.88	54.15	9839.4	1850
Aeq. nondurable consumption	126.36	126.86	10.59	1430.0	1850
Log(aeq. nondurable consumption)	4.539	0.7455	2.360	7.265	1850
Income ^a	162.00	555.01	-1089.0	6168.7	1850
Aeq. income	25.50	86.86	-162.12	934.66	1850
Wealth ^b	215.21	324.09	1.8	2,415.6	168
Aeq. wealth	34.55	49.87	0.313	403.6	168
Log(aeq. wealth)	2.838	1.246	-1.163	6.001	168
Land ^c	129.93	241.40	0	1573.43	168
Aeq. land	21.19	37.86	0	313.83	168
Log(aeq. land+1)	1.806	1.774	0	5.752	168
Household size	9.042	4.444	1	41	1850
Aeq. household size	6.417	2.376	1	16.89	1850
Age of head ^d	50.51	13.85	18	80	168
Gender ^e	0.4815	0.1202	0	0.8333	168
Education ^f	1.839	1.205	1	5	168

^aMeasured in 1986 Pakistani rupees per week. 16 rupees = 1 US dollar in 1986.

^bMeasured in thousands of 1986 Pakistani rupees.

^cValue of land owned at time 1, measured in thousands of 1986 Pakistani rupees.

^dAs of the time of the first interview.

^eProportion of women among adults.

^fCategorical variable describing schooling achievement of the head of the household.

on the final schooling achievement of the head of the household. In particular, `educ` equals 1 if the head is illiterate, 2 if he has been to primary school or learnt to read in some other way, 3 or 4 if he has attended middle or secondary school (including technical studies) respectively, and 5 if he went to college or university. Finally, village dummies are also included.

I delete observation *it* if consumption is missing, income is missing, or any income component is outside some reasonable range. I delete households whose head is over 80 years of age, or if the head changes over the three-year period of the interviews. Thereafter, we are left with a sample size of 168 households in 6 villages and 1850 observations for Faisalabad district.²¹ Table 1 presents descriptive statistics.

On average, daily adult equivalent nondurable consumption is just over one 1986 US dollar, which is about two 2009 US dollars. Measured income is only a fraction of measured consumption, which reflects general underreporting, and the fact that agricultural production was hit hard by bad weather conditions during the years of the survey. Another sign of income underreporting is that mean adult equivalent wealth is about 26 times mean yearly adult equivalent income. Almost 60 percent of household heads are illiterate, and about 14 percent have schooling achievement above middle school.

There is scope for risk sharing in a community only if incomes are not highly correlated. The correlation coefficient of any two households' income per adult equivalent is 20.7 percent on average in Faisalabad district, while in the six villages the correlation coefficients are 30.8, 12.2, 39.2, 23.0,

²¹The villages are Saddoana (27 households), Singpura (27), Jaranwala-Chak (30) Gojra-Subadarwala (28) Khalisabad (29), and Sumundri (27).

31.5, and 23.3 percent. The relatively low correlation of incomes in these small communities is not surprising, given that households have many sources of income. In particular, the share of total income (before subtracting medical and educational expenditures) from crop production is 36.2 percent on average. Income from livestock makes up 13.1 percent, craftwork 10.4, produce of orchard 8.0 percent. Wages account for only 1.8 percent, income from assets 18.2, and transfers from outside the community 12.3 percent.

For the structural estimations, I delete the 5 percent extreme consumption and income observations in the pooled panel. Finally, I rescale aggregate income to be equal to aggregate consumption in the community at each t , since consumption aggregate changes are taken as exogenous.

4.2 Existing evidence

The present data set has been used by a number of papers to examine risk sharing. Dubois (2000) constructs a nondirectional test of perfect risk sharing based on overidentifying restrictions implied by the model. Allowing for preference heterogeneity, he is able to reject households sharing risk perfectly. He also shows that sharecropping contracts are used to achieve better insurance. Dubois, Jullien, and Magnac (2008) also provide evidence that perfect risk sharing is not achieved, and both short-term formal and informal insurance contracts are important. Foster and Rosenzweig (2001) find that limited commitment is important in explaining transfers, and altruism helps to relax enforcement constraints, and thereby improves risk sharing.

On the other hand, Ogaki and Zhang (2001) are not able to reject perfect risk sharing for the vast majority of villages examined, when relative risk aversion is not constant. They use both the ICRISAT data from India and the data from Pakistan used in this paper. The authors argue that earlier tests of perfect risk sharing do not take into account the possibility of decreasing relative risk aversion. The present paper also allows risk aversion to differ across households, and compares the perfect risk sharing model to a well-specified alternative, namely, the model of risk sharing with limited commitment.

5 Structural estimation and model selection results

I consider seven models: risk sharing with limited commitment with and without preference heterogeneity and with heterogeneous and homogeneous income risk (LC^{u_i,r_i} , $LC^{u_i,r}$, LC^{u,r_i} , and $LC^{u,r}$), perfect risk sharing with and without heterogeneous preferences (PRS^{u_i} and PRS^u respectively), and autarky (AUT); and for each model, I take a community to be a village. That is, aggregate consumption in the village is assumed to be exogenous, and I look at each household's consumption relative to the village mean. Vuong's (1989) tests are performed to statistically compare the seven models pairwise. The computations have been done using the software R (see www.r-project.org).

Before turning to the main results, some preliminary steps are necessary (see subsection 3.3.4). Namely, we have to regress wealth on land owned and household characteristics (see equation (29)). This gives the following:

$$\begin{aligned} \text{wealth}_i = & 1.968^{***} + 0.608^{***} \text{land}_i - 0.011 \text{hhsz}_i \\ & (0.293) \quad (0.026) \quad (0.017) \\ & - 0.001 \text{age}_i - 0.493 \text{gender}_i + 0.089^{**} \text{educ}_i, \\ & (0.003) \quad (0.332) \quad (0.040) \end{aligned}$$

for 168 observations. Standard errors are in parentheses. *** indicates significantly different from zero at the 1% level, and ** at 5%. Village dummies are also included, but the estimated coefficients are not presented.

Land and household characteristics explain 83.8 percent of the variation in wealth. Land has a significant and positive relationship with wealth. Education also increases wealth significantly. In particular, the head of the household reaching a higher education level raises wealth by about 9 percent. Below I include predicted log wealth of household i from the above regression as z_i when estimating the structural models. Figure 2 presents the histograms of log wealth and predicted log wealth. We cannot capture well the variation at the lower tail of the distribution, where landless households are found.

