Heterogeneity in the Efficiency of Intrahousehold Resource Allocation: Empirical Evidence and Implications for Investment in Children*

Manuela Angelucci† and Robert Garlick‡

December 1, 2016

Abstract

We present some of the first evidence of within-sample variation in the efficiency of intrahousehold resource allocation. In a sample of rural, low-income Mexican households, observed consumption patterns are Pareto efficient for households with relatively old heads, but not households with relative young heads. This variation in efficiency has important welfare implications: younger households invest less in children's education and their education investments are less sensitive to the receipt of cash transfers. We believe this is the first empirical link between inefficient household resource allocation and lower investment in children's education, though this link follows naturally from standard models of investment in household public goods. Heterogeneity in efficiency by household head age may be due to cohort and/or lifecycle effects. The specific patterns of heterogeneity we find are more consistent with cohort than lifecycle effects, which means that average efficiency may decline over time in this population. From a policy perspective, these results highlight a limitation of cash transfers relative to other forms of poverty alleviation policy in settings where at least some households engage in inefficient resource allocation. Many papers have already established that the distributional effects of development policy depend on whether households behave as unitary actors; our results emphasize that it is also important to assess whether (some) households behave as efficient collective units.

*We are grateful for valuable comments from Erica Field, Andrew Foster, David Lam, Jesse Naidoo, Duncan Thomas, Alessandra Voena and seminar participants at Barcelona GSE Summer Forum, Duke University, Econometric Society World Congress, Northeast Universities Development Conference, Society of Labor Economists Conference, University of British Columbia, University of Connecticut, University of Michigan, University of South Carolina, and the World Bank. Monica Hernandez and Javier Romero Haaker provided valuable research assistance. All errors are our own.

†Assistant Professor of Economics, University of Michigan; mangeluc@umich.edu

‡Assistant Professor of Economics, Duke University; robert.garlick@duke.edu

1 Introduction

Intrahousehold resource allocation is an important and widely-studied topic in economics. An extensive literature studies how households accumulate resources and allocate these resources between members. Describing and modeling these decisions can provide important insights into households' use of productive resources (Udry, 1996), investment in children's human capital (Thomas, 1990), and labor supply decisions (Rangel, 2006). In particular, many papers test whether households can be accurately modeled as *unitary* actors with rational and consistent preferences and/or as *cooperative* units whose consumption, investment, and production decisions are Pareto efficient. Most empirical work rejects unitary models of household behavior but the evidence regarding cooperative models is inconclusive.[1]

We present a novel potential explanation for the range of results regarding the validity of cooperative models of household behavior. We show that in a sample of rural Mexican households, the cooperative model is strongly rejected for young households but not rejected for old households.[2] Our tests are based on empirical implications of static cooperative bargaining models that assume households always reach Pareto efficient consumption allocations. Our result demonstrates that younger households on average fail to achieve Pareto efficient consumption allocations and so "leave resources on the table." We believe that this is the first documentation of within-sample heterogeneity over observed demographic characteristics in the efficiency of household resource allocations.[3] This result is consistent with both cohort and life-cycle factors. Households may be more likely to attain Pareto efficient consumption allocations at older ages, perhaps because efficient bargaining is a learned skill or because those households unable to achieve Pareto efficiency are more likely to dissolve. Alternatively, efficient bargaining may be linked to some time-invariant characteristic that is more prevalent in older cohorts in these rural Mexican states. We present suggestive evidence that our results are more consistent with cohort factors, by comparing cross-sectional surveys at different points in time and by examining heterogeneity over cohort characteristics that vary minimally through the life cycle.

We make three contributions with this paper. First, and most important, our findings show that the efficiency of household decision-making, and hence the validity of cooperative models of the house-

[1] A substantial body of empirical work fails to reject the null hypothesis that household resource allocations are Pareto efficient (Akresh, 2005; Bobonis, 2009; Bourguignon, Browning, Chiappori, and Lechene, 1993; Browning, Bourguignon, Chiappori, and Lechene, 1994; Chiappori, Fortin, and Lacroix, 2002; Rangel and Thomas, 2005). However, rejections of Pareto efficient consumption decisions are not uncommon (Dercon and Krishnan, 2000; Djebbari, 2005) and several studies reject Pareto efficient allocation of productive resources (Udry, 1996; Duflo and Udry, 2004). Several recent papers explore possible reasons why Pareto efficiency is not achieved, including limited information (Ashraf, 2009), heterogeneous time preferences (Schaner, 2013), and limited intertemporal commitment (Mazzocco, 2007).

[2] We define "young" and "old" households as households whose is head is respectively younger or older than the sample median, 39. The results are robust to small changes in the value of this cutoff.

[3] In related work, Angelucci (2008) shows that behavioral responses to cash transfers vary across rural Mexican households with different baseline characteristics. Farfán (2013) documents differences in resource-sharing and expenditure for Mexican households depending on whether their members are co-resident in one household or split across multiple locations. Hoel (2013) shows that asymmetric information prevents some but not all households from playing Pareto efficient strategies in modified dictator games in Kenya. Kazianga and Klopper (2006) show that in polygynous Malian households, fertility and child mortality is consistent with a collective model for senior wives but not for junior wives.

hold, can differ even when households face the same geographic, institutional and socio-economic environment. This suggests that tests for "the validity of cooperative household models" may be misplaced; researchers should instead aim to identify which models best characterize different groups of households. For example, the disparate results of prior empirical tests may reflect differences in the age composition of samples across different studies. Standard approaches that test whether an entire sample's behaviour is consistent with a specific model do not take this into account. Second, we show that inefficiency is more consistent with cohort than lifecycle explanations. This suggests that average efficiency may fall over time, with important welfare implications. This may provide guidance for the active research program into models of non-cooperative household behavior.

Third, we show inefficiency is associated with lower investment in children's education and with spending a smaller fraction of cash transfers on children's education. This suggests that inefficiency is associated with direct welfare costs for members of inefficient households and with lower effectiveness of public anti-poverty policy. While cash transfers have become a popular policy measure in many developing and developed countries, their effectiveness may be limited in environments where some households bargain inefficiently and hence fail to take advantage of the potential welfare gains from these transfers. We believe that we are the first to document an empirical relationship between (in)efficient household bargaining and public goods provision. There is a clear theoretical prediction that in inefficient households, members will contribute less to public goods such as children's education. However, no previous work has identified two different groups of households whose aggregate behaviour respectively is and is not consistent with efficiency. This has made it impossible to test whether inefficiency is indeed associated with lower public goods provision. Our approach relies on aggregating households into groups whose aggregate behaviour is more or less consistent with efficiency. We thus cannot be sure that our results are not driven by unmodeled heterogeneity at the household level that is correlated with (in)efficiency. But our main result is robust to controlling flexibly for the most obvious measures of heterogeneity: differences in household resources and composition, and differences in local market conditions.

Studies of household decision-making have considerably broader implications for understanding group decision-making. Households may plausibly face more conducive conditions for achieving efficient outcomes than other groups. Household members have relatively extensive information on other members' behavior, face substantial costs of leaving the group, are more likely to have altruistic preferences toward each other, and may have relatively homogeneous preferences. If households are not able to achieve efficient outcomes, the scope for other groups to do so may be even more limited.

We begin by outlining a theoretical and empirical framework in section 2. This framework closely follows the canonical cooperative model of household behavior (Chiappori, 1992; Browning and Chiappori, 1998) and the alternative tests developed in Bourguignon, Browning, and Chiappori (2009). We describe the setting and data in section 3. The data come from surveys conducted in rural Mexico between 1997 and 2007 to evaluate the Progresa/Oportunidades program. We present our main results in section 4, documenting the outcomes of efficiency tests for the full sample and the subsamples of younger and older households. In section 5 we explore cohort and life-cycle explanations for the main findings. We discuss relative education investments in more and less efficient households in section 6 and conclude in section 7.

2 Theoretical and Empirical Framework: Testing the Cooperative Model

We lay out a simple model household consumption decisions to motivate the empirical work. This model follows closely from Bourguignon, Browning, Chiappori, and Lechene (1993) and Browning, Bourguignon, Chiappori, and Lechene (1994) and is fairly standard in the literature on household bargaining. Households maximize the weighted sum of the utilities of members A and B

$$u_A(q^A, q^B, Q; a) + \mu(y, z, p) \cdot u_B(q^A, q^B, Q; a), \tag{1}$$

subject to a budget constraint

$$p \cdot q \leq y$$

where q^i is a vector of consumption of private goods for household member i, Q is a vector of consumption of public goods, $q = (q^A + q^b, Q)$, p is a vector of prices, a is a vector of preference parameters, and μ is the relative weight attached to member B's utility. μ depends on a vector of distribution factors $z = (z^1, \ldots, z^M)$, which are typically interpreted as measures of bargaining power. We assume that the utility functions u^A and u^B satsify standard conditions but do not impose any parametric assumptions at this stage. The specification of the demand problem in equation 1 assumes that distribution factors influence equilibrium consumption only through the weight function μ, not through the utility functions u_A and u_B. Substantively, this rules out candidate distribution factors from which household members might derive direct utility.

The solution to this problem yields a reduced form household demand for good j

$$d^j = \bar{d}(y, \mu(y, p, z), p, a) = \tilde{d}(y, z, p, a),$$

This demand function implies that the household budget shares for goods $j \in \{1, \ldots, J\}$ are given by

$$w^j = \frac{d^j}{x} = d(x, z, p, a) \tag{2}$$

where $x = \sum_{j=1}^{J} w^j$ is total expenditure.[4]

If the household behaves as a unitary actor, then intrahousehold resource allocations should not depend on the bargining power of the individual members. Hence, conditional on total household consumption, the distribution factors should not influence the distribution of consumption between different goods and so $\frac{\partial w^j}{\partial z^m} = 0$ in equation 2 for all $j \in \{1, \ldots, J\}$ and $m \in \{1, \ldots, M\}$. This provides a testable restriction of the unitary model.

