Poverty and Time Preferences

Leandro S. Carvalho*

Princeton University

November, 2008

Abstract

This paper estimates the time preference of poor households in rural Mexico by combining experimental variation with simulation-based econometric methods. I use data from PROGRESA, a program that randomly assigned communities to treatment and control groups and paid cash transfers to poor households in treatment communities. The randomization implies that differences in the consumption profiles of control and treatment households are due to the program. A standard buffer stock model predicts how consumption will respond to cash transfers as a function of the discount factor. I use the method of simulated moments to estimate such a model by matching simulated treatment effects on consumption to sample treatment effects. My estimates indicate that, under a range of assumptions, households in the PROGRESA sample have very low discount factors. It is difficult to reconcile their behavior with the much higher discount factors that have been estimated for U.S. households. I conclude that either poor households are very impatient or a richer model is needed to describe the consumption behavior of poor households in developing countries.

1 Introduction

Empirical evidence shows that the poor and the non-poor behave differently. In particular, they make different decisions when it comes to choices involving intertemporal trade-offs. The poor have lower saving rates (Attanasio and Szekely 2000; Dynan, Skinner and Zeldes, 2004), invest less in their children's education (Behrman, Birdsall and Szekely 1998) and have more children (Schultz 2005). Economists have long debated the underlying reasons for these differences (Duflo 2006) and have suggested numerous explanations. One explanation is that the poor might face lower returns to their investments (Rosenzweig 1995, van de Walle 2003). Another explanation is that some investment opportunities might not be available

*I am grateful to David Lee, Chris Paxson and Sam Schulhofer-Wohl for their advice and support. I am also indebted to Silvia Helena Barcellos, Angus Deaton, Bo Honore, Alisdair McKay, Ashley Ruth Miller, Felipe Schwartzman and seminar participants at the Princeton University Development Lunch, the Princeton University Labor Lunch and the Princeton Public Finance Working Group. All remaining errors are mine. Email: lcarvalh@princeton.edu.

to the poor due to inefficiencies in credit markets or insurance markets (Banerjee 2001).[1] Finally, the poor and the non-poor might have different preferences (Lawrance 1991; Atkeson and Ogaki 1996, 1997). In particular, the poor might be more impatient than the non-poor.

The last explanation suggests that time preferences might be a mechanism that transmits poverty across generations. Becker and Mulligan (1997) argue that parents can spend resources to alter their children's time preferences. For example, schooling can increase children's appreciation of future utility. Moreover, theory predicts that parents with lower discount factors invest less in their children's human capital (Lang and Ruud 1986). These children, in turn, will develop time preferences that are more present-oriented. When they reach adulthood, they will be less educated and more impatient parents. To investigate whether this hypothesis has any merit, it is necessary to know whether the poor and the non-poor – as adults – have different time preferences.

Previous studies (Hausman 1979; Lawrance 1991; Atkeson and Ogaki 1997; Gourinchas and Parker 2002; Cagetti 2003; Harisson, Lau and Williams 2002; Tanaka, Camerer and Nguyen 2006; Stephens and Krupka 2006) have investigated differences in time preferences between poor and non-poor by comparing estimates for the two groups. These studies have produced mixed results.[2] However, one limitation of this research is that it is difficult to distinguish between differences in time preferences, and differences in the economic environments in which the poor and the non-poor live. For example, differences in intertemporal choices of the two groups might be partly explained by the fact that the poor are more likely to be liquidity constrained.[3]

In this paper, I estimate the discount factor of poor Mexican households living in rural communities. I explicitly model liquidity constraints and I present estimates for a wide range of interest rates and risk aversion parameters to address concerns that particular assumptions about the economic environment might drive my results. I also compare my estimates

[1]More recently, it has been suggested that deviations from rational behavior might have worse consequences for the poor given their circumstances (Bertrand, Mullainathan and Shafir 2004; Mullainathan 2006).

[2]Hausman (1979), Lawrance (1991), Harisson, Lau and Williams (2002) and Tanaka, Camerer and Nguyen (2006) find that the poor have higher discount rates. Gourinchas and Parker (2002), Cagetti (2003) and Stephens and Krupka (2006) yield inconclusive results while Atkeson and Ogaki (1997) find no differences in time preferences.

[3]Alternatively, Carroll (1997) argued that Lawrance's(1991) results might be partly explained by the faster labor income growth of more educated households.

of the discount factor with existing estimates for non-poor households from studies that use similar methods. I use data from PROGRESA, a program that randomly assigned communities to "treatment" and "control" groups, and extended conditional cash transfers to poor households in the treatment communities 18 months before those in the control communities. The random assignment implies that, in the absence of the program, there would not be any differences in the consumption behavior of control and treatment households. A model gives predictions on how treatment households would spend the experimental variation in wealth over time as a function of households' time preference. Specifically, I use the canonical buffer-stock model of consumption to simulate the effects of PROGRESA transfers on consumption at different discount factors. I estimate the discount factor using the method of simulated moments by matching simulated and sample treatment effects on log consumption.

My empirical approach bridges two methods that have been used to measure time preferences.[4] The first approach elicits discount rates from subjects by asking them to make intertemporal choices over hypothetical or real outcomes (Coller and Williams 1999; Harrison, Lau and Williams 2002; Andersen et al 2008).[5] The most common experimental procedure asks respondents to choose between a smaller, more immediate reward and a larger, more delayed reward. A benefit of this method is the researcher's ability to control the terms of the experiment. However, there are concerns about the external validity of these estimates. The second method uses observational data on individuals' intertemporal choices. For example, macroeconomists have estimated structural life-cycle consumption models using data on consumption and wealth (Gourinchas and Parker 2002; Cagetti 2003).[6] This approach entails simulating the consumption-age or wealth-age profiles as a function of the parameters of the model, which are estimated by matching the simulated and sample profiles. The limitation of this method is that it assumes changes in consumption or wealth over the life cycle

[4]See Frederick, Loewenstein and O'Donoghue (2002) for a review of the literature on estimating time preferences.

[5]Frederick, Loewenstein and O'Donoghue (2002) make the point that there is enormous variation in estimates of the discount rate. They attribute this to differences in procedures but also to a faulty assumption that there is a single discount rate that applies to different choices. The variation is particularly large for articles that elicit discount rates through experiments. Andersen et al (2008) show that estimates of the discount rate are much lower when risk and time preferences are elicited jointly.

[6]One strand of the literature estimates discount rates using data on the purchase of durable goods (e.g., Hausman 1979). Other studies estimate discount rates from wage-risk tradeoffs (Viscusi and Moore 1989).

are entirely due to optimizing behavior as predicted by the model. The approach I take – using observational data that embeds experimental variation in wealth – addresses concerns about confounding factors that arise in observational studies but avoids the concerns about external validity that arise in many experimental studies.

My results indicate that poor rural households in Mexico are very impatient. My benchmark estimates for the annual discount factor range from 0.08 to 0.70. These estimates are much lower than what is found for the much wealthier households in the United States. Three recent articles (Gourinchas and Parker 2002; Cagetti 2003; Laibson, Repetto and Tobacman 2007) use data from large U.S. household surveys to estimate the parameters of a stochastic life-cycle consumption model using the method of simulated moments. Their estimates of annual discount factors range from 0.92 to 0.99 and concentrate around 0.95. I take these estimates for the U.S. as reflecting the time preference of non-poor households and formally test whether households in my sample have an annual discount factor above 0.90.[7] For almost all model specifications, I am able to rule out at a 1% significance level any annual discount factor above 0.90.

The intuition for my result is as follows. PROGRESA transfers increased over the first few semesters of the program. If households faced no liquidity constraints, they would have smoothed consumption either by drawing down their own assets when the program started, or by borrowing against future transfers. This behavior would have produced treatment effects in consumption that declined over time, with larger increases in consumption in the earlier periods of the program. However, the data do not show such a pattern. Instead, the treatment effects on log consumption increased over time, implying an absence of consumption smoothing. The observed pattern of treatment effects is consistent with borrowing constraints working in concert with impatience. A key point is that patience governs the stock of assets that households can use for consumption smoothing. Patient households hold larger stocks of assets before the program, on average, and so can use their own assets to smooth consumption even in the presence of borrowing constraints. Impatient house-

[7]The fact that the different articles match different moments and still find estimates in the same ballpark suggests that, despite potential biases, these estimates are not too far from the true discount factor .

holds, who hold smaller stocks of assets, cannot. A high degree of impatience – implying lower initial assets – is necessary to reconcile the model with the rising treatment effects in consumption observed in the data.

The structure of the paper is as follows. Section 2 gives some background on PROGRESA. Section 3 introduces the model. Section 4 discusses the identification strategy and the method of simulated moments. Section 5 presents the empirical results, and section 6 presents robustness checks. Section 7 concludes with a discussion of alternative models that could be consistent with the data.

2 Background on PROGRESA

PROGRESA is a conditional cash transfer program in Mexico that pays cash transfers to households which comply with the program requirements. To evaluate the program, a randomized social experiment was conducted in its initial phase: 506 rural communities were randomly assigned to either participate in the program or serve as controls. A survey was conducted before the start of the program to identify eligible families. Eligible households in treatment communities started receiving transfers 18 months before eligible households in control communities.

PROGRESA's cash transfers have two major components: a scholarship conditional on 85% or higher attendance of school-age children, and a cash transfer for food that is conditional on mandatory visits of household members to health clinics.[8] Households are not required to spend the transfer on food or school-related expenses. Children over 7 and under 18 years are eligible to receive scholarships. The amount of the scholarship depends on the grade and sex of the child, and there are no payments during the July-August school holiday period. The cash grant for food is the same for all families. The payments are made every two months, and the total transfer per household is capped.[9] Hoddinott, Skoufias and Washburn (2000) report that average transfers corresponded to roughly 20% of average

[8]Households also get twice a year support for school supplies; the amount varies according to children's educational level.
[9]Benefits are adjusted every 6 months for inflation according to the national CPI.

expenditures of eligible households in control communities. I do not model the program conditionalities and the household's participation decision. A majority (approximately 90%) of eligible households in treatment communities enrolled in the program (Gertler, Martinez and Rubio-Codina 2006).

In 1997, PROGRESA conducted a census in treatment and control communities to identify eligible households. Households were classified as above or below a poverty line based on a measure of household income that excluded child income. Variables that best discriminated between poor and non-poor were then identified, and an equation was developed to compute an index – a discriminant score – that is used to classify households as poor/eligible or non-poor/noneligible.[10] The index is meant to reflect a household's permanent income rather than income reported at the baseline survey. In this article, I use the term "eligible" to refer to households in both control and treatment communities which were classified by PROGRESA as poor.

In the summer of 1998, government officials visited each community that had been randomized into the treatment group and announced in a public meeting that the program would be available to eligible households (Schultz 2004). At the public meeting, officials disclosed the list of eligible households and explained the goals and rules of the program to beneficiaries. Program officials informed households that benefits would be paid for three years and that their status would not change irrespective of changes in household income.[11] Treatment households started receiving transfers in May 1998.

Households in control communities were not incorporated into the program until November 1999 and started receiving transfers in December 1999. Control households were not notified about the program until two months before incorporation and surveyors were told to conceal that the survey was related to PROGRESA.[12]

[10]For more details, see Skoufias, Davis and Behrman (1999).

[11]The existence of the program could not be guaranteed beyond 2000 because of the 2000 general election. It was common in Mexico for each new administration to scrap existing programs and to start its own (Levy 2006).

[12]Interviewers reported an increasing number of complaints from households in control communities as the program expanded, sometimes incorporating households in localities close by the control communities (Parker and Teruel 2005). There is some controversy on whether control households anticipated participation in the program. See Attanasio et al (2004), Skoufias (2005) and Todd and Wolpin (2006).

3 The Model

3.1 Theoretical Model

The framework is a standard discrete-time intertemporal allocation problem (Deaton 1991a). Infinitely-lived households maximize their lifetime utility:

$$\max_{\{C_{i,t+k}\}_{k=0}^{\infty}} E_t \left[\sum_{k=0}^{\infty} \beta^k u\left(C_{i,t+k}\right) \right], \tag{1}$$

where u is instantaneous utility, β is the discount factor, C is consumption, t indexes time and i indexes the household. Every period, households receive stochastic labor income $Y_{i,t}$ – which is the only source of uncertainty – and some of them receive a positive cash transfer $B_{i,t}$. There is a single asset in the economy with a constant gross real interest return R, and the evolution of assets is given by

$$A_{i,t+1} = R\left(A_{i,t} + Y_{i,t} + B_{i,t} - C_{i,t}\right). \tag{2}$$

The consumer is subject to a borrowing constraint

$$A_{i,t} \geq 0. \tag{3}$$

The instantaneous utility function is of the constant relative risk aversion (CRRA) form with coefficient of relative risk aversion ρ:

$$u\left(C\right) = \frac{C^{1-\rho}}{1-\rho}. \tag{4}$$

I assume that transitory shocks to income are lognormally distributed:

$$\ln Y_{i,t} = \ln \omega_i + \varepsilon_{i,t}, \tag{5}$$

where ω_i is the household-specific permanent component of income and $\varepsilon_{i,t}$ represents transitory shocks to income.

The literature for the U.S. (Zeldes 1989; Carroll 1997; Gourinchas and Parker 2002; Cagetti 2003) models income as a nonstationary serially correlated process and adopts a finite horizon version of the model. Here, I assume that the income process is stationary because the available data are not sufficient to estimate the parameters of a nonstationary process. I initially assume that transitory shocks to income are independently identically distributed. However, I also present results when transitory shocks follow an AR(1) or MA(1) process. I assume that agents are infinitely lived, which means I do not have to specify the value function at retirement (Gourinchas and Parker 2002) or the utility from leaving bequests (Cagetti 2003). The majority of households in the sample are young couples with school-age children and are mostly saving for precautionary reasons rather than for retirement. Carroll (1997) shows that, for the U.S., the infinite-horizon model is a good approximation to the finite horizon model up to age 50.

I follow Zeldes (1989) and Deaton (1991a) in imposing a liquidity constraint, which seems particularly relevant in the context of rural Mexico. An alternative approach (Carroll 1997, Gourinchas and Parker 2002) is to set up the model such that the consumer chooses never to borrow by assuming that there is a positive probability that income will be equal to zero.[13] Deaton (1991b) questions whether this assumption should be taken seriously since households might always be able to guarantee some minimal consumption level with help from friends or family.

3.2 Modeling the Program

I explain how the introduction of a cash transfer such as PROGRESA affects consumption decisions within the framework of the canonical buffer stock model. I use the superscript 0 to denote control households and 1 to denote treatment households. To simplify the exposition, I present the case in which transitory shocks are i.i.d.[14] I observe consumption

[13] Schechtman (1976) showed that if (1) an Inada condition $\lim_{c\to 0} u'(c) = +\infty$ holds and 2) there is a strictly positive probability that labor income will be arbitrarily close to zero, then there is a strictly positive probability that cash on hand will be zero and the marginal utility infinity if the agent holds no assets. Thus, the agent finds optimal to always have strictly positive assets.

[14] If transitory shocks follow an AR(1), then the state vector is formed by cash on hand and current income. If transitory shocks follow a MA(1), the state vector is formed by cash on hand and the current innovation to income.

every 6 months, and my analysis uses data from three 6-months "semesters." To match the frequency of the data, I assume that households solve their intertemporal problem on a semi-annual basis. Henceforth, all discount factors presented are semi-annual discount factors unless noted otherwise.

I initially assume that control households did not anticipate their participation in the program, in which case their intertemporal allocation problem is the standard one.[15] (I will relax this assumption in section 6.) The value function of a control household is:

$$V^0\left(X_{i,t};\omega_i\right) = \max_{\left\{C_{i,t+k}\right\}_{k=0}^{\infty}} u\left(C_{i,t}\right) + \beta E_t\left[V^0\left(X_{i,t+1};\omega_i\right)\right], \tag{6}$$

subject to constraints (2) and (3) and where $X_{i,t} = A_{i,t} + Y_{i,t} + B_{i,t}$ is cash on hand, the sum of assets, income and transfers.

