



**UCD GEARY INSTITUTE
DISCUSSION PAPER
SERIES**

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

Dr. Marina Pavan

UCD Geary Institute, University College Dublin, Belfield, Dublin 4
marina.pavan@ucd.ie

Dr. Aldo Colussi

University of Western Ontario, USA
acolussi@uwo.ca

27th February 2008

The views expressed here do not necessarily reflect those of the Geary Institute.
All errors and omissions remain those of the author.

Geary WP/7/2008

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

February 22, 2008

Abstract

We adopt a structural approach to studying the effects of public transfers on consumption smoothing, risk sharing and welfare in small village economies. We calibrate 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 calibrated model correctly predicts the increase in consumption smoothing of transfers' recipients, and the decrease in risk sharing *between* beneficiaries and non beneficiaries of the program. Progresa transfers are found to crowd-out between 3% and 10% of the pre-existing private transfers, but the overall direct effect of the subsidy on consumption is welfare improving for all households. Last, we use our structural model to evaluate a counterfactual, fully funded, insurance scheme.

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

JEL CODES: H31,I38,D31,E21

*E-mail: marina.pavan@ucd.ie

†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 behaviour 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 calibrate and validate a structural behavioural model, which is then used to conduct counterfactual policies.

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 calibrate 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 commit-

¹Notable 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.

²Ligon, Thomas and Worrall (2002), for example, 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.

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

We use the data set gathered as part of the Mexican Progresa transfers program.³ Progresa is a large anti-poverty program begun in Mexico in 1997, whose main goal is to increase the 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 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 behavioural model should apply to both. Following an approach pioneered by Todd and Wolpin (2006), we calibrate the model 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 behaviour under the subsidies scheme) to the impact obtained under the experiment, observed in the treatment villages.

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. In particular, we find that in evaluating the effects of a public policy such as Progresa it is of paramount importance to take into consideration the heterogeneous responses in individual behaviour dictated by different preferences.

Further, our results suggest that the Progresa transfers crowd-out between 3% and 10% of the pre-existing private transfers, but nonetheless lead to an overall increase in consumption and welfare for all households.

Finally, the main advantage of our structural approach is that we can use the behavioural model to predict the impact of counterfactual policies meant to provide insurance against

³The program is still running in Mexico under the name of Oportunidades.

uncertainty. The dynamic model is used to perform an evaluation of a social insurance program against individual uncertainty in which transfers are conditional on idiosyncratic shocks, and such that the subsidies are fully financed by the population through a proportional income tax. The results from the calibrated model show that while there is some crowding-out of private transfers, overall consumption smoothing would improve as a result of the insurance policy. Nonetheless such a policy would not be welfare enhancing for all households: depending on preferences, some households would gain and some would lose from the introduction of the insurance scheme.

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 behavioural model and the observed micro-data. We use the data to calibrate the model and validate the calibration taking advantage of the experimental design of Progresa. Moreover, modelling explicitly household behaviour enables us to perform welfare analysis of the Progresa subsidy and of an alternative counterfactual insurance policy.

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 labour 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 labour 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 the 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 labour endowment $\alpha_g l_{it}$, where α_g denotes the deterministic group-specific fixed effect, while l_{it} is an idiosyncratic shock, $l_{it} \in L \equiv \{l_1, \dots, l_N\}$, that follows a Markov process with transition probabilities $\pi(l'/l) = \text{prob}\{l_{t+1} = l' / l_t = l\}$. 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 $l_i^t = \{l_{i0}, l_{i1}, \dots, l_{it}\}$ 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 $\sum_{\eta=1}^2 \pi_{\eta/g} = 1 \forall g$.

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

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

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 $a_t \in \mathfrak{R}_+$ 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

$$U^{aut}(a_t, l_t; g, \eta) = (1 - \beta) \max_{0 \leq a_{t+1} \leq \alpha_g l_t + R_t a_t} \left\{ u_\eta(\alpha_g l_t + R_t a_t - a_{t+1}) + \beta \sum_{l_{t+1}} \pi(l_{t+1}/l_t) U^{aut}(a_{t+1}, l_{t+1}; g, \eta) \right\} \quad (2)$$

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 $t + 1$.⁴ 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

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

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

$$U(c, l^t; g, \eta) \geq U^{aut}(0, l_t; g, \eta) - \psi_\eta \quad \forall l^t \quad (4)$$

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 depen-

⁴The price depends only on current period realization of earnings and not on the whole history because of the Markov process properties.

