Assessing the Impact of Public Transfers on Private Risk Sharing Arrangements: Evidence from a Randomized Experiment in Mexico

Marina Pavan*
The Geary Institute University College Dublin

Aldo Colussi†
University of Western Ontario

October 10, 2007

Abstract

We adopt a structural estimation approach to studying the effects of public transfers on consumption smoothing, risk sharing and welfare in small village economies. We estimate the key parameters of a dynamic limited commitment model using data gathered as part of the Mexican Progresa program, and take advantage of the randomized experimental design of the data to validate the model using the treatment sample. The limited commitment model enriched to allow for unobserved heterogeneity in preferences can reasonably well explain consumption dynamics and cross-sectional distributions. The estimated model correctly predicts the increase in consumption smoothing and risk sharing *within* transfers recipients, and the decrease in risk sharing *between* beneficiaries and non beneficiaries of the program.

KEYWORDS: Risk Sharing, Limited Commitment, Crowding-out, Transfers Programs

*E-mail: pavanma@bc.edu
†Email: acolussi@uwo.ca

1. Introduction

Village economies in developing countries are characterized by high volatility in individual incomes. Recent empirical work find that in these economies changes in household consumption are systematically related to changes in household income over time, rejecting the hypothesis of full risk-pooling or complete insurance. However, there does appear to be a certain level of risk sharing, even if far from complete (Townsend (1994), Udry (1994), Ligon, Thomas and Worrall (2002) among others).[1] Households in these small societies often lack access to a developed financial sector, nonetheless establish a system of informal credit or "quasi-credit" transactions as a means to smooth consumption over time.

Governments are often involved in designing policies to provide insurance against risk or in general to alleviate poverty in village economies. These policies, however, could affect the incentives of households to participate in the informal insurance arrangements that are in place before the public intervention. In evaluating the effect of public transfers programs, it is then of paramount importance to understand their influence on individual behavior and on the pre-existing private arrangements to share risk. In this paper, we study the effects of public transfers on consumption smoothing, risk sharing and welfare in small village economies. We take advantage of the observable experience of a transfers program to estimate and validate a structural behavioral model, which is then used to conduct policy counterfactuals.

Our starting point is a dynamic limited commitment model, which is considered to be an appropriate framework to describe the economies under consideration.[2] In the context of limited commitment environments, the degree of risk sharing that can be sustained in equilibrium differs depending on the parameters that characterize household preferences and uncertainty. The predicted effects of public transfers on private risk sharing arrangements also depend upon those same parameters.

A first goal of this paper is to estimate the key parameters of a limited commitment model using observed earnings and consumption patterns of households in poor small villages.

We incorporate in the model the main features of a small agricultural economy: idiosyncratic income uncertainty, public information and limited commitment with lack of enforcement, informal contingent contracts, and the possibility of defaulting on these contracts and reverting to self-insurance through a storage technology. Allowing for savings is important in order to measure the extent of risk sharing in the economy. Not only savings are a means to self-insure and thus smooth consumption over time, but they can also constitute a commitment

[1] Notable exceptions are Ogaki and Zhang (2001), who do not reject the full insurance hypothesis within Pakistani and Indian villages when assuming a Decreasing Relative Risk Aversion utility function. Mazzocco and Saini (2007) assume heterogeneous risk preferences and also do not reject efficient risk sharing at the caste level in rural India.

[2] Ligon, Thomas and Worrall (2002) find that a dynamic limited commitment model performs much better than both the full insurance and the autarky models to predict the response of consumption to idiosyncratic income shocks in three Indian villages.

device, when used as collateral (see Gobert and Poitevin (2002), Ligon, Thomas and Worrall (2000)).[3]

To estimate the model, we use the data set gathered as part of the Mexican Progresa transfers program. Progresa is a large anti-poverty program begun in Mexico in 1997, whose main goal is the developing of human capital of poor households in rural villages. It provides educational subsidies (contingent on children's regular schooling attendance), some cash transfers and nutritional supplements for infants and small children. The benefits that families receive are substantial relative to their income level, so that the program can be seen as both an insurance scheme aimed at limiting the impact of bad shocks to income and a long-term human capital development policy.

The fact that the Progresa villages are in general agricultural, small and isolated, and the fact that all village households are surveyed by Progresa allow us to clearly define the risk sharing pool as the households belonging to the same village.

We adopt the restrictions derived from the theory to investigate whether the estimated model is able to match key features of the observed data. We find that a version of the model enriched to allow for heterogeneous preferences can reasonably well fit the dynamic response of consumption to income and the cross-sectional variation of consumption given the income distribution. There does exist, however, as also found in Ligon et al. (2002), a trade-off between consumption time-series dynamics and cross-sectional distribution: the model cannot very precisely fit both at the same time.

A second goal of the paper is to determine the validity of the model exploiting the randomized experimental design of the data. In the Progresa social experiment the Mexican villages were randomly assigned to either participate in the program or to serve as controls. Being control and treatment groups randomly drawn by the same population, the relevant behavioral model should apply to both. The model is estimated using data on households in the randomized-out control group. We then assess its performance by comparing the impact of the program predicted by the model (obtained simulating the model behavior under the subsidies scheme) to the impact obtained under the experiment, observed in the treatment villages.

We find that the limited commitment model accurately predicts the impact of a public transfers program on the beneficiary households' consumption smoothing and on the risk sharing *between* recipients and non recipients in the villages. The model captures the fact that risk sharing improves *within* the group of beneficiary households, but fails to generate the reduction in risk sharing and the (small) decrease in consumption smoothing observed for the non recipients in the treatment.

Finally, the main advantage of our structural estimation approach is that we can use the

[3]The effects of storage on the level of risk sharing and welfare are ambiguous if savings are not used as collateral (i.e. can be kept upon breaching the contract). The higher utility from self-insurance can make mutual insurance more difficult to sustain, thus decreasing welfare. This is the main result in Ligon, Thomas and Worrall (2000).

behavioral model to predict the impact of alternative hypothetical policies meant to provide insurance against uncertainty. The estimated dynamic model can be used to perform an evaluation of policy interventions as a social insurance program against individual uncertainty in which transfers are conditional on idiosyncratic shocks, or a program which conditions the provision of public transfers to some positive action to be taken by the household (e.g. a minimum strictly positive amount of storage). Eventually, one could design and test alternative transfers schemes that could be implemented in order to improve efficiency in the presence of enforceability problems.

Many studies have addressed the issue of the potential crowding-out of private transfers by public ones (Cox (1987), Cox, Hansen and Jimenez (1999), Attanasio and Rios-Rull (2000b), Albarran and Attanasio (2001), among others). In particular, Attanasio and Rios-Rull (2000a) show that in environments with limited commitment a public insurance program can improve as well as deteriorate welfare, depending on the characteristics of the underlying economy. Their numerical exercises find that key elements to understand and evaluate these crowding-out effects are agents' preferences, the persistence of idiosyncratic and aggregate shocks and the relative variance of aggregate and idiosyncratic risks. Similarly, Krueger and Perri (2005) analyze the effect of a change in the progressivity of the tax system on the incentives private agents have to engage in private risk sharing arrangements. They calibrate their model to US income and tax data and find that a more redistributive tax system could indeed lead to less risk sharing and lower ex-ante welfare.

In the empirical literature, Albarran and Attanasio (2001, 2003) use the data sets from the Progresa program to reject the hypothesis of full risk sharing, and find that the program does crowd-out private transfers (the program decreases both the likelihood to receive a transfer and the amount received conditional on receiving one). Skoufias (2007) performs a test of full risk sharing using the Progresa data, and finds that consumption changes are significantly correlated with income changes, thus rejecting the complete insurance hypothesis. Moreover, his analysis reveals that Progresa does not affect risk sharing within the treated villages, even if the recipient households in these villages do smooth consumption better.

The novel approach in our research is the direct link we draw between a dynamic behavioral model and the observed micro-data. Moreover, we are able to use those data to validate the theory. The alternative reduced form approach does not allow taking into account the change in behavior induced by a public intervention, since it does not explicitly model that behavior. Also, welfare evaluation is not possible with a reduced form analysis. On the other hand, the existing theoretical results in the literature have not yet been tested empirically in a coherent and structured framework. We intend to fill that gap. This is the first attempt, to the best of our knowledge, to try and estimate the structural parameters of a dynamic limited commitment model with storage possibilities.[4]

[4] Ligon, Thomas and Worrall (2002) take into consideration an extension to their benchmark model of dynamic limited commitment in which they incorporate the possibility of storage. However, they estimate only the simpler

The next section describes the model, which builds on Krueger and Perri (2006). Section 3 briefly introduces the Progresa program and the data we use, while Section 4 describes the estimation procedure. In Section 5 we present our results and Section 6 concludes.