If the household's allocation of resources is Pareto efficient, then the relative effect of changes in different distribution factors on the budget share for each consumption good should be identical across all goods: $\frac{\partial w^j}{\partial z^m} \Big/ \frac{\partial w^j}{\partial z^n} = \frac{\partial w^k}{\partial z^m} \Big/ \frac{\partial w^k}{\partial z^n}$ for all $j, k \in \{1, \ldots, J\}$ and $m, n \in \{1, \ldots, M\}$. This is typically called the *joint proportionality* condition. Intuitively, this condition is easiest to understand in terms of a multi-stage budgeting process. Browning and Chiappori (1998) shows that the collective model is formally equivalent to a model where households members first pool their incomes, then divide incomes between members based on bargaining power, and members finally allocate their income shares in accordance with their individual preferences. If a change in the distribution factor z^m increases member A's bargaining power, her income share will increase in the division stage. She will adjust her expenditure on each element of q based on her individual utility function $u_A(q^A, q^B, Q; a)$, which does not depend on z. The household budget share vector w will then change, but this change will be due only to the change in the income share received by A, not on which distribution factors induced this change. Hence, the relative shifts in budget shares will be identical across distribution factors, yielding the joint proportionality condition. This provides a testable restriction of the collective or cooperative model. Note that this is a strictly weaker restriction on consumption behavior than that implied by the unitary model. Hence, consumption behavior may be consistent with the cooperative model and not the unitary model but not *vice versa.*

We impose some parametric structure on the budget share models in equation 2. While this structure is restrictive, it is standard in the literature and nonparametric estimation of household demand models is fairly rare (Cherchye, Rock, and Vermeulen, 2009). We begin by estimating Engel curves derived from the almost ideal demand system (Deaton and Muellbauer, 1980):

$$w_{i,t}^j = \alpha^j \ln(x_{i,t}) + \sum_{m=1}^{M} \phi_m^j z_{i,t}^m + \beta^j a_{i,t} + \eta_{i,t}^j + \epsilon_{i,t}^j \tag{3}$$

for each good j and for households $i \in \{1, \ldots, n\}$ in periods $t \in \{1, \ldots, T\}$. $x_{i,t}$ represents total consumption expenditure; $z_{i,t}^1, \ldots, z_{i,t}^M$ represent distribution factors; $a_{i,t}$ represents a vector of demographic characteristics which proxy for unobserved preferences a; $\eta_{i,t}^j$ are state-by-survey wave fixed effects which proxy for unobserved prices, and $\epsilon_{i,t}^j$ captures remaining unobserved preferences.[5] We

[4]Given that this is a static framework, we abstract away from saving and so treat income y and expenditure x as equal.

[5] The canonical almost ideal demand system includes price indices and interactions between total consumption

estimate the set of all budget shares are simultaneously using a seemingly unrelated regression and cluster standard errors at the village level. We use log total household income $\ln(y_{i,t})$ as an instrument for log total household expenditure to address possible measurement error.[6]

We also estimate a quadratic almost ideal demand system adapted from Banks, Blundell, and Lewbell (1997):

$$w_{i,t}^j = \alpha_1^j \ln(x_{i,t}) + \alpha_2^j \ln(x_{i,t})^2 + \sum_{m=1}^{M} \phi_m^j z_{i,t}^m + \beta^j a_{i,t} + \eta_{i,t}^j + \epsilon_{i,t}^j \tag{4}$$

which conditions on quadratic log total expenditure and so allows a more flexible Engel curve. In particular, this specification does not assume that preferences are homothetic and allows the budget share to vary with income in a more general form. We again estimate the set of all budget shares simultaneously using a seemingly unrelated regression, cluster standard errors at the village level, and instrument the linear and quadratic expenditure terms with linear and quadratic income terms. As a final robustness check, we estimate equations 3 and 4 with interactions between the log consumption expenditure terms and the state-by-survey wave fixed effects. This allows the slope of the Engel curve to vary with regional or temporal price variation, in the spirit of the full linear and quadratic almost ideal demand systems.

The unitary model predicts that $\phi_m^j = 0$ for all $j \in \{1, \ldots, J\}$ and $m \in \{1, \ldots, M\}$. The cooperative model predicts that $\phi_m^j / \phi_n^j = \phi_m^k / \phi_n^k$ for all $j, k \in \{1, \ldots, J\}$ and $m, n \in \{1, \ldots, M\}$. We implement these tests in section 4. We depart from the literature in estimating the tests for the entire sample and separately for younger and older households. Our key innovation is the finding that the efficiency of household resource allocations varies across households with different observed characteristics: younger households' allocations are not on average efficient while older households' allocations are on average efficient.

We also implement a second test based on Bourguignon, Browning, and Chiappori (2009). They note that if there is at least one good d^j and one observed distribution factor z^m such that $d^j(x, z, p, a)$ is monotone in z^m, then dj can be inverted and z^m written as a function of $(x, z^{-m}, p, a; d^j)$. This function can be plugged into the budget share function for every other good $k \neq j$ to yield a "z-conditional demand function." We therefore obtain

$$w^k = \theta_j^k(x, z^{-m}, p, a, ; d_j)$$

expenditure and prices. Our data contain no direct measures of prices and prices can be calculated only for food shares. We therefore omit the price terms except where noted below.

[6] This strategy follows Attanasio, Battisin, and Mesnard (2011), who emphasize that total consumption expenditure is typically calculated by aggregating over item-specific expenditure. Measurement errors in individual budget shares and total consumption expenditure will thus be correlated by construction. Total income is likely to be a valid instrument because it is reported separately on most surveys, so measurement error in this variable will be uncorrelated with measurement error in the budget shares. The instrument passes standard weak instrument tests, with F-statistics of at least 20 in all specifications.

from equation 2, where $z^{-m} = (z^1, \ldots, z^{m-1}, z^{m+1}, \ldots z^M)$. Without loss of generality, we denote the good and distribution factor used in the inversion by $j = 1$ and $m = 1$ respectively. We estimate this equation using the linear specification

$$w_{i,t}^j = \alpha^j \ln(x_{i,t}) + \sum_{m=2}^{M} \phi_m^j z_{i,t}^m + \beta^j a_{i,t} + \gamma^j w_{i,t}^1 + \eta_{i,t}^j + \epsilon_{i,t}^j, \tag{5}$$

where $w_{i,t}^j$, $x_{i,t}$, $a_{i,t}$, $\eta_{i,t}^j$, and $\epsilon_{i,t}^j$ are defined as above. The ϕ vector contains the coefficients of interest. Bourguignon, Browning, and Chiappori (2009) show that under the assumption of Pareto efficiency, the distribution factors should have no direct effect on $w_{i,t}^j$ for $j > 2$ after conditioning on w_1, the expenditure share on good 1. Intuitively, w_1 is a sufficient statistic for the effect of all the distribution factors on all the budget shares. Hence, testing the cooperative model amounts to testing the joint null hypothesis that $\phi_m^j = 0$ for all $m \in \{2, \ldots, M\}$ and for all $j \in \{2, \ldots, J\}$

We also estimate quadratic versions of the z-conditional demand functions

$$w_{i,t}^j = \alpha_1^j \ln(x_{i,t}) + \alpha_2^j \ln(x_{i,t})^2 + \sum_{m=2}^{M} \phi_m^j z_{i,t}^m + \beta^j a_{i,t} + \gamma^j w_{i,t}^1 + \eta_{i,t}^j + \epsilon_{i,t}^j \tag{6}$$

which allow the slope of the Engel curve to vary nonlinearly over the observed support of income. We again estimate equations 5 and 6 with a full set of interactions between the log consumption terms and state-by-year fixed effects.

The z-conditional demand system permits a simple and direct test of the cooperative model, based on a joint test of a vector of linear restrictions. In contrast, the joint proportionality test entails a test of a nonlinear cross-equation set of restrictions. Bourguignon, Browning, and Chiappori (2009) note that such tests have relatively low power and so may fail to reject the cooperative model even when it does not apply. The z-conditional test does, however, have two key restrictions. First, the test requires a monotonic relationship between at least one distribution factor and at least one budget share. In our setting, this relationship needs to hold for both subsamples in which we implement the cooperative test. Second, the test conditions on w^1 which is by assumption correlated with the unobserved preference term ϵ^j for all $j \in \{2, \ldots, J\}$. To address this problem, we use the excluded distribution factor z^1 as an instrument for w^1 in equations 5 and 6. This instrument is valid under the assumptions used to derive the inversion argument above: w^1 is a sufficient statistic for the effect of all the distribution factors on all the budget shares. Hence, z^1 should affect w^j for all $j \in \{2, \ldots, J\}$ only through w^1. Instrument strength is testable, so we require that the relationship between z^1 and w^1 is both monotonic and statistically significant.

3 Distribution Factors, Expenditure and Demographics in Progresa 1998-1999

We use data from the evaluation of Mexico's Progresa/Oportunidades program. This program, started in 1997, still ongoing, and currently reaching about one quarter of the Mexican population, provides conditional cash transfers to eligible households. To be eligible, a household must be sufficiently poor. The transfers, paid bimonthly, are largely in the form of scholarships to the last 4 grades of primary school and the three grades of secondary school. These transfers are contingent on (i) schoolchildren attending at least 85% of classes, (ii) household members undergoing periodical health checks, and (iii) transfer recipients (typically, the mothers of the schoolchildren) attending nutrition and health classes.

The evaluation data are collected from 506 rural villages across seven states whose school-aged children met minimum school attendance and clinic visitation conditions. We have a complete census of all village residents. Moreover, in our data, all households received a score on a wealth proxy-means test: households who score below a threshold score are eligible for the Progresa transfers.

