Empirical implications of limited commitment.

Evidence from Mexican villages[⊗]

Pedro Albarran[♣] and Orazio P. Attanasio[♥]

First Draft July 2001
This Draft February 2003

[⊗] A preliminary draft of this paper was presented at the SIEPR Conference "Effects of Credit Market Frictions on the Macroeconomy" at Stanford, August 2001 and was prepared while Attanasio was visiting the Department of Economics at the University of Chicago, whose hospitality is gratefully acknowledged. The authors thank Fernando Alvarez, Marcel Fafchamps, Narayana Kocherlakota, Victor Rios-Rull and Rob Townsend for useful discussions, as well as seminar audiences at UCL, Cambridge, Wharton, Kellogg, MIT, Princeton, Minnesota, Stanford, Chicago, Oxford (Nuffield and CSAE), Southampton, LSE for useful comments. Josep Pijoan-Mas helped with the simulations, while Ana Santiago helped with the Progresa data. Several people in the Progresa staff, including Daniel Hernandez, Patricia Muniz and Susan Parker helped us with information about the data and the program in the early part of the project. We also wish to acknowledge our debt to the late Jose Gomez de Leon who first stimulated our interest into Progresa and without whose vision none of this would have been possible.
[♣] CEMFI & Universidad Carlos III Madrid
[♥] University College London, IFS, and NBER

1 Introduction

In this paper we study the observable implications of a class of models of consumption smoothing characterized by less than full risk sharing. In particular, we focus on models where the first best allocation of resources is not achieved because of the presence of imperfect enforceability. That is, while the individuals within the economy we consider have full information about the endowments and the actions taken by the other individuals with whom they interact, they cannot convey this information to the external world. As a consequence, full enforceability of contracts might be exceedingly costly. We are therefore led to consider models where individuals only enter contracts that are self-enforceable.

The models we consider are not new. They consider constrained-efficient allocations, where the constraints are given by the inability of enforcing contract agreements. They have been studied by several authors, including Thomas and Worrall (1988, 1990), Kocherlakota (1996), Alvarez and Jerman (2000a,b), Ligon, Thomas and Worrall (2000), Kehoe and Perri (1999), Krueger and Perri (2001), Attanasio and Rios-Rull (2000), among others. In the class of models we consider, contracts are enforced by the threat to reverting to the worst sub-game perfect equilibrium: autarky. This equilibrium concept, proposed by Abreu et al. (1988) has proved very useful in this literature.

These models can generate very complex dynamics and allocations of resources that differ substantially from those implied by complete markets. In particular, they can replicate features of inter-households agreements that are reportedly prevalent in many village economies. Platteau (1987), for instance, claims that the risk sharing agreements often observed in fisherman communities are half way between credit and insurance, while Udry (1994) reports that interest rates and maturities on loans in Northern Nigeria seem to be state contingent and vary not only with shocks affecting the borrower but also with those affecting the lender. Moreover, the assumptions typically used in these models (complete information, lack of strong enforcement mechanisms, repeated interactions, large scope for risk sharing) seem to be appropriate to characterize poor and isolated rural villages in many parts of the world.

As we discuss below, the crucial state variable of these models, the ratio of marginal utilities, unlike in models with perfect risk sharing, moves over time to guarantee that the participation constraints are always satisfied. One can interpret movements in this state variable as movements in the (relative) Pareto weights that a social planner given to an individual: when the participation constraint of an individual is binding, the social planner will increase the weight given to that individual. It is the partial inertia created by this mechanism that gives to the efficient contract some of the features of a debt contract. However, when the constraint of another individual becomes binding, past history becomes irrelevant and contract resembles more an insurance one.

Not much evidence exists on the empirical relevance of models with imperfect enforceability. The only papers we are aware of are those by Foster and Rosenzweig (2001), Ligon, Thomas and Worrall (2001), Krueger and Perri (2002) and Albarran and Attanasio (2002).

In their imaginative and interesting paper, Foster and Rosenzweig (2001) consider the implications of the model for transfers. In particular, they notice that, conditioning on the current shock, the net transfers an individual receives from her partner in a self enforcing risk sharing agreement is negatively related to the cumulate of past transfers received. After showing with some simulations that such a relationship is well approximated by a linear one, they consider its first difference to relate changes in transfers to the change in shocks and the lagged level of transfers. Such an equation can be used on panel data that contain at least two observations per household on transfers to test that the coefficient on lagged transfers is negative. Foster and Rosenzweig (2001) also extend the model to consider the presence of altruism. Applying these tests to data from Pakistan and Bangladesh they find a pattern of coefficients that is in accordance with the predictions of the model.

Ligon, Thomas and Worrall (2001) take a more structural approach and estimate a version of the model with imperfect enforceability by maximum likelihood using the ICRSTAT data used, among others, by Townsend (1994) to test full insurance. Their estimation method involves, for each evaluation of the likelihood function, the numerical solution of the optimal contract. The numerical complexity of the problem forces them to some important simplifying assumptions. In particular, they assume the lack of a storage technology (so that aggregate income and consumption are also the same) and, instead of considering the contract among the N households in the village they consider N contracts between each household and the remaining N-1 households. While they convincingly justify the latter assumption on the basis of an aggregation argument, the former is likely to be more critical and hard to justify. However, even with these counterfactual assumptions, the model with imperfect enforceability seems to be able to fit the observed data better than a model that assumes perfect risk sharing.

In a very recent paper, Krueger and Perri (2001) consider the implications of the model for the relationship between the evolution of the cross sectional variance of income and consumption and present some evidence that is not inconsistent with the model.

In Attanasio and Albarran (2002), we propose a simple test that considers how the introduction of a public transfer program affects, under imperfect enforceability, private transfers. In that paper, we exploit the randomized adoption of the welfare program implemented in the villages that form our data base. We start by noticing that a welfare program that involves public transfers to some (or all) of the partners of an insurance agreement with imperfect enforceability is likely to reduce or 'crowd out', under many preference specifications, private transfers. This implication, however, is not unique to the models with imperfect enforceability: it might also occur in models with perfect risk sharing and certainly in models with altruism. However, within the set of models we consider, the amount by which private transfers are reduced is determined by other features of the economy, such as the variance of income and its persistence. It is this set of implications, that are unique to the models with imperfect enforceability that we consider.

The main scope of this paper is to propose some new tests that can be informative about the empirical relevance of such models. We propose two different approaches that focus on different features of the model and of the available data. Our approaches differ considerably from those mentioned above. The main difference of both our tests relative to those mentioned above is that our approach focuses on the implications of the model for intertemporal consumption allocations. We can therefore afford to be silent about the particular mechanisms that are used to achieve that particular allocation of resources. They are therefore in the spirit of Townsend (1994) test. This is particularly true for our first test, based on village level statistics.

We apply our tests to a unique data set from rural Mexico. The data, a 4 waves panel collected between 1997 and 1999, was gathered to evaluate a large welfare program started by the Zedillo administration in 1998, called PROGRESA, whose aim is to foster the accumulation of human capital in rural communities by providing financial incentives to the investment in health and education. To evaluate such a program the administration started collecting a large database that gathered information on all the households living in 506 localities in 7 states that qualified for the program. [1]

The first piece of evidence we consider can be interpreted as a test of full insurance and first best allocation. As we discuss below, and as it was noticed by Deaton and Paxson (1994), a straightforward implication of models with complete insurance markets is that the cross sectional variance of the marginal utility of consumption is constant over time. In what follows, we start from this implication and construct a measure of the deviations from first best allocations given by the magnitude of the changes in the cross sectional variance of marginal utility. We then relate this measure with various features of the economies we study, such as the variance and persistence of income, which would affect it under the imperfect enforceability of contracts. Our test predicts that these quantities relate in a particular way with the level of risk sharing: everything that increases the enforceability of contracts should result in a greater amount of risk sharing.

The second test we propose develops some of the ideas proposed by Kocherlakota (1996). It uses the first order conditions implied by the constrained optimization that characterizes the intertemporal allocation of resources under imperfect enforceability. In particular, it exploits the fact that cross sectional differences in the rate of growth of marginal utility of consumption can be informative about which households have a binding enforceability constraint. Moreover, the constrained efficiency nature of these contracts has specific implications for both the set of constrained and unconstrained households. We generalize the ideas in Kocherlakota to consider storage and measurement error in consumption.

The rest of the paper is organized as follows. In section 2, we sketch the theoretical framework that we will be investigating empirically. While most of this material is not new, we need to establish notation and state some results that can be translated into testable implications. In section 3, we propose our three tests of the model. In section 4 we discuss

[1] As the program is very large and was phased in over a period of two years, 186 randomly chosen localities of the evaluation sample were placed 'at the end of the queue' so that the program in those localities was started in December 1999 (rather than July 1998). This randomization scheme was explicitly implemented with the purpose of facilitating the evaluation of the effectiveness of the program. In what follows we do not use the variation induced by the program to identify our models. We followed that strategy in Albarran and Attanasio (2002).

the data set we use, while in Section 5, we present our results. Section 6 concludes the paper with some thoughts on future research.