2. The Model

We consider a pure exchange economy, populated by a continuum of infinitely lived agents of measure one. Time is discrete. Households face idiosyncratic uncertainty in their labor endowments, and there is no aggregate uncertainty. Agents can share risk by trading each period a complete set of Arrow securities which are contingent on the observable labor endowment realizations. However, contracts cannot be enforced, and individuals can renege on their payments, in which case they lose all savings and they are excluded from future trade in the insurance market. We first describe in detail the individual problem, and then define an equilibrium for this economy. We shall specify more precisely the functional forms of the model economy in a later section.

2.1. The Individual Problem

The population is divided into three different groups, with p_g indicating the fraction of agents belonging to group $g \in G \equiv \{1, 2, 3\}$. Each group is meant to capture the rich set of observed household characteristics that define the entitlement to program transfers in the empirical implementation. Groups are characterized by different levels of expected earnings. At each period t, household i belonging to group g receives a labor endowment $\alpha_g l_{it}$, where α_g denotes the deterministic group-specific fixed effect, while l_{it} is an idiosyncratic shock, $l_{it} \in L \equiv \{l_1, .., l_N\}$, that follows a Markov process with transition probabilities $\pi(l'/l) = prob\{l_{t+1} = l'/l_t = l\}$. We assume a law of large numbers, so that the fraction of agents in the population who face the shock l today and the shock l' tomorrow is equal to $\pi(l'/l) > 0$ $\forall l, l'$, with $\sum_{l' \in L} \pi(l'/l) = 1$ $\forall l \in L$. Let $\Pi(\cdot)$ be the unique invariant measure corresponding to $\pi(l'/l)$, and denote with $l_i^t = \{l_{i0}, l_{i1}, ..., l_{it}\}$ the history of idiosyncratic endowments of household i up to period t. We simplify notation and disregard individual's subscripts from now on.

Utility from the consumption stream $c = \{c_t\}_{t=0}^{\infty}$ is separable over time. To take into account the effect of household size, sex and age composition on current utility, consumption is expressed in adult equivalent units. We allow for unobserved heterogeneity assuming there are two "types" of agents, characterized by different preferences. Unobserved types are common knowledge in the village, but are not observed by the econometrician. The distribution of types depends on the same set of characteristics that define households' groups. In each group g of households, there are $\pi_{\eta/g}$ unobserved type η households, with $\sum_{\eta=1}^{2} \pi_{\eta/g} = 1$ $\forall g$.

version on the grounds of computational difficulties and because their main goal is to test its performance compared to that of the models of full insurance, autarky and static limited commitment.

Each type η household maximizes the normalized expected value of a discounted sum of instantaneous utility functions:

$$E_0 \left\{ (1-\beta) \sum_{t=0}^{\infty} \beta^t u_\eta (c_t) \right\} \tag{1}$$

where $\beta \in (0,1)$ is the common discount factor, E_0 denotes expectations at $t = 0$, and the function u_η is assumed to be strictly increasing, strictly concave, twice differentiable and satisfying Inada conditions.

Each agent has the possibility of self-insuring through a storage technology which pays (risk-free) R_t units of the consumption good for each unit saved. Let $a_t \in \Re_+$ denote savings at t. Given the nature of the stochastic process, it is possible to write the autarky problem in recursive formulation. The (normalized) continuation value U^{aut} for the agent in autarky is therefore

$$U^{aut}(a_t, l_t; g, \eta) = \tag{2}$$

$$(1-\beta) \max_{0 \le a_{t+1} \le \alpha_g l_t + R_t a_t} \left\{ u_\eta(\alpha_g l_t + R_t a_t - a_{t+1}) + \beta \sum_{l_{t+1}} \pi(l_{t+1}/l_t) U^{aut}(a_{t+1}, l_{t+1}; g, \eta) \right\}$$

Moreover, consumers have access to a full set of state-contingent contracts. At any period t, after observing the shock l_t, the household can buy an Arrow security $a_{t+1}(l') \in A$ for each possible $l' \in L$, which pays only conditional on l' being realized in $t+1$. These securities can be interpreted as informal credit arrangements, whereby households can borrow (or save) with repayments that are contingent on the future shocks. Let $q_t(l', l)$ be the price of one unit of consumption delivered in $t+1$ conditional on realizations l at t and l' at $t+1$.[5] Then each security costs $q_t(l', l)\, a_{t+1}(l')$. In equilibrium, the no arbitrage condition implies that it must be $q_t(l', l) = \frac{\pi(l'/l)}{R_t}$, with $1/R_t$ denoting the price for receiving one unit of consumption with certainty in the following period.

Define as

$$U\left(c, l^t; g, \eta\right) = (1-\beta) \sum_{\tau=t}^{\infty} \sum_{l^\tau / l^t} \beta^{\tau-t} \pi\left(l^\tau / l^t\right) u_\eta\left(c_\tau\left(l^\tau\right)\right) \tag{3}$$

the normalized continuation utility from an allocation $c = \{c_\tau(l^\tau)\}_{\tau=t}^{\infty}$, from event history l^t onwards.

We assume that each household has the option to renege on its obligation at any point in time. As in Krueger and Perri (2006), agents who default lose all their savings, are banned from future insurance markets, and can only self-insure. Moreover, default entails a utility cost (or stigma), ψ_η, which differs across types. The participation constraint that prevents individuals from defaulting at any time t and at any history l^t is

$$U\left(c, l^t; g, \eta\right) \ge U^{aut}(0, l_t; g, \eta) - \psi_\eta \qquad \forall l^t \tag{4}$$

[5] The price depends only on current period realization of earnings and not on the whole history because of the Markov process properties.

At period 0 households start with initial asset holdings a_0 and initial shock l_0. Let $\Phi_0(a, l, g, \eta)$ be the joint measure of initial assets, shocks, groups and unobserved types.

2.2. Competitive Equilibrium with Solvency Constraints

We use here the concept of competitive equilibrium with solvency constraints introduced by Alvarez and Jermann (2000). They define implicit borrowing constraints that are state dependent and are such that the individual who borrows at the limit is indifferent between defaulting and paying his debt (these constraints are called "not too tight solvency constraints").[6] The problem in recursive formulation can be expressed as follows:

$$V(a, l; g, \eta) = \max_{c,\{a'(l')\}_{l'}} \left\{ (1-\beta)u_\eta(c) + \beta \sum_{l'} \pi(l'/l)\, V(a'(l'), l'; g, \eta) \right\} \tag{5}$$

$$\text{s.t.} \qquad c + \sum_{l'} q(l', l)\, a'(l') = a + \alpha_g l$$

$$a'(l') \geq \overline{A}(l'; g, \eta) \qquad \forall l' \in \overline{L}$$

where we omit the time subscripts for simplicity, denoting with $a'(l')$ the choice of Arrow security that pays in $t+1$ conditional on l' being realized. The implicit borrowing limit $\overline{A}(l; g, \eta)$ satisfies

$$V\left(\overline{A}(l; g, \eta), l; g, \eta\right) = U^{aut}(0, l; g, \eta) - \psi_\eta \qquad \forall l, g, \eta \tag{6}$$

We restrict ourselves to the analysis of a stationary equilibrium, i.e. one in which the distribution Φ is constant over time.

Definition 2.1. *Given the initial distribution Φ_0, a stationary recursive competitive equilibrium with borrowing limits $\left\{\overline{A}(l; g, \eta)\right\}_{l,g,\eta}$ that are not too tight is defined by policy functions $c(a, l; g, \eta), \{a'(a, l, l'; g, \eta)\}_{l' \in L}$, a value function $V(a, l; g, \eta)$ for each $(a, l, g, \eta) \in A \times L \times G \times \{1, 2\}$, an interest rate R, and a measure Φ such that:*

1) Given prices, the allocations $c(a, l; g, \eta), \{a'(a, l, l'; g, \eta)\}_{l' \in L}$ solve the individual optimization problem (5), with $\overline{A}(l; g, \eta)$ defined as in (6);

2) The good market clears at every period:

$$\int_{A \times L \times G \times \{1,2\}} c(a, l; g, \eta)\, \partial\Phi = \sum_g \sum_{l \in L} \alpha_g l p_g \Pi(l) \tag{7}$$

or, equivalently, the assets market clears:

$$\sum_{l \in L} \int_{A \times L \times G \times \{1,2\}} a'(a, l, l'; g, \eta) \pi(l'/l)\, \partial\Phi = 0 \tag{8}$$

[6] We refer to Alvarez and Jermann (2000) and Krueger and Perri (2005, 2006) for a discussion of the decentralization of constrained efficient consumption allocations, and a proof of the existence of the competitive equilibrium with solvency constraints.