dent 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").⁵ The problem in recursive formulation can be expressed as follows:

$$\begin{aligned}
V(a, l; g, \eta) &= \max_{c, \{a'(l')\}_{l'}} \left\{ (1 - \beta)u_\eta(c) + \beta \sum_{l'} \pi(l'/l) V(a'(l'), l'; g, \eta) \right\} & (5) \\
\text{s.t.} \quad & c + \sum_{l'} q(l', l) a'(l') = a + \alpha_g l \\
& a'(l') \geq \bar{A}(l'; g, \eta) \quad \forall l' \in \bar{L}
\end{aligned}$$

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 $\bar{A}(l; g, \eta)$ satisfies

$$V(\bar{A}(l; g, \eta), l; g, \eta) = U^{aut}(0, l; g, \eta) - \psi_\eta \quad \forall l, g, \eta \quad (6)$$

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 $\{\bar{A}(l; g, \eta)\}_{l, g, \eta}$ 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 $\bar{A}(l; g, \eta)$ defined as in (6);
- 2) The good market clears at every period:

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

or, equivalently, the assets market clears:

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

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

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

⁵We 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.

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

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

with \tilde{A} and \tilde{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 $\bar{A}(l; g, \eta)$ and solve for $V(a, l; g, \eta)$ subject to the budget constraint and borrowing constraints, using policy function iteration. If the limits $\bar{A}(l; g, \eta)$ are not too tight, i.e. if $V(\bar{A}(l; g, \eta), l; g, \eta) = U^{aut}(0, l; g, \eta) - \psi_\eta$ 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 $\bar{A}(l; g, \eta)$ 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

Progreso 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.⁶ Within each village (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

⁶Behrman and Todd (2000) provide evidence that the randomization was accurately implemented, i.e. most variables are not significantly different between control and treatment groups.

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 53% to 79% of the sample (we will call these new beneficiaries "Densificados").⁷

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 centres. 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.⁸ 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 15% 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.

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 equivalent consumption dividing total household consumption by a sum of individual weights attributed to household components based on their sex and age.⁹ We then normalize each household's adult equivalent

⁷This new selection method added more elderly poor who no longer lived with their children (Skoufias, Davis, and de la Vega, 1999).

⁸Available at <http://evaluacion.oportunidades.gob.mx:8010/en/index.php>.

⁹The 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

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 labour earnings from the main job, secondary occupations and other informal work activities (cooking, sewing, repairs, etc. generating some income), plus pensions, for all family members. 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.¹⁰

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 20,550 households, of which about 7,990 are living in control and about 12,560 in treated villages. Some descriptive statistics is presented in Table 6 in the Appendix.

4. Specification of Model Functional Forms and Parameterization

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

Ideally, one would like to estimate the model parameters directly from the available data, using a method of moments or a maximum likelihood estimator (or their simulated counterparts). However, the computational solution of the model is very time intensive, and a formal estimation was not possible. As a consequence, we calibrate the model parameters to replicate some key features of the data.

We assume the insurance pool to be identified by the village, and use only control villages to calibrate the model. 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.¹¹

Our calibration exercise is actually conducted using the same algorithm we would need for a method of simulated moments. Given the parameters characterizing the economy (household preferences, earnings process, unobserved types distribution, etc.), we numerically solve for

gender; for children 4-6, 0.52; for toddlers 1-3, 0.32; and for infants 0.05 (Townsend (1994), Footnote 12, p.554).

¹⁰We deflate all variables using the National Agricultural Consumer Price Index for Mexico.

¹¹Since it would be impossible to calibrate the model for each village separately, the moments are averages across the control villages. The Appendix describes in detail how we construct these moments.

the optimal policy functions and the equilibrium consumption distribution. Our calibration procedure consists in searching for the parameter vector that obtains the best possible fit to the data, with the caveat that given the complexity of the problem the optimization algorithm does not reach convergence in a feasible time framework.

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. Given that our data are observed over a relatively short period of time (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.

We now introduce the model functional forms and their parameterization.

Instead than attributing a value to each entry of the transition matrix $\pi(l'/l)$ for the idiosyncratic shock to labour earnings, we assume that matrix to be the approximation of the following 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 $\sigma_l^2(1 - \rho^2)$. 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 $E(\alpha_{ig}) = 1$.

The variance of the idiosyncratic shock that best matches our moments is $\sigma_l^2 = 0.102$ (or a standard deviation of about 0.64 on a yearly basis), with a persistence coefficient ρ of 0.975 (or 0.95 yearly). The variance of the fixed effect is found to be equal to $\sigma_\alpha^2 = 0.025$.