2 The theoretical framework

In this section, we sketch a simple model of risk sharing with imperfect enforceability. While most of the material we present is not particularly new, we need it to establish notation and state formally some of the results we will be using in our empirical work. Moreover, we could not find the characterization of some of the properties of the equilibria we study in the existing literature. This is partly because the properties we study are targeted towards our empirical application. In other words, we are interested in 'observable' properties of the equilibria with imperfect enforceability.

For expositional simplicity we start our presentation considering the case of two households without storage opportunities. We then discuss how our results extend to the case of many households and storage.

2.1 The basic model

Let us consider two infinitely lived agents who, in each period receive an endowment $e^i,\ i=1,2$, which is a function of an aggregate shock z and an idiosyncratic shock υ^i: $e^i = e^i(\upsilon^i, z)$. Both υ^i and z have discrete support. The two idiosyncratic shocks are independent of the aggregate shock and of each other. As we start by assuming that there are no storage opportunities, the state of the world is fully described by the aggregate shock and the two idiosyncratic shocks. We denote the vector that contains these three variables with s_t and assume that it is Markov with a transition probability matrix Γ. We shall denote with $s^t = \{s_1, s_2, ..., s_t\}$ the history of the system up to time *t*.

The assumption of no storage makes the value of autarky, which we will denote by $\Omega^i(e^i(s^t))$, very easy to compute: it will be equal to the present discounted value of the utility of consuming the individual endowment minus some penalty P that the community could impose on the individual who does not comply with the terms of a contract:

$$(2.1)\qquad \Omega^i(e^i(s^t)) = u(e^i(s^t)) + E\left[\sum_{j=1}^{\infty}\beta^j u(e^i(s^{t+j}))\Big|e^i(s^t)\right] - P(s^t)$$

where u is a well-behaved utility function (continuous, concave, and satisfying the Inada conditions) and β is the discount factor assumed to be between zero and one. As the penalty will not play any important role in what we are going to do, we will set it to zero in what follows.

We assume that there is complete information between the two individuals: each observes completely the vector s_t and its history. As the two idiosyncratic shocks are uncorrelated there will be scope for risk sharing. A contract between the two individuals will specify the net transfer from individual 1 to individual 2 as a function of the current history of the world, $\kappa(s^t)$. The value of being in such a contract at time *t*, given the current history of the world, for individual 1 will therefore be:

$$(2.2)\qquad \begin{aligned} &U^1(s^t) = u(c_t^1(s^t)) + E\left[\sum_{j=1}^{\infty}\beta^j u(c_{t+j}^1(s^{t+j}))\Big|e^i(s_t)\right] \\ &c_{t+k}^1(s^{t+k}) = e^1(s_{t+k}) - \kappa(s^{t+k}),\ k = 0,1,2.... \end{aligned}$$

Notice that the contract $\kappa(s^t)$ is, in this framework, the only way the two individuals have to share risk. The evidence reported by Townsend (1994), Udry (1994) and many others, leads to believe that in village economies idiosyncratic shocks are important enough and storage technologies and access to credit limited enough so that even in a more general framework, these contracts constitute an important way to share risk.

In the absence of enforceability problems, the first best allocation of resources can be achieved and the two individuals will share idiosyncratic risk fully. Of course full risk sharing allocations are not unique and will depend on the resources controlled by each individuals, or, if one chooses to characterize them by considering a social planner problem, by the Pareto weights used by the planner.

The empirical implications of full risk sharing have long been recognized and studied. In particular, Townsend (1994) tested on data from some Indian villages the hypothesis that changes in the individual marginal utility of consumption are not affected by changes in individual incomes, which, as we will see below, is one of the implications of complete insurance.

If risk-sharing contracts are not fully enforceable, one has to restrict one's attention to contracts that are self-enforcing. It has become practice in this literature to focus on contracts that are constrained efficient and enforced by the threat to reverse to the worst sub-game perfect equilibrium, which is easily proven to be the autarkic equilibrium.[2] To characterize the self-enforceable risk sharing contracts we can therefore consider a modified social planner problem where, in addition to the resource constraint, the social planner faces two participation constraints. This can be written as follows:

$$
\begin{aligned}
& Max \quad \lambda_1 U^1(s^t) + \lambda_2 U^2(s^t) \\
& s.t. \\
(2.3) \qquad & \mu_t : \quad c_t^1(s^t) + c_t^2(s^t) \leq e^1(s_t) + e^2(s_t) \\
& k_t^1 : \quad U^1(s^t) \geq \Omega(e_t^1) \\
& k_t^2 : \quad U^2(s^t) \geq \Omega(e_t^2)
\end{aligned}
$$

where λ_1 and λ_2 are the Pareto weights assigned by the social planner to the two individuals, μ_t is the multiplier associated to the resource constraint, and k_t^i is the multiplier associated to the enforceability constraint of individual *i*.

Written as in (2.3), the program is not immediately recursive. However, it is possible to re-write it in a different way that makes it recursive and allows one to derive some useful expressions. This approach, proposed by Marcet and Marimon (1999) has been used, among others, by Kehoe and Perri (2002) and Attanasio and Rios Rull (2000).[3] For this purpose, it is useful to define recursively a quantity $K_t^i(s^t) = K_{t-1}^i(s^{t-1}) + k_t^i(s^t)$ with $K_0^i = \lambda^i$. $K_t^i(s^t)$ is, for each individual, the cumulate of the multiplier associated to the enforceability constraint for that individual, using as an initial value the Pareto weight for that individual. With this definition, it is straightforward to show that one can re-write the planner problem as:

$$
\begin{aligned}
& Max \quad E_t \sum_{i=1}^{2} \sum_{t=0}^{\infty} \beta^t \left\{ K_{t-1}^i(s^{t-1}) u(c_t^i(s^t)) + k_t^i [u(c_t^i(s^t)) - \Omega(e_t^i)] \right\} \\
(2.4) \qquad & s.t. \\
& \mu_t : \quad c_t^1(s^t) + c_t^2(s^t) \leq e^1(s_t) + e^2(s_t)
\end{aligned}
$$

[2] The equilibrium concepts used in this literature are those advocated by Abreu, Pearce and Stacchetti (1988).
[3] An alternative solution method is used by Ligon et al. (2001)

Given the representation (2.4), a first order condition for the efficient allocation of resources is given by the following equation:

$$\text{(2.5)} \qquad \frac{u'(c_t^1(s^t))}{u'(c_t^2(s^t))} = \frac{K_t^2(s^t)}{K_t^1(s^t)} = \frac{K_{t-1}^2(s^{t-1}) + k_t^2(s^t)}{K_{t-1}^1(s^{t-1}) + k_t^1(s^t)} \equiv x(s^t)$$

The variable x, which is equal to the ratio of marginal utilities of consumption of the two individuals, fully represents the evolution of the system. If the enforceability constraints are never binding, x is a constant and equal to the ratio of Pareto weights. This is the standard result one gets under full risk-sharing. When one of the enforceability constraints is binding, the ratio of marginal utilities will change. Notice that, as it is not possible that both constraints are binding at the same time, at each time t at least one of the two multipliers $k_t^i, i = 1,2$; is going to be equal to zero. Therefore, at each point in time, it will be possible to establish whether the variable x is increasing or decreasing. Moreover, equation (2.5) implies that the marginal utility of the constrained individual increases by less than the marginal utility of the unconstrained individual. This fact can be used to establish, on the basis of observable variables, which of the two consumers is constrained. One can interpret the ratio of marginal utilities as the relative weight that the social planner gives to the two individuals. In other words, the planner compensates an individual whose enforceability constraint is binding with an increase in her relative weight.

The model can generate some interesting dynamics. It is possible, for instance, not only that the amount of net transfers is (in absolute value) below the level implied by first best, but that it has the opposite sign. In other words, we can have situations in which the individual who is relatively less lucky makes a transfer to her luckier partner. This happens if the relative weight of an individual in the social planner problem has declined so much that even if she receives a smaller endowment than that of her partner, she is asked to make a transfer.[4] This might happen when an individual has received from her partner transfers for some time. This example illustrates well the sense in which the contracts we are considering are half-way between debt and insurance. There are states of the world in which the optimal contracts keeps track of past transfers and makes agent 'repay' their debts. However, when

[4] Notice that in a two-person context, it is always the person making a net positive transfer whose enforceability constraint might be binding. A very clear discussion of these issues can be found in Ligon et al. (2001).

the state of the world is such that the identity of the constrained individual changes, the optimal contract erases all past debts (see the discussion in Ligon et al. (2001) on this).

2.2 *Extensions*

There are several ways in which one can add storage to the model. We will assume that two agents in our model have access to a storage technology that allows them to transform consumption at time t into consumption at time t+1 at a rate R, which is assumed constant. In principle we could also add a 'liquidity constraint' stating that storage cannot be negative and capturing the fact that most of the agents in village economies have only very limited access to external credit. This would complicate the notation further.