I now discuss the expected sign of the parameter β , the coefficient of wealth. First of all, note that, if the coefficient of relative risk aversion is higher, then the marginal utility of consumption is lower. This also means that the household bears less of the uninsured risk in the community. There is agreement in the literature on the fact that absolute risk aversion decreases with wealth. However, the effect of wealth on relative risk aversion is debated. Both Arrow (1965) and Pratt (1964) hypothesized that relative risk aversion increases with wealth, while empirical evidence on the matter is mixed (see Halek and Eisenhauer, 2001, and references therein).

The discrete distribution from which income is drawn for each household i has to be estimated as well, before estimating (25). I estimate a household-specific mean income and quantiles over multiplicative risk faced by households. I do the latter under two alternative assumptions: (i) households are homogeneous in terms of the income risk they face, or (ii) they are heterogeneous. In the latter case I estimate quantiles separately for each decile based on the longitudinal coefficient of variation of household income. In the homogeneous case the estimated 7-state multiplicative risk input in the structural model is (0.406, 0.538, 0.717, 0.960, 1.173, 1.360, 1.758).

5.1 Main results

Table 2 shows the structural estimation results for nondurable consumption for all the models considering households sharing risk within villages in Faisalabad district. The second panel shows the model selection test statistics, conducting Vuong's (1989) tests for nested and non-nested models, as appropriate.

The first panel of Table 2 shows that the coefficient of relative risk aversion is increasing with wealth. The coefficient β is positive and significant in all models with preference heterogeneity. This result is consistent with the hypothesis of Arrow (1965) and Pratt (1964) that relative risk aversion is increasing with wealth. This means that wealthier households bear less consumption risk.

In the homogeneous preferences case with limited commitment, introducing income risk heterogeneity raises the discount factor from 0.87 to 0.91. Accounting for income risk heterogeneity, the incentive to default on the informal risk sharing contract is found to be less. Interestingly, the percentage of times when an enforcement constraint binds goes up from 30.5 percent to 34.6 percent when income risk heterogeneity is taken into account.

The estimated discount factor decreases to 0.79 when preference heterogeneity is introduced. When both parts of risk sharing are taken into account, namely, income pooling and that less risk averse households should bear more uninsured risk, more incentives to default are in line with the observed consumption allocation. In other words, the sharing rule is further away from the first best. This can also be seen from the average percentage of times a household's enforcement constraint is

Figure 2: Histograms of log wealth and predicted log wealth

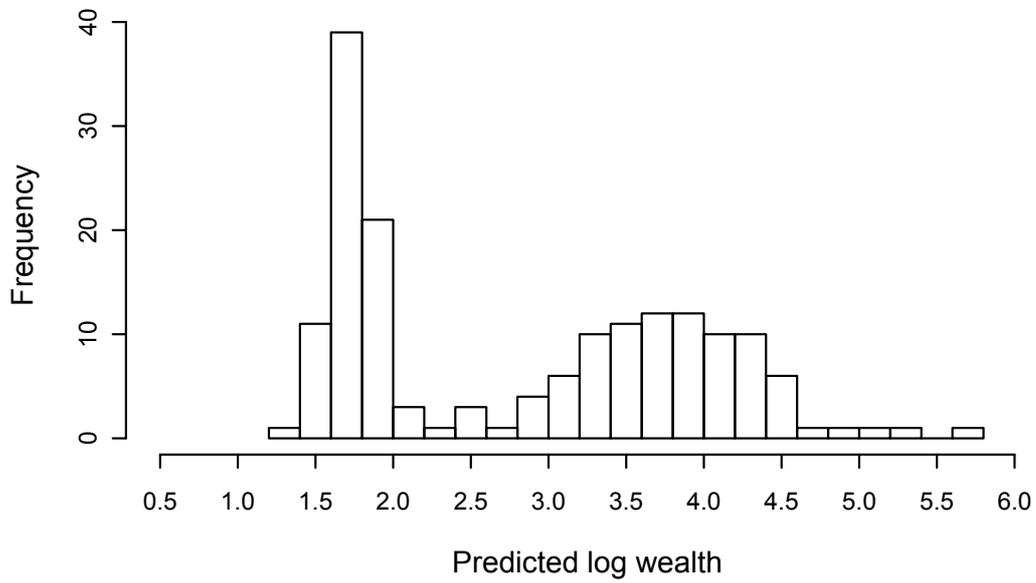
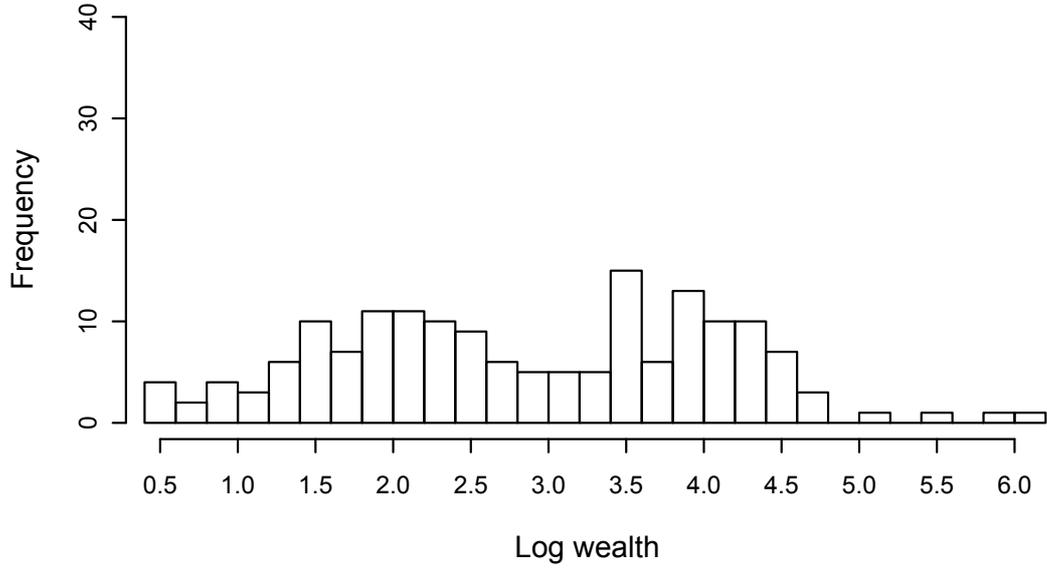


Table 2: Risk sharing within villages in Faisalabad district, adult-equivalent nondurable consumption