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 normalized to one, ξ_2 calibrated), 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.¹² We therefore decide to adopt

¹²With 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

a more general utility function than the CRRA typically used. Using the calibration procedure described above, however, the parameter γ is found to be negligible, equal to 0.003.

Only sharp differences between the two unobserved types of households can generate the kind of cross-sectional heterogeneity in consumption observed in the data.¹³ According to our calibration strategy, the first type ($\eta = 1$) has a coefficient of risk aversion equal to 2.25 and a negligible utility cost of default ($\psi_1 = 0.001$). The other type of household has a much lower risk aversion ($\sigma_2 = 0.65$), but also a lower marginal utility of consumption $\xi_2 = 0.45$, and a cost of default $\psi_1 = 0.25$ (or about 20% of his current utility for a unit of consumption). Given these preferences, the less risk averse type is the "borrower" in this economy, while the first type typically wants to save (having no stigma in default, he is more borrowing constrained than the second type). Households discount the future at a rate $\beta = 0.969$ per semester, which translates into an annual discount rate of 6.5%, in the range of estimated values for that parameter: in the United States the discount rate is usually estimated to lie between 4% and 15% (see for example Gourinchas and Parker, 2002, and Cagetti, 2003).

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 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 a possibly very complex set of characteristics in one variable, but also to easily conduct counterfactual experiments with the calibrated model. The distribution of these groups in the model economy corresponds to the observed distribution.

We assume the probability of being a certain type depends on the household being classified

full insurance hypothesis within each village (with a couple of exceptions), thus showing that results obtained with the assumption of CRRA can be misleading.

¹³As already observed in Ligon et al. (2002), this class of limited commitment models is not able to fit cross-sectional statistics and dynamics at the same time. Ligon et al. (2002) estimate (a simplified version of) the model using two different equations, one to match consumption levels, the other to fit changes in consumption shares. Neither of these equations provides estimates that simultaneously explain distribution and dynamics.

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

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 $\pi_{2/g} = 1 - \pi_{1/g}$. Our calibration procedure leads to a value of $\pi_{1/g}$ equal to 35%, $\pi_{2/g}$ equal to 43%, and $\pi_{3/g}$ of 27%. The prevalent type is therefore the second one, with lower risk aversion and higher stigma in default.

5. Results

5.1. Data Fit of the Model

In Table 1 and Figure 1 we present the data moments used in calibration, together with the corresponding moments predicted by the calibrated model. As mentioned above, we concentrate our attention on measures of consumption smoothing, consumption inequality, and risk sharing both within and between groups.

A first measure of how well people smooth consumption is the correlation between the percentage change in consumption shares and the percentage change in earnings shares over time. Under full insurance that correlation would be equal to zero: household consumption patterns would be independent of individual idiosyncratic shocks. These correlations in the data range between 10% for the Non Poor group to 13% for the Densificados, all significantly different from zero (and not significantly different from each other). Coherently with the findings of previous tests of the full insurance hypothesis in developing countries (Townsend, 1994, Ligon, Thomas and Worrall, 2002, and many others), insurance in small Mexican villages is not complete. However, there seems to be a certain degree of consumption smoothing over time, which is similar among the three groups of households. The calibrated model fits very well this feature of the data: the average responses of consumption shares to earnings shares are around 0.10 for all groups, as shown in the right column of Table 1.

Figure 1 compares the Lorenz curves of earnings and consumption computed for the calibrated model economy to the ones in the data. The model does a very good job in fitting earnings and consumption inequality within each group of households.

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 (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 0.94 (for Densificados).¹⁴ The between groups ratio, on the other end, is equal to 0.64. The calibrated model predicts a value of rs_g^w equal to 0.69 for Poor, 0.71 for Densificados, and 0.66 for Non Poor, thus explaining about 75% of the observed ratios. The equilibrium between groups ratio rs^b is calculated to be 0.93.

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.¹⁵ 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.

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

After having calibrated the model using the data from the control villages, we can assess its performance by comparing the impact of the Progresa program as predicted by the model to the impact observed in the data in the treatment villages.

To this extent, we simulate the calibrated model attributing to beneficiary 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, the more direct effect of the subsidy is that of increasing the capability of households to smooth consumption over time. 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

¹⁴These values of the ratios rs_g^w are significantly different from each other (by pair of groups) at the 1% level.

¹⁵Preference heterogeneity, however, seems to be an essential feature of any model that aims at explaining wealth and consumption inequality (Venti and Wise, 2001, among others).

cannot be pre-determined a-priori, and will depend on the specific parameters characterizing the economy.