As pointed out by Kehoe and Perri (2002), equation (2.5) will hold even in the presence of storage. The extent to which storage changes the nature of the equilibrium depends on how storage affects the value of autarky. In general, unless storage is communally held, the amount of private storage will affect the value of autarky, which will now denote with $\Omega(e_t^i, A_{t-1}^i)$. Adding storage to the model also modifies the resource constraint faced by the planner in equation (2.3) and (2.4). In particular, if we denote by A_t^i the amount of storage held by individual *i,* the aggregate resource constraint will be:

$$(2.6) \qquad \mu_t: \quad c_t^1(s^t) + c_t^2(s^t) + A_{t+1}^1 + A_{t+1}^2 \le e^1(s_t) + e^2(s_t) + (1+R)(A_t^1 + A_t^2)$$

Given this resource constraint, it will be possible to derive an Euler equation for each of the two consumers:

$$(2.7) \qquad u'(c_t^i(s^t)) = \beta E_t\left[(1+R)(1+\frac{k_{t+1}^i}{K_t^i})u'(c_{t+1}^i(s^{t+1})) - \frac{k_{t+1}^i}{K_t^i}\frac{\partial\Omega(e_{t+1}^i, A_t^i)}{\partial A_t^i}\right];\ i = 1,2$$

Notice that when the participation constraint is not binding for individual *i,* equation (2.7) reduces to a standard Euler equation of the kind analyzed in the consumption literature. When the multiplier k_{t+1}^i is positive, that is when the enforceability constraint for individual *i* is binding, it introduces a distortion in the intertemporal allocation of resources for such individual. The multiplier enters twice, reflecting two different effects. Moving resources to period *t+1* makes individual *i* better off and therefore relaxes the enforceability constraint.

On the other hand, an increase in the private storage of individual *i* also increases the value of autarky and therefore make the enforceability constraint worse. However, given that equation (2.5) holds, we are able to sign the difference between the two effects: for the unconstrained individual the expected rate of growth of marginal utility should be higher than for the constrained individual. Moreover, for the unconstrained individual the expected rate of growth of marginal utility is $1/\beta(1+R)$, implying that

$$E_t\left[\frac{k_{t+1}^i}{K_t^i}\left(\frac{u'(c_{t+1}^i)}{u'(c_t^i)}-\frac{\partial\Omega(e_{t+1}^i,A_t^i)}{\partial A_t^i}\right)\right]>0\,.$$

The case with two individuals is particularly simple because at most one of the two enforceability constraints can be binding at any given time. However, even when we have more than two individuals, there must be at least one for whom the constraint is not binding. In general, there will be a group of individuals for whom the constraint is not binding and whose marginal utilities of consumption change at the same rate, and another for whom the constraint is binding and for whom (if the first effect discussed in the previous paragraph is stronger) the marginal utility grows less than that of the unconstrained group.

Notice also that, as noted by Kocherlakota (1996), if an individual is given by the social planner a certain Pareto weight (implied by the value of its current marginal utility of consumption relative to the other individuals in the economy), and, on entering period t+1 the state of the world is such that she is constrained (in that her participation constraint is binding), the new allocation of consumption will be determined uniquely by such participation constraint. That is, the central planner will give to her enough consumption and continuation utility to make her current value at least as large as the value of autarky. Therefore, the value of consumption for such an individual will depend only on the current state of the world and on other individuals lagged marginal utility. The forgiveness of the system, which is crucial for the constrained efficiency of the contract, will imply that her current consumption will be independent of her lagged marginal utility. In section 3.3, we discuss how to exploit this property and that discussed in the previous paragraph to construct a test of the model.

Most results obtained in the two-individual model extend to the case in which we have many individuals, as long as the complete information structure is preserved. To consider many individuals, it is useful to modify the planner problem and write it as the maximization of the utility of an arbitrarily chosen individual given the weights given to the other individuals subject to participation and promise-keeping constraints. This is the framework used, for instance, by Ligon et al. (2001).

2.3 *Characterizing properties of the equilibria*

Two natural benchmarks to compare the allocation of resources implied by the equilibrium contract are the autarkic and the first best one. While the equilibrium contract can coincide with one of these two, in general will call for some limited amount of risk sharing. The effect of changing some of the features of an economy on the likelihood that the equilibrium allocation coincides with first best or autarky is reasonably intuitive and has been discussed in the literature. Proposition 4.9 in Alvarez and Jermann (2000) and Proposition 2 in Ligon et al. (2001) contain the main comparative static results. In particular, it is easy to prove that for values of the discount factor that are high enough, it is always possible to implement the first best allocation. The same is true for the coefficient of relative risk aversion: enough curvature in the utility function will deliver first best equilibria. The intuition behind the two results is reasonably obvious: as the punishment for deviating from the optimal contract is in the future and consists in the exclusion from some risk sharing agreement, more patient and more risk averse consumers will find such punishment harsher. It will therefore be easier to enforce risk-sharing agreements. At the other extreme, there will be values of the discount factor and of risk aversion that will imply the existence of no risk-sharing in equilibrium and no equilibrium other than the autarkic one.

Similarly intuitive results can be derived if one considers changes to the environment individuals face. In particular, a decrease (increase) in the variance of individual and aggregate income will make people better (worse) off in autarky and give more scope for risk sharing. It is therefore not surprising that low enough values of the income variance will yield only the autarkic equilibrium, while high enough value of the same parameter will yield full risk sharing. The intuition is, once again, in what happens to the value of autarky when one changes the variance of the income process. A higher variance implies stronger incentives to stay in a risk-sharing agreement.

It is also straightforward and intuitive to see what happens when one changes the persistence of idiosyncratic shocks. More persistent shocks are harder to insurance as an individual will resist sharing her permanent luck. In other words, it will be possible to find a high enough level of persistence that will cause autarky to be the only equilibrium outcome.

While most results in the literature are stated in terms of a how varying a given parameter can give rise to either full risk sharing or (on the other extreme) to autarkic equilibria, it is useful, especially if one wants to test the empirical relevance of the model, to consider measures of how far a given allocation is from first best and how this changes when changing taste and technology parameters.

In some situations, namely in models with two individuals and simple error structures, a measure of the amount of risk sharing is quite natural and the results we have mentioned so far generalize in that the convergence of the economy from autarky to first best when some parameters change is 'monotonic'.[5] (see Alvarez and Jermann (2001)). However, proving these results in more general situations or deriving the particular functional form between some parameters and the amount of risk sharing can be hard. For this reason, after considering a measure of the amount of risk sharing in an economy, in Appendix A, we use simulations to characterize these relationships.

In establishing a measure of the amount of risk sharing, we want to establish a metric that allows us to measure the 'distance' of a given observed intertemporal allocation from the benchmark of full risk sharing. In other words, we want to measure how much risk sharing happens in equilibrium compared to full risk sharing.

One possibility would be to compare equilibrium transfers with the transfers that one would observe under first best. While this measure has an intuitive appeal, it also has some important drawbacks. First, it ignores other mechanisms (such as saving) that individuals might be using at a given point in time, to get close to first best allocations. The relationship between individual savings (and borrowing) and private transfers could be quite complex, especially in models where the effect of the stock of (private) saving on the value of autarky can be important. Second, as we discuss above, even in the absence of storage, in our model, enforceable transfers might have the opposite sign of first best transfers, making the ratio measure less meaningful. Because of these two problems, we focus on a different measure.

In a symmetric first best equilibrium full risk sharing implies complete equality across individuals. On the contrary, autarky implies that a large fraction of shocks to endowments

[5] What we mean by 'monotonic' is that an increase in a given parameter either increases or decrease the amount of risk sharing up to the point where the equilibrium contract coincides with either first best or autarky.

or incomes are reflected into consumption. In the case in which individuals have no access to self insurance the cross sectional variance of consumption will actually be equal to the value of income. Following Attanasio and Rios Rull (2000), we can therefore take the ratio of the cross sectional variance of consumption to the cross sectional variance of income as our measure of the amount of risk-sharing achieved by a given economy. Under full risk-sharing (and with symmetric endowments), such a measure is zero. Under autarky (and without storage) such a measure is 1.

As we mentioned above, the first order conditions that one gets from the central planner problem, are silent about the determination of the Pareto weights. A competitive equilibrium is not necessarily symmetric. However, given a set of Pareto weights, which might generate a non-zero cross-sectional variability of consumption in equilibrium, under first best, the cross sectional variance of consumption will not change. Instead, under a self-enforceable risk sharing contract, the cross sectional variance of consumption will vary if the participation constraint become binding. Therefore, to accommodate the possibility of asymmetric competitive equilibria, we use the changes in the cross sectional variance of consumption as our measure of risk-sharing.