To evaluate the program, until the end of 1999 the transfers were offered only to eligible households in 320 randomly chosen treatment villages. The baseline data was collected in September-October 1997, while the first cash transfers were paid in March/April 1998 in the treated villages and between the end of 1999 and the beginning of 2000 in the control villages. We also have three follow-up waves, collected in October/November 1998, May/June 1999, and October/November 1999.[7]

Since the cash transfers were paid directly to mothers, women's bargaining power in the household plausibly increases among eligible households in treatment villages. Therefore, we follow the existing literature (Bobonis, 2009; Angelucci and Attanasio, 2013; Attanasio and Lechene, 2014) and use an indicator variable for Progresa eligibility as a distribution factor, denoted z^1. Since the Progresa cash transfer are conditional, they might change other features of the household which affect the shape of the Engel curves, besides bargaining power. One relevant change that Progresa causes is increased school attendance in households that would not have sent their children to school in the absence of the program. This, in turn, may affect household demand by changing food and non-food expenditures, as the children attending school may have school-related expenses and also change their and their household members' nutrition. Since primary school enrollment is virtually 100% in these villages, the program can have minimal changes in primary school enrolment and attendance. Conversely, secondary school enrollment in the absence of the program is about 65% and, indeed, the program has been shown to have large effects on secondary school enrollment (Schultz, 2004; Angelucci, de Giorgi, Rangel, and Rasul, 2010).

[7] By November 1999, some eligible households in control villages had received their first transfer. We return to this issue later.

To use living in a treated village as a distribution factor, we restrict the sample to households eligible for Progresa in both treatment and control villages. Moreover, avoid potential confounding effects of the program on households in which it caused a change in school attendance, we drop households with potentially eligible secondary schoolchildren at baseline. To do that, we restrict the sample to households without any child aged 10 to 16 with 5 or 6 maximum completed school grades in September 1997.[8] These are the children who may start secondary school in 1998 or 1999, i.e. the group whose enrollment is most likely affected by the program.[9][10] Lastly, we drop single-headed households and also all households for which we have missing observations. This provides a total of 25372 observations spread across 506 villages and three survey waves.

Our second distribution factor is the municipality-level ratio of female to total population in 1995, computed using data from the Mexican National System of Municipal Information database. This variable is approximately continuous and a higher female:male sex ratio is likely to reduce the relative bargaining power of women by skewing the marriage market in men's favor. The distribution factors are assumed to be independent of prices and preferences. The second part of this independence assumption may fail if there is selective migration out of villages in response to perceived marriage market conditions. However, cross-state migration for marriage is not common in this setting: approximately 99% of husbands and 95% of wives live in their state of birth. Preliminary results show that our main findings are robust to replacing the contemporaneous village sex ratio with the sex ratio from the Mexican census closest to the time that each cohort reached age 18. This also goes some way toward addressing the potential concern that the sex ratio at the time marriages are formed has a larger impact on bargaining power within relationships than the contemporaneous sex ratio. The latter measure is only appropriate if exit from marriages is a credible threat, as discussed in Mazzocco (2007).

We split our sample into "old" and "young" households, defined by whether the household head is above or below median age. Table 1 shows the means of socio-economic variables separately for the two sub-samples (columns 1 and 2) and their difference (column 3). Older households have a bigger spousal age gap and a higher likelihood of being indigenous, illiterate and less educated. They also have fewer children aged 0-9, more children aged 10-18, and more adults in the household. Conversely, neither the village characteristics – emigration share and marginalization index – nor the distribution factors – the treatment dummy and female ratio – vary by sub-group. We also find that the characteristics of these two groups are balanced by village type, with few exceptions, consistent with random assignment (columns 4 and 5).

[8] All results are robust to also dropping households with any children aged 10 to 16 irrespective of their level of schooling at baseline.

[9] Once enrolled in seventh grade, the first of the three years of secondary school, the likelihood of dropping out is low.

[10] The receipt of Progresa may also change health and nutrition for recipient households via changes in knowledge and habits. However, Angelucci and Attanasio (2013) provide evidence inconsistent with this hypothesis using a sample of urban program recipients.

Lastly, we regress female ratio on the socioeconomic variables by household sub-group (columns 6 and 7). As expected, female ratio is negatively correlated with spousal age gap (as one can marry within one's age range) and both positively correlated with village emigration (which occurs primarily among males) and negatively correlated with village marginalization (as marginalized villages are more geographically isolated, hence less likely to have emigrants). The other variables are generally not statistically significantly correlated with female ratio, besides adult males and female, as expected.

Next, we repeat this exercise for household expenditure, income, and budget shares. We use thirteen expenditure categories to construct the budget shares w^j: seven food types (dairy, fats, fruit, meat, starch, sugar, and vegetables) and six non-food commodities (children's clothing, children's school-related expenditures, female clothing and household goods (eg. kitchen equipment), health, male clothing and tobacco products, and and utilities).[11]

Table 2 shows that households are poor, as shown by their low income (about 1000 pesos at constant 1998 values, which equals approximately USD100), which is entirely consumed, and their budget is spent largely on food. The main differences in budget shares between old and young households are that the latter group, having younger children, spends more on child-related items. However, these differences are not large. The other budget shares are similar. Even when there are statistically significant differences, they are not very large.

When we look at the differences between treatment and control households, we find, as expected, that expenditure increases in the treatment group, more so for younger households, which have more school-age children and, therefore, larger cash transfers. Nevertheless, the effect of being in the treatment group on the budget shares is qualitatively similar for young and old households: the budget shares of starchy food decreases while the budget shares of fruit and meat increase, as well as child-related and women-related expenses, consistent with the evidence that women have more marked preferences than men for private goods such as women's clothing and household goods, and child-related goods. The budget share for school and health expenses decrease. This change may be partly caused by Progresa providing school- and health-related goods to its recipients.

The correlations between the female ratio and income, expenditures, and budget shares are also qualitatively similar for old and young households. Areas with larger budget shares are wealthier (higher income, consumption, and nonfood, meat, dairy, and utilities budgets shares; lower starches budget share) and have a higher budget share on women, children, and health goods.

In sum, the descriptive relationship between consumption behavior and both distribution factors is broadly consistent with the informal predictions of economic theory.

[11] We consider several alternative expenditure categories as robustness checks. In addition to the seven food shares used in the primary analysis, we consider four food shares based on the survey design (carbohydrates, fruit/vegetables, protein, and other) and five food types following Attanasio and Lechene (2014) carbohydrates, fruit/vegetables, protein, pulses, and other). In addition to the six non-food shares used in the primary analysis, we consider five shares with children's clothing and school-related expenditure combined together and four shares with all child-related expenditure combined with female clothing and household goods. Results are robust to all nine combinations of these food and non-food shares.

4 Tests of Efficiency in the Cooperative Model

4.1 Estimating Equations and Sample Definition

Recall that we estimate both conventional Engel curves of the form

$$w_{i,t}^j = \sum_{p=1}^{P} \alpha_p^j \ln(x_{i,t})^p + \alpha_2^j \ln(x_{i,t})^2 + \sum_{m=1}^{2} \phi_m^j z_{i,t}^m + \beta^j a_{i,t} + \eta_{i,t}^j + \epsilon_{i,t}^j$$

and z-conditional demand curves of the form

$$w_{i,t}^j = \sum_{p=1}^{P} \alpha_p^j \ln(x_{i,t})^p + \phi_2^j z_{i,t}^2 + \beta^j a_{i,t} + \gamma^j w_{i,t}^1 + \eta_{i,t}^j + \epsilon_{i,t}^j.$$

In each case we instrument log total expenditure $\ln(x)$ with log total income $\ln(y)$, estimate the budget shares $j = 1, \ldots, J$ jointly as a system, and cluster the standard errors at the village level. We estimate all models with four specifications: linear (P=1) and quadratic (P=2) terms in log expenditure and with and without interactions between the log expenditure terms and state-by-survey wave fixed effects.

We use binary Progresa eligibility and continuous municipality-level sex ratio as distribution factors. We estimate 13 budget shares: dairy, fats, fruit, meat, starch, sugar, vegetables, children's clothing, children's school-related expenditures, female clothing and household goods (eg. kitchen equipment), health, male clothing and tobacco products, and utilities. We include 18 covariates in the a vector to proxy for preferences that affect the budget shares: household head age, education, literacy, and indigenous status; counters for number of household members of each gender in the age brackets 0-5, 6-9, 10-12, 13-15, 16-18, and ≥ 19; the village-level marginality index; and the share of the village population who have migrated.

In the z-conditional demand estimation, we use the municipality sex ratio as the inverted distribution factor z^1 and utilities as the conditioning good w^1. The effect of the sex ratio on the budget share spent on utilities is more robust than for other expenditure categories. We instrument the budget share spent on utilities with the sex ratio. This exercise requires that (1) the sex ratio is a significant determinant of the budget share spent on utilities and that (2) the relationship is monotonic. Both assumptions are testable. The first stage F-statistic for the entire sample is at least 15.9 ($p < 0.001$) across all specifications and the F-statistics for the samples of younger and older households are at least 10.4 and 11.9 respectively ($p \leq 0.001$). We test for monotonicity by including a the sex ratio squared in the first stage and find that it is not statistically significant. We also note that the slope of the utilities-sex ratio relationship implied by the quadratic estimates is monotone over the observed support of the data. We cannot test whether the instrument is valid but this is implied by the maintained assumption of the model, that equilibrium w^1 is a sufficient statistic for the effect of the sex ratio on all budget shares.