Results are shown in Table 2. The first column shows the values of the data moments computed for the treatment villages: 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 6.5% and 5% respectively, significantly lower than the ones for the corresponding groups in control villages (10.6% and 12.5%). Households who receive the transfers are thus found to better smooth consumption over time.¹⁶ Instead there is not much difference for the Non Poor households, whose correlation is 11.7% in the treatment group and 10.4% in the control group. After the transfers are introduced in the model economy, the equilibrium correlations between consumption growth and income growth are predicted to be about 6% for all groups of households. The model therefore captures very well the change for beneficiary households but not so for those who do not receive the transfers.

In the data the ratio of the standard deviation of log consumption to the standard deviation of log earnings changes very slightly: it decreases from 0.92 to 0.89 for the Poor group, and from 0.94 to 0.93 for the Densificados, whereas it increases from 0.89 to 0.93 for the Non Poor. 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, and for those families who receive the transfers only later, but higher in the treatment villages for those who are not eligible to the transfers. The model does not fully capture the change in these ratios: for the Poor group, it predicts an increase from 0.69 to 0.75; for the Densificados, the model correctly predicts no change; and for the Non Poor group it predicts a decrease from 0.66 to 0.55.

In the model, these changes in the "within" ratios rs_g^w can be explained by the fact that after the introduction of public transfers households change their behaviour in different ways, depending not only on their eligibility to the benefits, but also on their type. Among the beneficiaries, while type one households (the "savers") just accumulate more assets when receiving the transfers, on the opposite less risk averse agents (the "borrowers") increase their borrowing. For the latter, indeed, the new borrowing constraint becomes looser upon the transfers receipt: the higher permanent income allows them to repay and thus borrow more, despite their lower incentive to commit. As a result of this higher borrowing demand, in fact, the equilibrium interest rate increases (from 3.16% to 3.19%). Given the opposite reactions of these two types of agents, the standard deviation of consumption within the recipients (especially the Poor) tends to increase.

On the other hand, probably due to the price effect of the interest rate increase, both types of households within the non beneficiaries (the Non Poor) increase their saving (or decrease their borrowing), so that the standard deviation of consumption decreases within this group.

¹⁶Analogous result is found by Skoufias (2007), who conducts a standard test of full insurance using the Progres data, and finds that beneficiary households are able to insulate their consumption from fluctuations in income better than their counterparts in control villages.

Even though the model does not exactly replicate the changes in the cross-sectional distribution of consumption observed in the data (the kind of preferences heterogeneity considered here is still too simplistic), from this exercise we do learn that households can react in potentially very different ways to the introduction of public transfers in the economy, depending on their preferences.

As for the between groups measure of risk sharing, the ratio of the standard deviation of log consumption to the standard deviation of log earnings is on average 70% higher in treatment than in control villages: this is because the standard deviation of log earnings once transfers are introduced decreases proportionally more than the standard deviation of log consumption.¹⁷ 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 equal to 1.44, or 55% higher than the ratio before the transfers, explaining about three quarters of the increase in the between groups ratio.

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 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.¹⁸

5.3. Risk Sharing and Welfare

As pointed out in Attanasio and Rios-Rull (2000a), in an environment with limited commitment a public program can reduce welfare through the crowding out of private risk sharing arrangements. To evaluate this possibility we can use the calibrated model to further analyze the changes in risk sharing and welfare induced by the Progresya transfers program.

In Table 3 we present an alternative measure of risk sharing: the average amount of private

¹⁷The between group standard deviation of log earnings is on average equal to 0.15 in both control villages and treatment villages. Adding the transfers to the earnings of recipients leads to a standard deviation of log income equal to 0.098. The standard deviation of log consumption is 0.10 in control villages, 0.08 in treatment villages. The ratios used in calibration 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).

¹⁸Clearly, the limited commitment model has also difficulty capturing these different effects.

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. Using the calibrated model, we found some evidence of crowding out: the public transfers program reduces between 3 and 10% this ratios with respect to those in the benchmark economy for those groups of households who receive the subsidies.

Nonetheless, 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 equivalent consumption variations, as if the utility function were CRRA. Type one Poor households are the ones that benefit the most from the program: they are willing to forgo 20% of their pre-program consumption. Before the introduction of the transfers, type one households in Densificados and Non Poor groups are willing to forgo 13% and 5% of their consumption, respectively, to participate to the public subsidies program. Independently of the group they belong to, type two households are willing to forgo about 10% of their consumption for their villages to participate in the program. Given the higher equilibrium interest rate in the economy with transfers, the "borrower" recipients benefit less than the "savers".