A certain change in the cross sectional variance of (log) consumption might mean very different things in terms of the amount of risk sharing that is achieved in a given economy. In a given economy (village) the cross sectional distribution of consumption might change very little even if there is no risk-sharing if the cross sectional distribution of incomes does not change. On the other hand, one could observe a village where a substantial proportion of idiosyncratic risk is diversified and yet the cross sectional variance of consumption changes substantially. In other words, to construct a proper measure of the degree of risk sharing of a given economy, changes in the cross sectional variance of consumption should be normalized by something that takes into account the need for risk sharing of a given economy. We propose to normalize it by the value of the variance of income. Simple simulations, whose results we report in Appendix A, show that such a ratio behaves is extremely correlated with the ratio of variances that we discussed in section 2.[6]

[6] We have also experimented with normalizing the changes in the variance of consumption with the *changes* in the variance of income. The results we obtained both in the simulations and in the empirical exercise were very similar.

We can therefore confidently take this ratio as a measure of the distance of a given economy's allocation from that implied by perfect risk sharing.

To obtain a relationship between our measure of risk sharing and the parameters of the model, we start from a baseline economy and shock randomly the parameters that characterize the endowment process to generate many different economies. We can then solve for the equilibrium contract and simulate each of these economies to study the relationship between efficient allocations and the endowment parameters. In particular, we focus on the measure of risk sharing we use in the empirical analysis (see the discussion in Section 3.1).

In Appendix A, for power utility, we show the following results:

(i) an increase in the variance of income (either idiosyncratic or aggregate) increases monotonically risk sharing;

(ii) an increase in the persistence of idiosyncratic income decreases monotonically risk sharing;[7]

(iii) an increase in the mean of the process (keeping the variance constant) decreases risk sharing;

In section 3.1, we discuss how to make our measure of risk sharing operational and we show how to use these results to propose a test of the model with imperfect enforceability.

3 The proposed tests

The main aim of this paper is to propose two different tests of the empirical relevance of models with imperfect enforceability. Our tests consider some explicit implications of the model we sketched above and differ in that the first is based on village level statistics, while the second relies on the explicit estimation of the Euler equations implied by the model.

3.1 Changes in the cross sectional variance of the marginal utility of consumption and village characteristics

The starting point of our first approach is a test of the perfect insurance model and first best allocation. As is well known, if one considers the optimization problem of a social planner

[7] Attanasio and Rios-Rull (2000b), using simulations, show that an increase in the persistence of aggregate shocks that leaves the mean and the variance of the process unaffected, will cause less risk sharing in equilibrium.

that maximizes expected utility of a set of individuals with a give set of Pareto weights, one can derive the following first order condition:

$$U_c(c_t^i(s^t), z_t^i(s^t))\lambda^i\beta^i = \mu(s^t) \tag{3.1}$$

where U_c is the marginal utility of consumption for individual *i,* which is assumed to depend on its non-durable consumption and, possibly, on a vector of other variables *z*, some of which might be unobservable. λ^i is the Pareto-weight given by the social planner to individual i in the maximization problem. Different sets of weights will correspond to different competitive equilibria with full risk sharing. The theory is silent about what determines these weights, except in saying that they are constant over time. β^i is the discount factor for individual *i,* and $\mu(s^t)$ is the Lagrange multiplier associated with the resource constraint at time *t.* Since the work of Townsend (1994), Mace (1991) and Cochrane (1991), equation (4.1) has been used extensively to test the empirical implications of a model with full risk sharing. If one takes the log of equation (3.1) and considers it at two different time periods, one can eliminate both the discount factor and the unobserved Pareto weight. In particular, one gets:

$$\log(U_c(c_t^i(s^t), z_t^i(s^t))) - \log(U_c(c_\tau^i(s^\tau), z_\tau^i(s^\tau))) = \log(\mu_t) - \log(\mu_\tau) \tag{3.2}$$

Notice that times t and τ need not be adjacent periods.[8] If they are k periods apart, one can write equation (3.2) as:

$$\Delta^k \log(U_c(c_t^i(s^t), z_t^i(s^t))) = \nu(s_t, s_{t-k}) \tag{3.3}$$

The main implication of equation (3.3) is that changes over time in the marginal utility of different individuals should be the same. Differencing has eliminated the Pareto weights. The changes in marginal utility should be unaffected by the idiosyncratic shocks received by individuals. Full insurance means that changes in the amount of resources available to an individual over and above the aggregate change should not be reflected in changes in marginal utility. The resource constraint can be taken into account by considering either time

[8] Moreover, if one considers many pairs of observations, the distance between the two time periods need not be the same.

dummies, or, as in Mace (1991), aggregate consumption. Any other variable, such as individual income, should therefore not enter equation (3.3).

In the absence of panel data that follow the same individuals over time, one can still test the implications of perfect risk sharing by aggregating equation (3.3) over individuals belonging to a given group whose membership is assumed to be fixed over time. Insurance across groups implies that the average marginal utility of consumption for different groups should change in the same way and should be unaffected by group level shocks. Attanasio and Davis (1996) used synthetic panels to test equation (3.3) by forming year of birth and education groups and following the averages for these groups over time. The synthetic panel approach has two big advantages and one disadvantage. The advantages are the possibility of testing equation (3.3) even in the absence of longitudinal data and the gains in power that might be obtained averaging measurement error in wages, income or whatever measures of individual resources are used over the members of a group. The disadvantage is the fact that one focuses only on the insurance across groups. By taking averages over the members of a group, one cannot say anything about the extent of risk sharing within a group.

An alternative test of the perfect insurance model can be constructed by considering, instead of the cross sectional average of the variables in equation (3.3), their cross sectional variance. Consider once again the log of equation (3.1) above, and re-write it as follows:

$$(3.4) \qquad \log(U_c(c_t^i(s_t), z_t^i(s_t))) = \log(\mu(s_t)) - \log(\lambda^i \beta^i)$$

Having defined groups with fixed membership one can compute the cross sectional variance of both sides of equation (3.4). As the resource constraint multiplier is common across individuals it does not contribute to the cross sectional variance. Under perfect risk sharing, the Pareto weights and the discount factors are constant over time: therefore an implication of the theory is that the cross sectional variance of the marginal utility of consumption is constant over time.

If we assume that the utility function is isoelastic, equation (3.4) can be written as:

$$(3.5) \qquad Var_v(\log(c_t^i)) = d_v$$

where the subscript v indicates the fact that the variance is computed within a group (village), the d are group dummy variables and reflect the variance of Pareto weights and discount factors. Taking first differences of equation (3.5) one gets:

$$\text{(3.6)} \qquad \Delta Var_v(\log(c_t^i)) = 0$$

Equations (3.5) can be tested by regressing the cross sectional variance of the marginal utility of consumption on group dummies and other variables (such as income shocks indicators or the variance of income) and test that the coefficients on these other variables are not different from zero. Analogously, equation (3.6) could be tested regressing the changes in the cross sectional variance of consumption on income shocks indicators or the changes in the income variance and test the hypothesis that these variables have no systematic effect.

While this implication had been noted in passing by Deaton and Paxson (1994), Attanasio and Szekely (2001) and Attanasio (2002)[9] propose and implement explicit tests of such a hypothesis. In this paper, we use this approach as a starting point to study the implications of the model with imperfect enforceability. Before moving on, however, one point that should be noticed here is that one can readily think of situations in which the null of full insurance is violated but in which the test based on variance changes fails to pick up this violation. Another implication of the null is that individuals maintain their relative position in the distribution of the marginal distribution of consumption. It is possible that individuals change their position (therefore violating the null), while leaving the cross sectional variance unaffected. In a recent paper, Jappelli and Pistaferri (2001) test the implication of the perfect insurance model of no mobility in the cross sectional distribution of consumption by constructing the Shorrocks mobility index in a panel of Italian households. While they present convincing evidence rejecting the implications of full risk sharing, they do not consider the implications of alternative models.

The test based on the evolution of the cross sectional variance of consumption is complementary to those based on means (such as Townsend (1994) and Attanasio and Davis (1996)). This is both because it focuses on the insurability of shocks within a group (or village) rather than across groups and because in some circumstances could be more

[9] See also Attanasio and Jappelli (2001).

powerful than the test based on means. An example of the latter situation is when large measurement in income reduces the power of the test based on income, while, if the variance of measurement error does not change over time, does not affect the test based on variances.

The expression on the left-hand side of equation (3.6) is similar to our measurement of the degree of risk-sharing we discussed in Section 2.3. There we pointed out that the farer away is the intertemporal allocation of consumption from first best, the larger are the changes in the cross sectional variance of (the log of the marginal utility of) consumption.

The next step is then to use this measure of the distance from perfect risk sharing,[10] to relate it to variables that, according to the model with imperfect enforceability are important determinants of the deviation from perfect risk sharing. In particular, we consider the variance of income, its autocorrelation and its mean. As mentioned in Section 2, economies with higher values of income variance should achieve greater degrees of risk sharing, while the opposite is true for economies where the income processes are highly persistent. The effect of the mean, keeping constant the variance of log income and its autocorrelation, depends on the form of the utility function.