3) The measure Φ is stationary, i.e. satisfies

$$\Phi = Q \cdot \Phi$$

where Q is the transition matrix defined by (with a slight abuse of notation, discretizing the state space):

$$Q\left(a', l', \widetilde{A}, \widetilde{L}; g, \eta\right) = \begin{cases} \sum_{l' \in \widetilde{L}} \pi\left(l'/l\right) & \text{if } a'\left(a, l, l'; g, \eta\right) \in \widetilde{A} \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

with $\widetilde{A}$ and $\widetilde{L}$ being subsets of the state space for Arrow securities and the state space L respectively.

The model is solved numerically. The problem is non-standard in that we have to compute not only the equilibrium prices, but also the endogenous borrowing constraints.

We first guess the gross interest rate R, and compute the autarky value function U^{aut} using policy function iteration. We guess borrowing limits $\overline{A}(l; g, \eta)$ and solve for $V\left(a, l; g, \eta\right)$ subject to the budget constraint and borrowing constraints, using policy function iteration. If the limits $\overline{A}(l; g, \eta)$ are not too tight, i.e. if $V\left(\overline{A}(l; g, \eta), l; g, \eta\right) = U^{aut}(0, l; g, \eta) - \psi_\eta$ for each l, g, and η, then we solved for the endogenous borrowing constraints corresponding to the interest rate R; if not so, then we update the guesses for $\overline{A}(l; g, \eta)$ and solve for the value functions again, until the above equalities hold.

Once we found the value and policy functions with related borrowing limits that are not too tight, we compute the stationary distribution of assets and earnings as the unique fixed point of the operator H:

$$\Phi = H(\Phi)$$

where $H(\Phi) = Q \cdot \Phi$ (and Q is defined as in (9) above).

Given the policy functions and the stationary distribution, we can compute aggregate wealth as defined in (8). If aggregate wealth is equal to zero, then we found the interest rate such that markets clear. If not, we update the guess for the interest rate and iterate until convergence. We approximate value functions with piece-wise linear approximations and use Brent's root finding method to find the equilibrium interest rate.

3. Data

Progresa is one of Mexico's major ongoing programs aimed at fostering the human capital of poor households in rural areas, through subsidies on health, nutrition and schooling investments in children. To evaluate the impact of the program a set of 506 rural villages in 7 Mexican states was randomly divided into 320 treatment and 186 control villages.[7] Within each village

[7] Behrman and Todd (2000) provide evidence that the randomization was accurately implemented, i.e. most variables are not significantly different between control and treatment groups.

(both control and treatment), in a first step households were classified as poor if their income fell under the poverty line, as non poor otherwise. A final list of beneficiaries (we will call them Poor) was determined in a second step, using an index obtained from a discriminant analysis which incorporated other household characteristics. The program subsidies were offered to all Poor households in the treatment villages starting in May 1998 (transfers were actually delayed by a few months for some households in these villages). In June 1999 the program added new households to the list of beneficiaries through a process called "densification", which increased the percent of transfers' recipients from 52% to 78% of the sample (we will call these new beneficiaries "Densificados").[8]

The information is collected from all beneficiary and non beneficiary households in control and treatment villages both before and after the implementation of the program. After two initial surveys in November 1997 and in March 1998, prior to the program introduction, data on 24,000 households from the treatment and control villages were collected every six months between October 1998 and November 1999. By 2000 the localities serving the role of a control group started receiving Progresa benefits as well.

Progresa gives benefits exclusively to mothers, and cash transfers and nutritional supplements are conditioned on children's regular school attendance and visits to health care centers. The size of the grants increases with the grade and, for secondary education, it is slightly higher for female students. In total, the benefits that families receive are substantial relative to their incomes. We gather information about the transfers received from an administrative source.[9] For our sample, the monthly payment to beneficiaries who start receiving transfers in 1998 amounts on average to 26% of our measure of monthly earnings, while the average transfer for beneficiaries that were added to the program only later (Densificados) is 16% of their monthly earnings. Despite the fact that the largest component of the benefits is not a pure transfer, but is conditional on school attendance, it can be said that this component possibly represents a substantial wealth effect. We therefore exploit this wealth effect to analyze the indirect impact of the program on consumption risk sharing.

Since we do not have complete information on all assets and saving and borrowing arrangements among families, to retrieve our measures of risk sharing and inequality we exploit available data on consumption and earnings only. The Progresa survey contains detailed information on several categories of food and non-food consumption, and income from several sources. The Appendix provides a detailed description of the creation of these variables.

For estimation, we transform consumption and earnings observed in the data in a way that is consistent with the theoretical model. We follow Ligon, Thomas and Worrall (2002) and define consumption as expenditures on food and clothing plus the consumption value of self-production of these goods, measured at the household level. We first compute adult

[8] This new selection method added more elderly poor who no longer lived with their children (Skoufias, Davis, and de la Vega, 1999).

[9] Available at http://evaluacion.oportunidades.gob.mx:8010/en/index.php.

equivalent consumption dividing total household consumption by a sum of individual weights attributed to household components based on their sex and age.[10] We then normalize each household's adult equivalent consumption by the mean of the same variable in the village at the observation period. We are mostly interested in analysing the distribution and dynamics of these consumption shares, especially in relation to the distribution and dynamics of earnings shares.

Our measure of earnings captures all sources of household revenues that are exogenous to the consumption and saving decisions of households, coherently with the theoretical model. We measure earnings as the sum of labor earnings from the main job, plus pensions, for all family members. Unfortunately the data on earnings from secondary jobs were not always available or coherent across waves, so that we found them not reliable.[11] To have a measure of earnings consistent with the definition of consumption, we compute adult equivalent earnings for each household in the sample, normalized by the average village earnings.[12]

We use data from the three waves from October 1998 to November 1999 (i.e. until the control group starts receiving the treatment). Some of the households were excluded from the sample because of missing or inconsistent information on some of the key variables. Our sample includes a total of about 20550 households, of which about 7990 are living in control and about 12560 in treated villages. Some descriptive statistics is presented in Table 7.1 in the Appendix.

4. Estimation

4.1. Specification of Model Functional Forms

Going from the model to the data, we need to make specific assumptions on preferences and the labor endowment process, as well as on the distribution of households' groups and unobserved types.

As discussed in the data section above, a discriminant analysis was used to classify households into beneficiaries and non beneficiaries of the program. That analysis took into account a variety of household characteristics to capture the "multi-dimensional" nature of poverty (Skoufias, Davis, and Vega, 1999). We take advantage of this classification in the data to define the three groups of households: "Poor" ($g = 1$), who receive the transfers from the beginning of the Program in 1998, "Densificados" ($g = 2$), who were added to the list of beneficiaries

[10] The weights are taken from Townsend (1994) and are the following: for adult males, 1.0; for adult females, 0.9. For males and females aged 13-18, 0.94 and 0.83 respectively; for children aged 7-12, 0.67 regardless of gender; for children 4-6, 0.52; for toddlers 1-3, 0.32; and for infants 0.05 (Townsend (1994), Footnote 12, p.554).

[11] As an example, in October 1998 the survey asks how many weeks the individual worked on a second job in the last six months, but a non response is classified as 8, making it impossible to retrieve the real time spent working. In waves where data seemed of a better quality, earnings from second jobs were negligible values, so that their exclusion does not seem to change results.

[12] We deflate all variables using the National Agricultural Consumer Price Index for Mexico.

in June 1999, and "Non Poor" ($g = 3$), who do not receive any transfer at any time. We differentiate households that were classified as transfers recipients from the beginning of the program ("Poor") from the so called "Densificados", who were included as beneficiaries only later, since we have reason to believe these households might come from demographic groups with different characteristics (Skoufias, Davis, and de la Vega, 1999), and participated to the program at different times. The definition of these types of households directly from the data allows us not only to capture in one variable a possibly very complex set of characteristics, but also to easily conduct counterfactual experiments with the estimated model. The distribution of these groups in the model economy corresponds to the observed distribution.