Treatment households behave differently from controls since treatment households receive transfers from the program and consequently face a different budget constraint. If the program was a one-period unanticipated intervention, then the policy function for control and treatment households would be the same – although it would be evaluated at different points because control and treatment households would have different levels of cash on hand. However, if treatment households anticipate receiving transfers for more than one period, control and treatment households have different policy functions. In particular, the consumption behavior of treatment households depends on their expectations about future transfers.

I initially assume that treatment households did not expect the introduction of the program but that, once the program was announced, households had no uncertainty about the levels of transfers in future periods. I model the seasonality in transfers – transfers are smaller in one semester because households are not entitled to scholarships during the school holiday period – by letting households have different expectations for each semester. Define the "summer" semester as the six-months period from May to October and the "winter"

[15]This choice corresponds to assuming that control households expected future transfers to be equal to zero: $E_t\left[B_{i,t+k}\right] = 0$ for any $k \geq 0$.

semester as the period from November to April of the following year. Let $t = 1$ denote the first semester of the program, from May to October of 1998, so odd periods denote summer semesters and even periods denote winter semesters.

I assume that treatment households perfectly forecast transfers two semesters ahead and expect transfers to remain constant from there into the future. Thus, in the summer semester of 1998, treatment household i expects to receive in future winter semesters $B_{i,2}$ (the winter semester of 1998 actual transfers) and in future summer semesters $B_{i,3}$ (the summer semester of 1999 actual transfers):

$$\text{future winters} \quad : \quad E_{t=1}\left[B_{i,2+2k}\right] = B_{i,2}, \tag{7}$$

$$\text{future summers} \quad : \quad E_{t=1}\left[B_{i,3+2k}\right] = B_{i,3} \text{ for } k \epsilon \left\{0, 1, 2, ...\right\}. \tag{8}$$

Similarly, in the winter semester of 1998 ($t = 2$), household i expects to receive in future summer semesters $B_{i,3}$ (the summer semester of 1999 actual transfers) and in future winter semesters $B_{i,4}$ (the winter semester of 2000 actual transfers):

$$\text{future summers} \quad : \quad E_{t=2}\left[B_{i,3+2k}\right] = B_{i,3}, \tag{9}$$

$$\text{future winters} \quad : \quad E_{t=2}\left[B_{i,4+2k}\right] = B_{i,4} \text{ for } k \epsilon \left\{0, 1, 2, ...\right\}. \tag{10}$$

Treatment households expect the program to last for 6 semesters, after which they expect to stop receiving cash transfers. The value function of treatment households who have τ periods remaining until the end of the program is:

$$V^{1,\tau}\left(X_{i,t}, B_{i,t+1}, B_{i,t+2}; \omega_i\right) = \max_{\left\{C_{i,t+k}\right\}_{k=0}^{\infty}} u\left(C_{i,t}\right) + \beta E_t\left[V^{1,\tau-1}\left(X_{i,t+1}, B_{i,t+2}, B_{i,t+3}; \omega_i\right)\right], \tag{11}$$

subject to constraints (2) and (3), where $\tau = \max\left\{7 - t, 0\right\}$. At the end of the six semesters – i.e, when $\tau = 1$ – the value function of a treatment household is equal to the value function of a control household:

$$V^{1,\tau}\left(X_{i,t},\cdot;\omega_i\right)=V^{0}\left(X_{i,t+1};\omega_i\right)\ \text{ for }\tau\leq 1. \tag{12}$$

The isoelastic utility function makes the problem homogeneous of degree 1 in the permanent component of income ω. I write the optimal consumption rule as a function of the state variables and transfers, all normalized by ω:

$$c_{i,t}=\frac{C_{i,t}}{\omega_i},\ x_{i,t}=\frac{X_{i,t}}{\omega_i}\text{ and }b_{it}=\frac{B_{i,t}}{\omega_i}. \tag{13}$$

The policy function of a control household is given by:

$$c_{it}=c^{0}\left(x_{i,t}\mid\Omega\right), \tag{14}$$

where Ω is the vector containing the key parameters – β,ρ,R and σ_ν^2, where σ_ν^2 is the variance of transitory shocks to income – and c^0 is the control policy function. The policy function of a treatment household is given by:

$$c_{i,t} = c^{1,\tau}\left(x_{i,t},b_{i,t+1},b_{i,t+2}\mid\Omega\right)\ \text{ for }\tau\geq 2, \tag{15}$$

$$c_{i,t} = c^{1,\tau}\left(x_{i,t}\mid\Omega\right)=c^{0}\left(x_{i,t}\mid\Omega\right)\ \text{ for }\tau\leq 1, \tag{16}$$

where $c^{1,\tau}$ is the policy function of a treatment household with τ periods remaining until the end of the program.

The general consumption rule can be written as:

$$c\left(s_{i,t},D_i|\Omega\right)=c^{0}\left(x_{i,t}|\Omega\right)\left(1-D_i\right)+c^{1,\tau}\left(s_{i,t}|\Omega\right)D_i, \tag{17}$$

where $s_{i,t}=(x_{i,t},b_{i,t+1},b_{i,t+2})$ and D_i is a treatment status dummy. The consumption of household i in period t is then given by:

$$C_{i,t}=c\left(s_{i,t},D_i|\Omega\right)\omega_i. \tag{18}$$

The consumption problems of control and treatment households are solved respectively by policy function iteration and backward recursion. The first step is to solve an infinite horizon dynamic program to find the policy and the value functions for control households. The second step involves solving a finite horizon dynamic program to find the policy function for treatment households, where the horizon is equal to six semesters, the expected duration of the program. The terminal value of the treatment household problem is given by the value function of control households.

3.3 The Income Process

The data generating process for observed household income is assumed to be:

$$\ln Y^*_{i,t} = \ln Y_{i,t} + \xi_{i,t} = \ln \omega_i + \varepsilon_{i,t} + \xi_{i,t}, \tag{19}$$

where $Y^*_{i,t}$ is observed household income, $Y_{i,t}$ is actual household income, ω_i is household permanent income, $\varepsilon_{i,t}$ captures transitory shocks to income and $\xi_{i,t}$ is measurement error.

I decompose permanent income into household size and per capita permanent income:

$$\ln \omega_i = \ln hhsize_i + \ln \mu_i, \tag{20}$$

and assume that $\ln \mu_{i,t}$ is normally distributed with mean $\overline{\ln \mu}$ and variance $\sigma^2_{\ln \mu}$.

The program identified eligible households by choosing the ones more likely to be poor based on a proxy for per capita permanent income. Thus, I assume that $\ln \mu$ is distributed among eligible households as a truncated normal with mean $\overline{\ln \mu}$ and variance $\sigma^2_{\ln \mu}$ bounded above by some threshold level b.

I assume transitory shocks and measurement error are independent of permanent income. Measurement error is serially uncorrelated and normally distributed with mean zero and variance σ_{ξ}^{2}. The distribution of the transitory shock $\varepsilon_{i,t}$ depends on the income process:

$$\begin{aligned} iid &: \varepsilon_{i,t} = \nu_{i,t}, \\ AR(1) &: \varepsilon_{i,t} = \eta\varepsilon_{i,t-1} + \nu_{i,t}, \\ MA(1) &: \varepsilon_{i,t} = \eta\nu_{i,t-1} + \nu_{i,t}, \end{aligned} \tag{21}$$

where

$$\nu_{i,t} \overset{iid}{\sim} N\left(0, \sigma_{\nu}^{2}\right). \tag{22}$$

4 Methodology

4.1 Identification Strategy

The data-generating process for observed consumption C^{*} for a family i in period t can be generally described by:

$$\ln C_{it}^{*} = \ln C_{it} + \epsilon_{it} = \ln\left(c\left(s_{it}, D_{i}|\Omega\right)\omega_{i}\right) + \epsilon_{it}, \tag{23}$$

where the first term on the right hand side is optimal consumption predicted by the model and ϵ is an error term that captures deviations of observed consumption from optimal consumption. The error term can be thought of as measurement error in consumption levels.

The goal is to estimate β by minimizing the distance between theoretical moments and sample moments. The equation above states that the mapping of s and D to consumption is a function of Ω, so it would be natural to match on moments that summarize co-movements between the right-hand-side variables and consumption. Let z be a vector including some of the right-hand-size variables. One possible set of matching moments would be the coefficient vector $\widehat{\gamma}_{1}$ from an OLS regression of log consumption on z. The probability limit of these coefficients is:

$$p\lim\widehat{\gamma}_{1} = Var\left(z\right)^{-1}Cov\left(\ln C, z\right) + Var\left(z\right)^{-1}Cov\left(\epsilon, z\right). \tag{24}$$

The model predicts the first term on the right hand side of the equation above. However, the equation makes it clear that the sample moments are not consistent estimators of the theoretical moments if the second term on the right hand side is different from zero. In other words, the theoretical moments differ from the observed moments at the true parameter vector if the variables that provide the source of variation used to identify Ω are correlated with the error term.

Gourinchas and Parker (2002) estimate the parameters of the intertemporal allocation problem by matching the theoretical consumption-age profile to its empirical counterpart, which is equivalent to having as z a vector of age dummies. Their identifying assumption is that the error term is uncorrelated with age. But if there are other explanations for why observed consumption evolves over the life cycle, the second term on the right hand side of (24) will differ from zero. For example, the error term will be correlated with age if there is time-varying measurement error in consumption or if omitted factors that determine consumption vary over the life cycle.

In this article, I exploit the fact that the consumption behavior of control and treatment households would have been the same in the absence of the program. In other words, any treatment-control differences in consumption must be due to the program. I use the model to predict program treatment effects on log consumption and match these to sample treatment effects. In this case, z is a vector of time-treatment dummies.

I estimate the time preference of sample households from the following moment conditions:

$$E^1\left[\ln C^*_{i,t}\right] - E^0\left[\ln C^*_{i,t}\right] = E^1\left[\ln c\left(s_{i,t}, D_i|\Omega_0\right)\omega_i\right] - E^0\left[\ln c\left(s_{i,t}, D_i|\Omega_0\right)\omega_i\right], \qquad (25)$$

where $E^d\left[\cdot\right] = E\left[\cdot|D = d\right]$. The expression on the left hand side is the sample treatment effect on log consumption at period t while the expression on the right hand side is the treatment effect predicted by the model. The identifying assumption is that the error term ϵ is orthogonal to treatment status:

$$E\left[\epsilon_{i,t} \mid D_i\right] = E\left[\epsilon_{i,t}\right]. \qquad (26)$$

The discussion above clarifies the advantage of using experimental variation to estimate time preferences. When using changes in consumption over the life cycle to identify the parameters of interest, the assumption is that the changes are entirely due to the optimizing consumption behavior described by the model. This rules out any unobservables that are correlated with age and with consumption. When using treatment-control differences in consumption, the nature of the experiment implies that differences in the consumption profiles of control and treatment households are due to the program. The model is used to describe how treatment households' consumption responds to the injection of PROGRESA transfers, relative to the consumption of control households. The identifying assumption is that the mean deviation of log observed consumption from log optimal consumption is the same for control and treatment households.

4.2 Method of Simulated Moments

This section justifies the use of simulation methods and explains the simulation procedure.

The dataset does not contain reliable data on households' permanent incomes nor on their wealth, so $s_{i,t}$ is unobservable. Thus, s must be integrated out of the moment conditions:

$$E^d \left[\ln c \left(s_{i,t}, D_i | \Omega\right)\right] = \int \ln c \left(s, D | \Omega\right) dF_t \left(s; \Omega \mid D = d\right), \tag{27}$$

where F_t is the cumulative distribution of s at date t conditional on treatment status D.

The consumption function does not have a closed form, so that (27) cannot be calculated analytically. Another option is to compute (27) numerically, but evaluating F is cumbersome, and I follow Gourinchas and Parker (2002) in using the method of simulated moments to overcome this difficulty. The idea of the method of simulated moments is to draw many realizations of s from $dF_t(\cdot)$, compute $\ln c$ for each s, and calculate the mean of $\ln c$ across the simulated draws.[16]

To simulate consumption for a given household, I need data on assets a, income y, transfers

[16] For earlier formulations of the method of simulated moments, see McFadden (1989), Pakes and Pollard (1989), and Duffie and Singleton (1993).

b and treatment status D. Recall that all lowercase variables are normalized by permanent income ω_i. For each household i, I generate $H = 50$ vectors ζ^h:

$$\zeta_i^h = \left(\boldsymbol{\nu}_i^h, \mu_i^h\right), \tag{28}$$

where μ_i is per capita permanent income and

$$\boldsymbol{\nu}_i^h = \left(\nu_{i,-99}^h, \nu_{i,-98}^h, ..., \nu_{i,3}^h\right) \tag{29}$$

is a history of income shocks going from 100 semesters before the program to 3 semesters into the program. The generated normalized income for household i in the i.i.d. case is:

$$\ln y_{i,t}^h = \nu_{i,t}^h. \tag{30}$$

The next step is to simulate assets. The randomization implies that the asset distribution for control and treatment households should have been the same before the start of the program. Thus, the goal is to simulate the pre-program asset distribution. With that in mind, I take advantage of the stationary of the asset distribution and use the policy function of control households and the history of income shocks to simulate the asset distribution at the eve of the program. I set $a_{i,-99}^h = 0\ \forall i, h$ and run history for 100 periods:

$$a_{i,t+1}^h = R\left(a_{i,t}^h + y_{i,t}^h - c_{i,t}^h\right), \tag{31}$$

$$c_{i,t}^h = c^0\left(a_{i,t}^h + y_{i,t}^h|\Omega\right), \tag{32}$$

where $t\epsilon\{-99, ..., 0\}$. Notice that the generated asset distribution is a function of the parameters of the model.

Finally, I simulate consumption after the program was introduced. At this time, treatment households are receiving transfers and therefore different policy functions apply for control and treatment households. I use survey data on household i's size to compute:

$$\omega_i^h = \mu_i^h * hhsize_i \tag{33}$$

and administrative data on transfers B to calculate:

$$b_{i,t}^h = \frac{B_{i,t}}{\omega_i^h}. \tag{34}$$

The data on transfers and household size come from the actual household-level data and are not simulated.

Consumption for a control household i is simulated as:

$$C_{i,t}^h = c^0 \left(a_{i,t}^h + y_{i,t}^h | \Omega \right) * \omega_i^h \tag{35}$$

while consumption for a treatment household i is simulated as:

$$C_{i,t}^h = c^{1,\tau} \left(a_{i,t}^h + y_{i,t}^h, b_{i,t+1}^h, b_{i,t+2}^h | \Omega \right) * \omega_i^h, \tag{36}$$

where $\tau = 7 - t$ for $t \epsilon \{1, 2, 3\}$.

The average log consumption for treatment households is simulated by:

$$\widetilde{E}^1 \left[\ln C_{i,t} \left(\Omega \right) \right] = \frac{1}{H} \frac{1}{N_1} \sum_{h=1}^{H} \sum_{i=1}^{N} D_i \ln C_{i,t}^h \tag{37}$$

and in a similar way for control households:

$$\widetilde{E}^0 \left[\ln C_{it} \left(\Omega \right) \right] = \frac{1}{H} \frac{1}{N_0} \sum_{h=1}^{H} \sum_{i=1}^{N} \left(1 - D_i \right) \ln C_{i,t}^h, \tag{38}$$

where $\widetilde{E}$ refers to the simulated moments, N_1 is the number of treatment households and N_0 is the number of control households. For any parameter Ω, the theoretical moment can be replaced by its simulated counterpart.

In estimating the model presented above, two complications emerge. First, the data available are not rich enough to identify all the parameters of the model. In particular, there is not enough variation in the data to estimate the risk aversion parameter and the interest rate. Thus, I choose to present results for different values of the interest rate and the risk aversion. Second, I follow Gourinchas and Parker (2002) in employing a two-stage estimation procedure to estimate the parameters of interest. In the first stage, I use income data to estimate the variance of log permanent income and the variance of transitory shocks. In the second stage, I use the treatment effects on log consumption to estimate the discount factor and the mean of permanent income.