Our results suggest that non beneficiary households are also better off with the Progresa program: as explained above, these agents choose to save more in the economy with transfers, which lead to an increase in their consumption and welfare.

5.4. An alternative transfers scheme

The calibrated model shows that all households benefit from the introduction of Progresa, just because the direct effect of the subsidy on consumption will dominate the indirect effects on private risk sharing agreements. To net out the effect of the subsidy, it is therefore interesting to examine a fully funded insurance scheme. In particular, we use the calibrated model to analyze the effects of a transfer equal to 18% of income (roughly the average transfer in the Progresa Program), paid when the household income is lower than 75% of the mean income in the economy: under this scheme all households (also those belonging to the Non Poor group) receive a transfer if their shock to earnings is sufficiently bad. These transfers are totally financed by those households with income higher than the threshold through a tax equal to 3.84% of their income. Under this policy then, at each point in time there would be a redistribution of resources from the lucky to the less lucky agents.

As Table 4 shows, this policy would also improve consumption smoothing: whereas the simulated policy slightly increases the ratio of the standard deviation of log consumption to the standard deviation of log earnings, the correlation between growth rates in consumption shares and growth rates in earnings shares is halved for all groups. Nonetheless, as presented in Table 5, the change in the ratios of private transfers to earnings shows that there would be crowding-out of private risk sharing. More interestingly, the policy would not be welfare improving for all households: as measured by the equivalent consumption variation, type two

households would lose from the introduction of this insurance scheme. The income effect of this kind of policy on the borrowing constraints is much less strong than in a "pure" transfer program, since permanent income increases by less (depending on the distribution of earnings shocks for each group of households). This together with the fact that the equilibrium interest rate increases (from 3.16% to 3.20%) lead to a reduction in the welfare of type two, who is typically a borrower.

6. Conclusions

Our structural approach allows us to take into account the direct and indirect effects of public policies on individual behaviour. We find that a limited commitment model can replicate some of the features of consumption smoothing and cross-sectional variation in the small Mexican villages surveyed as part of the Progresa program. In particular, we find that heterogeneity in preferences is a crucial factor for explaining the quite unequal cross-sectional distribution of consumption. Different preferences dictate dramatically different behaviours with respect to saving and consumption, and also dramatically different reactions to changes in the environment. Therefore, in evaluating the effects of a public transfers program on risk sharing and welfare, it is of paramount importance to allow for such heterogeneous response.

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

Our results suggest that private risk sharing agreements are partially, even if not totally, crowded-out by the introduction of the Progresa subsidy. Nonetheless, the direct effect of the subsidy on consumption prevails on the indirect effects on private risk sharing arrangements, and welfare increases for all households, even for the non beneficiaries, who save and consume more after the introduction of the public transfers.

A counterfactual, fully financed, insurance scheme is also evaluated. While overall improving consumption smoothing, the effect of this policy on welfare depends on preferences: some households would in fact lose from the introduction of such policy. This again stresses the importance of accounting for this dimension in the evaluation of public programs.

Our results therefore suggest that an important direction for future research would be that of finding tractable ways to model heterogeneity in risk and, possibly, time preferences when analyzing the equilibrium effects of public policies. In this, our paper is in line with the recent strand of the literature (Kurosaki, 1999, Mazzocco and Saini, 2007, Schulhofer-Wohl, 2007) that investigates (mostly in reduced form) the effects of individual preferences heterogeneity on risk sharing arrangements.

Appendix

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 labour income from the main job and other secondary occupations, informal work activities generating some income, plus pensions received by all members of the household. We impute semester earnings from daily, weekly and monthly earnings.¹⁹ 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. Informal work generating income includes the provision of transportation services, cooking, sewing, repairs, construction and similar, and its value is computed as the difference between the revenues from these activities and the costs incurred in performing them. Further, we include in the earnings

¹⁹Some 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).

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 225 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.²⁰

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 outliers. 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 lose 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 compute 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 levels 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 6 shows descriptive statistics of the sample used in estimation.