We report results here including only households where no child is age-eligible to attend secondary school. Progresa transfers depend in part on school attendance behavior, so the realized transfer will be endogeneous with respect to school attendance decisions. This concern is unlikely to be important at primary school level, where enrollment is near universal. Schultz (2004) and Angelucci, de Giorgi, Rangel, and Rasul (2010), amongst others, show that secondary school enrollment does respond to Progresa eligibility. All results are broadly robust to including households where one or more members is *age-eligible* to attend secondary school in 1998/9 but has not yet completed primary school and so is not *grade-eligible*.

4.2 Joint Proportionality Tests

Table 3 shows relevant results from estimating equation 3 with the full sample of households. Progresa eligiblity increases expenditure shares for fruit, meat and healthcare, while decreasing expenditure shares for education, utilities, children's clothing and adult men's clothing. This shift in expenditure may reflect the health and nutrition counselling required for Progresa recipients. Higher values of the female:total population ratio are associated with increases in the expenditure shares for vegetable and utilities, and with decreases in the expenditure shares for sugary foods and for adult men's and women's clothing. Recall that higher values of this ratio are plausibly linked to lower female bargaining power, so this pattern of results is perhaps surprising. Most non-food categories are luxury goods, whose expenditure share increases with income. This is consistent with a broad literature documenting that low-income households spend a higher share of their income on food. Protein-rich dairy and meat also appear to be luxury goods, while the expenditure shares of fats and starch are falling in log total consumption.

The distribution factors are jointly significant in 8 of the 13 Engel curves. It is thus unsurprising that the unitary model, which assumes that neither distribution factor have an effect on any budget share, is strongly rejected. The χ^2 test statistic for this set of exclusion restrictions is reported in column 1 of panel A of table 5. The test statistic for the joint proportionality conditions implied by the cooperative model is 20.07, which is significant at the 5% level. We therefore reject the hypothesis that consumption allocations within households are on average Pareto efficient. Both results are highly robust to a range of alternative model specifications. The test statistics are marginally larger when controlling for a quadratic term in log total expenditure (columns 2 and 4) or interacting the log total expenditure terms with state-by-survey wave fixed effects. This provides some reassurance that the results are not driven by unobserved prices that prevent us from estimating a fully specified almost ideal demand system.

Table 3 shows relevant results from estimating equation 3 separately for households whose head is below the median age (panel A) or above the median age (panel B). The relationship between budget shares and total expenditure is broadly consistent across the two subgroups: non-food shares are generally luxuries, as are dairy and meat. Expenditure shares on fats and starch fall particularly quickly as total expenditure rises.

The Progresa-eligible households in both age subsamples spend more on fruit, meat and healthcare, while decreasing expenditure on education and some categories of clothing. In both subsamples, a higher fraction of women in the population is associated with much higher expenditure on utilities and lower expenditure on sugary foods and adult men's clothing. The magnitudes of many of these relationships differ across the subsamples but the signs seldom differ and in very few cases can we reject statistical equality across the two subsamples.

However, results of the formal tests reports in table 5 panels B and C do differ across the subsamples. We reject the unitary model for both subsamples but by a considerably larger margin for the younger than older households (χ^2 test statistics 146 and 91 respectively, both with 12 degrees of freedom). We reject the cooperative model for the younger households only, not for the older households. These results are again robust to alternative model specifications, shown in columns 2-4. Some rejections for the younger sample are marginal, in the sense that the p-values are in the neighborhood of 0.1. We therefore turn to the z-conditional approach as an alternative, potentially more powerful, testing strategy.

4.3 Z-Conditional Demand Tests

We find consistent test results using z-conditional demand systems, which are reported in table 6. We reject the cooperative model for the entire sample across all specifications. The p-values for this test are consistently smaller in the z-conditional test than the joint proportionaly tests discussed above, reflecting the generally low power of cross-equation nonlinear tests. We reject the cooperative model for younger households but fail to reject the cooperative model for older households. This result is robust across all four specifications and, in results not reported here, to alternative definitions of the expenditure categories and to alternative conditioning categories. The higher power of the z-conditional tests is again visible in each subsample: the p-values for the cooperative test decrease in the older sample from 0.49 to 0.19 and in the younger sample from approximately 0.10 to 0.04 (averaging across all four reported model specifications).

The z-conditional deman systems also provide a direct test of our hypothesis that the younger and older subsamples differ in their consumption allocations. We test the statistical hypothesis that the coefficients on the included distribution factor, Progresa eligibility, are equal for older and younger households across all budget shares. This amounts to testing the economic hypothesis that two samples' consumption allocations across different goods responds in the same way to the resources that Progresa eligibility potentially brings. This hypothesis is rejected across most specifications, as shown in panel D of table 6. This direct test complements the indirect evidence of differences in behavior obtained by comparing test results for younger and older households (inefficient vs efficient).[12]

We interpret these results as robust evidence that the cooperative model better describes the intrahousehold bargaining process of households with older heads than households with younger heads. In the next section we explore whether this difference is more consistent with cohort or life-cycle factors.

4.4 Reconciling Results with Attanasio and Lechene (2014)

Attanasio and Lechene (2014) (hereafter AL) test the collective model using the same dataset and fail to reject the model. The difference between their and our results reflects several differences in our implementation of the test. First, they use two waves of data after Progresa began to distribute cash transfers; we use three waves of data. Second, they use the relative size of husbands' and wives family networks as a distribution factor; we use the local sex ratio. We prefer the local sex ratio because it is "more continuous" than relative network sizes. Relative network sizes in our data range from zero to one but 79.3% of the observed values are 0, 0.5, or 1. The local sex ratio ranges from 0.481 to 0.525 but no single value accounts for more than 7.2% of the data. Third, they model only food consumption and condition on observed food prices as controls; we model both food and non-food consumption and condition on state-by-survey wave fixed efects to proxy for unobserved non-food prices. Fourth, we exclude households with high school-eligible children, whose enrollment behavior determines households' realized cash transfers. AL include these households and use distance to the nearest secondary school as an instrument for enrollment. We have only recently obtained this distance data and have not yet incorporated it into our main results. Fifth, they use village wages as an instrument for household consumption; we use household income. In our sample the village wage instrument does not pass standard instrument strength tests.

We replicate the AL result by estimating the same z-conditional demand system with the same two waves of data, using the relative family network distribution factor, modeling only food consumption, instrumenting high school enrollment with secondary school distance, and instrumenting household expenditure with village wages. This yields substantively similar results: we fail to reject the collective model with a Wald test statistic of 3.7 ($p = 0.449$). This result comes from inverting on the expenditure share on carbohydrates, which has a first-stage F-statistic of 13.5 ($p < 0.001$) with respect to the relative family networks instrument. The first stages for secondary school enrollment and total household expenditure have first-stage F-statistics of 10.0 and 44.3 with respect to school distance and village wages respectively ($p < 0.001$ in both cases).

We also implement AL's specifications using separate old and young subsamples. We find that

[12] We are not aware of an analogous test available in the classic joint proportionality approach. This approach employs joint nonlinear tests across the individual budget share equations of the form $H_0 : \frac{\phi_1^j}{\phi_2^j} = \frac{\phi_1^k}{\phi_2^k} \forall j \neq k$. The resultant test statistic has a χ^2_{J-1} distribution under the null hypothesis. Taking the difference between the test statistics for younger and older households does not yield useful insights, as there is no known closed-form expression for the difference between two χ^2 statistics.

the relative sibling network variable is not robustly significant in the budget share equations for these subsamples. This makes the z-conditional testing framework infeasible when applying our subsample comparison to AL's specifications and datasets.

We interpret the difference between our and AL's headline results as broadly consistent with our observation that the validity of the collective model may be sample-specific. For example, including households with high school-eligible children yields a marginal older sample (mean age 43.3 as opposed to 42.7). AL's sample thus includes more the older households whose consumption decisions we find are more consistent with the collective model.

5 Is Inefficiency a Cohort or Life Cycle Phenomenon?

This difference in the efficiency of consumption decisions between older and younger households is consistent with both cohort and life-cycle effects. Households may reach more efficient outcomes over repeated rounds of bargaining, perhaps by learning about each others' preferences or because exit becomes a less credible threat for at least one partner. Alternatively, efficiency may be an essentially time-varying characteristic of households that is more common amongst the cohort of rural Mexican households with older heads. One obvious life-cycle explanations can be quickly tested: older individuals are more likely to be divorced or separated than younger individuals but this difference is relatively small (2.8% versus 1.9%). It is possible that households that reach inefficient allocations are more likely to separate over time but the relatively small number of separations in rural Mexico makes this explanation unlikely.

We attempt to separate cohort and life-cycle effects in two ways. First, we consider characteristics of the older and younger cohorts that may drive the difference in the efficiency of household allocations. If there exist characteristics that (1) do not vary through the life-cycle but (2) differ across cohorts and (3) are associated with cross-sectional differences in the efficiency of household allocations, this provides evidence in support of cohort factors. Second, we implement the young-old efficiency tests using data from a follow-up survey of the same villages in 2007. The age distribution of household heads remains approximately constant between 1998/9 and 2007 but there is some change in the distribution of birth cohorts. If the results are broadly consistent between 1998/9 and 2007, this provides evidence in support of life-cycle factors.

5.1 Efficiency Comparisons Across Time-Invariant Cohort Characteristics

Here we consider characteristics of the older and younger cohorts might potentially drive the difference in the efficiency of household allocations. If there exist characteristics that (1) do not vary through the life-cycle but (2) differ across cohorts and (3) are associated with cross-sectional differences in the efficiency of household allocations, this provides evidence in support of cohort ahead of life-cycle factors. We depart from the finding that in households with high scores on a "tradition" index, Progresa eligibility is more likely to cause rising domestic violence Angelucci (2008). This reflects anthropological evidence that increases in female income share within the household are associated with disruptions to men's perceptions of their traditional gender roles. Angelucci (2008) constructs an index of households' tradition-orientation by averaging three proxies: the age gap between spouses, the husband's education (which enters the index negatively), and the husband's age. We construct this index, split households into those with above- and below-median values, and implement the same analysis described above.