Define:

$$\psi = \left(\sigma_{\nu}^{2}, \sigma_{\ln\mu}^{2}\right), \qquad (39)$$

$$\theta = \left(\overline{\ln\mu}, \beta\right). \qquad (40)$$

The parameter vector ψ is estimated in the first stage while θ is estimated in the second stage. Notice that $\widehat{\theta}$ is a function of $\widehat{\psi}$ as well as ρ and R since $\widehat{\theta}$ will be estimated conditional on particular values for the interest rate and the risk aversion parameter.[17]

I proceed by defining g_t, the difference between simulated treatment effects and sample treatment effects on log consumption:

$$g_t\left(\theta, \widehat{\kappa}\right) = \left(E^1\left[\ln C_{it}^{*}\right] - E^0\left[\ln C_{it}^{*}\right]\right) - \left(\widetilde{E}^1\left[\ln C_{it}\left(\theta, \widehat{\kappa}\right)\right] - \widetilde{E}^0\left[\ln C_{it}\left(\theta, \widehat{\kappa}\right)\right]\right), \qquad (41)$$

where $\widehat{\kappa} = \left(\widehat{\psi}, \rho, R\right)$.

The Method of Simulated Moments estimator is given by:

$$\widehat{\theta}_{MSM} = \arg\min_{\theta} \begin{pmatrix} g_1\left(\theta, \widehat{\kappa}\right) \\ g_2\left(\theta, \widehat{\kappa}\right) \\ g_3\left(\theta, \widehat{\kappa}\right) \end{pmatrix}' W^{*} \begin{pmatrix} g_1\left(\theta, \widehat{\kappa}\right) \\ g_2\left(\theta, \widehat{\kappa}\right) \\ g_3\left(\theta, \widehat{\kappa}\right) \end{pmatrix}, \qquad (42)$$

where the three periods refer to the three semesters during which the program was random-

[17] The standard errors presented in section 5.4 do not correct for the first stage estimation.

ized and where the optimal weighting matrix W^* is calculated from survey data.

5 Empirical Analysis

In the next subsection, I describe the data used in the empirical analysis. I then discuss how I estimate the parameters of the income process and present the simulated treatment effects. Finally, I present estimates of the discount factor.

5.1 Data

The empirical exercise involves three steps. First, I use survey data on household consumption and treatment status to estimate the program treatment effects on log consumption. Second, I use survey data on household per capita wage income to estimate parameters of the income process. Finally, I simulate consumption in the model taking the estimated parameters of the income process as given, and find the discount factor that best matches the model to the data. To simulate the consumption of a household, one needs data on transfers, household size, assets, income and permanent income. I use household-level data on transfers and household size in the simulation exercise. Income – including permanent income – and assets are simulated as explained in Section 4.2.

The PROGRESA survey data consists of longitudinal data on approximately 24,000 households from 506 communities – 320 selected into the treatment group and 186 into the control group – that were interviewed in the November 1997 baseline survey. PROGRESA conducted follow-up surveys approximately every six months between May 1998 and November 2000. I restrict my sample to households in control and treatment communities who were classified by the program as poor/eligible. These households correspond to 52% of the original sample. I exclude households who had their eligibility status revised from non-eligible to eligible.[18]

[18] In an initial stage, PROGRESA's beneficiary selection method classified 52% of households in the sample as eligible. By July 1999, PROGRESA revised its method and extended eligibility to 54% of households that were initial classified as ineligible. In consequence, eligibility increased from 52% to 78% of the sample. However, about 60% of households living in treatment communities with a status revised from noneligible to eligible were not incorporated into the program until late 2000. It is unclear what were the expectations and the behavior of these households so I use the more stringent eligibility criterion and drop families with revised eligibility status from my eligible sample.

Data on consumption, household size and income come from PROGRESA surveys. Consumption data were collected in all survey rounds starting in May 1998 except for May 2000. I use the November 1998, May 1999 and November 1999 survey rounds, which correspond to the period after the introduction of the program but before controls started receiving transfers. The appendix contains more details about how I construct the measure of household consumption. Income data were collected in all survey rounds except for May 1998. The survey instrument used to collect individual and household income data changed twice during the period, so it is difficult to compare total income data across rounds. Instead, because the questions on wage income did not change across surveys, I use household per capita wage income to estimate the parameters of the income process.

The data on transfers come from administrative data compiled by PROGRESA. The dataset shows the amount of transfers each household was entitled to receive for each two-months period in which it participated in the program – transfers were paid every two months. The dataset does not have information on the dates of transfer payments. I assume that transfers corresponding to participation during months m and $m+1$ were paid sometime during months $m+4$ and $m+5$, a rule that PROGRESA has followed since 2000.[19] This assumption may be incorrect because before 2000, there were discrepancies across localities in the timing of payments.[20] Table A9 of the appendix shows results under alternative assumptions about the timing.[21]

5.2 Estimating the parameters of the Income Process

I assume the data generating process for observed household per capita income is:[22]

[19] I adjust transfer values to September-October 1998 figures using the monthly national CPI calculated by Banco de Mexico and for each household summed the values received in a given semester.

[20] Local governments are in charge of processing the forms necessary for authorizing the payments and it took some time for PROGRESA to harmonize the procedure across different communities. The forms served to verify whether households were meeting the program conditionalities.

[21] Hoddinott, Skoufias and Washburn (2000) suggest that there was some "lumpiness" in transfer payments with large payments in months such as December 1998, July 1999 and December 1999. In one of the robustness checks, I show results assuming that transfers that should have been paid during September-October 1998 – by the end of the first semester – were actually paid during November-December 1998. In a different exercise, I assume that it took 2 months for PROGRESA to make transfer payments.

[22] In previous sections, I used lowercase letters to denote variables which were normalized by ω. To avoid making the notation more cumbersome, in this section I use y to denote per capita income.

$$\ln y_{i,t}^{*} = \ln y_{i,t} + \xi_{i,t} = \ln \mu_i + \varepsilon_{i,t} + \xi_{i,t}, \tag{43}$$

where $y_{i,t}^{*}$ is observed household per capita income, $y_{i,t}$ is actual household per capita income, $\xi_{i,t}$ is measurement error, μ_i is permanent per capita income and $\varepsilon_{i,t}$ captures transitory shocks to income. The distribution of $\varepsilon_{i,t}$ depends on the income process, as described in equation (21).

The income process implies the following theoretical moments:

$$Var_E\left(\ln y_{it}^{*}\right) = Var_E\left(\ln \mu\right) + \varphi_0 \sigma_{\nu}^2 + \sigma_{\xi}^2, \tag{44}$$

$$Cov_E\left(\ln y_{it}^{*}, \ln y_{it-1}^{*}\right) = Var_E\left(\ln \mu\right) + \varphi_1 \sigma_{\nu}^2, \tag{45}$$

where the subscript E indicates that the moments are calculated for the eligible population and φ_0 and φ_1 depend on the process of transitory shocks:

$$i.i.d. \; : \; \varphi_0 = 1; \varphi_1 = 0, \tag{46}$$

$$AR(1) \; : \; \varphi_0 = \frac{1}{1-\eta^2}; \varphi_1 = \frac{\eta}{1-\eta^2}, \tag{47}$$

$$MA(1) \; : \; \varphi_0 = 1 + \eta^2; \varphi_1 = \eta. \tag{48}$$

In Sections 2 and 3.3, I discussed how households were classified as eligible or non-eligible based on a proxy for their permanent income. Thus, the log permanent income is assumed to be distributed among the eligible population as a truncated normal. The variance of the unconditional distribution is given by:

$$\sigma_{\ln \mu}^2 = Var\left(\ln \mu_i\right) = \frac{Var_E\left(\ln \mu_i\right)}{\left[1 - \alpha\gamma\left(\alpha\right) - \left[\gamma\left(\alpha\right)\right]^2\right]}, \tag{49}$$

where $\alpha = \Phi^{-1}\left(.52\right)$, .52 is the fraction of households who were deemed eligible, Φ is the standard normal cdf and γ is the Mills ratio.

The variance of transitory shocks to income σ_{ν}^2 and the variance of log permanent income $Var_E\left(\ln \mu\right)$ are estimated by minimizing the distance between theoretical and sample moments. The mean of log permanent income is estimated jointly with the discount factor in the second stage.[23] I use data from five survey rounds on household log per capita wage income.[24] I restrict the sample to eligible households and compute the sample counterparts of the theoretical variance and autocovariance of order 1, calculating separate moments for control and treatment households. The moments are shown in Table A1. I present the estimates of the parameters of the income process in Table 2. In the AR(1) and MA(1) cases, I estimate $Var_E(\ln \mu)$ and σ_ν^2 conditional on particular values of η: .25 and .5. The variance of measurement error is estimated by assuming a reliability ratio λ of 0.7:[25]

$$Var_E(\xi_{it}) = (1-\lambda)\, Var_E(\ln y_{it}^*) . \tag{50}$$

The U.S. literature (Carroll and Samwick 1997, Gourinchas and Parker 2002) has estimated income uncertainty from within-household changes in income over time. Data limitations prevent me from using the same strategy here. First, there is an intrinsic difficulty in measuring income of rural households in a developing country given that a large fraction of these households is engaged in subsistence agriculture or self-employment. Second, the survey instrument used to collect income data changed during the period of analysis. In particular, data on farming and livestock income were collected in only 2 survey rounds. Finally, in each survey round, a large fraction of households reports no income. Given these limitations, I use data on household wage income, whose survey instrument was not modified across surveys. The goal is to provide a ballpark estimate of the uncertainty faced by the households in the sample. Nevertheless, it is a real concern that the variability in wage income data might not accurately reflect the income uncertainty they face. Therefore, in the robustness section, I vary the parameters of the income process drastically and show how the results are affected.

[23] Estimating the mean of permanent income is important because in many household surveys in developing countries income seems to be underreported relative to expenditures.

[24] I trimmed household log per capita wage income which was below the 1st percentile or above the 99th percentile of the distribution of a given survey round.

[25] The reliability ratio is defined as $\lambda = Var_E(\ln y_{it}) / Var_E(\ln y_{it}^*)$. Bound and Krueger (1991) find that the reliability ratio in US cross-sectional data is 0.82 for men and 0.92 for women.

5.3 Simulated Treatment Effects

Before moving on to the empirical analysis, I present results from simulations of the model. These results provide intuition on the economics behind the empirical results, which I present in the next section.

In a buffer-stock model, risk-averse households accumulate assets to use as a buffer against income shocks. However, accumulating assets is costly because it reduces current consumption. The more impatient households are, the more willing they are to substitute current consumption for future consumption, and the less assets they will hold.

Figure 1 shows the pre-program (or control) policy function for three different discount factors – 0.5, 0.75 and 0.9 – assuming an annual interest rate of 5% and a coefficient of relative risk aversion of 2. Consumption is shown on the vertical axis and cash on hand – the sum of assets, income and transfers – on the horizontal axis. The policy function is steeper for lower discount factors, implying that more impatient households have a higher marginal propensity to consume. For low levels of cash on hand, households consume all their cash on hand. For higher levels, households with lower discount factors consume more for any given level of cash on hand.

Suppose some households were randomly chosen to receive a one-time unexpected cash transfer. According to the policy functions shown in Figure 1, more impatient households would consume a larger fraction of the transfer. In reality, treatment households in PROGRESA received transfers for more than one semester and had to decide how to allocate this windfall over time. The transfers increased over the first few semesters of the program. As shown in Table 1, transfers increased roughly 150% between the first and second semesters and 18% between the second and third semesters. If treatment households anticipated rising transfers, they would have tried to smooth consumption. In particular, impatience would have led them to run down their assets in the first semesters of the program and let consumption fall over time. As a consequence, treatment effects in consumption would have declined over time.

The previous intuition relies – given the borrowing constraints – on households having enough assets to smooth consumption. But in a buffer-stock model, households' asset holdings are a function of their time preferences. The more impatient households are, the less willing they are to forego current consumption to accumulate buffer wealth, and hence the lower their assets will be on the eve of the program. Figure 2 presents the simulated pre-program asset distribution for three different discount factors: 0.5, 0.75 and 0.9.[26] Patient households hold, on average, more assets. The relationship of asset holdings to time preferences has important implications. When the program is introduced, treatment households want to smooth consumption in anticipation of higher future transfers. However, because households cannot borrow, they must use their own assets to finance any increase in first-semester consumption that exceeds the first-semester transfer. A high degree of impatience implies that households cannot easily do so because they have a small stock of initial assets; in this case, treatment households' consumption will track transfers over time.

Figure 3 shows simulated average log consumption for the control and treatment groups when households' pre-program asset holdings are a function of time preferences. The simulation uses household-level data on actual transfers. Because patient households have higher initial assets, the average log consumption of control households is an increasing function of the discount factor. Similarly, for low discount factors the average log consumption of treatment households increases with the discount factor. For higher discount factors, the average treatment log consumption becomes a decreasing function of the discount factor because more patient households consume a lower fraction of the transfers. The curves shown in Figure 3 result in the simulated treatment effects shown in Figure 4, which plots the treatment effects at different discount factors. The distances between the curves tell how the windfall is distributed over the first three semesters of the program. The graph shows that the increase in treatment households' consumption is roughly evenly distributed over time if households have a high discount factor, but not if households are less patient.[27]

[26] For more details on how the asset distribution is simulated, see Section 4.2.

[27] In the second and third semesters, the treatment effect is decreasing in the discount factor because more patient households have a lower marginal propensity to consume out of transfers. The treatment effect is non-monotonic in the first semester because households want to increase consumption by more than the transfers in anticipation of higher future transfers. Households have to finance an increase in consumption which exceeds the first semester transfers using their pre-program asset holdings – which is as increasing function of the discount factor – given that they cannot borrow against future transfers.

Finally, I present the time profile of treatment effects. Figure 5 plots the simulated treatment effects against the program semester for three different discount factors: 0.5, 0.75 and 0.9. The profile is positively sloped. Because of borrowing constraints, treatment households' consumption – and the treatment effects – grow over time following the growth in transfers. Moreover, patient households have larger initial assets and are more able to smooth consumption over time. As a consequence, the time profile of simulated treatment effects is flatter for higher discount factors.

Figure 5 also shows the program treatment effects on log consumption estimated from survey data, which give rise to a steep profile. The comparison of the sample treatment effects with the simulated treatment effects suggests that the two can be matched only if households have a low discount factor.

5.4 Results

5.4.1 Treatment Effects

Table 3 presents the program treatment effects on log consumption estimated from the survey data. I will estimate the discount factor by matching simulated treatment effects to these sample treatment effects. The top panel of Table 3 reports the estimates of a regression of household log consumption on treatment status, by program semester.[28] The treatment effects rise over time, from 8.5% for the first semester of the program, to 14.4% in the second semester, to 16.4% for the third semester. These findings are consistent with other work (Stephens 2003; Shapiro 2004) which find that recipients of government transfers such as Social Security checks and food stamps in the U.S. do not smooth consumption between payment dates.

The second panel of Table 3 reports average PROGRESA transfers for the treatment group and average consumption in levels for the control group. The last row of the panel shows the ratio of average transfers to control average consumption, which can be compared

[28] To compute the optimal weighting matrix, I have to restrict my sample to households that had non-missing consumption in all three survey rounds. The sample contains 9,397 observations with non-missing consumption data in all three survey rounds. In the simulation, I choose the missing observations to be exactly the same as in the observed data.

to the treatment effects. If the treatment effect is higher than the ratio of transfers to control consumption, this suggests that treatment households increased their consumption by more than the transfer. The treatment effect was higher than the ratio of transfers to control consumption in the first semester, but lower in the second and third semesters.

5.4.2 Estimates of the Discount Factor

Table 4 presents my benchmark estimates of the discount factor.[29] All estimates of the discount factor are conditional on the interest rate and the risk aversion.[30] The table displays results for nine possible combinations of the interest rate and the coefficient of relative risk aversion. The interest rate varies across columns while the risk aversion varies across rows. I follow Laibson, Repetto and Tobacman (2007) and show results for three different risk aversion coefficients – 1, 2 and 3 – that are commonly used in lifecycle consumption models.[31] I vary the interest rate over a wide range, from -15% to 10%. Annual inflation for the period was approximately 15% and the real interest rate was 5%. The real interest rate would have been -15% if households received no interest on their savings. For each combination of risk aversion and interest rate, four figures are shown: the semi-annual discount factor, its standard error, the p-value from a hypothesis test that the semi-annual discount factor is equal to $\sqrt{0.9}$ (corresponding to an annual discount factor of .9) and the p-value from a test of overidentifying restrictions.