	LC^{u_i, r_i}	$LC^{u_i, r}$	LC^{u, r_i}	$LC^{u, r}$	PRS^{u_i}	PRS^u	AUT
β (preference heterogeneity)	0.510*** (0.002)	0.504*** (0.004)			0.305*** (0.038)		
δ (discount factor)	0.790*** (0.001)	0.785*** (0.004)	0.905*** (0.003)	0.865*** (0.003)			
γ^2 (var of measurement error)	0.235*** (0.010)	0.222*** (0.008)	0.169*** (0.007)	0.161*** (0.008)	0.153*** (0.006)	0.141*** (0.006)	0.276*** (0.011)
Log likelihood	-783.9	-790.9	-902.9	-904.2	-1122.2	-1152.9	-1136.8
Observations	1467	1467	1467	1467	1467	1467	1467
Vuong's tests							
$LC^{u_i, r}$	4.327*** (0.000)						
LC^{u, r_i}	235.5*** (0.000)	5.347*** (0.000)					
$LC^{u, r}$	5.322*** (0.000)	223.3*** (0.000)	1.226 (0.110)				
PRS^{u_i}	676.5*** (0.000)	662.6*** (0.000)	9.922*** (0.000)	10.13*** (0.000)			
PRS^u	738.0*** (0.000)	724.1*** (0.000)	502.6*** (0.000)	500.8*** (0.000)	61.51*** (0.000)		0.440 (0.330)
AUT	705.7*** (0.000)	691.8*** (0.000)	470.3*** (0.000)	468.5*** (0.000)	0.382 (0.351)		

Notes: LC^{u_i, r_i} ($LC^{u_i, r}$): equation (25) is estimated by maximizing the log-likelihood function (28) assuming heterogeneous (homogeneous) income risk when estimating F_{Y_i} , $\forall i$. LC^{u, r_i} ($LC^{u, r}$): equation (25) with $\beta = 0$ is estimated assuming that income risk is heterogeneous (homogeneous). PRS^{u_i} : equation (16) is estimated by maximizing the log-likelihood function (17). PRS^u : equation (16) with $\beta = 0$ is estimated. AUT: equation (19) is estimated by maximizing the log-likelihood function (20). In the first panel, standard errors are in parentheses. They have been calculated taken into account misspecification and serial correlation. In the second panel, p -values of Vuong's tests are in parentheses, indicating whether the model of the line can be rejected to be as close to the true generating process as the model of the column. In the case of nested models, the p -values are simulated. *** indicates significance at the 1% level.

binding: the constraint binds 38.4 percent of times for the $LC^{u_i, r}$ model, and 45.0 percent of times when households are heterogeneous in both risk preferences and income risk.

The second panel of Table 2 presents the results of the model selection tests. The test statistics presented indicate whether the model of the line can be rejected as being as close to the true generating process as the model of the column. The table shows that perfect risk sharing with homogeneous preferences (PRS^u) and autarky (AUT) perform equally badly. In terms of predicting today's consumption shares, using yesterday's consumption shares or today's income shares we achieve a similar fit to data. Considering risk sharing within villages, even the perfect risk sharing model with heterogeneous preferences (PRS^{u_i}) does not significantly outperform AUT.

Introducing limited commitment, we can explain the consumption allocation significantly better than the benchmark models. Further, the models of risk sharing with limited commitment with heterogeneous preferences (LC^{u_i, r_i} and $LC^{u_i, r}$) outperform all other models, including LC^{u, r_i} and $LC^{u, r}$. Therefore, I find evidence that both limited commitment and heterogeneity in risk preferences

Table 3: Robustness checks for the LC^{u_i, r_i} model

	200 simulations	8 income states	60 grid points ^a
β (preference heterogeneity)	0.511*** (0.002)	0.512*** (0.007)	0.507*** (0.001)
δ (discount factor)	0.788*** (0.001)	0.786*** (0.001)	0.799*** (0.001)
γ^2 (var of measurement error)	0.242*** (0.012)	0.238*** (0.012)	0.232*** (0.015)
Log likelihood	-784.5	-783.1	-774.6
Observations	1467	1467	1467

Notes: Equation (25) is estimated by maximizing the log-likelihood function (28) assuming heterogeneous income risk when estimating $F_{Y_i}, \forall i$. Standard errors are in parentheses. They have been calculated taken into account misspecification and serial correlation.*** indicates significance at the 1% level.

^aIn this case, I increase the precision in approximating the policy functions by a stricter convergence criterion as well.

are important in risk sharing in rural Pakistan. Accounting for heterogeneity in income risk further improves the model's fit to data.

LTW point out that the model of risk sharing with limited commitment with homogeneous preferences cannot explain the allocation of consumption well. Introducing heterogeneity improves the model's fit to data in this respect. In particular, the LC^{u_i, r_i} model explains 89.4 percent of the cross-sectional variation of consumption in the six villages on average. The percentage of variation explained varies between 86.1 percent and 94.0 percent across villages, and between 74.6 percent and 98.0 percent across villages and time periods, with 79.3 percent as the 5-percent quantile and 96.1 percent as the 95-percent one.

5.2 Robustness checks

I first look at whether our results are robust to changing the number of simulations J in the limited commitment case. I increase J to 200. Second, I check for the effects of approximation error when computing the solution of the risk sharing with limited commitment model. In particular, I increase the number of grid points to 60 when discretizing the state variable x_i . Third, I increase the number of income states for each household to 8. I only do the estimation with each of these changes for the LC^{u_i, r_i} model. The results are presented in Table 3. The parameter estimates are not sensitive to any of the three changes, but computing the policy functions more precisely improves the log likelihood value. Whether the same happens for the other models is to be checked.

6 Policy simulations

In this section, I examine the effects of counterfactual transfers targeting the poor. The poor are defined as having below median mean consumption, that is 109 rupees per week per adult equivalent. I look at the effects of (i) a one-time transfer and (ii) a permanent transfer to the poor on consumption by eligible and ineligible households.

I simulate the effects of the following two social policies: (i) the poor receive a one-time unconditional cash transfer of 20 rupees at time t , and (ii) the poor receive an unconditional cash transfer of

Table 4: The effect of a one-time transfer of 20 rupees to the poor on consumption by eligible and ineligible households (% of the transfer)

	LC^{u_i, r_i}	$LC^{u_i, r}$	LC^{u, r_i}	$LC^{u, r}$	PRS^{u_i}	PRS^u	AUT
Eligible households	79.4	75.5	74.5	76.2	53.0	49.3	100
Ineligible households	20.6	24.5	25.5	23.8	47.0	50.7	0

20 rupees in all periods starting from time t . Mean income of the poor is 80 rupees, thus the transfer increases income by 25 percent on average. This is similar in magnitude to the transfer received by households in the Progresa program in Mexico, see Gertler (2004), for example.²² I consider the introduction of the policy at each time $t = 2, \dots, 12$, and average over the predicted consumption changes at the time of the introduction of the program. I assume that the one-time transfer is seen as a shock, and that it increases permanent income so little that the decision power of recipients does not increase. In the case of a permanent transfer, there is a change in income processes, which are important in determining insurance transfers when commitment is limited. Since this means that decision powers change, the informal risk sharing contract is renegotiated.