The test results are shown in tables 7 and 8. The main patterns are fairly similar to those shown in tables 5 and 8. The unitary model is rejected for both subsamples but by a larger margin for the less than more traditional sample. When using the joint proportionality test (table 7), the cooperative model is rejected for the less traditional but not for the more traditional sample. However, the cooperative model is rejected for *both* subsamples when using the z-conditional tests (table 8), albeit by a significantly larger margin for the less traditional households.

The tradition and household head age splits are fairly strongly correlated ($\rho = 0.69$) but there is substantial variation in the former measure conditional on household head age. We also construct a tradition index using only the two time-invariant characteristics, the husband's education and the spousal age gap. The resultant split is less strongly correlated with the split on household head age ($\rho = 0.44$). We find similar results using this entirely time-invariant tradition index. The joint proportionality tests reject the cooperative model for less traditional households but not more traditional households ($p = 0.004$ and 0.609 respectively), while the z-conditional tests reject for both subsamples but by significantly different margins ($p = 0.000$ and 0.035 respectively, with $p = 0.049$ for equality across subsamples).

These results show that cohort-level characteristics that vary little over the life-cycle are also associated with differences in the extent to which consumption allocations are consistent with the cooperative model. However, the distinction is somewhat less clear for these time-invariant cohort characteristics than for the household head age. We do not believe that this provides strong evidence that the age pattern observed in the cross-section can be ascribed to either cohort or life-cycle factors. We therefore turn to an alternative testing strategy that employs a later survey of the same villages.

5.2 Efficiency Comparisons between 1998/9 and 2007

The main results are based on three sequential censuses of households in sampled villages in 1998 and 1999, from which we construct a household panel. The same villages were surveyed again in 2003 and 2007, albeit using different survey instruments. In preliminary results, we test the validity of the cooperative model using the 2007 wave and compare these results to the 1988/9 wave. The 2007 sample is slightly older than the 1998/9 sample and was born slightly later.[13] If older households attain more efficient consumption allocations due to life-cycle factors, then allocations should be more efficient for the sample in 2007, which is relatively older. If older households attain more efficient consumption allocations due to cohort factors, then allocations should be less efficient for the sample in 2007, which was born in later cohorts.

The main empirical strategy requires some modifications when applied to the 2007 data. In particular, there is experimental variation in village-level Progresa eligiblity in 1998/9 due to random rollout of the program. By 2007, all surveyed villages are eligible for transfers via Progresa. We therefore construct an alternative distribution factor based on the share of total household income controlled by the wife: her labor income and transfers paid to her, including from Progresa. This provides a proxy for the wife's bargaining power in the household and ranges from 0 to 1 with mean 0.61 and interquartile range 0.26-1.00. Progresa transfers constitute by far the largest portion wives' income share. The cross-household variation in income share is explained by differences in Progresa eligibility (grants depend on the number and gender composition of relevant children) and take-up, as well as differences in labor force participation. This variation may be correlated with unobserved household-level preferences and so wife's income share in the 2007 data is on balance a less compelling distribution factor than the randomly assigned Progresa eligibility in the 1998/9 data. We use municipality sex ratios in 2005 as the other distribution factor. We once again use the sex ratio for inverting the z-conditional demand functions, expenditure share on utilities as the conditioning good, and the wife's income share to conduct the cooperative model test.

In addition, the survey only asks education levels for a subset of respondents so it is impossible to obtain a comprehensive measure of husbands' education. We therefore report test results for subsamples defined by husbands' age and by the spousal age gap in tables 9 and 10 respectively. Finally, the survey asks about school enrollment status for a subset of households. We therefore cannot identify households with children who are enrolled or are eligible to enroll in secondary school and so cannot remove them from the estimation sample.

The tables present three key patterns in the 2007 data. First, the cooperative model is strongly rejected for the entire sample. The test statistics are comparable in magnitude to those estimated using 1998/9 data, despite the substantially smaller sample size (18834 households in 2007 versus 25392 households in 1998/9).[14] The cooperative model is also rejected for households above and below the median head age and, in most specifications, above and below the median spousal age gap. The key dimension of heterogeneity observed in the 1998/9 is substantially less prominent in the 2007, though we continue to reject equality of effect of wife's income share on budget shares across the younger and

[13] The mean ages in 2007 and 1998/9 are respectively 46.7 and 43.8 and the interquartile ranges are [35,58] and [32,53]. The mean birth cohorts in 2007 and 1998/9 are respectively 1962 and 1957 and the interquartile ranges are [1949,1972] and [1945,1967].

[14] The test statistics are asymptotically χ^2-distributed so it is unclear whether their finite sample values should be adjusted to reflect the differences in sample size.

older households.[15]

The 2007 results show stronger evidence against the cooperative model than the 1998/9 results and less evidence of heterogeneity. We interpret this preliminary cross-period result as evidence that intrahousehold consumption allocations may be less efficient amongst older households who are members of later cohorts than amongst younger households who are members of earlier cohorts. This shifts our interpretation of the main results, which employ within-period comparisons between younger and older households, toward a cohort effect rather than a life-cycle effect. Households headed by individuals born in later years appear less likely to achieve Pareto efficient intrahousehold consumption allocations.

6 Public Good Provision: Children's Education

Inefficient resource allocation within the family can result in the under-provision of public goods (Lundberg and Pollak, 1993). We proceed to test the hypothesis that children's education level and participation are lower in the subsample of younger households, for which we reject collective rationality. We argue that children's education is plausibly a public good: parents' psychic benefits from children's well-being are non-rival and largely non-excludeable, and children may not be able to condition transfers that they make to parents in later life on their recollection of which parent argued more strongly for investing in the child's education. We document that children in younger households have both lower levels of education before receiving cash transfers and a lower increase in school participation when eligible for cash transfers, despite higher parental aspirations for their children's educational attainment.

We estimate the parameters of the following equation for household i at time t:

$$y_{it} = \alpha_0 + \alpha_1 t_t + \alpha_2 P_i + \alpha_3 o_i + \alpha_4 t_t P_i + \alpha_5 t_t o_i + \alpha_6 P_i o_i + \alpha_7 t_t P_i o_i + \beta X_{it} + \epsilon_{it}, \tag{7}$$

where the indicator variables t, P, and o equal 1 for post-treatment periods, eligible households in Progresa villages, and households with older than median husbands. This is a triple-difference model that compares the time change in the outcome variables for eligible and ineligible and older and younger households. We consider four outcome variables y: (1) the share of offspring of the household head aged 11 to 16 with at least completed primary education, (2) the share of offspring of the household head aged 11 and older with at least completed secondary education, (3) the share of children aged 6-11 enrolled in school, and (4) the share of children aged 11-16 who are enrolled in school. We choose the age brackets 6-11 and 11-16 because primary school lasts 6 years and children often first enroll at age 5. Therefore, children as young as 11 can complete primary school. Moreover, at baseline, a large

[15] This result may not be robust to using joint proportionality tests, but those results were not yet available at the time this draft was prepared.

fraction of children without complete primary school are not enrolled by age 16.

We include in X the share of children in the household aged 6 to 16 (eleven variables in total, one per year of age) and the village marginalization index. The parameters of interest are $\alpha_3 + \alpha_6 \bar{P}$, which measures baseline differences in the educational attainment of children from old and young households, and α_7, which measures whether the average treatment effect of being an eligible household in a Progresa village differs for children from old and young households.

We report the estimates of these parameters in Table 11. We report three sets of estimates for each outcome. The first drops households with any child aged 10 to 16 with 5 or 6 maximum completed school grades in September 1997. This is the sample we use in the rest of the analysis. The second column keeps the same sample and adds the household wealth index and the education and indigenous status of the household head, measured at baseline, to see whether any differences between old and young households are explained by differences in these variables, as older households are less educated and more likely to be indigenous. In this way, we control for potential differences between young and old households caused by different preferences or credit constraints. Lastly, the third column shows the results for the sample without dropping households with any child aged 10 to 16 with 5 or 6 maximum completed school grades in September 1997. With minor exceptions, the results are qualitatively the same in the three sets of estimates.

The main finding from Table 11 is that investment in children's schooling is lower in younger households, for which we reject collective rationality. Children in older households have both higher stocks of education (proxied by completion of primary and secondary education, in the top panel) and a higher increase in school enrollment when eligible for Progresa. We proceed to describe our results in details.

Rows 1 and 2 in columns 1 to 6 show that, at baseline, children's educational attainment is higher in older households: the share of children with at least completed primary (secondary) school is between 1.9 and 3.4 (1.8 and 2.5) percentage points higher in old than in young households and this difference is not explained by differences in wealth, education, and indigenous status (columns 2 and 5). These differences are economically significant, especially for secondary education, as the baseline means in the control group are 0.56-0.51 for primary school and 0.07-0.16 for secondary school in the restricted and non restricted samples.

Rows 3 and 4 in columns 1 to 6 look at the effect of Progresa eligibility on educational attainment and generally find no statistically significant effects for both primary and secondary education in both younger and older households, despite some weak evidence that Progresa may increase primary school attainment by 3 percentage points more among older households thank younger households. This lack of effects is not unexpected, since we know that the strongest effects of Progresa on education are on continuing to attend school beyond primary education. Since it takes at least three years to complete secondary school, the children who are induced to stay in school beyond 6th grade in 1998 and 1999 will complete secondary school (9th grade) not earlier than 2000 and 2001, years which are not in our data.