The point estimates of the semi-annual discount factor in Table 4 suggest that households are very impatient. The figures in the table have to be squared to obtain annual discount factors. The highest semi-annual discount factor is 0.84, which corresponds to an annual discount factor of 0.7 and to an annual discount rate of 43%. The lowest semi-annual discount factor is 0.28, which corresponds to an annual discount factor of 0.08 and to an annual discount rate of 1156%. Although the estimates vary according to the values of the risk

[29] Estimates of $\overline{\ln \mu}$ range from 7.29 to 7.33, which corresponds to an average ratio of transfers to income of about .2. Hoddinott, Skoufias and Washburn (2000) estimate that the ratio of average transfers to average expenditures of the control group is 20%.

[30] As discussed in Section 4.2, I do not estimate the risk aversion parameter because I cannot separately identify the discount factor from the risk aversion parameter.

[31] See Laibson, Repetto and Tobacman (2007) for a discussion on the values of the risk aversion parameter.

aversion parameter and the interest rate, all the estimates imply a high degree of impatience and are well below the discount factors typically estimated for the U.S. For comparison, Gourinchas and Parker (2002), Cagetti (2003) and Laibson, Repetto and Tobacman (2007) apply similar methods to U.S. data and obtain estimates of the annual discount factor ranging from 0.921 to 0.989.

Overall, it is hard to reconcile the behavior of poor rural households in my sample with discount factors which have been estimated for the wealthier households in the U.S. Table 4 reports the p-value for the null hypothesis that the semi-annual discount factor is $\sqrt{.9}$ against a one-sided alternative. I can reject the null at 1% for all combinations of risk aversion and interest rate except one: the case with a negative interest rate of 15% and a risk aversion of 1. Despite the large standard errors, I can generally bound the estimates away from annual discount factors above 0.9.[32]

The estimates are sensitive to the risk aversion parameter but change little with the interest rate. As explained in Section 5.3, households decide how much to consume by striking a balance between their impatience and their precautionary motives. When households are more risk averse, the precautionary motives for savings are stronger. If, for example, households were risk neutral, they would have no assets at the eve of the program and treatment households' consumption would track transfers over time. This intuition explains why there is a negative relationship between the discount factor that the data imply and the risk aversion parameter. Similarly, households have more incentives to save when the interest rates are higher, giving rise to a negative relation between the estimated discount factor and the interest rate.

Finally, Table 4 reports a measure of model fit. I estimate two parameters using three moments so there is one overidentifying restriction, which can be tested using a chi-squared test. Table 4 reports the p-value from this test; I cannot reject the overidentifying restriction at any conventional significance level.

[32]The standard errors are large because the estimates are in a region in which the derivatives of the simulated treatment effects with respect to the discount factor are close to zero.

6 Robustness Checks

The results shown in the previous section relied on assumptions about the income process and the households' expectations. Here, I examine how sensitive the results are to changes in some of these assumptions. Tables 5 and 6 present robustness checks assuming an annual interest rate of 5% and a coefficient of relative risk aversion parameter of 2. The appendix replicates the robustness checks for different combinations of the interest rate and the risk aversion parameter.

6.1 Income Process

Table 5 shows results under alternative assumptions about the income process. The first row of the table, labeled row 0, repeats the results for the benchmark case. The benchmark assumes that transitory shocks are i.i.d. with standard deviation of .36, that the standard deviation of log permanent income is .46 and that control and treatment households have the same income process. Rows 1 to 8 present results for alternative assumptions. The first column says which assumption is varied. The remaining columns present the semi-annual discount factor estimate, its standard error and the p-value from a one-sided test of the null hypothesis that the semi-annual discount factor is $\sqrt{0.9}$.

I begin by varying the parameters of the income process. Row 1 reduces the standard deviation of transitory shocks from .36 to .2, which is the estimate of Gourinchas and Parker (2002) from U.S. data. I take their estimate as a lower bound for the income risk rural households in Mexico face. Households hold fewer assets if the income process is less risky, which suggests that treatment effects in consumption would track transfers more closely. The estimate of the discount factor is, as expected, higher, but it remains low relative to the U.S. Row (2) reduces the standard deviation of log permanent income from .35 to .15, which is the variance of log permanent income estimated by Gourinchas and Parker (2002).[33] This change does not affect the policy function but changes the importance of transfers in

[33] In my model, the variance of log permanent income refers to the time-invariant cross-section variance of log permanent income. Gourinchas and Parker (2002) estimate a within-household time variance of log permanent income, which I take to be a lower bound of the cross-sectional variance of my model.

households' budget. The estimate of the discount factor remains roughly the same.

The estimates discussed so far assume that the control and treatment households continued to have the same income process after the program was introduced. However, PROGRESA may have reduced the income (excluding PROGRESA transfers) of treatment households because the program reduced income from child labor or because households enrolled in PROGRESA were not entitled to receive government transfers from other programs. Row 3 of Table 5 shows results when I assume that treatment households' permanent incomes were on average 5% lower than control households' permanent incomes. If treatment households' incomes were lower, they would have less cash on hand to finance an increase in current consumption in anticipation of higher future transfers. Treatment effects in consumption would then track transfers more closely, and the estimated discount factor would be higher. Indeed, the estimate for the semi-annual discount factor raises to 0.66 and the annual discount factor to 0.43 so that the annual discount rate falls to 130%. To check whether PROGRESA affected the income process of treatment households, Table A1 shows the mean, variance and autocovariance of household per capita wage income for control and treatment households and a test of whether the moments are different across the two groups.[34] The evidence suggests that control and treatment households did not have different income processes. Other work finds similar evidence.[35]

Even if control and treatment households have the same income process, control and treatment households might respond differently to shocks. In Row 4 of Table 5, I investigate whether macro shocks can explain the growth of treatment effects over time. I estimate the discount factor assuming that there was an unexpected aggregate shock of -5% to income of control and treatment households in the second and third semesters of the program. When hit by a negative income shock, treatment households use the transfers to buffer the

[34] Table A2 shows moments of total income for the 2 survey rounds when data on farming and livestock income were collected.

[35] Parker and Skoufias(2000) find no effect of PROGRESA on adult labor supply, but they do find effects on children's labor force participation. Gertler, Martinez and Rubio-Codina (2006) compare ineligible households in control and treatment communities and find no differences in consumption or investments in micro-enterprises and farm activities. They also find no differences in the agricultural wages of control and treatment communities. Angelucci and De Giorgi (2006) find no differences in wage income between treatment and control communities for the ineligible. For the eligible, they find no clear pattern. Albarran and Attanasio (2002) Gertler, Martinez and Rubio-Codina (2006) find evidence that PROGRESA had a negative effect on income from transfers.

shock. As a consequence, the consumption of treatment households changes less than the consumption of control households. Thus, treatment effects are expected to be larger in semesters with negative aggregate shocks. Indeed, the estimate of the annual discount factor rises to 0.41 and the annual discount rate falls to 145%.

Finally, I investigate how results change when I relax the assumption that transitory shocks to income are i.i.d. If the income process is positively serially correlated, negative shocks tend to be followed by negative shocks. In this case, households hold more assets and have a higher marginal propensity to save. Thus, in comparison to the i.i.d. case, the profile of treatment effects over time should be flatter and the discount factor smaller. Rows 6 to 9 of Table 5 show results for autoregressive and moving average processes of order 1 for η=.25 and η=.5, where η is the parameter that governs the serial correlation. The results show that estimated discount factors are lower than in the i.i.d. case and that they are a decreasing function of the persistence of the shocks. Although I present estimates of the discount factor conditional on particular values of η, I have also estimated η using second-order autocovariances, and the results suggest that income shocks are positively serially correlated but with very low persistence.

6.2 Expectations

The consumption behavior of control and treatment households also depends on their expectations about the program. In Table 6, I show how results change when I make alternative assumptions about expectations. For comparison, the first row of the table shows the results for the benchmark case.

I first explore alternative assumptions regarding treatment households' expectations about the path of future transfers. As explained in section 5.3, the results are driven by the growth in treatment effects over time. We would expect to observe a similar pattern if households did not anticipate the growth in transfers but instead revised their consumption plans upward as they received higher-than-expected transfers. I examine two ways in which households could fail to foresee the growth in transfers: 1) treatment households could have myopic expectations and expect transfers to remain constant at the current level, and 2) treatment households could be pessimistic about the continuation of the program and expect to receive zero transfers in the future. Rows 1 and 2 of Table 6 show results for these two cases.[36] The estimate for myopic expectations is higher than for the benchmark case. For pessimistic expectations, the estimate is slightly lower. The null of a semi-annual discount factor of $\sqrt{0.9}$ is rejected for myopic expectations at the 5% level but not at 1%. In Table A5, I show that both alternative models fit the data worse than the benchmark model.[37]

Households' consumption behavior also depends on whether they anticipate the program's arrival. Row 4 of Table 6 shows results when control households anticipate that they will be enrolled in the program by the end of 1999, as indeed they were. If households were not liquidity constrained, we would expect the time-profile of treatment effects to be steeper when control households anticipate enrollment than when they do not – because the behavior of control households would have produced a downward sloping (control) consumption-time profile. If households were liquidity constrained, anticipation would have no effect on their behavior since control households would not have been able to increase consumption before receiving transfers. The estimate actually changes very little when I assume control households anticipated the program.

Row 5 of Table 6 assumes treatment households anticipated participation in the program one semester before they started receiving transfers. In this case, treatment households would have increased consumption at the time of the announcement by reducing their buffer wealth. Treatment households' consumption would have tracked transfers more closely during the first three semesters of the program, implying a higher estimate of the discount factor. The estimates increase to an annual discount factor of 0.62, corresponding to an annual discount rate of 61%.

Finally, in row 6 of the table, I assume that households expected the program to last

[36]For myopic and pessimistic expectations, I cannot identify the discount factor separately from the mean log permanent income. Hoddinott, Skoufias and Washburn (2000) estimate that the ratio of average transfers to average expenditures of the control group correspond to 20%. The results I present are estimated conditional on $\overline{\ln \mu} = 7.3$, which corresponds to an average ratio of transfers to income of .2. Higher values of $\overline{\ln \mu}$ correspond to higher estimates of the discount factor but imply implausible low values of average income given that expenditures provide a lower bound for income.

[37]I cannot reject the overidentifying restriction for myopic expectations, but I can reject it at 5% for pessimistic expectations.

for four semesters rather than for six semesters. The longer is the expected duration of the program, the more households will be inclined to substitute current consumption for future consumption because the transfers reduce future uncertainty. The intuition suggests that the estimate of the discount factor should be lower. In fact, the estimate changes very little.

7 Conclusion

My benchmark results indicate that poor households in rural Mexico have very low discount factors, much lower than those found in U.S. data using similar methods. My result is robust to a wide variety of alternative assumptions about households' expectations and the economic environment they face. To obtain discount factors similar to those in the U.S., I would need to assume either that households have an unusually low degree of risk aversion or that the households had mistaken beliefs about future transfers – which seems unlikely because, as explained in section 2, government officials visited each community and explained the program at a public meeting.

There are three possible interpretations of the results. The first is that the poor are very impatient, which has implications both for policy design and for future research.[38] First, poverty reduction programs that try to increase poor households' investments in human capital will be more effective when part of the returns to these investments accrues in the short run. Second, future research should investigate the underlying reasons for differences in time preferences of poor and non-poor households.

The second interpretation is that the model does not contain a sufficiently rich description of households' information and the assets they can hold. In particular, the model abstracts from uncertainty about future transfers and does not allow for the possibility that households hold illiquid assets. When there is uncertainty about future transfers, households are reluctant to deplete their assets to increase consumption in anticipation of higher future transfers: they might not receive any transfers in the future and they would have no buffer

[38] Shapiro (2004) finds similar evidence. He shows that the caloric intake of food stamp recipients falls over a food stamp month and calibrates the annual discount factor of an exponential model to be 0.23.

wealth to shield future consumption from income shocks. Another explanation is that households failed to smooth because they hold mostly illiquid assets, which they could not use to substitute current consumption for future consumption.[39] Households would choose to hold few liquid assets if the returns to available liquid assets were sufficiently lower than the returns to illiquid assets.

The third interpretation is that the very low estimated discount factors suggest that a different model may be needed to describe the consumption behavior of poor households in developing countries. The co-movement between treatment effects in consumption and transfers observed in the data is, for example, consistent with sophisticated hyperbolic preferences, which induce households to keep most of their savings in illiquid assets to cope with their self-control problems (Laibson 1997). Time-inconsistent preferences imply that it might be desirable to adopt policy interventions that induce higher savings. Laibson (1996) shows that there would be a Pareto improvement if households could commit to a savings rate. Ashraf, Karlan and Yin (2006) provide evidence of demand for savings commitment devices.

This article estimated the discount factor of poor households in rural Mexico by combining experimental variation with a canonical buffer stock model. Future research could use a similar approach – using models together with experimental variation – to give us new insights into the consumption behavior of the poor.

8 Appendix A. Construction of Consumption Measure

I use the Nov98, May99 and Nov99 survey rounds which collected data on consumption. The module included questions on 36 food items and households reported for each of these items expenditures, quantity consumed and quantity purchased in the last 7 days. The respondent was given an option to report quantities in 3 units of measure: Kilos, Liter or Units. I

[39] This explanation is consistent with Paxson (1992) who finds that farmers in Thailand save a large fraction of their transitory income. Rosenzweig and Wolpin (1993) note that most of the assets held by farmers are inputs to production but also serve as buffer against income shocks. They propose a model in which investment assets are used both for production and to smooth consumption. Duncan, Rubacalva and Teruel (2008) suggest that rural Mexican households commonly save by owning livestock.

exclude data on consumption of alcohol, for which the data seems unreliable.

Household-specific prices were calculated by dividing expenditures by quantity purchased. I trimmed prices which were below the 1st percentile or above the 99th percentile. If a household reported consuming a give food item but the price was missing – either because of trimming, no quantity had been purchased or because the measure unit of purchase was different from the measure unit of consumption – then the household-specific price was set equal to the local price. Median prices at the village, municipality, state and national level were calculated for each measurement unit and the local price was set equal to the lowest level of aggregation with at least 20 household-specific price observations. The value of consumption for each item was calculated as the product of prices times quantity consumed, where prices and quantities were quoted in the same unit of measure. Household food consumption was calculated as the total value of consumption over the 35 items and converted to a semi-annual figure.

Households were also asked to report their expenditures with 24 non-food items. I exclude data on expenditures with weddings, funerals and ceremonies. The length of the reporting period was different for different items: transportation costs (last 7 days); personal and house hygiene, drugs and prescriptions, doctor visits, heating (i.e., wood, gas, oil), electricity (last month); clothing and shoes, school related expenses (last 6 months). I converted the expenditures to semi-annual figures and summed over all expenditures to calculate consumption of other goods.

Household total consumption is calculated as the sum of food consumption and consumption of other goods. Finally, I trimmed the total consumption distribution in each survey round in 1%.

9 Appendix B. Policy Function

I first solve the consumption problem for control households by using policy function iteration. I create a grid for cash on hand x with 250 points. For each value of cash on hand on the grid I find the level of consumption c which maximizes lifetime utility. The expected value of the value function one period ahead is calculated by discretizing the shock and using a three-node Gaussian quadrature scheme. I use cubic spline interpolation to approximate the policy function for levels of cash on hand between grid points.

The second step is to solve the consumption problem for treatment households using backward recursion. Again, I create a grid for cash on hand identical to the one created for control households. For treatment households, there is heterogeneity in the amount of transfers households receive. Thus, I proceed to create a 7 by 7 grid corresponding to different levels of summer semesters and winter semesters transfers – b_1 and b_2. Starting at period $T-1$, for each value of cash on hand on the x grid and for each value of transfers on the (b_1, b_2) grid I find the optimal consumption level – where period T is the last period in which households expect to receive transfers. The terminal value function is given by the value function of control households. The optimal consumption for periods before $T-1$ are found recursively. I use cubic spline interpolation to approximate the policy function.