Given these considerations, we consider the following relationship:

$$\frac{\left|\Delta Var_v(\log(c_t^i))\right|}{Var_v(\log(y_t^i))} = \alpha_0 + \alpha_1 st.dev.(\log(y_t^i)) + \alpha_2 \operatorname{var}(\log(y_t^i)) + \alpha_3 \rho(\log(y_t^i)) + \alpha_4 \rho^2(\log(y_t^i)) + \alpha_5 \mu(\log(y_t^i)) + \alpha_6 \mu^2(\log(y_t^i)) \tag{3.7}$$

where ϱ and μ are the autocorrelation and mean of (log) income respectively. As the relationship between our measure of deviations from first best and the features of the village we consider is not necessarily linear, we consider a flexible functional form and estimate equation 3.7 by LAD. The model with imperfect enforceability implies that the overall effect of income variance is negative while that of the autocorrelation is positive. In section 4 and in the Appendix, we describe in detail how we obtain estimates of the various quantities in equation (3.7).

[10] Such a measure is not perfect. As we mention above, there might be situations where the hypothesis of perfect risk sharing is violated and yet our measure is equal to zero.

3.2 A test based on Euler equations

The second test we consider is based on the Euler equation that we derived in section 2.2 for consumers who participate in a risk sharing agreement subject to enforceability constraints. We first describe the test considering the model with two individuals and no measurement error. We then relax these two assumptions. Assuming a CRRA utility function, we can re-write equation (2.7) above as:

$$(3.8) \qquad 1=\beta E_t\left[(1+R)(1+\frac{k_{t+1}^i}{K_t^i})\left(\frac{c_{t+1}^i}{c_t^i}\right)^{-\gamma}-\frac{k_{t+1}^i}{K_t^i}\frac{\partial\Omega(e_{t+1}^i,A_t^i)}{\partial A_t^i}\right];$$

where k_{t+1}^i is the multiplier associated with the participation constraint of individual i and $K_t^i=\sum_{\tau=0}^{t}k_\tau^i$. As mentioned in Section 2, if in period $t+1$ the enforceability constraint for individual i is not binding, (3.8) reduces to a standard Euler equation for the intertemporal allocation of consumption of an individual with a storage technology that pays an interest rate of R, and a discount factor of β. If there are only two individuals, only one of the two enforceability constraints can be binding at any point in time and at any state of the world. Therefore, at each point in time, for at least one of the two consumers, there will be an Euler equation holding. Moreover, if we have access to perfect data on consumption, we can easily establish which is the consumer who is potentially facing a binding constraint. If none of the two constraints is binding, the first best allocation corresponding to a given set of Pareto weights can be implemented and the marginal utility of the two consumers will grow at the same rate. However, if the constraint of consumer one is binding, her marginal utility will be growing less and, therefore, her consumption will be growing faster. The opposite is true if the constraint of the second consumer is binding.

If at time $t+1$ the constraint for the first consumer is binding, not only her consumption at time $t+1$ will be growing faster, but it will be fully determined by the upgrade in the continuation utility determined by the social planner. As noted by Kocherlakota (1996), an implication of this is that for the constrained individual, consumption at time $t+1$ will be a function of her shock, her partner's shock and her partner lagged marginal utility of consumption, but not of her own lagged marginal utility of consumption.

This discussion implies that we can write a switching regression system for the consumption growth of the two individual. Each consumer is in one of two regimes. If we denote with z_{t+1}^1 the consumption growth of consumer 1 minus that of consumer 2, with ε_{t+1}^i the income shock of individual *i,* we have:

$$\begin{array}{ll} \Delta\log(c_{t+1}^1) = const. + r + u_{t+1}^1; & if\ \ z_{t+1}^1 \leq 0 \\ \log(c_{t+1}^1) = g(\varepsilon_{t+1}^1, \varepsilon_{t+1}^2, \log(c_t^2)); & if\ \ z_{t+1}^1 > 0 \end{array} \tag{3.9}$$

where *g* is a function and *r=log(1+R).* The strong testable implication of equation (3.9) is that $\log(c_{t+1}^1)$ does not depend, after controlling for current shocks, on $\log(c_t^1)$. This is a strong implication of constrained efficiency stressed by Kocherlakota (1996) and follows from the lack of memory of the system once a constraint becomes binding. One problem in testing the switching regression system in (3.9) is that we do not have, in general, a closed form solution for the function *g*. A possibility, therefore, is to approximate *g* by a flexible function of its arguments, such as a polynomial.

To make this idea operational and apply it to our context, we have to extend the test to consider many individuals and the possibility of measurement error. We start with the latter. In the presence of errors in the measurement of consumption, it is not possible to exactly classify the observations in the two regimes in equation (3.9). However, if one writes down an explicit model of measurement error, one can establish, for each observation, the probability of each of the two regimes and therefore identify the model. In particular, let us assume a classical multiplicative measurement error in the levels of consumption, so that measured consumption $C_t^h = c_t^h v_t^h$, where v is a lognormal measurement error. We further assume that the measurement error process is i.i.d. across individuals and over time with a standard deviation of σ_v. Denoting with Z_{t+1}^1 the measured equivalent of z_{t+1}^1, we have:

$$\begin{array}{ll} \Delta\log(C_{t+1}^1) = const. + \log(1+r) + u_{t+1}^1 + \Delta v_{t+1}^1 & with\ \ prob.\ 1-\Phi(Z_{t+1}^1 / 4\sigma_v) \\ \log(C_{t+1}^1) = g(\varepsilon_{t+1}^1, \varepsilon_{t+1}^2, \log(c_t^2), v_{t+1}^2) & with\ \ prob.\ \Phi(Z_{t+1}^1 / 4\sigma_v) \end{array} \tag{3.10}$$

where Φ is the normal cdf.

Next, we can generalize this approach to the case in which we have many consumers. In this case, at least one consumer will not be constrained. More generally, we will have a set of unconstrained consumers, whose marginal utility of consumption will be growing at the same rate and the remaining consumers, who will have binding participation constraints and marginal utility of consumption will be growing at a rate lower than that of the unconstrained consumers (but not constant). Once again, with perfect information about consumption it will be possible to classify each consumer in one of two sets and to write down equations similar to (3.9). With measurement error, we have to generalize equation (3.10). We have an additional parameter to estimate that is the unobserved rate of growth of consumption for the unconstrained consumers. For a generic consumer *h* in village *v,* we have the following system:

$$(3.11)\qquad \begin{array}{ll} \Delta\log(C^{h,v}_{t+1}) = const. + r + u^2_{t+1} + \Delta v^2_{t+1} & w.\ prob.\ 1-\Phi((\log(C^{h}_{t+1})-\mu_{v,t})/2\sigma_v) \\ \log(C^{h,v}_{t+1}) = g(\varepsilon^{h,v}_{t+1}, \underline{\varepsilon^{v}_{t+1}}, \underline{\underline{\log(c^{v}_{t})}}) & w.\ prob.\ \Phi((\log(C^{h}_{t+1})-\mu_{v,t})/2\sigma_v) \end{array}$$

where $\mu_{v,t}$ is the unobserved mean consumption growth for the unconstrained households in village *v* at time *t*, and $\underline{\varepsilon^{v}_{t+1}}, and\ \underline{\underline{\log(c^{v}_{t})}}$ are the vector of time *t+1* shocks for all the other households and the vector of time t log consumption for the unconstrained households in village *v*. Once again notice that the testable restriction implied by equation (3.11) is the fact that *t+1* consumption for the constrained households does not depend on the lagged value of their own consumption.

The empirical implementation of (3.11) implies two additional problems. First, $\underline{\varepsilon^{v}_{t+1}}, and\ \underline{\underline{\log(c^{v}_{t})}}$ are potentially highly dimensional objects, as villages can have a few hundreds households. Second, the vector of parameters $\mu_{v,t}$ can also be of very high dimension, as we have several hundreds villages and several time periods.

To tackle the first problem, we summarize the distribution of shocks in the households, and lagged marginal utilities of the currently unconstrained households by a few distributional statistics. These statistics, the mean, variance, skewness and kurtosis of the distributions, can be estimated along with the other parameters of the system. Moreover, we can choose to implement this exercise in each village separately or pooling all the villages and make the distributional statistics function of observable village characteristics in order to keep the dimension of parameters manageable.

To solve the second problem, we notice that, at a point in time and for each village, the average consumption growth in the village constitutes an upper bound to $\mu_{v,t}$, which is the rate of consumption growth of unconstrained individuals. By how much the village aggregate growth of consumption exceeds $\mu_{v,t}$ depends on how many consumers are constrained and how binding are the constraints. This in turn depends on the distribution of income. We therefore make the following assumption:

$$\text{(3.12)} \qquad \mu_{v,t} = \varphi(\underline{y_t^v}, \underline{c_{t-1}^v}) gc_t^v$$

where φ is a function that takes values between 0 and 1 and gc_t^v is the aggregate growth of consumption in village *v* at time *t*. $\underline{y_t^v}$ *and* $\underline{c_{t-1}^v}$ are the income shock and lagged marginal utility distribution in village *v*. We approximate the function φ as a logistic function in the moments of the income and marginal utilities distributions.