Were households in autarky, or, under the policymaker's assumption the transfer increases consumption directly, eligible households' consumption should increase by 20 rupees, while consumption by ineligible households should be unaffected. On the other hand, if households share risk perfectly, the transfers to the poor become part of the common pool, and consumption by each household should increase by about 10 rupees. Any variation in consumption changes should come from differences in risk preferences.

In the case of a one-time transfer, the model of risk sharing with limited commitment predicts that consumption by eligible households increases by 15.88 with heterogeneity in both risk preferences and income risk, while consumption by ineligible households increases by 4.12. This means that eligible households' share of the increase is about four-fifths, and the share of ineligible households is one-fifth. The other models ($LC^{u_i, r}$, LC^{u, r_i} , and $LC^{u, r}$) predict that the share of ineligible households is about one-fourth, because enforcement constraints bind less often. A possible mechanism behind these results is that the transfer of the program crowds out some transfers previously provided by richer households in the villages. The simulation results are summarized in Table 4.

If the transfer of 20 rupees is expected to be received by the poor in all subsequent periods, the models with limited commitment predict that less is shared (see Table 5). This is because the transfer makes eligible households' expected income higher, and multiplicative income risk lower. Therefore, the value of their outside option increases, which implies that their decision power in determining informal insurance transfers increases. The model with heterogeneity in both risk preferences and income risk predicts that consumption by eligible households increases by 18.02 rupees on average, while consumption by ineligible households by 1.98 rupees. In other words, ineligible households' share is only one-tenth in the case of a permanent transfer to the poor. The other three models also predict a lower share going to ineligible households than for a one-time transfer, between 15 and 20 percent. However, more is shared in these cases compared to the model with heterogeneity in both risk preferences and income risk, as before. When risk sharing is perfect, the shares do not change, because I assume that Pareto-weights in the social planner's objective stay the same.

²²Progresa is a conditional cash transfer program. In this exercise, I cannot take such conditionality into account. In other words, I do not look at changes in household behavior other than with respect to insurance transfers. Neither do I deal with how poor households could be identified in practice.

Table 5: The effect of a permanent transfer of 20 rupees to the poor on consumption by eligible and ineligible households (% of the transfer)

	LC^{u_i, r_i}	$LC^{u_i, r}$	LC^{u, r_i}	$LC^{u, r}$	PRS^{u_i}	PRS^u	AUT
Eligible households	90.1	84.1	80.6	82.4	53.0	49.3	100
Ineligible households	9.9	15.9	19.4	17.6	47.0	50.7	0

These results draw attention to the fact that policy effects are likely to be miscalculated, if we ignore informal risk sharing arrangements. Policy makers should neither assume that a transfer given to the poor will be fully consumed by the recipients, nor should they ignore the positive spill-over effects on ineligible households. Recent empirical evidence by [Angelucci and De Giorgi \(2009\)](#) provides support for this argument.

7 Extensions

In this section, I develop a more general empirical model. In particular, I allow preferences to depend on (i) unobservable individual effects and (ii) time-varying household characteristics; further, (iii) preference shocks are introduced, and (iv) income is measured with error as well as consumption. I show that the predicted consumption allocation is not affected by the individual effects, and argue that the other extensions are possible in theory. They are, however, infeasible due to prohibitively long computation time.

Instead of (1), let us now specify the utility function as

$$u_{it}(c_{it}) = \exp(\xi_{it}) \frac{c_{it}^{1-\sigma_{it}} - 1}{1 - \sigma_{it}}, \quad (30)$$

where

$$\xi_{it} = \eta_i + w'_{it}\alpha + \varepsilon_{it}^\eta,$$

and the coefficient of relative risk aversion is

$$\sigma_{it} = 1 + z'_{it}\beta,$$

where η_i is a time-constant unobservable individual effect, w_{it} and z_{it} are vectors of observable characteristics of household i at time t , α and β are parameter vectors to be estimated, and ε_{it}^η is a normally distributed preference shock with mean 0 and variance γ_η^2 . Note that σ_{it} is only allowed to depend on observable covariates. Finally, let y_{it}^* be income observed by the econometrician. Measurement error in income is assumed to be multiplicative and log-normally distributed, that is,

$$y_{it}^* = \exp(\varepsilon_{it}^y) y_{it},$$

and $\varepsilon_{it}^y \sim N(0, \gamma_y^2)$.

7.1 Perfect risk sharing

The first order condition for household i , sharing risk with the rest of the community, can now be written as

$$\sigma_{it} \log c_{it} - \xi_{it} - \log c_{kt} = \log \lambda_i.$$

Replacing for ξ_{it} and σ_{it} , in terms of measured consumption we have

$$(1 + z'_{it}\beta) \log c_{it}^* - (1 + z'_{it}\beta) \varepsilon_{it} - \eta_i - w'_{it}\alpha - \varepsilon_{it}^\eta - \log c_{kt} = \log \lambda_i.$$

First differencing and rearranging give

$$\begin{aligned} \log \left(\frac{c_{it}^*}{c_{kt}} \right) &= \log \left(\frac{c_{i,t-1}^*}{c_{k,t-1}} \right) - z'_{it}\beta \log c_{it}^* + z'_{i,t-1}\beta \log c_{i,t-1}^* + w'_{it}\alpha - w'_{i,t-1}\alpha \\ &\quad + \varepsilon_{it}^\eta - \varepsilon_{i,t-1}^\eta + (1 + z'_{it}\beta) \varepsilon_{it} - (1 + z'_{i,t-1}\beta) \varepsilon_{i,t-1}. \end{aligned}$$

Thus η_i drops out along with $\log \lambda_i$. The error term is now distributed as

$$N \left(0, 2\gamma_\eta^2 + (1 + z'_{it}\beta)^2 \gamma^2 + (1 + z'_{i,t-1}\beta)^2 \gamma^2 \right).$$

7.2 Autarky

In the autarky case, preferences do not play a role. Thus only the additional measurement error in income has to be taken into account. Instead of (19), the equation to be estimated is

$$\log \left(\frac{c_{it}^*}{c_{kt}} \right) = \log \left(\frac{y_{it}^*}{y_{kt}} \right) + \varepsilon_{it} - \varepsilon_{it}^y,$$

and γ^2 and γ_y^2 are not jointly identified.