Rows 1 and 2 in columns 7 to 12 show that, at baseline, the shares of children aged 6-11 (12-16) enrolled in school are up to 2.0 (4.6) percentage points lower in old than in young households. This is consistent with the previous results, showing a higher stock of education for older households: as children are more likely to complete primary and secondary education in older households, they are, therefore, less likely to be enrolled in school.

Rows 3 and 4 in columns 7 to 12 further show that Progresa eligibility has a statistically and economically larger effect on school enrollment in older households, increasing up to 3.8 percentage points more than in younger households, although this difference is statistically significant only for enrollment for the younger children. These results suggest that Progresa is likely to have a higher effect on completion of secondary school for older households, although it is too early to pick it up in our data.

These results are consistent with the hypothesis that inefficient households – the younger households in our sample – under-invest in public goods despite facing similar prices and institutional environments as older households. Moreover, when offered the same financial incentives to increase school attendance, they are less likely to increase enrollment.

These findings are useful for the design of public policy. Differential household responses to the same policy can have important welfare effects that might require more tailored policies. Conditional cash transfers have become a very popular policy tool in many countries but their effectiveness will be reduced if inefficient household bargaining leads to under-investment in education. At the margin, this might increase the attractiveness of policies that provide transfers to households that cannot be lost due to inefficient household bargaining. For example, school tuition fee reductions can produce similar enrollment effects to conditional-on-enrollment cash transfers, without providing a direct pecuniary transfer to households that might be inefficiently spent (Garlick, 2013). Note that this issue applies to the Pareto efficiency of how cash transfers are spent, not to the adherence of household spending with policymakers' preferences.

We also rule out that these differences are driven by differential parental preferences. In unreported regressions, we tested whether parental aspirations regarding their children's educational attainment vary for older and younger headed households. Since data on aspirations is collected for the first time in March 1998, when the program had just about started, we restrict this analysis to control villages, in which the responses are likely unaffected by the treatment. We find that older households have statistically lower aspirations than younger ones, and that this gap widens when we condition on parental wealth, education, indigenous status, village marginalization index, and share of children by age. This finding rules out that younger households invest less in their children's education because they have lower aspirations. In fact, our results suggest that, based on these variables alone, investment in children's schooling should be higher among younger households, which is in clear contrast with our findings.

7 Conclusion

We present a novel empirical finding regarding the validity of the cooperative model of household resource allocation. This model predicts that households containing members with potentially heterogeneous preferences will engage in a process of bargaining that yields a Pareto efficient set of resource allocations. Prior empirical evidence regarding this model is mixed: most but not all tests fail to reject the key testable prediction. We show that within a sample of rural Mexican households, Pareto efficiency of household expenditure is rejected for younger but not older households. This result is robust to a range of model specifications, sample definitions, and to using more aggregated expenditure categories, though the latter results are not presented here.

Our finding suggests that tests for "the validity of cooperative household models" may be misplaced; researchers should instead aim to identify which models best characterize different groups of households. For example, the disparate results of prior empirical tests may reflect differences in the composition of samples. Standard approaches that test whether an entire sample's behaviour is consistent with a specific model do not take this into account. We present preliminary evidence that this finding may be driven by a time-invariant cohort-level differences between younger and older houseolds. This will ultimately help to identify the factors that explain why some households achieve Pareto-efficient resource allocations and others do not. This contributes to an active research program into the best way to model non-cooperative household behavior.

References

Akresh, R. (2005): "Understanding Pareto Inefficient Intrahousehold Allocations," Discussion Paper 1858, IZA.

Angelucci, M. (2008): "Love on the Rocks: Domestic Violence and Alcohol Abuse in Rural Mexico," *The B.E. Journal of Economic Analysis and Policy: Contributions*, 8(1).

Angelucci, M., and O. Attanasio (2013): "The Demand for Food of Poor Urban Mexican Households: Understanding Policy Impacts using Structural Models," *American Economic Journal: Policy*, 5(1), 146–178.

Angelucci, M., G. de Giorgi, M. Rangel, and I. Rasul (2010): "Family Networks and School Enrollment: Evidence from a Randomized Social Experiment," *Journal of Public Economics*, 94(3-4), 197–221.

Ashraf, N. (2009): "Spousal Control and Intra-Household Decision Making: An Experimental Study in the Philippines," *American Economic Review*, 99(4), 1245–1277.

Attanasio, O., E. Battisin, and A. Mesnard (2011): "Food and Cash Transfers: Evidence from Colombia," *Economic Journal*, 122(559), 92–124.

Attanasio, O., and V. Lechene (2014): "Efficient Responses to Targeted Cash Transfers," *Journal of Political Economy*, 122(1), 178–222.

Banks, J., R. Blundell, and A. Lewbell (1997): "Quadratic Engel Curves and Consumer Demand," *Review of Economics and Statistics*, 79(4), 527–539.

Bobonis, G. (2009): "Is the Allocation of Resources within the Household Efficient? New Evidence from a Randomized Experiment," *Journal of Political Economy*, 117(3), 453–503.

Bourguignon, F., M. Browning, and P.-A. Chiappori (2009): "Efficient Intra-Household Allocations and Distribution Factors: Implications and Identification," *Review of Economic Studies*, 76(2), 503–528.

Bourguignon, F., M. Browning, P.-A. Chiappori, and V. Lechene (1993): "Intra Household Allocation of Consumption: A Model and Some Evidence from French Data," *Annales d'Economie et de Statistique*, 29, 137–156.

Browning, M., F. Bourguignon, P.-A. Chiappori, and V. Lechene (1994): "Income and Outcomes: A Structural Model of Intrahousehold Allocation," *Journal of Political Economy*, 102(6), 1067–1092.

Browning, M., and P.-A. Chiappori (1998): "Efficient Intra-household Allocation: A General Characterization and Empirical Tests," *Econometrica*, 66(6), 1241–1278.

Cherchye, L., B. D. Rock, and F. Vermeulen (2009): "Opening the Black Box of Intrahousehold Decision Making: Theory and Nonparametric Empirical Tests of General Collective Consumption Models," *Journal of Political Economy*, 117(6), 1074–1104.

Chiappori, P.-A. (1992): "Collective Labor Supply and Welfare," *Journal of Political Economy*, 100(3), 437–467.

Chiappori, P.-A., B. Fortin, and G. Lacroix (2002): "Marriage Market, Divorce Legislation and Household Labor Supply," *Journal of Political Economy*, 110(1), 37–72.

Deaton, A., and J. Muellbauer (1980): "An Almost Ideal Demand System," *American Economic Review*, 70(3), 213–326.

Dercon, S., and P. Krishnan (2000): "In Sickness and in Health: Risk Sharing within Households in Rural Ethiopia," *Journal of Political Economy*, 108(4), 688–727.

Djebbari, H. (2005): "The Impact on Nutrition of the Intrahousehold Distribution of Power," Discussion Paper 1701, IZA.

Duflo, E., and C. Udry (2004): "Intrahousehold Resource Allocation in Cote d'Ivoire: Social Norms, Separate Accounts and Consumption Choices," Working Paper 10498, NBER.

Farfán, G. (2013): "Extended Families across Mexico and the United States," Duke University.

Garlick, R. (2013): "How Price Sensitive is Primary and Secondary School Enrollment? Evidence from Nationwide Tuition Fee Reforms in South Africa," Mimeo, University of Michigan.

Hoel, J. (2013): "When the Average Obscures What is True for Most: A Within-Subject Test of Household Model Assumptions," International Food Policy Research Institute.

Kazianga, H., and S. Klopper (2006): "The Intra-household Economics of Polygyny: Fertility and Child Mortality in Rural Mali," World Bank.

Lundberg, S., and R. Pollak (1993): "Separate Spheres Bargaining and the Marriage Market," *Journal of Political Economy*, 101(6), 988–1010.

Mazzocco, M. (2007): "Household Intertemporal Behavior: A Collective Characterization and a Test of Commitment," *Review of Economic Studies*, 74(3), 857–895.

Rangel, M. (2006): "Alimony Rights and Intrahousehold Allocation of Resources: Evidence from Brazil," *Economic Journal*, 116(513), 627–658.

Rangel, M., and D. Thomas (2005): "Out of West Africa: Evidence on the Efficient Allocation of Resources within Farm Households," University of Chicago.

Schaner, S. (2013): "Do Opposites Detract? Intrahousehold Preference Heterogeneity and Inefficient Strategic Savings," Dartmouth College.

Schultz, P. (2004): "School Subsidies for the Poor: Evaluating the Mexican PROGRESA Poverty Program," *Journal of Development Economics*, 74(1), 199–250.

Thomas, D. (1990): "Intra-household Resource Allocation: An Inferential Approach," *Journal of Human Resources*, 25(4), 635–664.

Udry, C. (1996): "Gender, Agricultural Production, and the Theory of the Household," *Journal of Political Economy*, 104(5), 1010–1046.