Construction of the Moments

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 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

²⁰The percentages are computed for households belonging to the Poor group. For Densificados, the food support accounts for 81% of the transfers.

villages in the control and thirty villages in the treatment sub sample. 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 shows that the mean standard deviation of log earnings does not differ between control and treatment villages (except than for the Non Poor group: however, the difference is only by one basis point). We 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 case. Table 8 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.²¹ 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 than the corresponding households in the control villages, 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*, vol. 113 (March), pp. C77-C85.
- [3] Alvarez F. and Jermann U.J. (2000). “Efficiency, Equilibrium, and Asset Pricing with Risk of Default”, *Econometrica*, vol. 68, pp. 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 Rios Rull V. (2000a). “Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?”, *European Economic Review*, vol. 44, pp. 1225-1258.

²¹In the regression, control variables are age, sex and education dummies, and a time dummy for the last wave.

- [6] Attanasio O. and Rios Rull V. (2000b). “Consumption Smoothing and Extended Families”, *Advances in Economic Theory: Eighth World Congress*.
- [7] Behrman J. and Todd P. (2000). “Randomness in the Experimental Samples of Progresa”, *International Food Policy Research Institute*, Washington D.C.
- [8] Cagetti M. (2003). “Wealth Accumulation Over the Life Cycle and Precautionary Savings”, *Journal of Business and Economic Statistics*, vol. 21 (3), pp. 339-353.
- [9] Cox D. (1987). “Motives for Private Income Transfers”, *Journal of Political Economy*, vol. 95, pp. 508-546.
- [10] Cox D., Hansen B. and Jimenez E. (2004). “How Responsive are Private Transfers to Income? Evidence from a Laissez-Faire Economy”, *Journal of Public Economics*, vol. 88, pp. 2193-2219.
- [11] Gourinchas P.O. and Parker J.A. (2002). “Consumption over the Life Cycle”, *Econometrica*, vol. 70, pp. 47-89.
- [12] Krueger D. and Perri F. (2005). “Public versus Private Risk Sharing”, Mimeo, NYU.
- [13] Krueger D. and Perri F. (2006). “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory”, *Review of Economic Studies*, vol. 73, pp. 163-193.
- [14] Kurosaki T. (1999). “Risk Sharing Arrangements and the Structure of Risk and Time Preferences: Theory and Evidence from Village India”, Mimeo, Institute of Economic Research, Hitotsubashi University.
- [15] Ligon E., Thomas J.P. and Worrall T. (2000). “Mutual Insurance, Individual Savings, and Limited Commitment”, *Review of Economic Dynamics*, vol. 3, pp. 1-47.
- [16] Ligon E., Thomas J.P. and Worrall T. (2002). “Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies”, *Review of Economic Studies*, vol. 69, pp. 209-244.
- [17] Mazzocco, M. and Saini S. (2007). “Testing Efficient Risk Sharing with Heterogeneous Risk Preferences”, Mimeo, UCLA.
- [18] Ogaki M. and Zhang Q. (2001). “Decreasing Relative Risk Aversion and Tests of Risk Sharing”, *Econometrica*, vol. 69 (2), pp. 515-526.
- [19] Schulhofer-Wohl S. (2007), “Heterogeneity, Risk Sharing and the Welfare Cost of Risk”, Mimeo, Princeton University.
- [20] Skoufias E. (2005). “PROGRESA and its Impacts on the Welfare of Rural Households in Mexico”, *International Food Policy Research Institute*, Washington, D.C.

- [21] Skoufias E. (2007). "Poverty Alleviation and Consumption Insurance: Evidence from PROGRESA in Mexico", *The Journal of Socio-Economics*, vol. 36, pp. 630-649.
- [22] 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.
- [23] Todd P. and Wolpin K. (2006). "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility", *The American Economic Review*, vol. 96 (5), pp. 1384-1417.
- [24] Townsend R., (1994). "Risk and Insurance in Village India", *Econometrica*, vol. 62, pp. 539-591.
- [25] Udry C., (1994). "Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria", *Review of Economic Studies*, vol. 61, pp. 495-526.
- [26] Venti S.F. and Wise D.A. (2001). "Choice, Chance, and Wealth Dispersion at Retirement", in Ogura S., Tachibanaki T. and Wise D.A. (eds.), *Aging Issues in the United States and Japan*, Chicago, University of Chicago Press, pp. 25-64.