7.3 Risk sharing with limited commitment

Instead of (23), the first order condition now is

$$\log \left(\frac{c_{it}^*}{c_{kt}} \right) = -z'_i\beta \log c_{it}^* + \eta_i + w'_{it}\alpha + \varepsilon_{it}^\eta + \log g(y_t, x_{i,t-1}, Z_i, W_i; \theta, \eta_i) + (1 + z'_i\beta) \varepsilon_{it}, \quad (31)$$

where Z_i and W_i are matrices of observable covariates of household i at different times, and θ now includes the vector α as well.

Let us first deal with the individual effect η_i , assuming $z_{it} = z_i, \forall t$, and ignoring the other time-varying components of the utility function (w_{it} and ε_{it}^η) for the moment. Suppose that we know the realization of measurement error in household i 's consumption at time $t - 1$, denoted $\varepsilon_{i,t-1}^j$, drawn from the distribution of $\varepsilon_{i,t-1}$, $N(0, \gamma^2)$. I show that

$$\begin{aligned} g \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta, \eta_i \right) &= \exp(-\eta_i) g \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right) \\ &= \exp(-\eta_i) \frac{\hat{c}_{it} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right)^{1+z'_i\beta}}{\hat{c}_{kt} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right)}, \end{aligned}$$

where $\hat{c}_{it} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right)$ and $\hat{c}_{kt} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right)$ are consumption by households i and k respectively, predicted by the model, normalizing $\eta_i = 0$. Replacing this in equation (31), η_i drops out. The fact that the function $g(\cdot)$ is homogeneous of order one in η_i is the direct consequence of the following claim.

Claim 1. $\hat{c}_{it} \left(y_t, x_{i,t-1}^*, \varepsilon_{i,t-1}^j, z_i; \theta \right) = \hat{c}_{it} \left(y_t, c_{i,t-1}^*, z_i, \varepsilon_{i,t-1}^j; \theta, \eta_i \right)$. That is, the consumption allocation predicted by the model of risk sharing with limited commitment does not depend on the individual effects.

Proof. To see this, let us take a closer look at the enforcement constraint of some household i , that can be written as (7) in general. Replacing the utility function (1) in (7) gives

$$\begin{aligned} & \exp(\eta_i) \frac{c_{it} (s^t)^{1-\sigma_i} - 1}{1-\sigma_i} + \sum_{r=t+1}^{\infty} \sum_{s^r} \delta^{r-t} \pi(s^r) \exp(\eta_i) \frac{c_{ir} (s^r)^{1-\sigma_i} - 1}{1-\sigma_i} \geq \\ & \geq \exp(\eta_i) \frac{y_{it} (s_t)^{1-\sigma_i} - 1}{1-\sigma_i} + \sum_{r=t+1}^{\infty} \sum_{s_r} \delta^{r-t} \pi(s_r) \exp(\eta_i) \frac{y_{ir} (s_r)^{1-\sigma_i} - 1}{1-\sigma_i}. \end{aligned} \quad (32)$$

Both sides can be divided by $\exp(\eta_i)$, thereby eliminating the individual effects. When no enforcement constraint is binding, we are back to perfect risk sharing, where $\exp(\eta_i)$ appears multiplicatively on both sides of $x_{it} = x_{i,t-1}$. \square

Then, with the utility function (30), a typical enforcement constraint can be written as

$$\begin{aligned} & \exp(w'_{it}\alpha + \varepsilon_{it}^{\eta}) \frac{c_{it} (s^t)^{-z'_{it}\beta} - 1}{-z'_{it}\beta} + \sum_{r=t+1}^{\infty} \sum_{s^r} \delta^{r-t} \pi(s^r) \exp(w'_{ir}\alpha + \varepsilon_{ir}^{\eta}) \frac{c_{ir} (s^r)^{-z'_{ir}\beta} - 1}{-z'_{ir}\beta} \geq \\ & \geq \exp(w'_{it}\alpha + \varepsilon_{it}^{\eta}) \frac{y_{it} (s_t)^{-z'_{it}\beta} - 1}{-z'_{it}\beta} + \sum_{r=t+1}^{\infty} \sum_{s_r} \delta^{r-t} \pi(s_r) \exp(w'_{ir}\alpha + \varepsilon_{ir}^{\eta}) \frac{y_{ir} (s_r)^{-z'_{ir}\beta} - 1}{-z'_{ir}\beta}. \end{aligned}$$

Note that future household characteristics and preference shocks enter into today's enforcement constraint. Therefore, some assumptions have to be made on households' expectations about their future characteristics and preference shocks.

If we assume that households are myopic, in the sense that they do not expect their characteristics to change relative to the rest of the village, and further that they do not expect preference shocks to differ from today's, then we may divide the above equation by $\exp(w'_{it}\alpha + \varepsilon_{it}^{\eta}) = \exp(w'_{ir}\alpha + \varepsilon_{ir}^{\eta})$, $\forall r > t$. In this case computation time is only multiplied by the number of time periods. However, the assumption that households are always surprised by a change in any of their characteristics is very strong.

Alternatively, we may assume that households form rational expectations, and their expectations about their characteristics next period are the observed values. Preference shocks can be integrated out using simulation, if we make some assumption about their distribution. To keep things tractable, I assume that preference shocks are i.i.d. over time and across households. Given household characteristics and a realization of the preference shock today, the enforcement constraint can be written in a recursive form as

$$\exp(w'_{it}\alpha + \varepsilon_{it}^{\eta}) \left[\frac{\hat{c}_i(s_t, x_{t-1})^{-z'_{it}\beta} - 1}{-z'_{it}\beta} - \frac{y_i(s_t)^{-z'_{it}\beta} - 1}{-z'_{it}\beta} \right] \geq \delta \sum_{s_{t+1}} \pi(s_{t+1}) \times$$

$$\int \left[V_i^{aut} \left(s_{t+1}, w_{i,t+1}, z_{i,t+1}, \varepsilon_{i,t+1}^{\eta} \right) - V_i \left(s_{t+1}, x_t(s_t, x_{t-1}), w_{i,t+1}, z_{i,t+1}, \varepsilon_{i,t+1}^{\eta} \right) \right] f \left(\varepsilon_{i,t+1}^{\eta} \right) d\varepsilon_{i,t+1}^{\eta}.$$

Thus the model should be solved on an extended state space that includes household characteristics. In practice, a grid has to be defined over each characteristic. Today's preference shock also has to be integrated out by simulation. Computation time thus becomes prohibitively long.