10 Appendix C. Simulation

In the simulation exercise, each household i is characterized by its treatment status D_i, its history of transfers $\{B_{i,t}\}_{t=1}^{5}$ and its household size $hhsize_i$ – all of which come from household-level data. I generate for each household $H = 50$ fictitious histories of income shocks $\nu_{i,t}^{h}$ and $H = 50$ fictitious log of permanent (per capita) income $\ln \mu_i^h$. Notice that both of these variables are assumed to be independent of household size, treatment status or transfers. Using the first 100 periods of each history and the control policy function, I simulate for each household-simulation draw (i, h) its initial asset level – i.e., the household's financial wealth at the beginning of the first semester of the program: $a_{i,t=1}^{h}$.

In the post-program period, control and treatment households have different policy functions. Consumption for control households (normalized by permanent income) is simulated as:

$$c_{it}^h = c^0\left(a_{i,t}^h + y_{i,t}^h|\Omega\right).$$

Consumption for treatment households (normalized by permanent income) is simulated as:

$$c_{it}^h = c^{1,\tau}\left(a_{i,t}^h + y_{i,t}^h + b_{i,t}^h, b_{i,t+1}^h, b_{i,t+2}^h|\Omega\right),$$

where:

$$b_{i,t}^h = \frac{B_{i,t}}{\omega_i^h}, \omega_i^h = hhsize_i * \mu_i^h.$$

The household size and the history of transfers $\{B_{i,t}\}_{t=1}^5$ are kept constant for a given household i across simulation draws h. I then average log consumption across simulation draws h and households i for each group. The treatment effect is the difference between the simulated average log consumption of treatment and control households.

References

Andersen, Steffen, Glenn W. Harrison, Morten I. Lau, and E. Elisabet Rutstrom, "Eliciting Risk and Time Preferences," *Econometrica*, 2008, *76* (3), p583 – 618.

Angelucci, Manuela and Giacomo De Giorgi, "Indirect Effects of an Aid Program: The Case of Progresa and Consumption," *American Economic Review*, forthcoming.

Ashraf, Nava, Dean Karlan, and Wesley Yin, "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines," *Quarterly Journal of Economics*, 20060501, *121* (2), p635 – 672.

Atkeson, Andrew and Masao Ogaki, "Wealth-Varying Intertemporal Elasticities of Substitution: Evidence from Panel and Aggregate Data," *Journal of Monetary Economics*, *38* (3), p507 – 534.

Attanasio, Orazio, Costas Meghir, and Ana Santiago, "Education choices in Mexico: using a structural model and a randomised experiment to evaluate Progresa," 2004.

Attanasio, Orazio P. and Miguel Szekely, "Household Saving in Developing Countries - Inequality, Demographics and All That: How Different are Latin America and South East Asia?," RES Working Papers 4221, Inter-American Development Bank, Research Department Jul 2000.

Becker, Gary S. and Casey B. Mulligan, "The Endogenous Determination of Time Preference," *Quarterly Journal of Economics*, 1997, *112* (3), p729 – 758.

Behrman, Jere R., Nancy Birdsall, and Miguel Szekely, "Intergenerational Schooling Mobility and Macro Conditions and Schooling Policies in Latin America," RES Working Papers 4144, Inter-American Development Bank, Research Department Sep 1998.

Bertrand, Marianne, Sendhil Mullainathan, and Eldar Shafir, "A Behavioral-Economics View of Poverty," *American Economic Review*, May 2004, *94* (2), 419–423.

Bound, John and Alan B. Krueger, "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?," *Journal of Labor Economics*, 1991, *9* (1), p1 – 24.

Cagetti, Marco, "Wealth Accumulation over the Life Cycle and Precautionary Savings," *Journal of Business and Economic Statistics*, 2003, *21* (3), p339 – 353.

Carroll, Christopher D., "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, 1997, *112* (1), 1–55.

_ **and Andrew A. Samwick**, "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 1997, *40* (1), p41 – 71.

Coller, Maribeth and Melonie B. Williams, "Eliciting Individual Discount Rates," *Experimental Economics*, 1999, *2* (2), p107 – 127.

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica*, 1991a, *59* (5), 1221–1248.

_ , *Understanding consumption*, Clarendon Lectures in Economics., 1991b.

Duffie, Darrell and Kenneth J. Singleton, "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, 1993, *61* (4), p929 – 952.

Duflo, Esther, "Poor but Rational?," *Understanding Poverty*, 2006.

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes, "Do the Rich Save More?," *Journal of Political Economy*, 2004, *112* (2), p397 – 444.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue, "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 2002, *40* (2), 351–401.

Gertler, Paul, Sebastian Martinez, and Marta Rubio-Codina, "Investing cash transfers to raise long term living standards," 20060101.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption over the Life Cycle," *Econometrica*, 2002, *70* (1), 47–89.

Harrison, Glenn W., Morten I. Lau, and Melonie B. Williams, "Estimating Individual Discount Rates in Denmark: A Field Experiment," *American Economic Review*, 2002, *92* (5), p1606 – 1617.

Hausman, Jerry A., "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables," *The Bell Journal of Economics*, 1979, *10* (1), 33–54.

Hoddinott, John, Emmanuel Skoufias, and Ryan Washburn, "The Impact of PROGRESA on Consumption," *International Food Policy Research Institute Working Paper*, 2000.

Laibson, David, "Hyperbolic Discount Functions, Undersaving, and Savings Policy," *NBER Working Paper 5635*, 1996.

_ , "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 19970501, *112* (2), p443 – 477.

_ **, Andrea Repetto, and Jeremy Tobacman**, "Estimating Discount Functions with Consumption Choices over the Lifecycle," *NBER Working Paper 13314*, 2007.

Lang, Kevin and Paul A. Ruud, "Returns to Schooling, Implicit Discount Rates and Black-White Wage Differentials," *Review of Economics and Statistics*, 1986, *68* (1), p41 – 47.

Lawrance, Emily C., "Poverty and the Rate of Time Preference: Evidence from Panel Data," *The Journal of Political Economy*, 1991, *99* (1), 54–77.

Levy, Santiago, *Progress against Poverty: Sustaining Mexico's Progresa-Oportunidades Program*, Washington, D.C.:, 2006.

McFadden, Daniel, "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration," *Econometrica*, 1989, *57* (5), p995 – 1026.

Melvin, Jr. Stephens, "'3rd of tha Month': Do Social Security Recipients Smooth Consumption between Checks?," *American Economic Review*, 20030301, *93* (1), p406 – 422.

_ **and Erin L. Krupka**, "Subjective Discount Rates and Household Behavior," *Carnegie Mellon University*, 2006.

Mullainathan, Sendhil, Francois Bourguignon, and Boris Pleskovic, *Development Economics through the Lens of Psychology*, Harvard U and Poverty Action Lab, Cambridge, MA: Washington, D.C.:,

Ogaki, Masao and Andrew Atkeson, "Rate of Time Preference, Intertemporal Elasticity of Substitution, and Level of Wealth," *Review of Economics and Statistics*, 1997, *79* (4), p564 – 572.

Pakes, Ariel and David Pollard, "Simulation and the Asymptotics of Optimization Estimators," *Econometrica*, 1989, *57* (5), p1027 – 1057.

Parker, Susan W. and Graciela M. Teruel, "Randomization and Social Program Evaluation: The Case of Progresa," *Annals of the American Academy of Political and Social Science*, 2005, *599*, p199 – 219.

Paxson, Christina H., "Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand," *American Economic Review*, 1992, *82* (1), p15 – 33.

Rabault, Guillaume, "When Do Borrowing Constraints Bind? Some New Results on the Income Fluctuation Problem," *Journal of Economic Dynamics and Control*, 2002, *26* (2), p217 – 245.

Rosenzweig, Mark R., "Why Are There Returns to Schooling?," *American Economic Review*, 1995, *85* (2), p153 – 158.

_ **and Kenneth I. Wolpin**, "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investment in Bullocks in India," *Journal of Political Economy*, 1993, *101* (2), p223 – 244.

Rubalcava, Luis, Graciela Teruel, and Duncan Thomas, "Investment, time preferences and public transfers paid to women," *Economic Development and Cultural Change*, forthcoming.

Schechtman, Jack, "An Income Fluctuation Problem," *Journal of Economic Theory*, 1976, *12* (2), p218 – 241.

Schultz, T. Paul, "School subsidies for the poor: evaluating the Mexican Progresa poverty program," *Journal of Development Economics*, June 2004, *74* (1), 199–250.

Schultz, T. Paul, "Fertility and Income," Working Papers 925, Economic Growth Center, Yale University Oct 2005.

Shapiro, Jesse M., "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle," *Journal of Public Economics*, 2005, *89* (2-3), p303 – 325.

Skoufias, Emmanuel, "PROGRESA and its impacts on the welfare of rural households in Mexico," *International Food Policy Research Institute Working Paper*, 2005.

_ **and Susan W. Parker**, "Conditional cash transfers and their impact on child work and schooling," *International Food Policy Research Institute Working Paper*, 2001.

_ **, Benjamin Davis, and Jere R. Behrman**, "An Evaluation of the Selection of Beneficiary Households in the Education, Health, and Nutrition Program (PROGRESA) of Mexico," *International Food Policy Research Institute Working Paper*, 1999.

Tanaka, Tomomi, Colin Camerer, and Quang Nguyen, "Poverty, politics, and preferences: Field Experiments and survey data from Vietnam," 2006.

Todd, Petra E. and Kenneth I. Wolpin, "Assessing the Impact of a School Subsidy

Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility," *American Economic Review*, 2006, *96* (5), p1384 – 1417.

van de Walle, Dominique, "Are Returns to Investment Lower for the Poor? Human and Physical Capital Interactions in Rural Vietnam," *Review of Development Economics*, November 2003, *7* (4), 636–653.

Viscusi, W. Kip and Michael J. Moore, "Rates of Time Preference and Valuations of the Duration of Life," *Journal of Public Economics*, 1989, *38* (3), p297 – 317.

Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 1989, *104* (2), p275 – 298.

Table 1. TRANSFERS

Panel A. Average Transfers

	Unadjusted		Adjusted	
	Control	Treatment	Control	Treatment
May98-Oct98	-	386	-	392
Nov98-Apr99	-	1,056	-	970
May99-Oct99	-	1,305	-	1,147
Nov99-Apr00	716	1,495	592	1,244
May00-Oct00	1,415	1,682	1,140	1,355
Nov00-Apr01	1,555	1,586	1,202	1,226
May01-Oct01	1,820	1,870	1,381	1,419

Panel B. Average Positive Transfers

	Unadjusted		Adjusted	
	Control	Treatment	Control	Treatment
May98-Oct98	-	428	-	434
Nov98-Apr99	-	1,147	-	1,054
May99-Oct99	-	1,436	-	1,261
Nov99-Apr00	718	1,605	593	1,335
May00-Oct00	1,532	1,834	1,234	1,478
Nov00-Apr01	1,719	1,793	1,329	1,386
May01-Oct01	2,042	2,159	1,550	1,638

Note: The Unadjusted column shows average nominal transfers per household in pesos. The Adjusted column shows average real transfers per household deflated bimonthly to Sep/Oct 98 prices using national CPI. Panel A shows average transfers for the entire sample. Panel B shows average transfers for households that received transfers.

Table 2. PARAMETERS of INCOME PROCESS

	Income Process				
Variable	i.i.d.	AR(1) η = .25	AR(1) η = .5	MA(1) η = .25	MA(1) η = .5
Std. Deviation of Transitory Shocks	0.4594	0.5086	0.5509	0.5047	0.5253
Std. Deviation of Permanent Component	0.3494	0.2822	0	0.2880	0.1929

Note: The table shows the estimated parameters of the income process for the cases in which transitory shocks are either i.i.d., autoregressive of order 1 or follow a moving average of order 1. For AR(1) and MA(1), the parameters are estimated conditional on the value of η, the parameter that governs the serial correlation.

Table 3. TREATMENT EFFECTS

	Log Consumption		
Variable	May-Oct98	Nov98-Apr99	May-Oct99
Treatment Effect	0.085	0.144	0.164
	(0.011)***	(0.011)***	(0.010)***
Control	8.565	8.532	8.584
	(0.009)***	(0.009)***	(0.008)***

Variable	May-Oct98	Nov98-Apr99	May-Oct99
Mean Consumption of Controls	6,024	5,870	6,059
Mean Transfer	386	1,056	1,305
Mean Transfer / Mean. Consumption of Controls	0.06	0.18	0.22

Note: The first panel reports treatment effects on log household consumption. Each column presents results from a separate regression by semester. The number in parentheses are robust standard errors. The sample, which includes 9,397 observations, is restricted to eligible households surveyed at the baseline survey with nonmissing consumption data in all three survey rounds. The second panel reports average PROGRESA transfers among treatment households, average consumption of control households and the ratio of average transfers to average consumption. All figures are nominal values in pesos. *** Significant at 1% level

Table 4. MAIN RESULTS

		Annual Interest Rate			
Risk Aversion		-15%	0%	5%	10%
1	Semi-annual Disc. Factor	0.84	0.77	0.75	0.74
	Standard Error	0.08	0.07	0.07	0.07
	p-value Ho: β = √.9	0.0787	0.0073	0.0028	0.0010
	p-value O.R. test	0.9305	0.9224	0.9203	0.9180
2	Semi-annual Disc. Factor	0.56	0.52	0.50	0.49
	Standard Error	0.13	0.12	0.12	0.11
	p-value Ho: β = √.9	0.0015	0.0002	0.0001	0.0000
	p-value O.R. test	0.9314	0.9232	0.9211	0.9188
3	Semi-annual Disc. Factor	0.32	0.30	0.29	0.28
	Standard Error	0.11	0.10	0.10	0.10
	p-value Ho: β = √.9	0.0000	0.0000	0.0000	0.0000
	p-value O.R. test	0.9372	0.9287	0.9265	0.9244

Note: The table presents estimates of the semi-annual discount factor for different combinations of annual real interest rate and risk aversion parameter. The mean of permanent income is estimated jointly with the discount factor and the estimates range from 7.29 to 7.33. The exercise assumes that families perfectly foresee transfers for the next year. The sample containing transfers and family size is the balanced sample of 9,397 eligible households and I draw 50 simulation draws for each household. The income process is assumed to be i.i.d. with the standard deviation of shocks being equal to .46 and the standard deviation of log permanent income equal to .35. The first p-value reported refers to the one-sided test that the semi-annual discount factor is √.9. The second p-value refers to the overidentifying restrictions test.

Table 5. INCOME PROCESS

	Semi-annual Discount Factor	Standard Error	p-value Ho: $\beta = \sqrt{.9}$
(0) **Benchmark:** i.i.d. shocks, $r = 5\%$, $\rho = 2$, $\sigma_v = .46$, $\sigma_{\ln\mu} = .35$ and $\eta = 0$	0.50	0.12	0.0001
(1) **Std. Deviation of Transitory Shocks** $\sigma_v = .2$	0.61	0.07	0.0000
(2) **Std. Deviation of Permanent Income** $\sigma_{\ln\mu} = .15$	0.50	0.12	0.0001
(3) **Income Differences** Assume aver. income treatment 5% lower than aver. income control	0.66	0.06	0.0000
(4) **Macro Shocks** Assume aver. income 5% lower in 2nd and 3rd semesters	0.64	0.03	0.0000
(5) **AR(1)** AR(1) shocks, $\sigma_v = .51$, $\sigma_{\ln\mu} = .28$ and $\eta = .25$	0.48	0.10	0.0000
(6) **AR(1)** AR(1) shocks, $\sigma_v = .46$, $\sigma_{\ln\mu} = 0$ and $\eta = .5$	0.47	0.09	0.0000
(7) **MA(1)** MA(1) shocks, $\sigma_v = .50$, $\sigma_{\ln\mu} = .29$ and $\eta = .25$	0.47	0.09	0.0000
(8) **MA(1)** MA(1) shocks, $\sigma_v = .52$, $\sigma_{\ln\mu} = .19$ and $\eta = .5$	0.44	0.10	0.0000

Note: The table presents estimates of the semi-annual discount factor for an annual interest rate of 5% and risk aversion of 2.. The mean of permanent income is estimated jointly with the discount factor. The exercise assumes that families perfectly foresee transfers for the next year. The sample containing transfers and family size is the balanced sample of 9,397 eligible households and I draw 50 simulation draws for each household. The p-value reported refers to the one-sided test that the semi-annual discount factor is $\sqrt{.9}$.