We assume the within period utility function to belong to the Hyperbolic Absolute Risk Aversion (HARA) class:

$$u_\eta(c) = \xi_\eta \frac{(c-\gamma)^{1-\sigma_\eta}}{1-\sigma_\eta}$$

where $\sigma_\eta > 0$ is a measure of risk aversion for type η, ξ_η denotes type η's marginal utility of consumption (ξ_1 is fixed to be equal to one, ξ_2 is estimated) , and γ is the parameter that governs how relative risk aversion varies with the level of consumption. A positive γ is typically interpreted as subsistence consumption, and implies a Decreasing Relative Risk Aversion (DRRA). Ogaki and Zhang (2000, 2001) find that it is important to allow for DRRA in testing the full risk-sharing hypothesis in economies with low-income households.[13] We therefore decide to adopt a more general utility function than the CRRA typically used, estimating the subsistence parameter γ, assumed equal across households.[14]

As mentioned above, unobserved types are common knowledge in the village, but are not observed by the econometrician, who can only estimate their distribution conditional on observed characteristics. We assume the probability of being a certain type depends on the household being classified as Poor, Densificado, or Non Poor in the data according to the following inverse logit function:

$$\pi_{1/g} = \frac{\exp(\phi_1 + \sum_{s=2,3} I\{g=s\}\,\phi_s)}{1+\exp(\phi_1 + \sum_{s=2,3} I\{g=s\}\,\phi_s)}$$

where $I\{g=s\}$ is an indicator function equal to one if the household belongs to group s, and $\pi_{2/g} = 1 - \pi_{1/g}$.

Instead than estimating each entry of the transition matrix $\pi(l'/l)$ for the idiosyncratic shock to labor earnings, we assume that matrix to be the approximation of the following

[13] With the assumption of a positive subsistence level in HARA preferences, consumption growth rates are more volatile for rich individuals than for poor ones, even under full insurance. This is in contrast with the standard result under Constant Relative Risk Aversion (CRRA) that consumption growth is identical for all agents. Ogaki and Zhang (2001) test the full risk-sharing hypothesis allowing for HARA preferences with the Pakistani ICRISAT data. They find strong evidence against risk sharing across villages, but cannot reject the full insurance hypothesis within each village (with a couple of exceptions), thus showing that results obtained with the assumption of CRRA can be misleading.

[14] We restrict γ to be positive and small compared to the average labor endowment.

autoregressive process of order one, common across groups of households:

$$\log(l_{it}) = \rho \log(l_{i,t-1}) + \varepsilon_{it}$$

where ε_{it} is an i.i.d. shock normally distributed with zero mean and variance $\sigma_l^2\left(1-\rho^2\right)$. For the numerical solution, we then discretize this process into nine grid points using Tauchen's method. The mean of the two components of the individual productivity endowments is normalized to one for simplicity: $E(l_{it}) = 1$ at every t, and $E(\alpha_{ig}) = 1$.

Let y_{it} be the adult equivalent labor earnings for household i in period t. We then have the following:

$$\log(y_{it}) = \log(\alpha_{ig}) + \log(l_{it})$$

and $var(\log(y_{it})) \equiv \sigma_y^2 = var(\log(\alpha_{ig})) + var(\log(l_{it})) \equiv \sigma_\alpha^2 + \sigma_l^2$.

We estimate the persistence coefficient ρ, the variance σ_l^2, and the variance of the fixed effect component σ_α^2 together with the other preference parameters and the coefficients characterizing the probability distribution for unobserved types. An alternative approach would have been to estimate the exogenous earnings process parameters in a first stage, using data on earnings only. However, since we wouldn't get consistent estimates of the variance and autocorrelation of earnings having to rely on three waves of data only, we opted to estimate all parameters fitting all data moments at the same time.

The length of the model period is six months, since that's the interval at which we observe household consumption and earnings in the data. Moreover, given that our data are observed over a relatively short period (one and a half years), we abstract from business cycle considerations, coherently with the assumption of no aggregate shocks. However, when computing moments, we remove time effects from both consumption and earnings data. We also remove differences due to education levels, sex and age of the household head. These differences wouldn't be accounted for in the theoretical model.

4.2. Estimation Method

We estimate our model using a Method of Simulated Moments (MSM). Given the parameters characterizing the economy (household preferences, earnings process, unobserved types distribution, etc.), we can numerically solve for the optimal policy functions and the equilibrium consumption distribution. The estimation procedure consists in searching for the parameter vector that obtains the best fit to the data.

Let $\ln(\widehat{v}_{it}) \equiv [\ln(\widehat{c}_{it}), \ln(\widehat{y}_{it})]$ be the vector of observed logarithm of consumption and earnings for household i in period t. Underlying assumption is that these data are generated by the following processes:

$$\ln(\widehat{c}_{it}) = \ln(c_{it}(x_{it}, \Psi)) + u_{it}^c$$

$$\ln(\widehat{y}_{it}) = \ln(y_{it}(l_{it}, g, \Psi)) + u_{it}^y$$

where $c_{it}(x_{it};\Psi)$ is the optimal consumption given the state vector $x_{it} \equiv (a_{it}, l_{it}; \eta_i, g_i)$ and the parameter vector Ψ, $y_{it}(l_{it}, g_i, \Psi)$ is the (exogenous) earnings for household i belonging at group g at time t, while u_{it}^c and u_{it}^y are normally distributed i.i.d. shocks which we interpret as measurement errors to consumption and income respectively.

We seek to estimate the model from a set of moment conditions, $E[m(\ln(\widehat{v});\Psi_0)] = 0$, where Ψ_0 is the true parameter vector and m indicates a vector of moments as a function of the observed variables. The unconditional expectation defined above would be a very complex object to compute analytically, since we would need to evaluate the unconditional distribution of the state variables. To overcome this difficulty, we then approximate that distribution computationally. Let $m(\Psi)$ be the N_m-dimensional vector of approximated moments corresponding to a parameter vector Ψ. Our method of simulated moments estimator solves the following problem:

$$\min_{\Psi}[G(\Psi)' W G(\Psi)]$$

where $G(\Psi) = [m(\Psi) - m(\ln(\widehat{v});\Psi_0)]$ and W is a $N_m \times N_m$ weighting matrix, chosen so to minimize the asymptotic variance of the estimator.

5. Results

5.1. Estimated Parameters and Data Fit of the Model

We consider the insurance group to be identified by the village, and use only control villages in estimation. Since it would be impossible to estimate the model for each village separately, moments used in estimation are averages across the control villages.[15] We are interested in measures of consumption smoothing, risk sharing and inequality. We therefore try to match the ratio of the standard deviation of (log) consumption to the standard deviation of (log) earnings, both within and between groups, the Lorenz curves for consumption and earnings, and the correlation between consumption growth and earnings growth for each group of households.

A challenge of the limited commitment model is that of fitting cross-sectional statistics and dynamics at the same time, as already noticed in Ligon et al. (2002).[16] We find that whenever the model parameters adjust to match the relatively low correlation between consumption growth and earnings growth in the data, the predicted cross-sectional variation in consumption is much lower than the actual one, probably due to too much "smoothing". Only extreme differences between the two unobserved types of households can generate the kind of cross-sectional heterogeneity in consumption observed in the data. We therefore had to restrict some of the parameter values and compromise the goodness of fit in order to avoid those

[15] The Appendix describes in detail how we construct these moments.

[16] Ligon et al. (2002) estimate their model using two different equations, one to match consumption levels, the other to fit changes in shares of consumption. Neither of these equations provide estimates that simultaneously explain distribution and dynamics.

Model Parameters	Estimated Value*
σ_α^2 - variance of earnings fixed effect	0.024
σ_l^2 - variance of earnings persistent part	0.204
ρ - earnings persistence coefficient	0.984
β - discount factor	0.968
σ_1 - risk-aversion coefficient, type 1	2.250
σ_2 - risk-aversion coefficient, type 2	0.697
ξ_2 - marginal utility of consumption, type 2	0.454
γ - subsistence level	0.003
ψ_1 - utility cost of default, type 1	0.001
ψ_2 - utility cost of default, type 2	0.249
$\pi_{1/1}$ - fraction of type 1's in group 1 (Poor)	34.6%
$\pi_{1/2}$ - fraction of type 1's in group 2 (Densificados)	43.9%
$\pi_{1/3}$ - fraction of type 1's in group 3 (Non Poor)	27.8%

Table 5.1: Estimated Parameters. *Standard Errors in parenthesis TBA

extreme differences between types and obtain more sensible estimates.[17] In particular, we restrict the marginal utility of the second type to lie between 0.45 and 1.3, while the utility cost of default parameters must be lower than 0.25. We also set an upper limit of 2.5 for the risk aversion coefficient.