Note that, once we have computed the consumption allocation predicted by the model, $\hat{c}_{it}(y_t, c_{i,t-1}^*, \varepsilon_{i,t-1}^j, Z_i, W_i; \theta)$ and $\hat{c}_{kt}(y_t, c_{i,t-1}^*, \varepsilon_{i,t-1}^j, Z_i, W_i; \theta)$, then

$$\log g(y_t, x_{i,t-1}, Z_i, W_i; \theta, \eta_i) = \exp(-\eta_i - w'_{it}\alpha - \varepsilon_{it}^\eta) \frac{\hat{c}_{it}(y_t, c_{i,t-1}^*, \varepsilon_{i,t-1}^j, Z_i, W_i; \theta)^{1+z'_i\beta}}{\hat{c}_{kt}(y_t, c_{i,t-1}^*, \varepsilon_{i,t-1}^j, Z_i, W_i; \theta)}.$$

Thus, along with η_i , $w'_{it}\alpha$ and ε_{it}^η drop out from equation (31). Therefore, adding measurement error in consumption is necessary to account for the residual in the estimating equation.

Finally, let us look at what changes if income is measured with error, as well as consumption. First, given the distribution from which income is drawn, measurement error does not affect the optimal intervals that characterize the constrained efficient risk sharing contract, since the model is solved using a grid on income. Second, the introduction of measurement error in income may plague the estimation of the income process. We could perturbate observed income, and then recompute the model given a matrix of draws of measurement error. Computation time is then proportional to the number of matrices of draws, which would make the estimation take prohibitively long. Third, today's income observation directly affects consumption predicted by the model. Once again, simulation is a simple way to deal with this problem, and the solution of the model does not have to be recomputed, thus computation time increases only moderately due to this third point.

8 Concluding remarks

This paper first performed statistical tests to compare seven models of risk sharing. Structural estimation results suggest that limitations in the enforcement of informal insurance contracts and heterogeneity in preferences and in income risk are important in explaining the consumption allocation in six villages in rural Pakistan. The model of limited commitment with heterogeneous preferences and income risk is able to account for 89 percent of the cross-sectional variation in consumption.

Using structural estimation results, this paper then simulated the effects of a simple social policy. The model predicts that consumption by both eligible and ineligible households should increase, consistently with the empirical findings of [Angelucci and De Giorgi \(2009\)](#). Research on the structural modeling of how consumption is allocated between households in poor communities can serve as an input for policy evaluation and design. Policy makers and members of non-governmental organizations could have a better understanding of the effects of their programs, such as redistributive policies or micro-insurance programs, by taking into account existing informal arrangements to share risk.

Several interesting extensions are possible. First, whether heterogeneity in the discount factor across households is important should also be addressed. Second, other models of risk sharing could be incorporated into the analysis, like the model of risk sharing with private information ([Wang, 1995](#)). In a recent paper, [Kinnan \(2010\)](#) finds that asymmetric information about income realizations is important in accounting for partial insurance in Thai villages. Whether such a model is useful for predicting the consumption allocation and quantitative policy effects is to be studied. Another important task for future work is to allow for individual savings, as in [Ligon, Thomas, and Worrall \(2000\)](#). Fourth, it is to be examined whether introducing more heterogeneity across households would mean that the model of risk sharing with limited enforcement could better capture the amount of risk sharing in other contexts, in particular, in the United States, since the homogenous

model predicts too much insurance (Krueger and Perri, 2006). Finally, when complete markets do not exist to insure against income fluctuations, households are expected to choose safer jobs, or safer production technologies in agriculture. In other words, they smooth income, not just consumption (Morduch, 1995). These ideas could be formalized in the context of this paper, endogenizing income by allowing households to choose between several income generating processes. Then, the cost of imperfect consumption insurance in terms of lower expected incomes could be quantified.

Acknowledgements: This paper has greatly benefitted from discussions with Árpád Ábrahám, Orazio Attanasio, Richard Blundell, Russell Cooper, Pierre Dubois, Johannes Gierlinger, Luigi Guiso, Ethan Ligon, Thierry Magnac, Ramon Marimon, Alessandro Menzuni, David Pacini, Alessandro Tarozzi, and Judit Vall Castello. I also thank participants of the Panel Data Conference in Bonn, the Summer Workshop in Budapest, and ESEM in Barcelona, and seminar participants at the EUI, TSE, Queen Mary, and the University of Alicante for useful comments and suggestions. All errors are mine. I thank the International Food Policy Research Institute (IFPRI) for making the data available. Generous funding from Agence Nationale de la Recherche (ANR-06-BLAN-0104) is gratefully acknowledged.

References

- Abreu, D. (1988). On the Theory of Infinitely Repeated Games with Discounting. *Econometrica* 56(2), 383–396.
- Ackerberg, D., J. Geweke, and J. Hahn (2009). Comments on “Convergence Properties of the Likelihood of Computed Dynamic Models” by Fernández-Villaverde, Rubio-Ramírez and Santos. *Econometrica* 77(6), 2009–2017.
- Alderman, H. and M. Garcia (1993). Poverty, Household Food Security, and Nutrition in Rural Pakistan. Research report 96, International Food Policy Research Institute.
- Ales, L. and P. Maziero (2009). Non-exclusive Dynamic Contracts, Competition, and the Limits of Insurance. Mimeo.
- Alvarez, F. and U. J. Jermann (2000). Efficiency, Equilibrium, and Asset Pricing with Risk of Default. *Econometrica* 68(4), 775–797.
- Angelucci, M. and G. De Giorgi (2009). Indirect Effects of an Aid Program: How Do Cash Transfers Affect Ineligibles’ Consumption? *American Economic Review* 99(1), 486–508.
- Arrow, K. (1965). *Aspects of the Theory of Risk-Bearing*. Yrjö Jahnssonin Saatio, Helsinki.
- Attanasio, O. and J.-V. Ríos-Rull (2000). Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare? *European Economic Review* 44(7), 1225–1258.
- Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption Inequality and Partial Insurance. *American Economic Review* 98(5), 1887–1921.
- Borch, K. (1962). Equilibrium in a Reinsurance Market. *Econometrica* 30(3), 424–444.
- Coate, S. and M. Ravallion (1993). Reciprocity without Commitment: Characterization and Performance of Informal Insurance Arrangements. *Journal of Development Economics* 40, 1–24.
- Cochrane, J. H. (1991). A Simple Test of Consumption Insurance. *Journal of Political Economy* 99(5), 957–976.
- Dercon, S. and P. Krishnan (2003a). Food Aid and Informal Insurance. UNU-WIDER Research Paper, World Institute for Development Economic Research.
- Dercon, S. and P. Krishnan (2003b). Risk Sharing and Public Transfers. *Economic Journal* 113, 86–94.
- Dixit, A., G. M. Grossman, and F. Gul (2000). The Dynamics of Political Compromise. *Journal of Political Economy* 108(3), 531–568.
- Dubois, P. (2000). Assurance complète, hétérogénéité des préférences et métayage au Pakistan. *Annales d’Economie et de Statistique* 59, 1–36.
- Dubois, P., B. Jullien, and T. Magnac (2008). Formal and Informal Risk Sharing in LDCs: Theory and Empirical Evidence. *Econometrica* 76(4), 679–725.
- Fafchamps, M. (1999). Risk Sharing and Quasi-Credit. *Journal of International Trade and Economic Development* 8(3), 257–278.