Table 6. EXPECTATIONS

	Semi-annual Discount Factor	Standard Error	p-value Ho: $\beta = \sqrt{.9}$
(0) **Benchmark:** Agents perfectly forecast transfers for next year and expect to remain constant; Control and treatment do not anticipate the program; Treatment expects to receive transfers for 3 years	0.50	0.12	0.0001
(1) **Myopic** Agents expect future transfers to remain constant at current levels	0.83	0.05	0.0122
(2) **Pessimistic** Agents expect to receive no future transfers	0.45	0.11	0.0000
(3) **Anticipation Control** Control anticipates participation as treatment is enrolled	0.49	0.10	0.0000
(4) **Anticipation Treatment** Treatment anticipates participation one semester before receive transfers	0.79	0.05	0.0009
(5) **Duration** Treatment expects to receive transfers for 2 years	0.49	0.09	0.0000

Note: The table presents estimates of the semi-annual discount factor for an annual interest rate of 5% and risk aversion of 2.. The mean of permanent income is estimated jointly with the discount factor. The income process is assumed to be i.i.d with the standard deviation of shocks being equal to .46 and the standard deviation of the permanent component equal to .35. The sample containing transfers and family size is the balanced sample of 9,397 eligible households and I draw 50 simulation draws for each household. The p-value reported refers to the one-sided test that the semi-annual discount factor is $\sqrt{.9}$.

Figure 1: Control Policy Function by Discount Factor

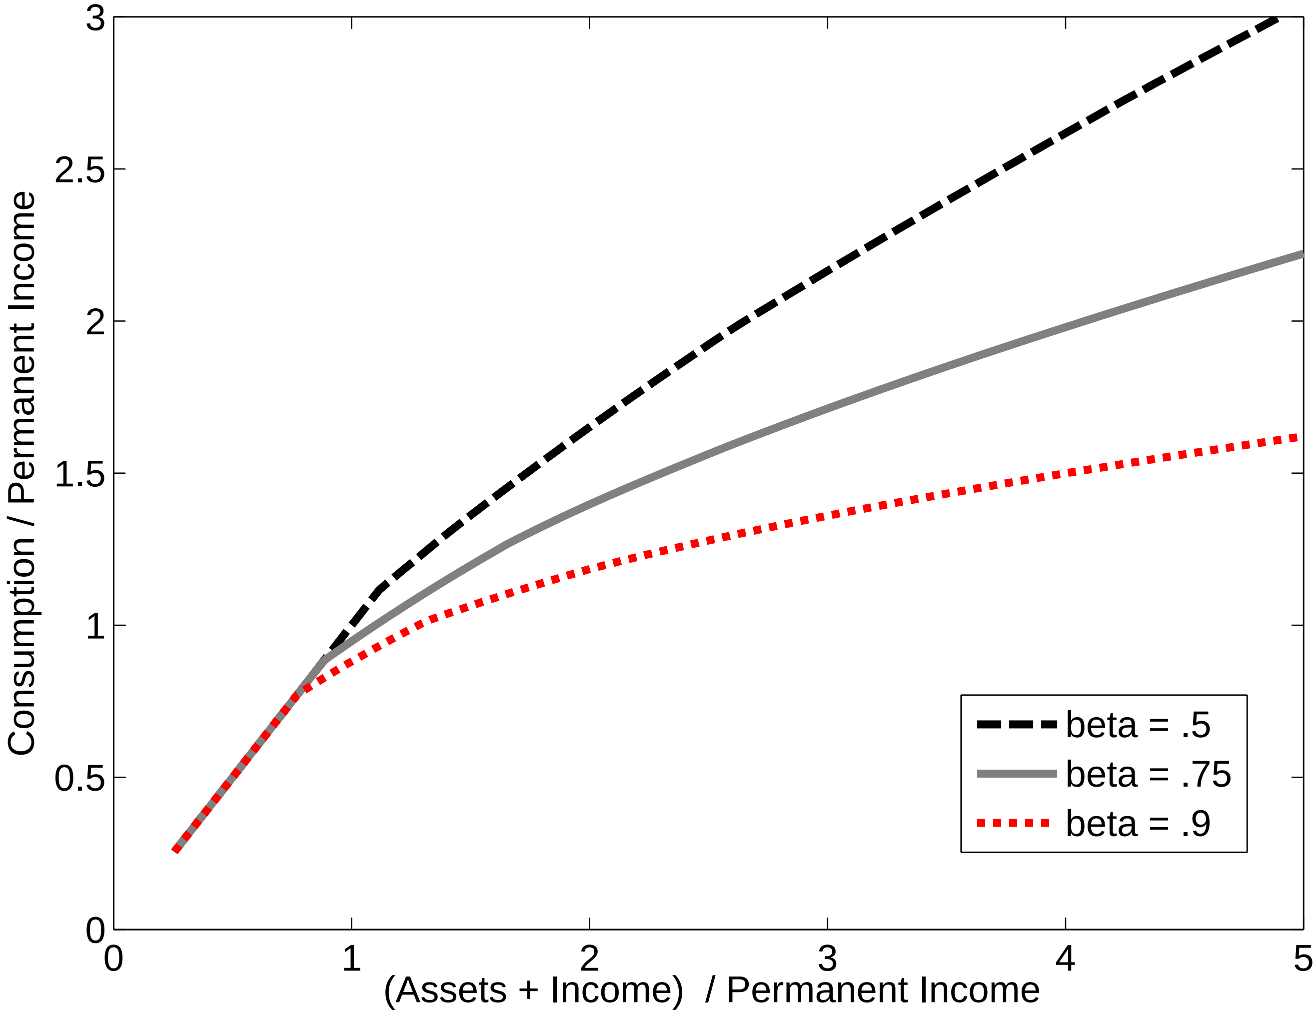

Note: The figure shows the control policy function for different semi-annual discount factors, assuming an annual interest rate of 5% and a risk aversion parameter of 2.

Figure 2: Simulated Initial Asset Distribution

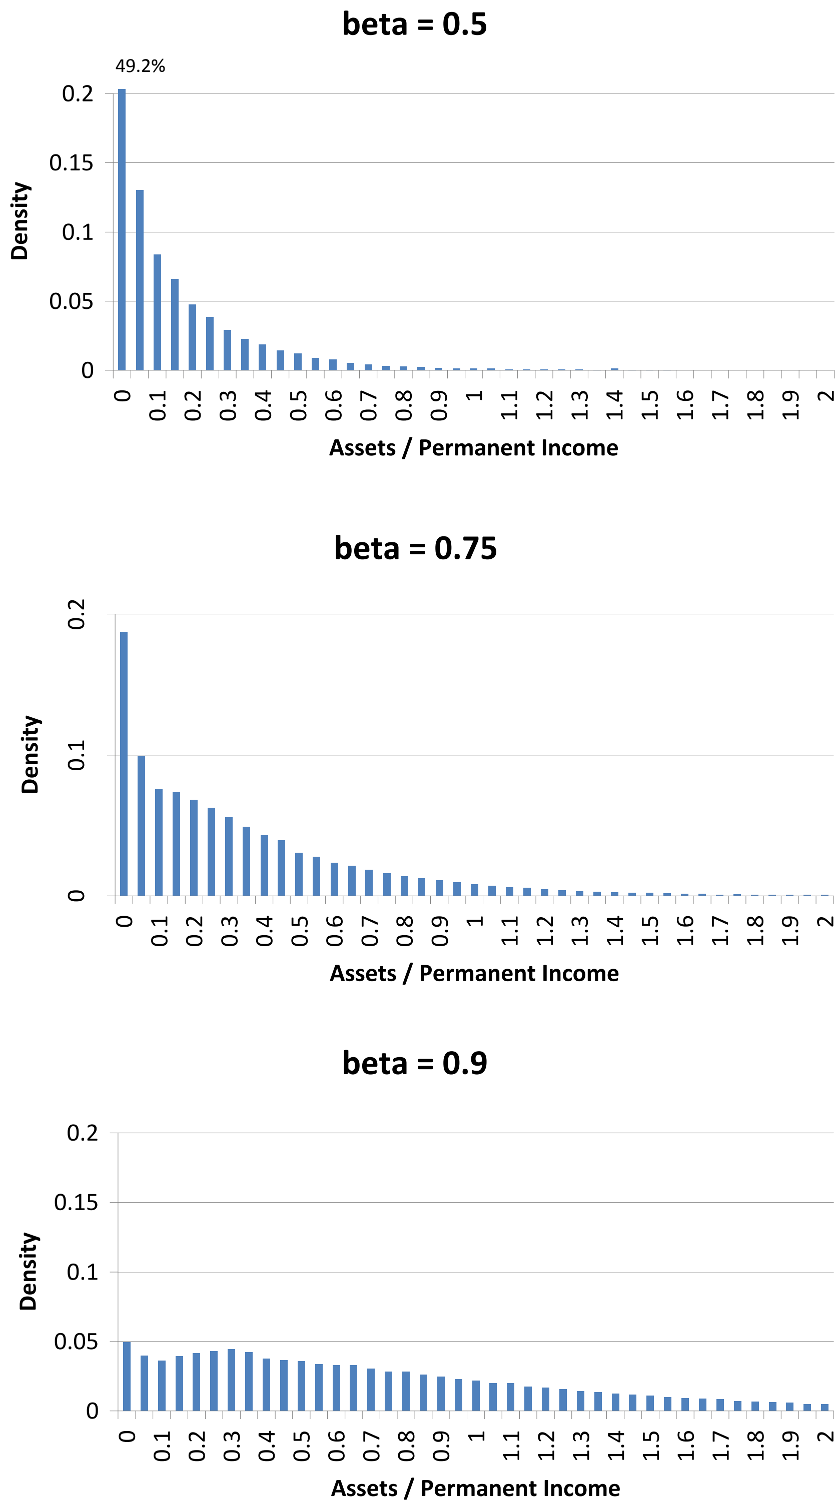

Note: The figure shows the simulated pre-program asset distribution for different semi-annual discount factors, assuming an annual interest rate of 5% and a risk aversion parameter of 2. The distribution is generated by running history for 100 periods using the control policy function, where households have zero assets in the first period.

Figure 3: Simulated Average Log Consumption by Discount Factor

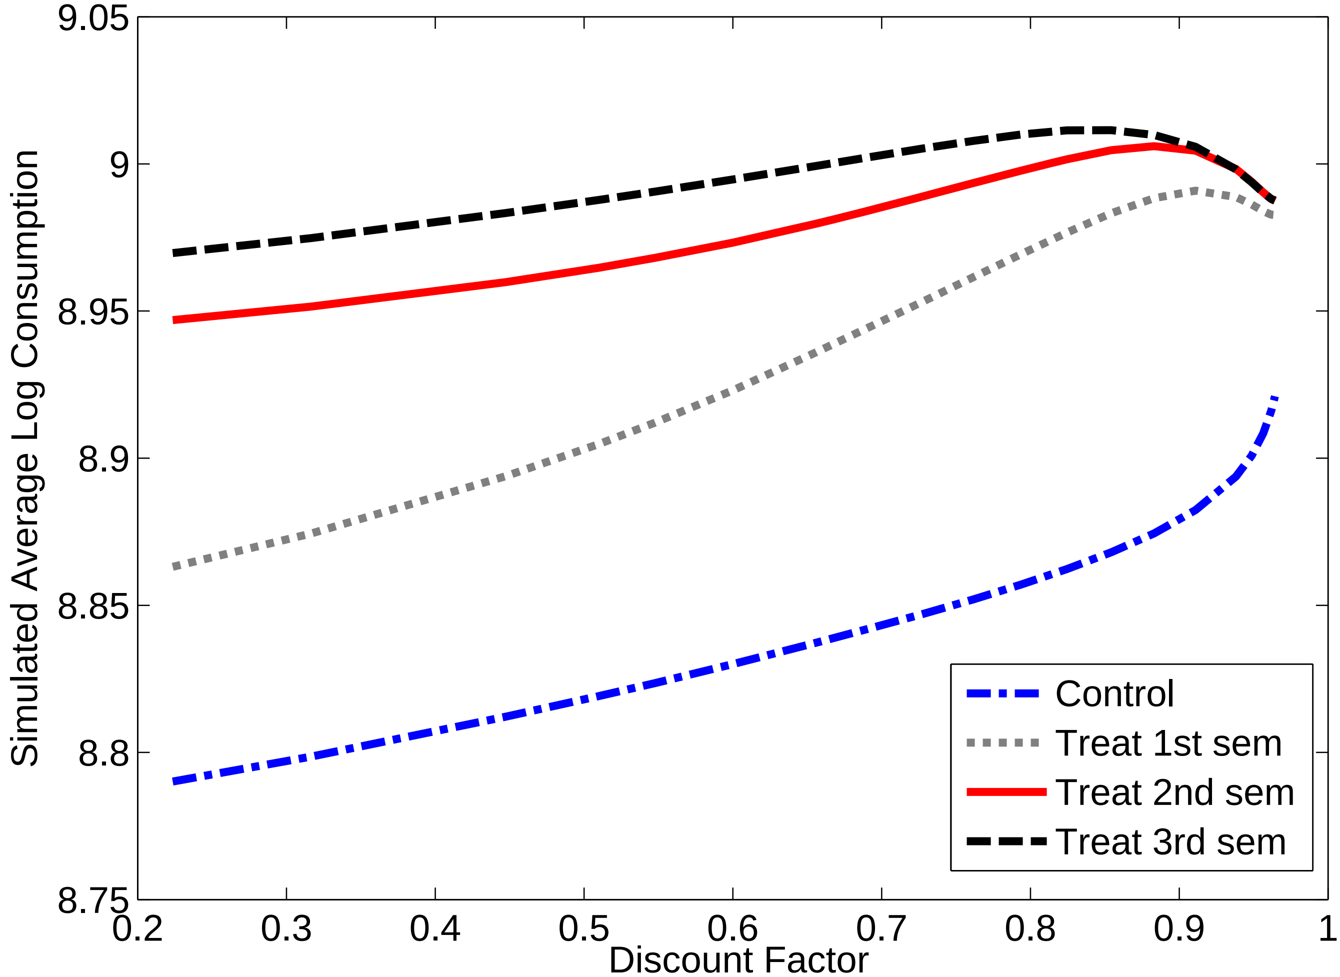

Note: The figure shows simulated average log consumption as a function of the semi-annual discount factor, assuming r = 5%, rho = 2 and lnmu = 7.3. For treatment households, there are 3 curves - one for each program semester. The simulation uses household-level data on PROGRESA transfers.