An alternative way to write equation (3.11), which incorporates the assumptions and approximations discussed so far, is therefore the following:

$$\text{(3.13)} \quad \Delta\log(C_{t+1}^{h,v}) = \begin{cases} const + r + u_{t+1}^2 + \Delta v_{t+1}^2 & w.\ prob.\ \ 1 - p_{v,t} \\ const + r + u_{t+1}^2 + \Delta v_{t+1}^2 + f(\varepsilon_{t+1}^{h,v}, \underline{\varepsilon_{t+1}^v}, \underline{\underline{\log(c_t^v)}}) & w.\ prob.\ \ p_{v,t} \end{cases}$$

where $p_{v,t} = \Phi((\log(C_{t+1}^h) - \mu_{v,t})/2\sigma_v)$, $\mu_{v,t} = \dfrac{1}{1+\exp(\theta m_{v,t})} gc_{v,t}$, *m* a vector of moments d $gc_{v,t}$ is the observed growth of consumption in village *v* at time *t*. Notice that the second line of equation (3.13) has the same left-hand side of the first. The *f* function is simply the multiplier associated to the participation constraint and is strictly positive for constrained individuals. For equation (3.13) to reduce to equation (3.11) for the constrained individual, the *f* function has to be of the form:

$$\text{(3.14)} \quad f(\varepsilon_{t+1}^{h,v}, \underline{\varepsilon_{t+1}^v}, \underline{\underline{\log(c_t^v)}}) = g(\varepsilon_{t+1}^{h,v}, \underline{\varepsilon_{t+1}^v}, \underline{\underline{\log(c_t^v)}}) - const - r + \pi \log(c_t^v)$$

with the coefficient π=1. This restriction, which can easily be tested, is required by the fact that lagged consumption (marginal utility) has to cancel out from equation (3.13) to get (3.11). As we mentioned above, the fact that lagged own marginal utility does not determine current consumption for constrained individuals is a very strong restriction that is at the heart of the limited enforceability model, as it reflects the constrained efficiency of the self-enforcing mechanism.

An additional restriction that can be tested is that current income does not enter the equation for the unconstrained individuals. This restriction is analogous to the excess sensitivity tests in the Consumption Euler equation literature.

4 Data

The tests proposed in the previous section are carried out in this paper using a unique data set: the evaluation sample of a large welfare program in rural Mexico, called PROGRESA. As we discuss below, PROGRESA is a program aimed at fostering the accumulation of human capital by increasing school enrolment, and improving nutrition and health practices. In order to evaluate this program, the administration collected information on all the households living in 506 localities (about 25,000 households were surveyed) before and after the implementation of the program. Furthermore, in a subset of 186 randomly chosen localities the program implementation was delayed for almost two years. This randomization scheme was explicitly implemented with the purpose of facilitating the evaluation of the effectiveness of the program.

4.1 The PROGRESA program and the evaluation data set

In 1997 the Mexican government decided to start the Education, Health and Nutrition Program, called PROGRESA (for its Spanish acronym). PROGRESA is a new and large welfare program targeted to rural communities. It aims to raise the living standards of very poor families by means of those three closely linked components (health, food and education); the underlying idea is that inter-actions amongst them enhance the effectiveness of an integrated program.

The health component consists of a number of initiatives oriented to provide better information about vaccination, nutrition, contraception and hygiene and of a program of visits for children and women to health centres. Participation into the health component is a pre-condition for participating into the nutrition component; this gives, in addition to a basic monetary subsidy received by all beneficiary households, some in kind transfers (nutrition supplements) to households with very young infants and pregnant women. The largest component of the program is the education one. Beneficiary households with school age children receive grants conditional on school attendance. The size of the grant increases with the grade and, for secondary education, is slightly higher for girls than for boys. In addition to the payments, beneficiaries with children in school age receive a small annual grant for school supplies. Finally, the mother in the household receives every two months all the transfers. The benefits of the program constitute a substantial help: over 20% of the beneficiaries average income.

The selection of eligible households is a multi-stage process. The Program first targeted the poorest communities in rural Mexico. Roughly speaking, the two criteria communities had to satisfy to qualify for the program were a certain degree of poverty (as measured by what is called an 'index of marginalization') and access to certain basic structures (schools and health centers). Once a locality qualifies, following a census in the relevant localities in 1997, individual households could qualify or not for the program, depending on a single indicator that is affected by a number of poverty variables (income, house type and so on).

The program was phased in slowly and is currently very large: at the end of 1999 its budget was US$777m and was implemented in more than 50,000 localities. At that time, about 2.6 million households, or 40% of all rural families and one ninth of all households in Mexico, were included in the program. The cost of the program is about 0.2% of Mexican GDP. The program has received a considerable amount of attention and publicity and similar programs are currently being implemented in Honduras, Nicaragua and Argentina. (See IFPRI (2000) for additional details on the program and its evaluation).

The agency running the program used the fact that, for logistic reasons, the program could not be started everywhere simultaneously, to start an evaluation sample. Among the beneficiaries localities, 506 where chosen randomly and included in the evaluation sample. Among these, 320 randomly chosen were assigned to the communities where the program started early, while 186 were assigned to the communities where the program started almost in December 1999 rather than May 1998.

The randomization seems well executed: Behrman and Todd (1999) present evidence in this respect. In particular, most variables seem to be not statistically different between the treatment and control villages.[11]

4.2 The data.

The data we use comes from the Survey of Household Socio-Economic Characteristics (ENCASEH) and from the Evaluation Survey (ENCEL) of PROGRESA. The Survey of Household Socio-Economic Characteristics was carried out before the program began at the end of 1997. In fact, this census was used to select which households in the eligible communities would participate in PROGRESA. On the other hand, the Evaluation Survey was specifically designed for the evaluation. First, the about 25,000 households of the evaluation sample (from 506 villages) were interviewed in May 1998 for a baseline survey (which complemented the ENCASEH census). Then, follow-up surveys were carried out every six months: November 1998, March 1999, November 1999 and April 2000. Within each village in the evaluation sample, the survey covers all the households and collects extensive information on consumption, income, labour supply, school enrolment, transfers, and a variety of other issues. While each instrument contained a core questionnaire, some of them also contained some additional modules.[12] For instance, the question on inter-household transfers were only asked in the November surveys. The household surveys are supplemented by a locality questionnaire that provides information on several village specific variables, including, among other things, information on prices of various commodities and village level shocks.

We use data on household consumption from the waves of November 1998, March 1999 and November 1999. We have added up across the different groups of commodities on which the surveys provide information: food, transportation, hygiene and health, education, clothing and durable goods. As, depending on the category, households are asked about expenditures made during the previous week, month, semester or year, we first converted all

[11] Of course one would expect a 5% of rejections. It is a bit worrying that one of the few variables for which a rejection was obtained was pre-program school enrollment.

into weekly basis. For food, we have included in our measure an estimate of the consumption of goods produced by the household valued at market prices.

In the surveys, each household member is asked to report separately her/his income from several sources. We build our measure of household income by adding-up the following sources for each household member: wages (both from a primary and from a secondary occupation), net profits (revenues minus expenditures) from businesses or income generating activities, communitary earnings, and income from pensions, from interests, and from rental (of land, animals or machinery).[13] As the respondent can report daily, weekly, biweekly, monthly or annual income, all reported income is first converted into weekly income; we used information on the number of days worked in the previous week to compute this weekly income when daily earnings were reported.

For reasons that will come clear later, we focus on households providing information on income and consumption for all the rounds available; this leaves us with a bit more than ten thousands households. In Table 1, we present some summary statistics relating income and consumption:

Table 1. Summary Statistics for Weekly Household Income and Consumption

Wave		Mean	Median	Standard Deviation
November 1998	Consumption	157.59	136	111.24
	Income	267.73	201	236.70
March 1999	Consumption	159.02	136	112.28
	Income	277.62	210	232.69
Number of Observations		10318		

To implement some of the second of the tests we discussed in the previous section, we need to compute several moments of the distribution of consumption and income in each of the villages. First, we need to compute the cross sectional variance of (log) consumption and income, as they will form our measure of the amount of risk sharing. Second, we need to compute the time series variance and autocorrelation of idiosyncratic income in each village (in addition to its mean), as they will constitute the explanatory variables of our regressions.

[12] Nevertheless, it is also true that some questions are slightly different from one wave to another.

In what follows we use the two central waves to compute the changes in the variance of consumption and income, while we use all four waves to compute the variance, mean and autocorrelation of income.

The cross sectional variance of income is informative about the risk faced by an individual income only under special circumstances. First, heterogeneity across individuals is not necessarily an indication of risk. Second, movements over time might be predictable and therefore could not be defined as risk. In an attempt to keep these issues into account, we first regress income on a set of observable and predictable factors. We then compute the time series variance and autocorrelation of the estimated residuals for each individual. Finally we compute the average of these individual variances and autocorrelations in each village. These are the measures we use in our regressions. Appendix B provides additional details on this procedure.