- Fernández-Villaverde, J., J. F. Rubio-Ramírez, and M. S. Santos (2006). Convergence Properties of the Likelihood of Computed Dynamic Models. *Econometrica* 74(1), 93–119.
- Foster, A. D. and M. R. Rosenzweig (2001). Imperfect Commitment, Altruism, and the Family: Evidence from Transfer Behavior in Low-Income Rural Areas. *Review of Economics and Statistics* 83(3), 389–407.
- Gertler, P. (2004). Do Conditional Cash Transfers Improve Child Health? Evidence from PROGRESA's Control Randomized Experiment. *American Economic Review* 94(2), 336–341.
- Gouriéroux, C. and A. Monfort (1997). *Simulation-based Econometric Methods*. Oxford University Press, Oxford.
- Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the Life Cycle. *Econometrica* 70(1), 47–89.
- Grimard, F. (1997). Household Consumption Smoothing Through Ethnic Ties: Evidence from Côte d'Ivoire. *Journal of Development Economics* 53(2), 391–422.
- Halek, M. and J. G. Eisenhauer (2001). Demography of Risk Aversion. *Journal of Risk and Insurance* 68(1), 1–24.
- Kehoe, P. J. and F. Perri (2002). International Business Cycles with Endogenous Incomplete Markets. *Econometrica* 70(3), 907–928.
- Kinnan, C. (2010). Distinguishing Barriers to Insurance in Thai Villages. Mimeo.
- Kocherlakota, N. R. (1996). Implications of Efficient Risk Sharing without Commitment. *Review of Economic Studies* 63(4), 595–609.
- Krueger, D. and F. Perri (2006). Does Income Inequality Lead to Consumption Inequality? Evidence and Theory. *Review of Economic Studies* 73(1), 163–193.
- Kurosaki, T. and M. Fafchamps (2002). Insurance Market Efficiency and Crop Choices in Pakistan. *Journal of Development Economics* 67(2), 419–453.
- Ligon, E. (1998). Risk Sharing and Information in Village Economies. *The Review of Economic Studies* 65(4), 847–864.
- Ligon, E., J. P. Thomas, and T. Worrall (2000). Mutual Insurance, Individual Savings, and Limited Commitment. *Review of Economic Dynamics* 3(2), 216–246.
- Ligon, E., J. P. Thomas, and T. Worrall (2002). Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies. *Review of Economic Studies* 69(1), 209–244.
- Mace, B. J. (1991). Full Insurance in the Presence of Aggregate Uncertainty. *Journal of Political Economy* 99(5), 928–956.
- Marcet, A. and R. Marimon (2009). Recursive Contracts. Mimeo.
- Mazzocco, M. (2007). Household Intertemporal Behaviour: A Collective Characterization and a Test of Commitment. *Review of Economic Studies* 74(3), 857–895.

- Mazzocco, M. and S. Saini (2009). Testing Efficient Risk Sharing with Heterogeneous Risk Preferences. Mimeo.
- Morduch, J. (1995). Income Smoothing and Consumption Smoothing. *Journal of Economic Perspectives* 9(3), 103–114.
- Ogaki, M. and Q. Zhang (2001). Decreasing Relative Risk Aversion and Tests of Risk Sharing. *Econometrica* 69(2), 515–526.
- Pratt, J. W. (1964). Risk Aversion in the Small and in the Large. *Econometrica* 32(1/2), 122–136.
- Renkow, M. (1993). Land Prices, Land Rents, and Technological Change: Evidence from Pakistan. *World Development* 21(5), 791–803.
- Schechter, L. (2007). Theft, Gift-Giving, and Trustworthiness: Honesty Is Its Own Reward in Rural Paraguay. *American Economic Review* 97(5), 1560–1582.
- Schulhofer-Wohl, S. (2007). Heterogeneity, Risk Sharing and the Welfare Costs of Idiosyncratic Risk. Mimeo.
- Thomas, J. and T. Worrall (1988). Self-Enforcing Wage Contracts. *Review of Economic Studies* 55(4), 541–554.
- Todd, P. E. and K. I. Wolpin (2006). Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility. *American Economic Review* 96(5), 1384–1417.
- Todd, P. E. and K. I. Wolpin (2008). Ex Ante Evaluation of Social Programs. Mimeo.
- Townsend, R. M. (1994). Risk and Insurance in Village India. *Econometrica* 62(3), 539–591.
- Vuong, Q. H. (1989). Likelihood Ratio Tests for Model Selection and Non-nested Hypotheses. *Econometrica* 57(2), 307–333.
- Wang, C. (1995). Dynamic Insurance with Private Information and Balanced Budgets. *Review of Economic Studies* 62(4), 577–595.
- Wilson, R. (1968). The Theory of Syndicates. *Econometrica* 36(1), 119–132.
- Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, Massachusetts.

Appendix

This appendix details how to compute the consumption allocation predicted by the model. That is, we want to find \hat{c}_{it} relative to mean consumption in the community \hat{c}_{kt} , for each household i at each time t , taking preference parameters and the discrete distributions from which incomes are drawn as given.²³ The model is solved for each household separately, since observables entering the utility function and household specific random income Y_i are taken as given.

Consider some household i sharing risk with the rest of the community, or, household k . Take preference parameters β , the discount factor δ , and the discrete distributions from which incomes are drawn F_{Y_i} and F_{Y_k} as given. The aim is to solve for the policy functions, that is, (true) consumption $c_{it}(s_t, x_{i,t-1})$, and the (true) relative weight $x_{it}(s_t, x_{i,t-1})$.