Table 5.1 shows the values of the estimated parameters. Households reasonably discount the future at a factor of 0.968 per semester, which translates into an annual discount rate of 6.7%. The two types of households are distinctly different, for the reasons underlined above. The prevalent type (constituting about 65% of the Poor and 72% of the Non Poor households) has a coefficient of risk aversion equal to 0.70, a lower marginal utility of consumption (very close to the lower limit) and a utility cost of default that is equal to about 17% of the current utility for a consumption unit (also at the limit). The other type of household has a much higher risk aversion (σ is about 2.3), and practically no cost of default. The subsistence level is found to be very close to zero, so that Constant Relative Risk Aversion seems to be a quite appropriate approximation of household preferences. Not surprisingly, the idiosyncratic shock is quite volatile ($\sigma_l^2 = 0.2$ or a standard deviation of about 0.89 on a yearly basis) and persistent ($\rho = 0.984$ for six months, i.e. 0.97 yearly).

In Table 5.2 and Figure 1 we present the data moments used in estimation, together with the corresponding moments predicted by the estimated model. As mentioned above, we concen-

[17] We could have increased the number of unobserved types in order to better fit cross-sectional moments. However, that alternative possibility is ruled out at the moment because computationally unfeasible.

trate our attention on measures of risk sharing both within and between groups. A commonly used measure of risk sharing is the cross-sectional variance of consumption, normalized by the cross-sectional variance of earnings to take into account the "need" for risk sharing in the economy (Albarran and Attanasio (2002), Attanasio and Rios-Rull (2000a)). Denote with rs_g^w and rs^b the ratios of the cross-sectional standard deviation of log consumption to the cross-sectional standard deviation of log earnings where the standard deviations are measured within each group g and between groups, respectively. In a complete markets equilibrium, assuming the only difference between groups were the fixed effects in productivity α_g, households within each group would end up consuming the same share of aggregate consumption over time, proportional to their life time human wealth. Under full risk sharing the ratio rs_g^w would then be zero within each group, but equal to one between groups (since the variance of consumption between groups would be equal to the variance of the fixed effects α_g).

The average value of the ratio rs_g^w we observe in the data is actually closer to one than to zero, being in the range of 0.89 (for the Non Poor group) to 0.95 (for the Poor).[18] The between groups ratio, on the other end, is equal to 0.96, only slightly lower than the one predicted by complete markets. The estimated model predicts a value of rs_g^w equal to 0.84 for Poor, 0.83 for Densificados, and 0.77 for Non Poor, thus explaining almost 90% of the observed ratios. The equilibrium between groups ratio rs^b is calculated to be 1.03, very close to the observed one.

The reasons why we do not observe these ratios to assume the values implied by full insurance could be that risk sharing is far from complete and/or that people do not differ only in their life time human wealth (or other observable factors which can be accounted for), but also in unobserved characteristics (taste, patience, etc.), which make them choose different consumption patterns for a same earnings volatility. Generally, we cannot quantify the importance of unobserved heterogeneity in determining consumption inequality (given income inequality), and we do not claim to be able to do that here. However, if unobserved factors (or their distributions) are thought to be constant over time, and not affected by changes in policy, a *change* in the ratios rs_g^w and rs^b induced by the introduction of a transfers program would only be consistent with incomplete insurance. We will exploit that idea for the validation of the model, using the effects of the Progresa transfers program on risk sharing in the treatment villages.

Another measure of how well people smooth consumption over time is the correlation between the percentage change in consumption shares and the percentage change in earnings shares. Under full insurance, this correlation would be equal to zero. These correlations in the data range between 8% for the Non Poor group to 12% for the Densificados, all significantly different from zero, but not significantly different from each other. Insurance is not complete, but there is a certain degree of consumption smoothing over time, which is similar among the three groups of households. The predicted average responses of consumption shares to

[18] These values of the ratios rs_g^w are significantly different from each other (by pair of groups) at the 1% level.

Statistical Properties	Data (Control Villages)	Model
Stdev($\ln c$)/Stdev($\ln l$) - Poor	0.95	0.84
Stdev($\ln c$)/Stdev($\ln l$) - Densificados	0.92	0.83
Stdev($\ln c$)/Stdev($\ln l$) - Non Poor	0.89	0.77
Stdev($\ln c$)/Stdev($\ln y$) - b/w groups	0.96	1.03
Corr($\Delta(\ln c), \Delta(\ln y)$) - Poor	0.11	0.20
Corr($\Delta(\ln c), \Delta(\ln y)$) - Densificados	0.12	0.19
Corr($\Delta(\ln c), \Delta(\ln y)$) - Non Poor	0.08	0.20

Table 5.2: Model Fit of Data Statistics

earnings share are higher than the actual ones, at around 0.20, as shown by the estimated correlations in the right column of Table (5.2). These higher values are due to the parameters restrictions described above: the goodness of fit in this case was compromised by the need to obtain economically sensible parameter values.

Last, Figure 1 compares the Lorenz curves computed for the estimated model economy to the ones in the data. The model does a good job in fitting earnings and consumption inequality.

5.2. Model Validation: the Effects of a Transfers Program on Risk Sharing

After having estimated the model using the data from the control villages, we can assess its performance by comparing the impact of the program predicted by the model to the impact obtained under the experiment, observed in the treatment villages.

We thus simulate the estimated model attributing to households the median value of the actual transfers received: 19% and 11% of monthly income for Pobre and Densificados, respectively. Assuming these transfers are perceived as permanent, they increase the expected life time earnings and therefore autarky values of the beneficiaries. While earnings volatility is not affected by the transfers, higher average wealth should possibly lead to more consumption smoothing over time (there is less need to borrow, and households have the capability to accumulate higher savings). However, the limited commitment model also predicts that households who receive the transfers have less incentive to participate in the risk sharing arrangements with other households in the village, given their higher autarky values. The overall result on consumption smoothing and risk sharing cannot be pre-determined a-priori, and will depend on the specific parameters characterizing the economy.

Results are shown in Table (5.3). The first column shows the values of the same moments used in estimation, but computed for the treatment villages. The ratio of the standard deviation of log consumption to the standard deviation of log earnings decreases from 0.95 to

0.90 for the Poor group, and increases for the two other groups, from 0.92 to 0.94 for the Densificados, and from 0.89 to 0.95 for the Non Poor. These changes are quite small, even if all significantly different from zero at the 1% level. Consumption variability as a fraction of earnings variability is thus lower among the Poor households in the treatment than among the Poor in the control villages, but higher in the villages that are part of the program for those families who receive the transfers only later, and for those who are not eligible to the transfers. The model captures the change in rs_1^w for the Poor group, predicting a decrease from 0.84 to 0.75. For the other two groups the computed change in rs^w has the wrong sign: the model predicts a decrease in the standard deviation of consumption with respect to the standard deviation of earnings within the Densicados and the Non Poor, even if these changes, especially for the latter group, are smaller than the effect on the Poor.

As for the between groups measure of risk sharing, the ratio rs^b is on average one third higher in treatment than in control villages: the standard deviation of log earnings once transfers are introduced decreases proportionally more than the standard deviation of log consumption.[19] The lower inequality between groups in terms of income does not lead to a comparably lower consumption inequality. The estimated model predicts a post transfers between groups ratio rs^b equal to 1.27, or 23% higher than the ratio before the transfers, explaining about two thirds of the increase in the between groups ratio.

In the data, the correlation between growth rates in consumption shares and growth rates in earnings shares for the treated Poor and Densificados groups are 7% and 5% respectively, significantly lower than the ones for the corresponding groups in control villages (11% and 12%). Households who receive the transfers are thus found to better smooth consumption over time. Although it would appear that the opposite be true for the Non Poor group of households, whose correlation is 12% in the treatment group and 8% in the control group, this difference is not significant.[20] After transfers are introduced in the model economy, the equilibrium correlations between consumption growth and income growth are predicted to be 19% for Poor, and 18% for Densificados and Non Poor. Again, the model captures the correct sign of the change for beneficiary households but not for those who do not receive the transfers.