Moments	Data (Control Villages)	Model
Corr($\Delta \ln c$, $\Delta \ln l$) - Poor	0.106	0.108
Corr($\Delta \ln c$, $\Delta \ln l$) - Densificados	0.125	0.104
Corr($\Delta \ln c$, $\Delta \ln l$) - Non Poor	0.104	0.112
Stdev($\ln c$)/Stdev($\ln l$) - Poor	0.917	0.694
Stdev($\ln c$)/Stdev($\ln l$) - Densificado	0.937	0.708
Stdev($\ln c$)/Stdev($\ln l$) - Non Poor	0.888	0.663
Stdev($\ln c$)/Stdev($\ln l$) - b/w groups	0.637	0.929

Table 1: Model Fit - Control Villages

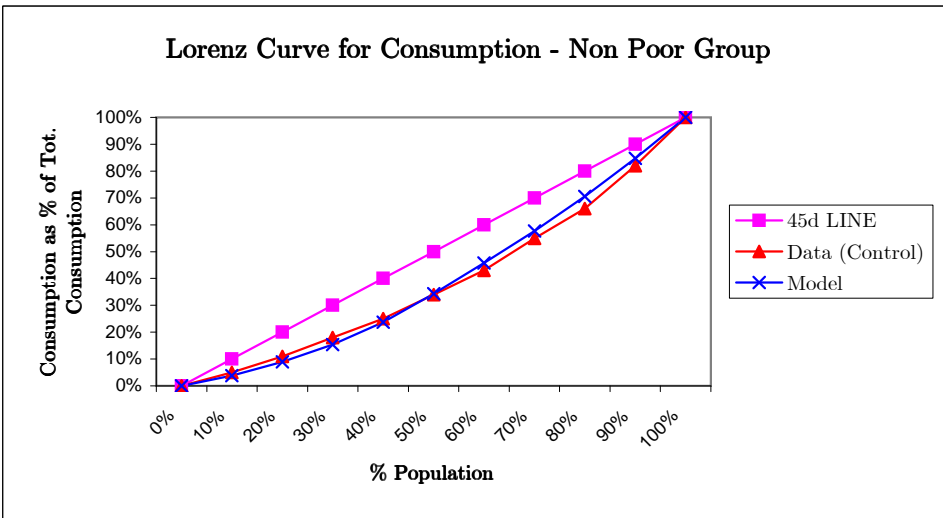

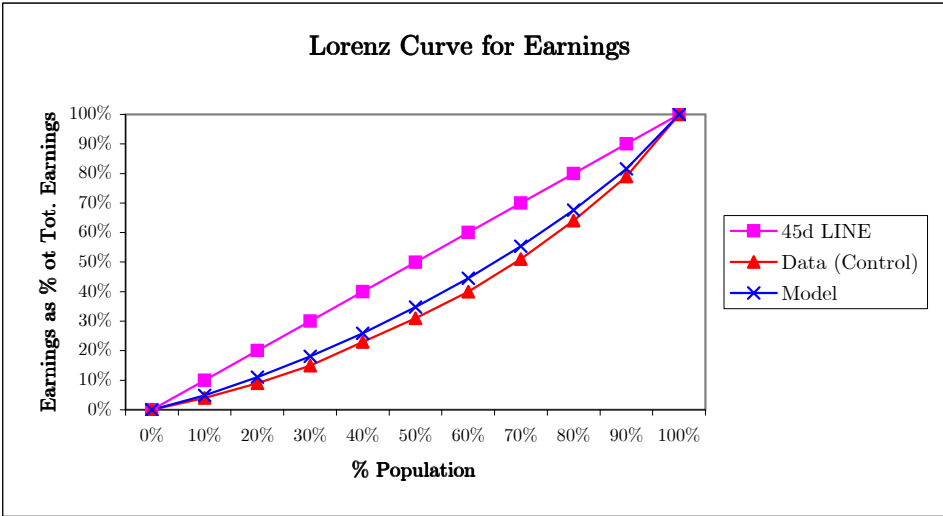


Figure 1: Lorenz Curves for Earnings and Consumption, Model vs. Data

Moments	Data		Model	
	Treatment Villages	% Change wrt Control	Model with Transfers	% Change wrt Benchmark
Corr($\Delta \ln c$, $\Delta \ln I$) - Poor	0.065	-0.384	0.062	-0.424
Corr($\Delta \ln c$, $\Delta \ln I$) - Densificados	0.049	-0.609	0.060	-0.425
Corr($\Delta \ln c$, $\Delta \ln I$) - Non Poor	0.117	0.129	0.062	-0.443
Stdev($\ln c$)/Stdev($\ln I$) - Poor	0.887	-0.033	0.758	0.093
Stdev($\ln c$)/Stdev($\ln I$) - Densificados	0.925	-0.013	0.706	-0.003
Stdev($\ln c$)/Stdev($\ln I$) - Non Poor	0.928	0.045	0.548	-0.175
Stdev($\ln c$)/Stdev($\ln I$) - b/w groups	1.039	0.631	1.444	0.554