To solve for the optimal intervals on the relative Pareto weight of household i , proceed as follows. First, define a grid over the continuous variable x_i for each value of s_t (I define the same points for all s_t). The support of the grid is the range of ratios of marginal utilities of household k and i given the income and consumption observations. I define an equidistant grid on $\log x_i$ of 30 points. Second, guess a solution for the value functions, that is, guess $V_i^0(s_t, x_{i,t-1})$, for each grid point. The algorithm does not converge from any initial guess for the value functions, but the value of perfect risk sharing will do.²⁴

Then, proceed to update the guess. Suppose we are at the h^{th} iteration. Let us look at grid point $(\tilde{s}_t, \tilde{x}_{i,t-1})$. Three cases have to be distinguished: (a) neither enforcement constraint binds, (b) the enforcement constraint for household i binds, and (c) the enforcement constraint for household k binds. Note that the two enforcement constraints cannot bind at the same time, because only one of the two households may be called upon to make a positive net transfer.

We first suppose that the enforcement constraints do not bind, that is, we try to keep x_i constant. This means setting $\hat{x}_{it}^h = \tilde{x}_{i,t-1}$ at state \tilde{s}_t , where the upper index h refers to iteration h . Then, using the first order condition and the resource constraint, we get the consumption allocation $(\hat{c}_{it}^h, \hat{c}_{kt}^h)$. Now the enforcement constraints have to be checked. This means verifying whether

$$\frac{(\hat{c}_{it}^h)^{-z'_i\beta} - 1}{-z'_i\beta} + \delta \sum_{s_{t+1}} \pi(s_{t+1}) V_i^{h-1}(s_{t+1}, \hat{x}_{it}^h) \geq U_i^{\text{aut}}(\tilde{s}_t) \quad (33)$$

and

$$\log \hat{c}_{kt}^h + \delta \sum_{s_{t+1}} \pi(s_{t+1}) V_k^{h-1}(s_{t+1}, \hat{x}_{kt}^h) \geq U_k^{\text{aut}}(\tilde{s}_t). \quad (34)$$

Note the upper index $h - 1$ for V_i and V_k , that is, we use the value function from the previous iteration.

- (a) *The enforcement constraints (33) and (34) do not bind.* This is the easy case, since we have already computed \hat{x}_{it}^h and the consumption allocation assuming that the enforcement constraints

²³The first step in estimation involves determining these distributions (see main text), while the last step is the maximization over the remaining structural parameters, which is done using a standard optimization algorithm available in R (function `optim()` with method BFGS with bounds (L-BFGS-B), which is a quasi-Newton method). See www.r-project.org. Here, I am talking about the computation between these steps.

²⁴Characterizing the convergence properties of the algorithm is left for future research. However, we know that the algorithm does not converge to the constrained-efficient solution from any initial guess for the value functions. For example, if we set the guesses for $V_i^0(s_t, x_{i,t-1})$ equal to the autarkic values, every iteration yields these same autarkic values. This is natural, since autarky is also a subgame perfect Nash equilibrium (SPNE).

do not bind. What remains to be done is to set

$$V_i^h(\tilde{s}_t, \tilde{x}_{i,t-1}) = \frac{(\hat{c}_{it}^h)^{-z'_i\beta} - 1}{-z'_i\beta} + \delta \sum_{s_{t+1}} \pi(s_{t+1}) V_i^{h-1}(s_{t+1}, \hat{x}_{it}^h)$$

and

$$V_k^h(\tilde{s}_t, \tilde{x}_{i,t-1}) = \log \hat{c}_{kt}^h + \delta \sum_{s_{t+1}} \pi(s_{t+1}) V_k^{h-1}(s_{t+1}, \hat{x}_{it}^h).$$

- (b) *The enforcement constraint (33) is binding.* In theory, we may compute \hat{c}_{it}^h and \hat{x}_{it}^h using (33) with equality and the first order condition. However, we do not know $V_i^{h-1}(s_{t+1}, \hat{x}_{it}^h)$ for any value of \hat{x}_{it}^h , only for the points on the grid. Therefore, we look for the point \hat{x}_{it}^h for which \hat{c}_{it}^h is such that (33) is satisfied and the left-hand side is closer to the right-hand side than for any other grid point. Finally, we update the value function as in case (a).
- (c) *The enforcement constraint (34) is binding.* We proceed similarly as in case (b).

Now we are done with grid point $(\tilde{s}_t, \tilde{x}_{i,t-1})$. We have to do the above steps at all other grid points as well. Then the h^{th} iteration is complete. I continue iterating until the policy function converges, that is, the optimal intervals on x_{it} do not change. More precisely, the solution has been found, if the length of the difference between the endpoints of the optimal intervals at iteration $h - 1$ and h is less than 0.001. In practice I allow for maximum 20 iterations.²⁵ In the end, we have the solution in the form $[\underline{x}_{it}(s_t, x_{i,t-1}), \bar{x}_{it}(s_t, x_{i,t-1})]$.

Computing the consumption of household i at time t , relative to mean consumption in the community, as predicted by the model is then done as follows. Remember that $c_{i,t-1}^*$ is the observed consumption by household i at time $t - 1$, and $\varepsilon_{i,t-1}^j$ is a realization of measurement error drawn from $N(0, \gamma^2)$, and γ^2 is given. I compute $x_{i,t-1} = \left(\exp(-\varepsilon_{i,t-1}^j) c_{i,t-1}^*\right)^{1+z'_i\beta} / c_{k,t-1}$, and check whether it is in the optimal interval for today's state s_t . Since only a discrete number of income states have been considered, I map observed incomes into the income states of the model by picking the closest point for each household. We have to consider the above three cases.

- (a) If $x_{i,t-1} \in [\underline{x}_{it}(s_t, x_{i,t-1}), \bar{x}_{it}(s_t, x_{i,t-1})]$, then I set $x_{it} = x_{i,t-1}$.
- (b) If $x_{i,t-1} < \underline{x}_{it}(s_t, x_{i,t-1})$, then I determine it from (33) with equality, using a linear interpolation of the value functions from the last iteration.
- (c) If $x_{i,t-1} < \bar{x}_{it}(s_t, x_{i,t-1})$, then I use (34) with equality, and proceed similarly as in case (b).

Finally, I use the first order condition and the resource constraint to determine the predicted consumption allocation. We may then write the likelihood of observation it , given $\varepsilon_{i,t-1}^j$, by plugging the \hat{c}_{it} and \hat{c}_{kt} computed here into (27).

²⁵A robustness check is performed with 60 grid points, requiring the endpoints of the optimal intervals not to change with a 0.00001 tolerance, and allowing for maximum 40 iterations (see Table 3).