Figure 2 shows the consumption Lorenz curves after the subsidies program is introduced in the model, comparing them to the actual Lorenz curves in the treatment villages. The fit is very good, and there is actually almost no change with respect to the same inequality measure for the control group.

We conclude that the model does quite a good job at predicting some of the qualitative and quantitative features of the effects of a transfers program on risk sharing and inequality

[19] The between group standard deviation of log earnings is on average equal to 0.147 in control villages and 0.143 in treatment villages. Adding the transfers to the earnings of recipients leads to a standard deviation of log income equal to 0.096. The standard deviation of log consumption is 0.11 in control villages, 0.09 in treatment villages. The ratios used in estimation are different from the ratios of numbers given here because they are the mean of the ratios (which is different from the ratio of the means).

[20] The analysis is presented in the Appendix at the end.

Statistical Properties	**Data (Treatment Villages)**	**% Change w.r. to Control**	**Model**	**% Change w.r. to Benchmark**
Stdev($\ln c$)/Stdev($\ln l$) - Poor	0.90	-5%**	0.75	-11%
Stdev($\ln c$)/Stdev($\ln l$) - Densificados	0.94	$+2\%$**	0.76	-8%
Stdev($\ln c$)/Stdev($\ln l$) - Non Poor	0.95	$+7\%$**	0.73	-5%
Stdev($\ln c$)/Stdev($\ln y$) - b/w groups	1.28	$+33\%$**	1.27	$+23\%$
Corr($\Delta(\ln c), \Delta(\ln y)$) - Poor	0.06	-36%*	0.19	-3%
Corr($\Delta(\ln c), \Delta(\ln y)$) - Densificados	0.04	-58%*	0.18	-7%
Corr($\Delta(\ln c), \Delta(\ln y)$) - Non Poor	0.12	$+50\%$	0.18	-10%
% Change in Consumption, Poor	14%		19%	
% Change in Consumption, Dens.dos	2%		11%	
% Change in Consumption, Non Poor	4%		3.5%	

Table 5.3: Model Fit of Data Statistics for the treatment villages. *significant at 5 percent, **significant at 1 percent

in a small economy. Note that the complete insurance model would imply no changes in the moments considered here. An alternative model we could have used to describe the economy would be the self-insurance/incomplete markets model, in which each household would have access to only one or a finite number of securities to buffer its consumption against income volatility. Our aim is not that of proving the superiority of a model over another. However, the self-insurance model would have quite different predictions from the limited commitment one. The injection of liquidity in that model would lead to a decrease in interest rates and more consumption smoothing by *all* groups of households. The evidence suggests that beneficiary and non beneficiary households are affected by the introduction of transfers in the economy in quite different ways, thus justifying the adoption of a model that at least in theory can predict these heterogeneous impacts.[21]

We can use the estimated model to further analyze the changes in risk sharing and welfare induced by the public transfers program. We present in Table (5.4) an alternative measure of risk sharing: the average amount of private transfers to earnings ratios for each group, where transfers are computed as the absolute value of the difference between earnings and consumption at each point in time (and are normalized by earnings). We also express these average transfers ratios as a fraction of the average efficient transfers to income ratios computed under complete markets with commitment. The public transfers program is found to improve

[21] Clearly, the limited commitment model has also difficulty capturing these different effects at the moment.

	Benchmark	**After Program**
Average Private Transfers ($\mid c - y \mid$) over Earnings:		
Poor	0.57	0.44
Densificados	0.56	0.49
Non Poor	0.48	0.51
Average Private Transfers as a fraction of Efficient Transfers:		
Poor	0.92	0.98
Densificados	0.93	0.98
Non Poor	0.80	0.84
Average Value Function Type 1/Type 2		
Poor	$-0.63/1.30$	$-0.51/1.39$
Densificados	$-0.53/1.37$	$-0.47/1.42$
Non Poor	$-0.39/1.50$	$-0.37/1.52$
Av. Difference b/w Limit. Commit. Value Function and Value Function in Complete Markets:		
Poor	0.34	0.23
Densificados	0.34	0.28
Non Poor	0.24	0.26

Table 5.4: Computed Statistics for the Model Economy

efficiency for each of the three groups of households: with the subsidies, ratios of private transfers to efficient transfers increase with respect to those in the benchmark economy.

Welfare under the public transfers program increases for all groups, especially, as expected, for the beneficiaries. Since the subsistence level is very close to zero, we compute differences in terms of consumption equivalent variations, as if the utility function were CRRA. Before the introduction of the transfers, type one households in Poor, Densificados and Non Poor groups are willing to forgo 18%, 11% and 3% of their consumption, respectively, to participate to the public subsidies program. Type two Poor households are the ones that benefit the most from the program: they are willing to forgo 22% of their pre-program consumption. Densificado and Non Poor type two households are willing to forgo 13% and 4% of their consumption, respectively, for their villages to participate in the program.

Overall, there does not seem to be crowding-out of private transfers by public ones, and actually the subsidies program is found to be welfare improving for all groups in the population, both recipients and not.

6. Conclusions

Our structural estimation approach allows us to take into account the direct and indirect effects of public policies on individual behavior. We find that a limited commitment model can replicate some of the features of consumpion smoothing and cross-sectional variation in the small Mexican villages surveyed as part of the Progresa program. However, we had to assume unobserved heterogeneity and restrict some of the preference parameters in order to get a better fit of the dynamics and cross-sectional distributions at the same time.

Using data from the treatment villages, we validate the model with a different sample than the one used in estimation. The model accurately predicts the increase in consumption smoothing and in risk sharing for the recipient households, and replicates the observed decrease in risk sharing *between* beneficiaries and non beneficiaries.

7. Appendix

7.1. The Data

Household food consumption is reported by food item (in detail for different kinds of fruits and vegetables, meats, cereals, etc.) in quantity consumed and/or value bought for the whole household in the previous week, and it includes both purchased and self-produced goods. Comparable information is available only in the surveys conducted after the beginning of the program, so we use the three waves in October/November 1998, May/June 1999 and November 1999.

Whenever the expenditure on a certain item is missing, but the quantity consumed is available, we impute the consumption value using the price of that item at the village level. The latter is computed as in Angelucci and De Giorgi (2006). We first construct household-specific prices for each item as the ratio of expenditures to the quantity bought of that item. If we have at least 20 observations of household-specific prices for a village, we use the median of these observations in the village to impute consumption values (multiplying the quantity consumed by the price). Otherwise, we use similarly computed median prices at the municipality or state level. If none of these prices are available, we use national prices. All home-produced consumption values are imputed in this way. In November 1999 the survey does not give direct information on home production, but does contain quantities consumed and quantities purchased for each item. In that wave consumption values are set equal to consumption expenditures if the quantity consumed is lower or equal to the quantity bought. We apply median prices as described above to evaluate the difference between quantity consumed and quantity purchased. Weekly consumption is transformed into semester values by multiplying it by 24.

Consumption on clothing is reported in values spent for general items (shoes, textiles, etc.) in the previous six months.

The earnings variable is the sum of labor income from the main job and pensions received by all members of the household. We impute semester earnings from daily, weekly and monthly earnings.[22] We assume there are six working days per week, and four weeks per month. When the observation on earnings is missing but the individual has worked in the week before the interview, we impute earnings by multiplying the number of hours worked by the median hourly wage rate for individuals with the same age, sex and education level.

As mentioned in the text, information on earnings from a second job as well as from work production appears not to be very reliable nor comparable across waves. We instead include in the earnings variable all pension income related to retirement, old age, invalidity or widowhood.

Progresa transfers to treated households are available in a separate administrative database and are reported every two months. They are the sum of monetary transfers for food support, the educational grants, and other money to be spent on school material (books, etc.). The average monthly transfer is around 200 pesos per beneficiary household (expressed in real October 1998 pesos). On average, about 64% of this is money for food support for the whole household, 33% is for the educational grant, and the rest is allocated to school supplies.[23]

We convert all variables in adult equivalent values as described in the text, and deflate them using the Agricultural Consumer Price Index for Mexico (as reported by the Central Bank of Mexico).

We trim the bottom and top percentile of the earnings and consumption variables to eliminate ouliers. Our variables are very likely to be measured with error, not only because they are self-reported, but also because we impute some values (in order not to loose a large part of the data). However, measurement error distributions shouldn't differ between control and treatment villages, so our comparisons should still be valid. Moreover, the measurement error problem should be mitigated by the fact that most of the moments we use in estimation are given by ratios of standard deviations and correlations between variables.