Figure 4: Simulated Treatment Effects by Discount Factor

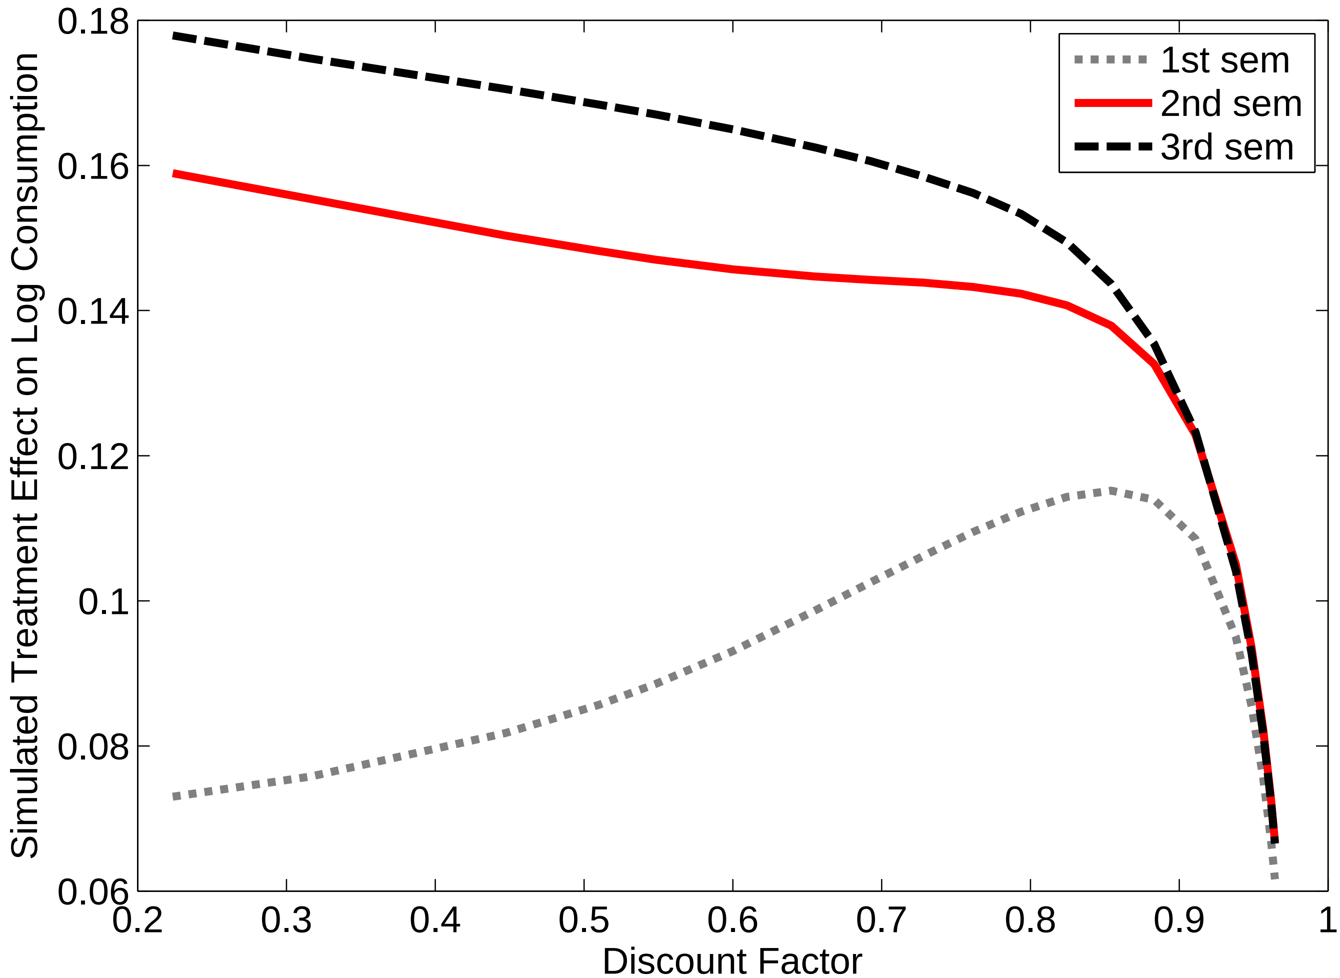

Note: The figure shows the simulated treatment effects on log consumption as a function of the semi-annual discount factor, assuming r = 5%, rho = 2 and lnmu = 7.3. The different curves correspond to different program semesters. The simulation uses household-level data on PROGRESA transfers.

Figure 5: Treatment Effects Over Time

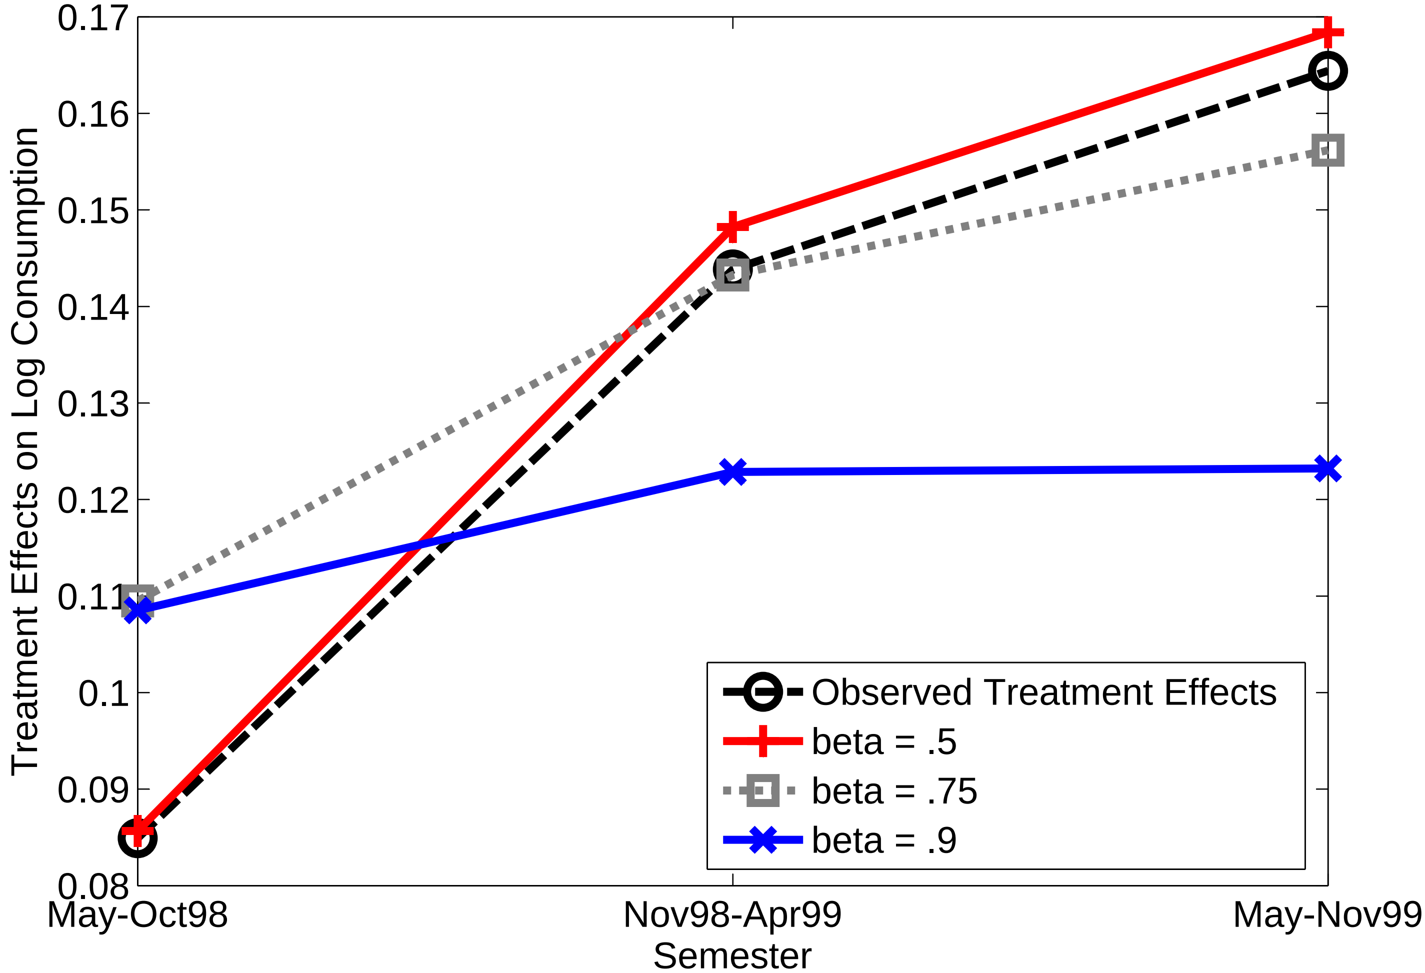

Note: The figure shows simulated and sample treatment effects on log consumption by program semester. The simulated treatment effects are shown for 3 semi-annual discount factors: 0.5, 0.75 and 0.9. The simulation assumes r=5%, rho=2, lnmu=7.3 and uses household-level data on PROGRESA transfers.

Table A1. Wage Income Moments

	Mean					
Variable	Nov-98	May-99	Nov-99	May-00	Nov-00	All periods
All	6.416	6.403	6.547	6.567	6.815	6.559
Control	6.422	6.379	6.539	6.582	6.834	6.561
Treatment	6.413	6.418	6.551	6.558	6.803	6.558
Treatment-Control	-0.009	0.039	0.012	-0.024	-0.031	-0.003
Std. Deviation	(0.02)	(0.019)**	(0.02)	(0.02)	(0.015)**	(0.01)

	Variance					
Variable	Nov-98	May-99	Nov-99	May-00	Nov-00	All periods
All	0.605	0.632	0.553	0.573	0.532	0.599
Control	0.606	0.636	0.545	0.578	0.532	0.603
Treatment	0.604	0.629	0.558	0.571	0.532	0.597
Treatment-Control	-0.002	-0.007	0.013	-0.007	0.000	-0.006
Std. Deviation	(0.019)	(0.021)	(0.017)	(0.018)	(0.017)	(0.009)

	Autocovariance of order 1					
Variable	Nov-98	May-99	Nov-99	May-00	Nov-00	All periods
All	0.208	0.198	0.186	0.200	-	0.206
Control	0.226	0.203	0.186	0.188	-	0.210
Treatment	0.197	0.194	0.186	0.207	-	0.203
Treatment-Control	-0.029	-0.009	0.000	0.019	-	-0.007
Std. Deviation	(0.017)*	(0.016)	(0.014)	(0.014)	-	(0.008)

Note: The table reports by survey round the mean, variance and autocovariance of order 1 for log household per capita wage income. The sample only includes eligible households surveyed at the baseline survey. The moments are computed for the entire sample and for control and treatment households separately. The fourth row in each panel reports the treatment-control difference and and the fifth row reports the standard error of the estimate.

Table A2. Total Income Moments

Variable	Mean		
	Nov-98	May-99	2 periods
All	6.686	6.445	6.574
Control	6.695	6.476	6.592
Treatment	6.680	6.425	6.562
Treatment-Control	-0.015	-0.051	-0.030
Std. Deviation	(0.02)	(0.022)**	(0.014)**

Variable	Variance		
	Nov-98	May-99	2 periods
All	0.738	0.992	0.871
Control	0.723	0.924	0.830
Treatment	0.747	1.033	0.896
Treatment-Control	0.024	0.109	0.066
Std. Deviation	(0.025)	(0.034)***	(0.022)***

Variable	Autocovariance of order 1		
	Nov-98	May-99	2 periods
All	0.184	0.171	0.177
Control	0.179	0.167	0.173
Treatment	0.187	0.174	0.179
Treatment-Control	0.008	0.007	0.006
Std. Deviation	(0.021)	(0.023)	(0.016)

Note: The table reports by survey round the mean, variance and autocovariance of order 1 for log household per capita total income. The sample only includes eligible households surveyed at the baseline survey. The moments are computed for the entire sample and for control and treatment households separately. The fourth row in each panel reports the treatment-control difference and and the fifth row reports the standard error of the estimate.

Table A3. INCOME PROCESS 1

Risk Aversion		Annual Interest Rate: 5% (1) Benchmark	(2) variance of transitory shocks	(3) variance of permanent income	(4) Treat. Income 5% lower than Control Income	(5) Macro Shocks
1	Semi-annual Disc. Factor	0.75	0.83	0.75	0.84	0.84
	Standard Error	0.07	0.03	0.07	0.03	0.02
	p-value Ho: $\beta = \sqrt{.9}$	0.0028	0.0003	0.0026	0.0001	0.0000
	p-value O.R. test	0.9203	0.8795	0.9161	0.8282	0.9748
2	Semi-annual Disc. Factor	0.50	0.61	0.50	0.66	0.64
	Standard Error	0.12	0.07	0.12	0.06	0.03
	p-value Ho: $\beta = \sqrt{.9}$	0.0001	0.0000	0.0001	0.0000	0.0000
	p-value O.R. test	0.9211	0.8754	0.9165	0.7515	0.8514
3	Semi-annual Disc. Factor	0.29	0.39	0.29	0.46	0.41
	Standard Error	0.10	0.07	0.10	0.08	0.03
	p-value Ho: $\beta = \sqrt{.9}$	0.0000	0.0000	0.0000	0.0000	0.0000
	p-value O.R. test	0.9265	0.8777	0.9215	0.702	0.809

Note: The table presents estimates of the semi-annual discount factor for different risk aversion parameters and for an annual interest rate of 5%. The first column shows the benchmark results. Column 2 shows results if the std. deviation of transitory shocks to income is equal to .2 and column 3 when the std. deviation of permanent income is assumed to be .15. The fourth column shows results when the average income of treatment households is 5% lower than of controls. Column 5 shows results assuming that there were unexpected aggregate shocks of -5% in the second and third semesters. The exercise assumes that families perfectly foresee transfers for the next year. The sample containing transfers and family size is the balanced sample of 9,397 eligible households and I draw 50 simulation draws for each household. The income process is assumed to be i.i.d. with the standard deviation of shocks being equal to .46 and the standard deviation of the permanent component equal to .35 unless mentioned otherwise. The first p-value reported refers to the one-sided test that the semi-annual discount factor is $\sqrt{.9}$. The second p-value refers to the overidentifying restrictions test.

Table A4. INCOME PROCESS 2

Risk Aversion		Annual Interest Rate: 5%				
		i.i.d.	AR(1) η =.25	AR(1) η =.5	MA(1) η =.25	MA(1) η =.5
1	Semi-annual Disc. Factor	0.75	0.74	0.75	0.74	0.72
	Standard Error	0.07	0.06	0.05	0.07	0.07
	p-value Ho: $\beta = \sqrt{.9}$	0.0028	0.0008	0.0000	0.0013	0.0003
	p-value O.R. test	0.9203	0.9236	0.7598	0.9072	0.7927
2	Semi-annual Disc. Factor	0.50	0.48	0.47	0.47	0.44
	Standard Error	0.12	0.10	0.09	0.09	0.10
	p-value Ho: $\beta = \sqrt{.9}$	0.0001	0.0000	0.0000	0.0000	0.0000
	p-value O.R. test	0.9211	0.9393	0.7879	0.9236	0.8139
3	Semi-annual Disc. Factor	0.29	0.25	0.24	0.25	0.22
	Standard Error	0.10	0.08	0.06	0.08	0.08
	p-value Ho: $\beta = \sqrt{.9}$	0.0000	0.0000	0.0000	0.0000	0.0000
	p-value O.R. test	0.9265	0.9502	0.8121	0.9357	0.8329

Note: The table presents estimates of the semi-annual discount factor for different risk aversion parameters and for an annual interest rate of 5%. The different columns show results for different income processes for transitory shocks to log income. For the AR(1) and MA(1) cases, I show estimates conditional on η - the parameter which governs the serial correlation - being equal to .25 or .5. The mean of permanent income is estimated jointly with the discount factor. The estimates of the mean of permanent income range from 7.30 to 7.31 for the i.i.d. case, from 7.1 to 7.28 for the AR(1) case and from 7.22 to 7.29 for the MA(1) case. The exercise assumes that families perfectly foresee transfers for the next year. The sample containing transfers and family size is the balanced sample of 9,397 eligible households and I draw 50 simulation draws for each household. The first p-value reported refers to the one-sided test that the semi-annual discount factor is $\sqrt{.9}$. The second p-value refers to the overidentifying restrictions test.