Table 1: Summary Statistics and Balance Tests for Households by Age Group: Socio-Economic Status

	Old	**Young**	**O - Y**	**T-C | O**	**T-C | Y**	**Female ratio**	
						Old	**Young**
Variable	**Mean**		**Coeff.**	**Coeff.**		**Coeff.**	
	(st. dev.)		**[st.err.]**	**[st.err.]**		**[st.err.]**	
spousal age gap	6.192	2.838	3.354	-0.283	-0.092	-42.048	-16.247
	(7.651)	(4.382)	[0.154]***	[0.309]	[0.159]	[15.100]***	[8.360]*
husband age	58.084	29.256	28.828	-0.148	-0.033	8.819	-8.327
	(12.368)	(5.166)	[0.246]***	[0.467]	[0.209]	[23.967]	[9.997]
husband education	1.665	4.102	2.438	0.187	-0.192	-5.429	-12.104
	(1.923)	(2.381)	[0.074]***	[0.122]	[0.162]	[6.672]	[8.088]
husband indigenous	0.432	0.356	0.076	0.007	0.009	-3.866	-4.643
	(0.495)	(0.479)	[0.017]***	[0.058]	[0.057]	[2.890]	[2.750]*
husband literate	0.569	0.836	-0.267	-0.067	0.018	0.555	-0.096
	(0.495)	(0.37)	[0.013]***	[0.030]**	[0.022]	[1.633]	[1.112]
# relatives	2.771	2.777	-0.006	-1.237	0.121	-21.924	0.711
	(4.837)	(3.162)	[0.212]	[0.505]**	[0.204]	[17.411]	[12.533]
#males aged ≤ 5	0.31	0.765	-0.456	-0.018	0.04	-1.172	-2.178
	(0.619)	(0.782)	[0.017]***	[0.023]	[0.028]	[1.085]	[1.358]
# females aged ≤ 5	0.321	0.783	-0.463	0.025	-0.044	0.077	-0.729
	(0.644)	(0.791)	[0.017]***	[0.025]	[0.027]	[1.256]	[1.352]
# males aged 6 - 9	0.263	0.395	-0.132	0.003	0.038	-2.477	-0.207
	(0.53)	(0.617)	[0.012]***	[0.019]	[0.019]**	[0.945]***	[0.953]
# females aged 6 - 9	0.249	0.369	-0.120	0.009	-0.014	-0.632	-0.335
	(0.515)	(0.608)	[0.012]***	[0.018]	[0.020]	[0.935]	[1.044]
# males aged 10-12	0.184	0.129	0.055	-0.009	-0.001	-1.116	0.825
	(0.433)	(0.366)	[0.009]***	[0.015]	[0.013]	[0.873]	[0.573]
# females aged 10-12	0.169	0.123	0.046	0.004	-0.017	-0.782	-0.68
	(0.414)	(0.356)	[0.008]***	[0.014]	[0.012]	[0.718]	[0.591]
# males aged 13-15	0.188	0.064	0.123	0.02	-0.003	-0.788	0.686
	(0.424)	(0.27)	[0.008]***	[0.014]	[0.009]	[0.784]	[0.391]*
# females aged 13-15	0.165	0.061	0.104	-0.034	-0.017	-0.928	-0.217
	(0.408)	(0.257)	[0.007]***	[0.015]**	[0.009]**	[0.670]	[0.403]
# males aged 16-18	0.18	0.042	0.138	0.027	-0.015	-1.169	-0.243
	(0.42)	(0.213)	[0.008]***	[0.014]*	[0.007]**	[0.734]	[0.369]
# females aged 16-18	0.174	0.072	0.102	0.002	0.006	0.524	-0.168
	(0.418)	(0.267)	[0.007]***	[0.013]	[0.008]	[0.744]	[0.422]
# males aged ≥ 19	1.474	1.041	0.433	-0.002	0.015	-3.19	-0.733
	(0.73)	(0.266)	[0.015]***	[0.032]	[0.009]*	[1.780]*	[0.452]
# females aged ≥ 19	1.466	1.047	0.419	-0.001	0.009	2.465	1.918
	(0.737)	(0.423)	[0.015]***	[0.030]	[0.018]	[1.587]	[0.632]***
Village emigration share	0.048	0.043	0.005	0.008	0.001	0.581	1.118
	(0.073)	(0.063)	[0.006]	[0.009]	[0.008]	[0.581]***	[0.430]**
Village marginalization	0.465	0.464	0.001	-0.020	0.011	-11.376	-13.360
index	(0.685)	(0.781)	[0.066]	[0.094]	[0.101]	[4.987]**	[4.876]***
Treatment	0.661	0.607	0.055	1	1	2.500	-4.261
	(0.474)	(0.490)	[0.043]			[3.417]	[3.111]
Female ratio 1995	0.501	0.501	0.001	0.001	-0.002	1	1
	(0.009)	(0.010)	[0.001]	[0.001]	[0.001]		

Notes: The sample consists of 25372 household-by-survey wave observations in 506 villages.

Table 2: Summary Statistics and Balance Tests for Households by Age Group: Socio-Economic Status and Demographics

	Old	Young	O - Y	T-C | O	T-C | Y	Female ratio	
						Old	Young
Variable	**Mean**		**Coeff.**	**Coeff.**		**Coeff.**	
	(st. dev.)		**[st.err.]**	**[st.err.]**		**[st.err.]**	
income	1037.016	948.599	88.417	44.706	-121.403	3,629.25	4,732.06
	(2829.693)	(5927.015)	[60.741]	[69.806]	[149.072]	[3,364.515]	[2,991.726]
expenditure	1002.558	997.293	5.265	49.869	116.335	4,709.38	4,475.08
	(932.802)	(1069.168)	[18.539]	[29.552]*	[34.328]***	[1,757.944]***	[1,937.912]
food share	82.688	81.529	1.159	0.442	0.376	-151.464	-217.041
	(12.717)	(11.944)	[0.233]***	[0.533]	[0.636]	[28.056]***	[30.206]***
nonfood share	17.312	18.471	-1.159	-0.442	-0.376	151.464	217.041
	(12.717)	(11.944)	[0.233]***	[0.533]	[0.636]	[28.056]***	[30.206]***
starches share	42.292	41.636	0.656	-0.097	-1.401	-118.085	-156.753
	(14.948)	(14.48)	[0.276]**	[0.637]	[0.643]**	[30.250]***	[32.117]***
vegetable share	7.72	7.839	-0.119	-0.036	0.156	3.315	9.359
	(5.519)	(5.537)	[0.092]	[0.181]	[0.172]	[10.129]	[8.543]
fruit share	2.688	2.764	-0.076	0.384	0.583	5.679	-8.247
	(4.271)	(4.333)	[0.067]	[0.157]**	[0.160]***	[8.532]	[9.283]
meat share	7.353	6.96	0.393	0.509	1.24	31.418	22.41
	(9.084)	(8.425)	[0.173]**	[0.396]	[0.354]***	[17.502]*	[16.935]
dairy share	6.137	6.76	-0.623	0.055	0.209	29.874	30.941
	(6.281)	(6.581)	[0.103]***	[0.215]	[0.247]	[10.824]***	[13.344]**
sugar share	10.757	9.839	0.919	-0.259	-0.117	-94.526	-101.814
	(8.407)	(7.764)	[0.140]***	[0.354]	[0.344]	[18.027]***	[16.254]***
fats share	5.74	5.732	0.009	-0.114	-0.295	-9.139	-12.938
	(4.846)	(5.223)	[0.080]	[0.151]	[0.150]*	[7.512]	[7.509]*
women share	1.17	1.185	-0.015	0.135	0.145	14.461	9.206
	(2.241)	(2.115)	[0.034]	[0.072]*	[0.068]**	[3.604]***	[3.185]***
men share	1.244	1.097	0.146	0.104	0.054	5.133	7.316
	(2.763)	(2.247)	[0.036]***	[0.083]	[0.070]	[4.898]	[3.290]**
kids share	1.408	2.644	-1.236	0.338	0.604	4.551	28.919
	(2.566)	(3.283)	[0.061]***	[0.084]***	[0.130]***	[4.394]	[5.981]***
school share	1.242	1.364	-0.123	-0.23	-0.26	-3.062	15.757
	(4.144)	(3.86)	[0.064]*	[0.116]**	[0.127]**	[5.907]	[5.867]***
health share	8.109	7.953	0.158	-0.422	-0.579	66.88	52.244
	(9.167)	(7.598)	[0.122]	[0.257]	[0.249]**	[12.747]***	[13.126]***
utilities share	4.139	4.228	-0.089	-0.367	-0.34	63.5	103.6
	(5.094)	(5.042)	[0.104]	[0.254]	[0.283]	[13.112]***	[13.781]***

Notes: The sample consists of 25372 household-by-survey wave observations in 506 villages.

Table 3: Engel Curves for Full Sample

	Dairy	Fats	Fruit	Meat	Starch	Sugar	Vegetables	Children	Health	Men	School	Utilities	Women
Log HH	3.804***	-3.025***	0.023	5.076***	-13.069***	-0.330	0.826	-2.985***	2.234***	2.262***	0.866*	2.632***	1.687***
expenditure	(0.719)	(0.578)	(0.500)	(1.068)	(1.929)	(0.953)	(0.682)	(0.930)	(0.320)	(0.272)	(0.455)	(0.643)	(0.260)
Progresa	-0.176	0.043	0.476***	0.590**	0.204	0.117	-0.036	-0.397**	0.260***	-0.122**	-0.292***	-0.651***	-0.015
eligibility	(0.153)	(0.110)	(0.122)	(0.268)	(0.479)	(0.250)	(0.140)	(0.175)	(0.061)	(0.056)	(0.101)	(0.170)	(0.050)
Sex ratio	13.946	8.210	0.341	-15.743	-40.614	-28.584**	14.256*	13.798	-0.818	-7.451**	4.957	41.806***	-4.015*
	(9.095)	(6.199)	(8.268)	(14.972)	(31.703)	(12.587)	(7.540)	(9.553)	(4.151)	(2.943)	(5.505)	(9.822)	(2.474)
χ^2 test stat. for both dist. factors	3.702	1.852	15.611***	6.065**	2.002	6.291**	3.757	7.740**	18.166***	8.989**	9.747***	28.160***	2.756
Mean of dep. var.	6.453	5.736	2.727	7.153	41.959	10.291	7.780	8.030	2.036	1.169	1.304	4.184	4.517

Notes: The sample consists of 25372 household-by-survey wave observations in 506 villages. Standard errors in parentheses are clustered at the village level. *, **, and *** denote significance at the 10, 5, and 1% levels respectively. The system of equations is estimated simultaneously in a seeming unrelated regression model. Log household expenditure instrumented by log household income with first stage F-statistic 225.95 ($p < 0.001$). All equations include state-by-survey wave fixed effects; controls for household head age, education, literacy, and indigenous descent; controls for the number of household members of each gender in ages brackets 0-5, 6-9, 10-12, 13-15, 16-18, and 19 or older; and controls for the village-level marginalization index and migrant share.