Last, the surveys contain also information on several demographic variables. Whenever age, sex or education level are missing for an individual in a certain wave, we try to recover them from other waves. Still, we need to drop 2,740 observations from the original sample because of missing or inconsistent information on any of these variables. Table (7.1) shows descriptive statistics of the sample used in estimation.

7.2. Construction of the Moments used in Estimation

In order to construct the ratios rs_g^w of the standard deviation of log consumption to the standard deviation of log earnings within each group g of households, we first compute numerator

[22] Some individuals report also fortnight or annual wages. However, we found the monthly mean of these wages is almost twice the monthly mean of wages reported daily, weekly, or monthly (which is the majority), so we drop these observations (losing 1,920 households in total over the three waves, or 2.7% of the sample).

[23] The percentages are computed for households belonging to the Poor group. For Densificados, the food support accounts for 81% of the transfers.

Variable	**Treatment**	**Control**
Number of Households:		
October/November 1998	13,207	8,277
May/June 1999	12,063	7,575
November 1999	12,414	8,115
Mean Age of Household Head	47.9	48.4
Gender of Household Head (% Male)	90.3	89.9
Mean Family Size (Townsend Weight)	4.61	4.67
Education (Head):		
No School	31.9%	30.8%
Some Primary	46.5%	48.9%
Primary	15.0%	14.0%
Secondary	3.9%	3.8%
% Poor in Sample	54.2	51.5
% Non Poor	21.8	21.9
Mean Adult Equivalent Monthly Earnings (real pesos 1998):		
All sample	272	273
Poor	229	226
Densificados	320	318
Non Poor	332	333
Mean Adult Equivalent Monthly Consumption (Food + Clothing):		
All sample	189	175
Poor	177	156
Densificados	203	199
Non Poor	200	192

Table 7.1: Data Descriptive Statistics

Mean Standard Deviation of Log Earnings*	**Poor**	**Dens.dos**	**Non Poor**	**All**
CONTROL	0.48	0.53	0.53	0.50
	(0.001)	(0.002)	(0.002)	(0.001)
Observations	3502	978	1173	13759
TREATMENT	0.49	0.53	0.52	0.50
	(0.001)	(0.002)	(0.002)	(0.000)
Observations	4897	1650	1317	18710
p-value of Difference b/w St. Dev.s	0.000	0.458	0.009	0.112

Table 7.2: Standard Deviation of Log Earnings, by groups and treatment. *Standard errors in parenthesis. We control for age, dummies on sex and on education of the household head, and a time dummy

and denominator separately for each group in each village. More precisely, we first divide each variable (earnings or consumption) by the mean of that variable in the village at the observation period, so that we consider "shares" instead than levels. We then compute the standard deviation of the residuals obtained from regressing the log of the shares on time dummies and demographic characteristics of the household head (age and dummies for sex and education). We construct the ratios rs_g^w as the mean across villages of the ratios of standard deviations of consumption residuals to standard deviations of earnings residuals, for each group, for control and treatment villages. To compute moments for the first group (the Poor), we consider villages with at least 30 households in that group. For the other two groups of households, we require at least 20 observations per village. Doing this allows us to have at least thirty villages in the control and thirty villages in the treatment subsample. There are on average 48 observations per village in the first group, 33 in the Densificados group, and 29 in the Non Poor group.

Similarly, to construct the between groups measure of risk sharing, rs^b, we consider the mean of the ratios for villages with at least 30 observations. There are 87 villages in the control group and 133 villages in the treatment group involved in this calculation.

Table (7.2) shows that the mean standard deviation of log earnings does not differ between control and treatment villages for the whole sample and for the Densificados group. However, the standard deviation is slightly higher for the Poor households in the treatment villages if compared to the same type of households in the control, and slightly lower for the Non Poor households in the treatment with respect to the Non Poor in the control. Despite statistically these numbers result to be significantly different, we do not think a one basis point variation in the standard deviation to be economically relevant. We therefore consider the groups in control and treatment villages to be facing the same kind of risk.

Correlations between the growth rate of consumption shares and the growth rate of earnings shares are computed pooling all individuals, since only the time dimension is important in this

Dependent Var.: % Change in Cons. Share*	Poor	Dens.dos	Non Poor	All
CONTROL: Coeff. on % Change in Earn. Share	0.103	0.110	0.067	0.096
	(0.016)	(0.021)	(0.023)	(0.009)
Observations	3576	1556	1315	6447
TREATMENT: Coeff. on % Change in Earn. Share	0.056	0.037	0.098	0.061
	(0.013)	(0.021)	(0.020)	(0.010)
Observations	4807	2049	1838	8694
p-value of Difference between Coefficients	0.0220	0.0141	0.3176	0.021

Table 7.3: Sensitivity of Consumption Changes to Changes in Earnings, by groups and treatment. *Robust standard errors in parenthesis. Control variables are age, dummies on sex and on education of the household head, and a time dummy

case. Table (7.3) shows the coefficients of a regression of the growth rate of consumption shares on the growth rate on earnings shares for the three groups in the treatment and in the control villages, and the p-values of the test that compares these coefficients between treatment and control. For both the groups of households who receive the transfers in the treatment villages the sensitivity of consumption to income is significantly lower (at the 5% level) than for the corresponding groups in the control sample. Households who do not benefit from the transfers in the treatment villages result to be more sensitive to changes in earnings, but this difference is not significant.

References

[1] Albarran P., and Attanasio O., (2001), "Do Public Transfers Crowd Out Private Transfers? Evidence from a Randomized Experiment in Mexico", Mimeo, UCL

[2] Albarran P., and Attanasio O., (2003), "Limited Commitment and Crowding Out of Private Transfers: Evidence from a Randomized Experiment", *The Economic Journal*, 113, C77-C85

[3] Alvarez F., and Jermann U.J. (2000), "Efficiency, Equilibrium, and Asset Pricing with Risk of Default", *Econometrica*, 68, 775-797

[4] Angelucci M., and De Giorgi G. (2006), "Indirect Effects of an Aid Program: the Case of Progresa and Consumption", IZA Discussion Paper No. 1955

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

[6] Attanasio O., and V. Rios Rull, (2000b), "Consumption Smoothing and Extended Families", *Advances in Economic Theory: Eighth World Congress*

[7] Behrman J., and P. Todd, (2000), “Randomness in the Experimental Samples of Progresa”, *International Food Policy Research Institute*, Washington D.C.

[8] Cox D., (1987), “Motives for Private Income Transfers”, *Journal of Political Economy*, 95, 508-546

[9] Cox D., Hansen B., and E. Jimenez (2004), “How Responsive are Private Transfers to Income? Evidence from a Laissez-Faire Economy”, *Journal of Public Economics*, 88, 2193-2219

[10] Krueger D., and F. Perri (2005), “Public versus Private Risk Sharing”, Mimeo, NYU

[11] Krueger D., and F. Perri (2006), “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory“, *Review of Economic Studies*, 73, 163-193

[12] Ligon E., J.P. Thomas, and T. Worrall (2000), “Mutual Insurance, Individual Savings, and Limited Commitment”, *Review of Economic Dynamics*, 3, 1-47

[13] 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, 209-244

[14] Mazzocco, M., and Saini, S. (2007), "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences", Mimeo, UCLA

[15] Ogaki M., and Zhang Q. (2001), "Decreasing Relative Risk Aversion and Tests of Risk Sharing", *Econometrica*, 69 n.2, 515-526

[16] Skoufias, E. (2005), "PROGRESA and its Impacts on the Welfare of Rural Households in Mexico", International Food Policy Research Institute, Washington, D.C.

[17] Skoufias, E. (2007), "Poverty Alleviation and Consumption Insurance: Evidence from PROGRESA in Mexico", *The Journal of Socio-Economics*, 36, 630-649

[18] Skoufias E., Davis B., and de la Vega S. (1999), "Targeting the Poor in Mexico: an Evaluation of the Selection of Households into Progresa", International Food Policy Research Institute, Washington, D.C.