Table 2: Model Fit - Treatment Villages

Average Private Transfers over Earnings ($ c-y /y$)		Benchmark	Economy with Transfers	% Change
Type 1 Household	Pobre	0.513	0.457	-0.109
	Densificado	0.513	0.496	-0.033
	Non Pobre	0.512	0.572	0.117
Type 2 Household	Pobre	0.271	0.247	-0.089
	Densificado	0.271	0.248	-0.085
	Non Pobre	0.271	0.284	0.048
Equivalent Consumption Variation*				
Type 1 Household	Pobre		20.0%	
	Densificado		13.4%	
	Non Pobre		4.7%	
Type 2 Household	Pobre		9.5%	
	Densificado		9.5%	
	Non Pobre		9.6%	

Table 3: Risk Sharing and Welfare Measures Before and After the Introduction of the Transfers

* The percentage of consumption to give to the agent at each period in the benchmark model economy in order to make him indifferent between the benchmark and the economy with transfers

Moments	Benchmark	Policy
Corr($\Delta \ln c$, $\Delta \ln I$) - Poor	0.108	0.05
Corr($\Delta \ln c$, $\Delta \ln I$) - Densificados	0.104	0.048
Corr($\Delta \ln c$, $\Delta \ln I$) - Non Poor	0.112	0.052
Stdev($\ln c$)/Stdev($\ln I$) - Poor	0.694	0.725
Stdev($\ln c$)/Stdev($\ln I$) - Densificados	0.708	0.768
Stdev($\ln c$)/Stdev($\ln I$) - Non Poor	0.663	0.679
Stdev($\ln c$)/Stdev($\ln I$) - b/w groups	0.929	1.03

Table 4: Fully Funded Insurance Policy

Average Private Transfers over Earnings ($c-y /y$)		Benchmark	Economy with Insurance	% Change
Type 1 Household	Pobre	0.513	0.459	-0.105
	Densificado	0.513	0.515	0.004
	Non Pobre	0.512	0.557	0.088
Type 2 Household	Pobre	0.271	0.227	-0.164
	Densificado	0.271	0.222	-0.182
	Non Pobre	0.271	0.245	-0.094
Equivalent Consumption Variation*				
Type 1 Household	Pobre		4.1%	
	Densificado		3.9%	
	Non Pobre		3.3%	
Type 2 Household	Pobre		-2.4%	
	Densificado		-3.6%	
	Non Pobre		-2.7%	

Table 5: Risk Sharing and Welfare Measures Before and After the Introduction of the Insurance Policy

* The percentage of consumption to give to the agent at each period in the benchmark model economy in order to make him indifferent between the benchmark and the economy with transfers

Variable	Treatment	Control
Number of Households:		
October/November 1998	13,206	8,277
May/June 1999	12,060	7,574
Nov-99	12,414	8,115
Mean Age of Household Head	47.9	48.4
Gender of Household Head (% Male)	90.3	89.8
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	281	282
Poor	236	232
Densificados	329	328
Non Poor	345	344
Mean Adult Equivalent Monthly Consumption (Food + Clothing):		
All sample	189	175
Poor	177	156
Densificados	203	199
Non Poor	201	192

Table 6: Data Descriptive Statistics

Mean Standard Deviation of Log Earnings*	Poor	Dens.dos	Non Poor	All
CONTROL	0.492 (0.081)	0.532 (0.079)	0.529 (0.075)	0.51 (0.065)
Observations	3502	1650	1173	13799
TREATMENT	0.487 (0.071)	0.533 (0.080)	0.524 (0.066)	0.51 (0.068)
Observations	4897	978	1317	18708
p-value of Difference b/w St. Dev.s	0.003	0.662	0.074	0.271

Table 7: Standard Deviation of Log Earnings, by groups and treatment.

*Standard errors in parenthesis. We control for age, dummies on sex and education of the household head, and a time dummy.

Dependent Var.: % Change in Cons. Share*	Poor	Dens.dos	Non Poor	All
CONTROL: Coeff. on % Change in Earn. Share	0.095 (0.015)	0.11 (0.022)	0.081 (0.021)	0.095 (0.011)
Observations	3576	1554	1315	6445
TREATMENT: Coeff. on % Change in Earn. Share	0.056 (0.012)	0.042 (0.019)	0.097 (0.020)	0.062 (0.009)
Observations	4804	2047	1838	8689
p-value of Difference between Coefficients	0.048	0.0183	0.5766	0.024

Table 8: 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 education of the household head, and a time dummy.