5 Results

5.1 Changes in the cross sectional distribution of marginal utilities and imperfect enforceability

In this subsection, we implement the first test described in Section 3. In particular, we estimate a flexible relationship between a measure of the amount of risk sharing observed in a village (compared to the amount of risk sharing that would occur under perfect risk sharing) and some features of the villages that can explain this departure under imperfect enforceability: the variance of income, its persistence and its mean (see equation 3.7). Following the discussion in subsection 3.2, our measure of how far a village is from perfect risk sharing is given by the absolute value in the change of the cross sectional variance of consumption (normalised by the cross sectional variance of income). We call this ratio *rs*. We have calculated the ratio *rs* using three different definitions of consumption: a) food, b) food and services (hygiene, health, education, and energy), and c) food, services, transportation, and durable goods (including clothing). As commented in section 4.2, we used data from November 1998, March 1999 and November 1999.

As we do not have a closed form solution for the relationship between our measure of risk sharing and the village level moments of idiosyncratic income, we regress the ratio *rs* on a

[13] Our measure of household income does not include transfers (neither the private transfers nor the

quadratic polynomial in the standard deviation, autocorrelation and mean of the process of individual income. In Table 2, we report the estimates of such a relationship (equation 3.7) obtained by LAD using our data set of Mexican villages. Columns (1)-(3) are different only because of the definition of consumption used to compute the ratio *rs*.

Several observations are in order. First, these results reject the hypothesis of full risk sharing. Remember that, according to this hypothesis, the left hand side of our regression equation should be zero. We observe, instead, that changes in the cross sectional variance of consumption are related to several features of the villages in our sample. Moreover, we find that the pattern of correlations present in the data is not inconsistent with some of the implications of the imperfect enforceability model. The results are very similar across the three columns corresponding to the three definitions of consumption.

Table 3. LAD estimates of

$$\frac{\left|\Delta Var_v(\log(c_t^i))\right|}{Var_v(\log(y_t^i))} = \alpha_0 + \alpha_1 st.dev.(\log(y_t^i)) + \alpha_2 \text{ var}(\log(y_t^i)) + \alpha_3 \rho(\log(y_t^i)) + \alpha_4 \rho^2(\log(y_t^i)) + \alpha_5 \mu(\log(y_t^i)) + \alpha_6 \mu^2(\log(y_t^i))$$

	(1)	(2)	(3)
st.dev.(log(y_t^i))	-0.9193	-1.0983	-1.1175
	(0.3965)	(0.3538)	(0.2298)
var(log(y_t^i))	0.4056	0.5247	0.5601
	(0.2508)	(0.2317)	(0.1359)
ρ(log(y_t^i))	0.6295	0.3425	0.2111
	(0.2747)	(0.2236)	(0.2350)
ρ^2(log(y_t^i))	2.0209	1.2716	1.0790
	(0.8941)	(0.5290)	(0.6303)
μ(log(y_t^i))	-0.0960	-0.0906	-0.0814
	(0.0313)	(0.0302)	(0.0262)
μ^2(log(y_t^i))	0.1807	0.1330	0.0602
	(0.0457)	(0.0495)	(0.0413)
Constant	0.6313	0.6634	0.6417
	(0.1422)	(0.1279)	(0.0889)
Joint significance: F(2, 951)			
α_1, α_2	12.120	14.880	19.590
(p-value)	(0.000)	(0.000)	(0.000)
α_3, α_4	2.690	5.480	3.890
(p-value)	(0.069)	(0.004)	(0.021)
α_5, α_6	9.160	5.730	5.060

program support).

(p-value)	(0.000)	(0.003)	(0.007)
Observations	958: 479 villages, two waves (March 1999 and November 1999)		

Note: standard errors in parentheses.

In particular, we notice that the standard deviation, the autocorrelation and the mean of income are all significant determinants of *rs*. The coefficient on the standard deviation is negative, while the coefficient on the variance is positive. In Figure 1, we plot the overall that is, $\square_1$st.dev.(log(y^i_t))+ $\square_2$var(log(y^i_t)), against st.dev.(log(y^i_t)) for the estimated coefficients in the three columns of Table 2. The overall effect is negative and mostly declining.[14] This evidence is consistent with the prediction of the model with imperfect enforceability that relates the amount of risk sharing to the size of the variance of income: the larger the latter, the closer should the equilibrium be to full risk sharing and therefore the smaller the movement in the cross-sectional variance of marginal utility.

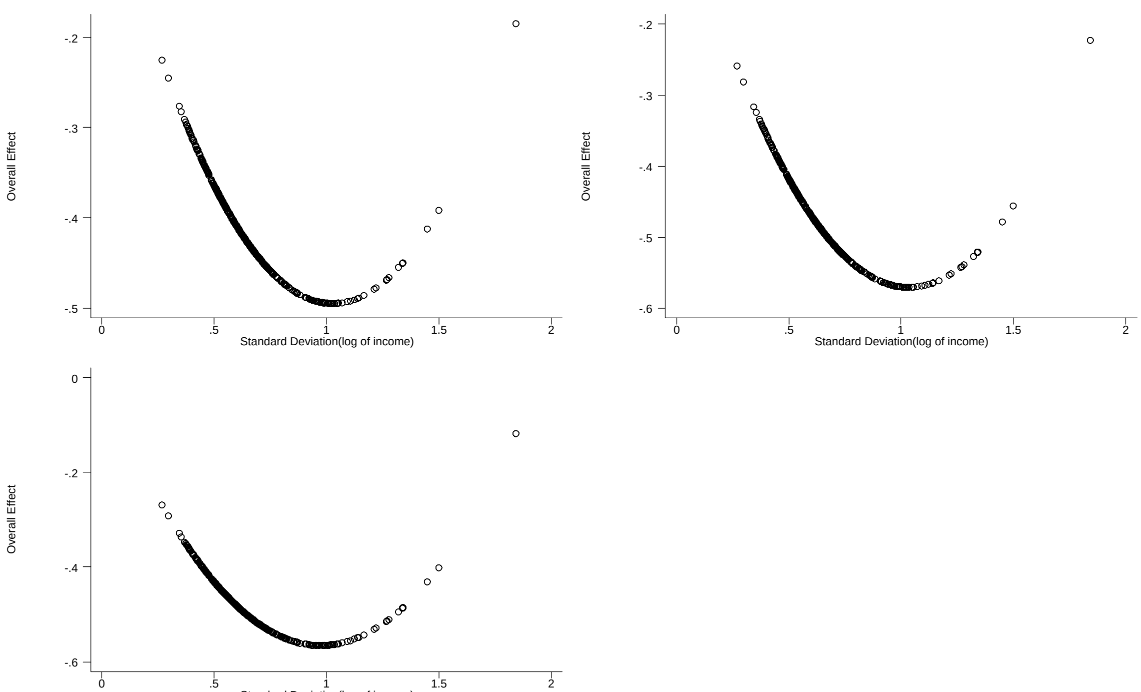

Figure 1. Overall Estimated Effect of Income Variance

Both coefficients associated with the autocorrelation of income are positive and significant, indicating an overall positive effect of autocorrelation. Once again, this is consistent with the implications of the model: individuals will be more reluctant to share persistent shocks.

Finally, the coefficient on the linear mean income term is negative, while the coefficient on the square term is positive. Over the relevant range, the overall effect which we plot in Figure 2, is positive for low values of income and negative for intermediate/ high values, and positive again. The prediction of the model in terms of the effect of mean income is not as strong as for the variance and the autocorrelation. With power utility and conditional on the variance of income, one would have a negative effect of income on risk sharing, and therefore we would expect a positive effect on our measure. However, it is easy to change the model so to modify this prediction. If, for instance one makes the penalty the village can impose on deviants, that is the enforcement technology, a function of average income, one can expect a negative effect in our regression.