[19] Townsend, R., (1994), “Risk and Insurance in Village India“, *Econometrica*, 62, 539-591

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

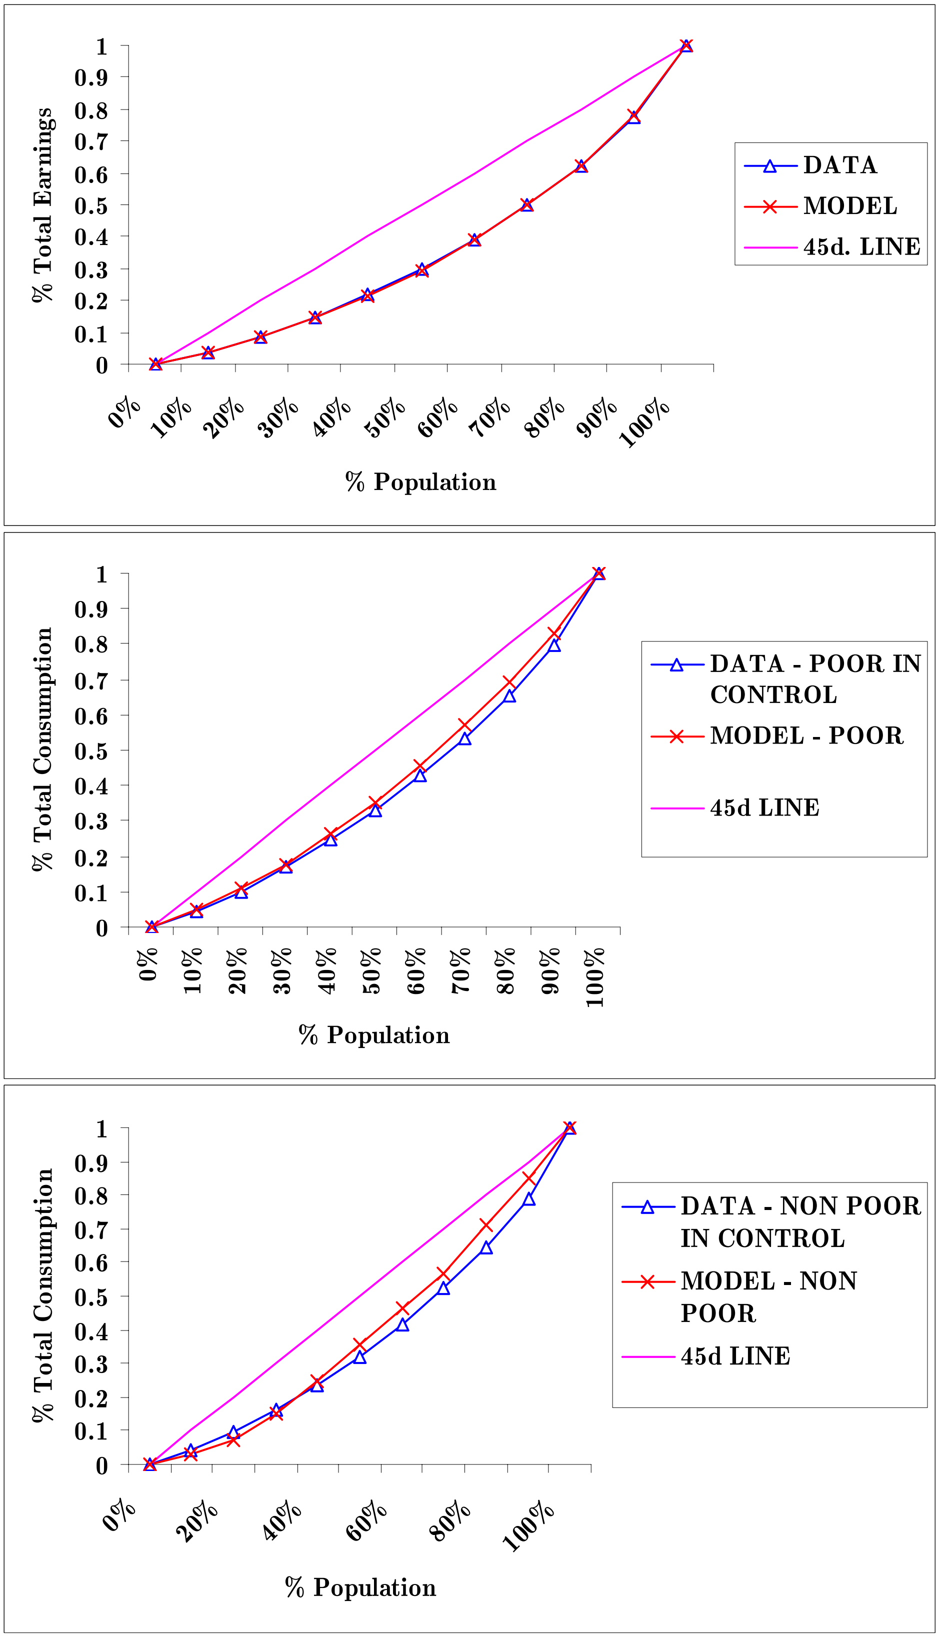

Figure 1: Lorenz Curves for Earnings and Consumption: Model vs. Data

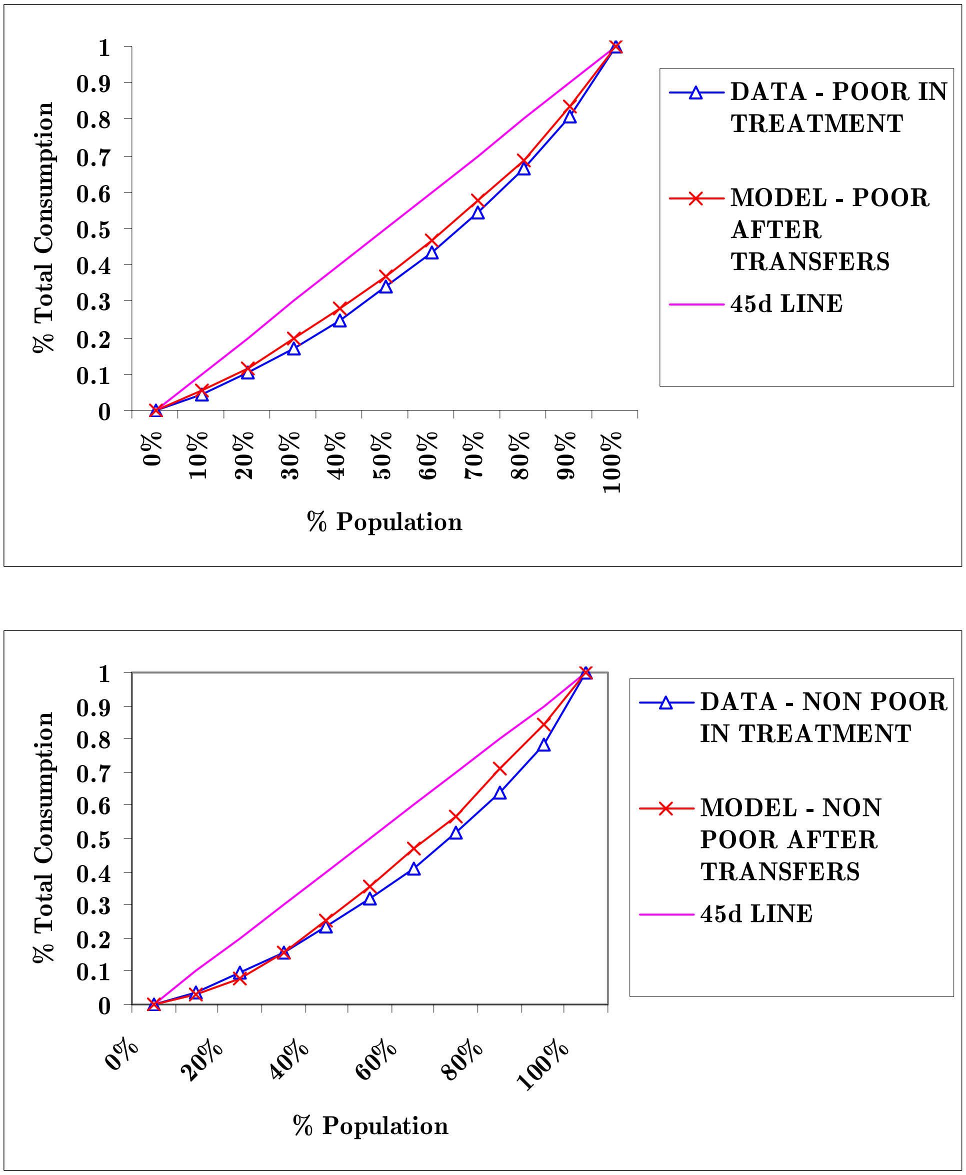

Figure 2: Lorenz Curves for Consumption Post-Transfers, Poor and Non Poor: Model vs. Data