Table A5. EXPECTATIONS

Risk Aversion		Annual Interest Rate: 5% (1) Benchmark	(2) Myopic	(3) Pessimistic	(4) Anticipation Control	(5) Anticipation Treatment	(6) Duration 2 years
1	Semi-annual Disc. Factor	0.75	0.90	0.71	0.74	0.89	0.73
	Standard Error	0.07	0.03	0.08	0.07	0.03	0.07
	p-value Ho: $\beta = \sqrt{.9}$	0.0028	0.0645	0.0010	0.0019	0.0148	0.0011
	p-value O.R. test	0.9203	0.1589	0.0197	0.9078	0.8436	0.7731
2	Semi-annual Disc. Factor	0.50	0.83	0.45	0.49	0.79	0.49
	Standard Error	0.12	0.05	0.11	0.10	0.05	0.09
	p-value Ho: $\beta = \sqrt{.9}$	0.0001	0.0122	0.0000	0.0000	0.0009	0.0000
	p-value O.R. test	0.9211	0.1641	0.0188	0.9305	0.8038	0.8245
3	Semi-annual Disc. Factor	0.29	0.76	0.25	0.28	0.67	0.28
	Standard Error	0.10	0.07	0.11	0.10	0.07	0.09
	p-value Ho: $\beta = \sqrt{.9}$	0.0000	0.0026	0.0000	0.0000	0.0001	0.0000
	p-value O.R. test	0.9265	0.1656	0.0182	0.9609	0.7589	0.8881

Note: The table presents estimates of the semi-annual discount factor for different risk aversion parameters and for an annual interest rate of 5%. The assumptions made in each column are: 1) households perfectly forecast transfers one year ahead and expect them to remain constant, 2) households expect transfers to remain constants at the current level, 3) households expect to receive no transfers in the future, 4) control anticipates participation as treatment starts receiving transfers, 5) treatment anticipates program one semester before receiving first transfers and 6) treatment expects to receive transfers for 2 years. The mean of permanent income is estimated jointly with the discount factor. The sample containing transfers and family size is the balanced sample of 9,397 eligible households and I draw 50 simulation draws for each household. The income process is assumed to be i.i.d. with the standard deviation of shocks being equal to .46 and the standard deviation of the permanent component equal to .35. The first p-value reported refers to the one-sided test that the semi-annual discount factor is $\sqrt{.9}$. The second p-value refers to the overidentifying restrictions test.

Table A6. INCOME PROCESS 1

Risk Aversion		Annual Interest Rate: -15% (1) Benchmark	(2) variance of transitory shocks	(3) variance of permanent income	(4) Treat. Income 5% lower than Control Income	(5) Macro Shocks
1	Semi-annual Disc. Factor	0.84	0.92	0.84	0.94	0.94
	Standard Error	0.08	0.04	0.08	0.03	0.02
	p-value Ho: $\beta = \sqrt{.9}$	0.0787	0.2204	0.0768	0.3375	0.2425
	p-value O.R. test	0.9305	0.8997	0.9261	0.8495	0.9881
2	Semi-annual Disc. Factor	0.56	0.68	0.56	0.73	0.71
	Standard Error	0.13	0.07	0.13	0.06	0.03
	p-value Ho: $\beta = \sqrt{.9}$	0.0015	0.0001	0.0013	0.0002	0.0000
	p-value O.R. test	0.9314	0.8927	0.927	0.763	0.8853
3	Semi-annual Disc. Factor	0.32	0.43	0.32	0.52	0.46
	Standard Error	0.11	0.07	0.11	0.08	0.03
	p-value Ho: $\beta = \sqrt{.9}$	0.0000	0.0000	0.0000	0.0000	0.0000
	p-value O.R. test	0.9372	0.8933	0.9319	0.7108	0.8446

Note: The table presents estimates of the semi-annual discount factor for different risk aversion parameters and for an annual interest rate of -15%. The first column shows the benchmark results. Column 2 shows results if the std. deviation of transitory shocks to income is equal to .2 and column 3 when the std. deviation of permanent income is assumed to be .15. The fourth column shows results when the average income of treatment households is 5% lower than of controls. Column 5 shows results assuming that there were unexpected aggregate shocks of -5% in the second and third semesters. The exercise assumes that families perfectly foresee transfers for the next year. The sample containing transfers and family size is the balanced sample of 9,397 eligible households and I draw 50 simulation draws for each household. The income process is assumed to be iid with the standard deviation of shocks being equal to .46 and the standard deviation of the permanent component equal to .35 unless mentioned otherwise. The first p-value reported refers to the one-sided test that the semi-annual discount factor is $\sqrt{.9}$. The second p-value refers to the overidentifying restrictions test.

Table A7. INCOME PROCESS 2

Risk Aversion		Annual Interest Rate: -15%				
		i.i.d.	AR(1) η =.25	AR(1) η =.5	MA(1) η =.25	MA(1) η =.5
1	Semi-annual Disc. Factor	0.84	0.83	0.83	0.82	0.80
	Standard Error	0.08	0.07	0.06	0.08	0.07
	p-value Ho: $\beta = \sqrt{.9}$	0.0787	0.0437	0.0167	0.0493	0.0232
	p-value O.R. test	0.9305	0.9192	0.7546	0.9045	0.7919
2	Semi-annual Disc. Factor	0.56	0.53	0.52	0.52	0.48
	Standard Error	0.13	0.11	0.10	0.10	0.10
	p-value Ho: $\beta = \sqrt{.9}$	0.0015	0.0001	0.0000	0.0000	0.0000
	p-value O.R. test	0.9314	0.9348	0.7862	0.9199	0.8143
3	Semi-annual Disc. Factor	0.32	0.28	0.26	0.27	0.24
	Standard Error	0.11	0.09	0.07	0.09	0.08
	p-value Ho: $\beta = \sqrt{.9}$	0.0000	0.0000	0.0000	0.0000	0.0000
	p-value O.R. test	0.9372	0.9442	0.8098	0.931	0.832

Note: The table presents estimates of the semi-annual discount factor for different risk aversion parameters and for an annual interest rate of -15%. The different columns show results for different income processes for transitory shocks to log income. For the AR(1) and MA(1) cases, I show estimates conditional on η - the parameter which governs the serial correlation - being equal to .25 or .5. The mean of permanent income is estimated jointly with the discount factor. The estimates of the mean of permanent income range from 7.30 to 7.31 for the i.i.d. case, from 7.1 to 7.28 for the AR(1) case and from 7.22 to 7.29 for the MA(1) case. The exercise assumes that families perfectly foresee transfers for the next year. The sample containing transfers and family size is the balanced sample of 9,397 eligible households and I draw 50 simulation draws for each household. The first p-value reported refers to the one-sided test that the semi-annual discount factor is $\sqrt{.9}$. The second p-value refers to the overidentifying restrictions test.

Table A8. EXPECTATIONS

Risk Aversion		Annual Interest Rate: - 15%					
		(1) Benchmark	(2) Myopic	(3) Pessimistic	(4) Anticipation Control	(5) Anticipation Treatment	(6) Duration 2 years
1	Semi-annual Disc. Factor	0.84	1.01	0.78	0.82	0.98	0.81
	Standard Error	0.08	0.02	0.08	0.08	0.03	0.08
	p-value Ho: $\beta = \sqrt{.9}$	0.0787	0.9918	0.0205	0.0498	0.8184	0.0324
	p-value O.R. test	0.9305	0.172	0.0258	0.8904	0.8789	0.7453
2	Semi-annual Disc. Factor	0.56	0.94	0.50	0.54	0.85	0.53
	Standard Error	0.13	0.04	0.12	0.11	0.06	0.11
	p-value Ho: $\beta = \sqrt{.9}$	0.0015	0.4321	0.0001	0.0001	0.0625	0.0001
	p-value O.R. test	0.9314	0.1772	0.0247	0.9131	0.8291	0.8043
3	Semi-annual Disc. Factor	0.32	0.88	0.27	0.31	0.71	0.30
	Standard Error	0.11	0.05	0.11	0.11	0.09	0.10
	p-value Ho: $\beta = \sqrt{.9}$	0.0000	0.0748	0.0000	0.0000	0.0039	0.0000
	p-value O.R. test	0.9372	0.1769	0.0238	0.9459	0.7762	0.8616

Note: The table presents estimates of the semi-annual discount factor for different risk aversion parameters and for an annual interest rate of -15%. The assumptions made in each column are: 1) households perfectly forecast transfers one year ahead and expect them to remain constant, 2) households expect transfers to remain constants at the current level, 3) households expect to receive no transfers in the future, 4) control anticipates participation as treatment starts receiving transfers, 5) treatment anticipates program one semester before receiving first transfers and 6) treatment expects to receive transfers for 2 years. The mean of permanent income is estimated jointly with the discount factor. The sample containing transfers and family size is the balanced sample of 9,397 eligible households and I draw 50 simulation draws for each household. The income process is assumed to be i.i.d. with the standard deviation of shocks being equal to .46 and the standard deviation of the permanent component equal to .35. The first p-value reported refers to the one-sided test that the semi-annual discount factor is $\sqrt{.9}$. The second p-value refers to the overidentifying restrictions test.

Table A9. TIMING of TRANSFERS

Risk Aversion		Annual Interest Rate: -15%			Annual Interest Rate: 5%		
		(1) Benchmark	(2) Participation+2	(3) Delayed	(4) Benchmark	(5) Participation+2	(6) Delayed
1	Semi-annual Disc. Factor	0.84	0.75	0.94	0.75	0.67	0.85
	Standard Error	0.08	0.14	0.03	0.08	0.14	0.03
	p-value Ho: $\beta = \sqrt{.9}$	0.0787	0.0732	0.3840	0.0028	0.0126	0.0001
	p-value O.R. test	0.9305	0.0941	0.2903	0.9203	0.0925	0.3082
2	Semi-annual Disc. Factor	0.56	0.43	0.75	0.50	0.38	0.68
	Standard Error	0.13	0.15	0.06	0.12	0.16	0.06
	p-value Ho: $\beta = \sqrt{.9}$	0.0015	0.0003	0.0007	0.0001	0.0002	0.0000
	p-value O.R. test	0.9314	0.081	0.3165	0.9211	0.0792	0.3269
3	Semi-annual Disc. Factor	0.32	0.20	0.54	0.29	0.18	0.50
	Standard Error	0.11	0.11	0.09	0.10	0.11	0.08
	p-value Ho: $\beta = \sqrt{.9}$	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
	p-value O.R. test	0.9372	0.0694	0.3383	0.9265	0.0677	0.3451

Note: The table presents estimates of the semi-annual discount factor for different risk aversion parameters. Columns 1 and 4 assume that there was a 4 months lag between participation and transfers payments. Columns 2 and 5 assume that there was a 2 months lag between participation and transfers payments. Columns 3 and 6 assume that transfers which should have been paid during Sep-Oct 98 were paid during Nov-Dec 98. The mean of permanent income is estimated jointly with the discount factor. The exercise assumes that families perfectly foresee transfers for the next year. The sample containing transfers and family size is the balanced sample of 9,397 eligible households and I draw 50 simulation draws for each household. The income process is assumed to be i.i.d. with the standard deviation of shocks being equal to .46 and the standard deviation of the permanent component equal to .35. The first p-value reported refers to the one-sided test that the semi-annual discount factor is $\sqrt{.9}$. The second p-value refers to the overidentifying restrictions test.

Figure A1: Objective Function

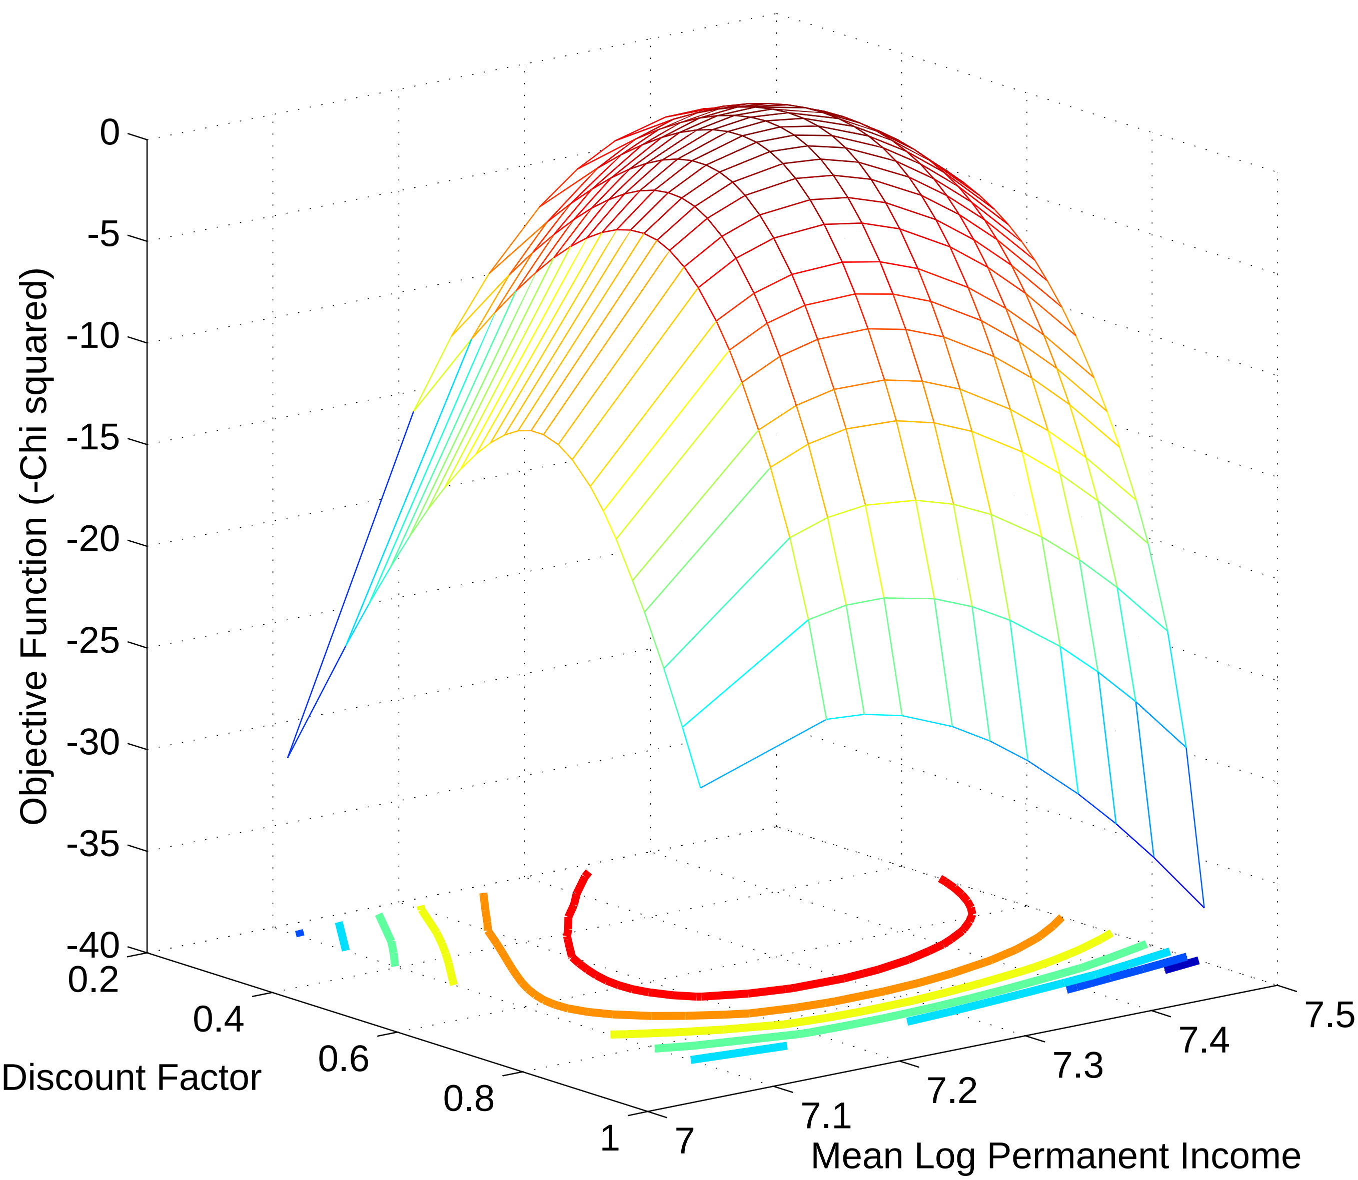

Note: The figure shows the inverse of the objective function, which is distributed as a chi-squared variable with one degree of freedom, in the beta - lnmu space. The objective function shown assumes r = 5% and rho = 2.