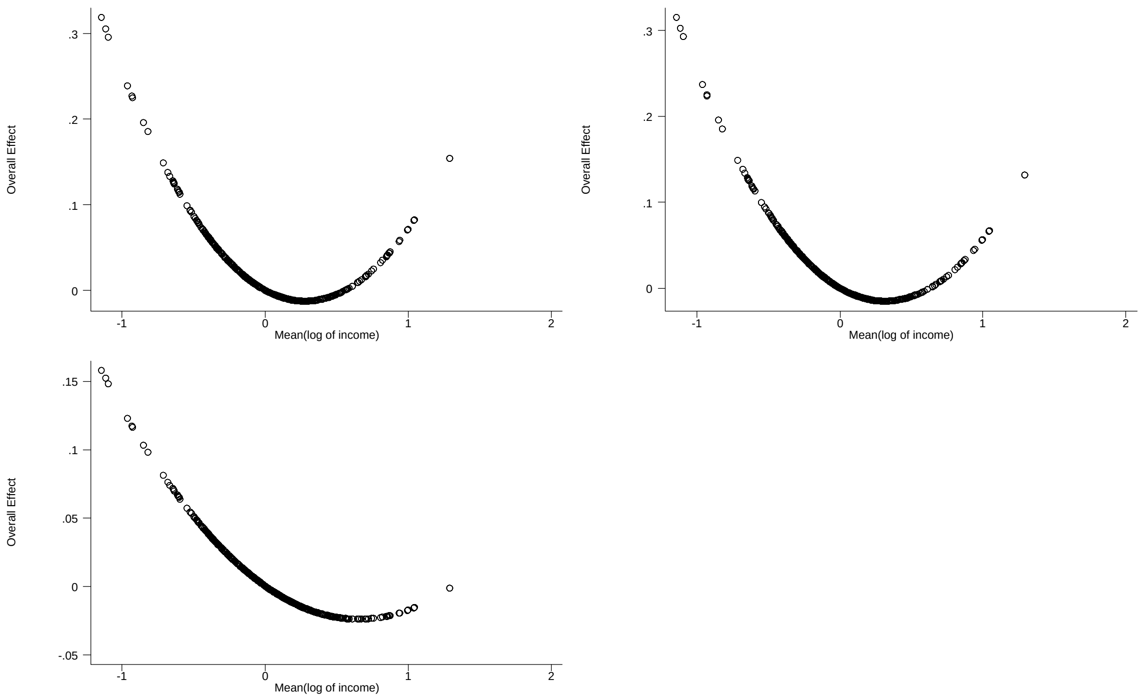

Figure 2. Overall Estimated Effect of Income Mean

[14] Notice that the coefficient of the standard deviation is always larger than that of the variance, and the standard deviation of the *logarithm* of income should be, in general, smaller than one (that is, income varies over time less than one hundred per cent).

5.2 *Euler equations*

In this section we discuss the econometric specification and the results of the second test we propose. To estimate the parameters of the model, we use a GMM approach. In particular, we consider the following moments:

$$
\text{(5.1)} \qquad \begin{aligned} E[\Delta v_{t+1}^{h,v} w_t^{h,v}] = E\left[\Delta v_{t+1}^{h,v} w_t^{h,v} \middle| constrained\right] \Pr(constrained) + \\ E\left[\Delta v_{t+1}^{h,v} w_t^{h,v} \middle| unconstrained\right] \Pr(unconstrained) \end{aligned}
$$

where $\Delta v_{t+1}^{h,v}$ is implicitly defined in equation (3.13), Pr(*constrained)* = $\Phi((\log(C_{t+1}^{h,v}) - \mu_{v,t})/2\sigma_v)$ and $\mu_{v,t}$ was defined in (3.12) and after equation (3.13).

As the number of time periods we can use is limited, the implicit assumption we need to obtain constistent estimator from the GMM procedure is that the expectational error average to zero across villages. To relax somewhat this assumption, we add state dummies.

To be completed

6 Conclusions

In this paper we have proposed some new tests of models with limited risk sharing and imperfect enforceability of contracts. Our evidence is, overall, in favour of these models, while rejecting strongly the hypothesis of full risk sharing. Much work remains to be done. In particular, we think it particularly valuable to extend the models we have been considering to include both imperfect information and imperfect enforceability.

References

Abreu, D. Pearce, D. and E. Stacchetti (1990) : "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring" *Econometrica* 58, 1041-63.

Albarran P. and O.P. Attanasio (2001) "Do public transfers crowd out private transfers? Evidence from a randomized experiment in Mexico", WIDER discussion paper No. 2000/6 http://www.wider.unu.edu/publications/dps/dps2002/dp2002-06.pdf.

Albarran P. and O.P. Attanasio (2002) "Limited commitment and crowding out of private transfers: Evidence from a randomized experiment", *Economic Journal,* forthcoming.

Alvarez, F. and U. Jermann (2000) "Asset Pricing when Risk Sharing is Limited by Default", *Econometrica, I68, 775-98.*

Alvarez, F. and U. Jermann (2001) "Quantative Asset Pricing Implications of Endogenous Solvency Constraints", forthcoming in the *Review of Financial Studies.*

Attanasio, O.P. (2002): "Consumption and Income Inequality: what we know and what we can learn from it", UCL Mimeo, paper presented for a plenary lecture to the Society of Economic Dynamics, New York, June2002.

Attanasio, O.P. and M. Mazzocco (in progress): "Intrahousehold allocation with limited commitment: empirical emplications".

Attanasio, O.P. and M. Székely (2001): "Wage shocks and Consumption Variability in Mexico During the 1990s", Mimeo, UCL

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

Attanasio, O.P. and J.V. Rios Rull (2000b): "Consumption Smoothing and Extended Families", forthcoming in Dewatripoint and Hansen (eds) *Advances in Economic Theory- The World Congress of the Econometric Society- Seattle 2000.*

Cochrane, J. (1991): "A Simple Test of Consumption Insurance", *Journal of Political Economy*, 99, 957-76.

Deaton, A., and C. Paxson. 1994. "Intertemporal Choice and Inequality." *Journal of Political Economy* 102: 384-94.

Foster, A. and M. Rosenzweig (1999):"Imperfect Commitment, Altruism, and the Family: Evidence from Transfer Behavior in Low-Income Rural Areas", *forthcoming in Review of Economics and Statistics.*

Kehoe P. and F. Perri (2002): "International Business Cycles with Endogenous Incomplete Markets, *Econometrica,* 70, 907-28.

Krueger D. and F. Perri (2001): "Does Income Inequality Lead to Consumption Inequality? Empirical Findings and a Theoretical Explanation", Mimeo, Stanford University.

Kocherlakota, N. (1996): "Implications of Efficient Risk Sharing without Commitment" *Review of Economic Studies,* 63, 595-609.

Ligon, E., Thomas, J. and T. Worrall (2000) "Informal Insurance Arrangements in Village Economies" forthcoming in the *Review of Economic Studies.*

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

Marcet, A. and R. Marimon (1992) "Communication, Commitment and Growth", *Journal of Economic Theory, 58, 219-49.*

Platteau, J.P. and A. Abraham (1987): "An inquiry into quasi-credit contracts: The role of reciprocal credit and inerlinked deals in small-scale fishing communities" *Journal of Development Studies,* 23, 461-90.

Thomas, J. and T. Worrall (1988): "Self-enforcing Wage Contracts", *Review of Economic Studies,* 55, 541-54.

Thomas, J. and T. Worrall (1990): "" *Journal of Economic Theory*

Townsend, R.M. (1994) "Risk and Insurance in Village India" *Econometrica,* 62(3), 539-91.

Udry, C. (1994): "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria" *Review of Economic Studies,* 61(3), 495-526

Appendix B

All the relevant variables in equation (3.7) involve some moments of the distribution of consumption and income in each of the villages; in particular, we need to obtain the cross-sectional variance of both consumption and income and the time series mean, variance and autocorrelation of income. We must pay special attention to measure properly the statistical features that are relevant according to the theoretical model in section 2 (see also simulations in Appendix A). As far as the ratio *rs* in the right-hand side of (3.7), it is apparent from discussion in section 3.1 that the changes in *cross-sectional variance of (log) consumption* (properly normalized by the variance of income) is a measure of the amount of risk-sharing. These cross-sectional variances can be consistently measured as the number of individuals in a cross-section for each village goes to infinity. We have used data from the three central consecutive Evaluation Surveys (ENCEL) of PROGRESA (November 1998, March 1999 and November 1999). We have around 30 observations on average in each wave for computing these variances, although if we restrict the analysis to villages with more than 30 observations, the results displayed in Table 3 did not change substantially.

As well, we need some estimates of the time series volatility and persistence of idiosyncratic shocks to income for the left hand side of (3.7). In principle, consistent estimates of the variance and autocorrelation can be obtained from a long time series. However, we only have five waves in our data set for this purpose: the three previously mentioned rounds and those carried out in March 1998 and in April 2000. As a consequence the approach we take is the following: First, we regress household income on a set of observable and predictable controls: the number of household members, educational attainment and age of the household head and his spouse, share of agricultural workers in the village and time dummies. The residuals of this regression can be regarded as a good approximation of the idiosyncratic shocks, since aggregate shocks and time-invariant individual-specific characteristics have been removed. We obtained an Adjusted-R^2 in this regression of 12.92%.

Then, we computed time series variances and autocorrelations (over only four waves) for each household in our sample using those residuals. Our final measures are obtained by averaging within each village these individual variances and autocorrelations. Indeed these are basically cross-section moments, although we can be confident that they are approximating properly the features of the idiosyncratic shocks we are interested in, since the residuals are not "contaminated" with cross-sectional variability. Some summary statistics of our measures of variability and persistence are displayed in Table B1.

Table B1. Summary statistics

	Variance	Autocorrelation
Percentiles		
5%	0.4212	-0.3672
10%	0.4548	-0.3155
25%	0.5263	-0.2458
Median	0.6128	-0.1718

75%	0.7366	-0.1078
90%	0.9414	-0.0449
95%	1.0306	-0.0042
Mean	0.6568	-0.1759
Std. Dev.	0.1939	0.1097
Observations	479	