Table 4: Engel Curves for Sample Split by Age of Household Head

Panel A: Households with Younger Heads	Dairy	Fats	Fruit	Meat	Starch	Sugar	Vegetables	Children	Health	Men	School	Utilities	Women
Log HH	5.759***	-2.750**	-0.754	6.289***	-19.410***	-2.194	0.461	-1.804	2.802***	3.430***	1.286*	4.702***	6.271***
expenditure	(1.335)	(1.155)	(0.868)	(1.985)	(3.453)	(1.577)	(1.186)	(1.596)	(0.657)	(0.496)	(0.726)	(1.053)	(1.308)
Progresa	-0.384	0.031	0.578***	0.488	0.885	0.408	0.068	-0.481*	0.209**	-0.328***	-0.392***	-0.958***	-0.305
eligibility	(0.107)	(0.164)	(0.150)	(0.343)	(0.581)	(0.306)	(0.194)	(0.250)	(0.101)	(0.068)	(0.135)	(0.209)	(0.207)
Sex ratio	8.807	8.018	-5.388	-21.649	-42.260	-25.310*	13.906	6.846	2.587	-7.197**	14.818**	53.847***	10.362
	(12.708)	(7.349)	(9.523)	(17.200)	(34.003)	(14.319)	(8.842)	(11.368)	(5.826)	(2.961)	(6.243)	(11.356)	(10.904)
χ^2 test stat. for both dist. factors	3.099	1.192	14.787***	3.780	4.261	6.431**	2.503	4.477	4.321	25.256***	20.800***	41.295***	3.724
Mean of dep. var.	6.760	5.732	2.764	6.960	41.654	9.839	7.839	7.953	2.644	1.097	1.364	4.228	5.193
Panel B: Households with Older Heads	Dairy	Fats	Fruit	Meat	Starch	Sugar	Vegetables	Children	Health	Men	School	Utilities	Women
Log HH	2.964***	-2.982***	0.151	4.231***	-8.584***	0.440	1.097	-4.135***	1.590***	1.553***	0.605	1.659**	3.606***
expenditure	(0.802)	(0.607)	(0.592)	(1.147)	(2.120)	(1.133)	(0.770)	(1.127)	(0.288)	(0.326)	(0.584)	(0.679)	(0.788)
Progresa	-0.145	0.018	0.427	0.589***	0.013*	0.006	-0.051	-0.396*	0.264***	-0.008	-0.230*	-0.534***	0.093
eligibility	(0.158)	(0.140)	(0.130)	(0.327)	(0.545)	(0.284)	(0.160)	(0.218)	(0.057)	(0.072)	(0.120)	(0.180)	(0.567)
Sex ratio	19.808**	8.598	9.216	-12.028	-34.052	-30.622*	13.878	21.272*	-5.293	-9.064**	-4.995	25.030**	-12.035
	(9.242)	(8.167)	(8.874)	(17.214)	(33.957)	(15.666)	(9.221)	(11.899)	(3.501)	(4.180)	(6.197)	(9.939)	(9.163)
χ^2 test stat. for both dist. factors	6.446**	1.157	12.704***	3.700	1.065	3.845	2.354	6.867**	25.995***	4.847*	4.895*	12.738***	1.988
Mean of dep. var.	6.137	5.740	2.688	7.353	42.292	10.757	7.720	8.110	1.408	1.244	1.242	4.139	3.819

Notes: Sample consists of 25372 household-by-survey wave observations in 506 villages. The young and old subsamples consist of 12888 and 12484 household-by-survey wave osbervations respectively. Standard errors in parentheses are clustered at the village level. *, **, and *** denote significance at the 10, 5, and 1% levels respectively. The system of equations, for young and old households, is estimated simultaneously in a seeming unrelated regression model. Log household expenditure instrumented by log household income with first stage F-statistics 80.32 ($p < 0.001$) and 185.27 ($p < 0.001$) for the young and old subsamples respectively. All equations include state-by-survey wave fixed effects; controls for household head age, education, literacy, and indigenous descent; controls for the number of household members of each gender in ages brackets 0-5, 6-9, 10-12, 13-15, 16-18, and 19 or older; and controls for the village-level marginalization index and migrant share.

Table 5: Unitary and Cooperative Test Results by Household Head Age

	(1)	(2)	(3)	(4)
Panel A: All Households				
Unitary test statistic	133.23	133.43	141.71	147.66
	[0.000]	[0.000]	[0.000]	[0.000]
Cooperative test statistic	20.07	20.10	20.22	20.39
	[0.082]	[0.081]	[0.078]	[0.074]
Panel B: Households with Younger Heads				
Unitary test statistic	145.88	145.97	153.84	130.65
	[0.000]	[0.000]	[0.000]	[0.000]
Cooperative test statistic	19.97	20.08	19.86	17.72
	[0.085]	[0.082]	[0.088]	[0.165]
Panel C: Households with Older Heads				
Unitary test statistic	91.22	91.90	90.84	94.37
	[0.000]	[0.000]	[0.000]	[0.000]
Cooperative test statistic	12.43	12.28	12.87	13.48
	[0.519]	[0.531]	[0.482]	[0.434]
Linear in log expenditure	×		×	
Quadratic in log expenditure		×		×
Expenditure interacted with state-by-survey wave fixed effects			×	×

Notes: Test results in column 1 are based on estimates reported in tables 3 and 4. Test results in columns 2-4 are based on estimates from augmented regressions (not shown) that include quadratic log expenditure terms (columns 2 and 4) and interactions between the log expenditure and the state-by-survey wave fixed effects (columns 3 and 4). These specifications follow the spirit of the linear and quadratic almost ideal demand systems discussed in section 2. Standard errors are clustered at the village level. *, **, and *** denote significance at the 10, 5, and 1% levels respectively. The system of equations, for young and old households, is estimated simultaneously in a seeming unrelated regression model. All equations include state-by-survey wave fixed effects; controls for household head age, education, literacy, and indigenous descent; controls for the number of household members of each gender in ages brackets 0-5, 6-9, 10-12, 13-15, 16-18, and 19 or older; and controls for the village-level marginalization index and migrant share.

Table 6: Z-Conditional Cooperative Test Results by Household Head Age

	(1)	(2)	(3)	(4)
Panel A: All Households				
Cooperative test statistic	52.2	59.4	55.3	66.4
	[0.000]	[0.000]	[0.000]	[0.000]
Panel B: Households with Younger Heads				
Cooperative test statistic	57.2	60.0	61.8	47.6
	[0.000]	[0.00]	[0.00]	[0.00]
Panel C: Households with Older Heads				
Cooperative test statistic	14.6	18.3	15.0	17.2
	[0.265]	[0.107]	[0.242]	[0.144]
Panel D: Comparing Age-Specific Curves				
Test statistic for equal slopes over subsamples	31.3	20.9	31.8	18.0
	[0.002]	[0.052]	[0.001]	[0.114]
Linear in log expenditure	×		×	
Quadratic in log expenditure		×		×
Expenditure interacted with state-by-survey wave fixed effects			×	×

Notes: Test results in columns 1 and 3 are from estimating models 5 and 6 respectively. Test results in columns 2 and 4 include interactions between the log consumption terms and state-by-survey wave fixed effects. Panel D reports the results of testing the hypothesis that the coefficients on the included distribution factor, Progresa eligibility, are equal for older and younger households across all budget shares. Standard errors are clustered at the village level. p-values are shown in brackets. The system of equations, for young and old households, is estimated simultaneously in a seeming unrelated regression model. All equations include state-by-survey wave fixed effects; controls for household head age, education literacy, and indigenous descent; controls for the number of household members of each gender in ages brackets 0-5, 6-9, 10-12, 13-15, 16-18, and 19 or older; and controls for the village-level marginalization index and migrant share.

Table 7: Unitary and Cooperative Test Results by Household Tradition

	(1)	(2)	(3)	(4)
Panel A: All Households				
Unitary test statistic	133.23	133.43	141.71	147.66
	[0.000]	[0.000]	[0.000]	[0.000]
Cooperative test statistic	20.07	20.10	20.22	20.39
	[0.082]	[0.081]	[0.078]	[0.074]
Panel B: Households with Low Tradition Scores				
Unitary test statistic	145.14	145.65	156.09	151.03
	[0.000]	[0.000]	[0.000]	[0.000]
Cooperative test statistic	25.65	26.09	25.35	25.03
	[0.010]	[0.008]	[0.012]	[0.013]
Panel C: Households with High Tradition Scores				
Unitary test statistic	87.93	88.65	87.14	90.64
	[0.000]	[0.000]	[0.000]	[0.000]
Cooperative test statistic	11.52	11.54	11.86	12.41
	[0.593]	[0.592]	[0.566]	[0.521]
Linear in log expenditure	×		×	
Quadratic in log expenditure		×		×
Expenditure interacted with state-by-survey wave fixed effects			×	×

Notes: Test results are generated by estimating equations 3 and 4 (columns 1 and 2) augmented with interactions between the log expenditure terms and the state-by-survey wave fixed effects. The tradition index is defined as the simple average of three standardized variables: the age gap between spouses, the husband's education and the husband's age. Standard errors are clustered at the village level. p-values are shown in brackets. The system of equations, for young and old households, is estimated simultaneously in a seeming unrelated regression model. All equations include state-by-survey wave fixed effects; controls for household head age, education literacy, and indigenous descent; controls for the number of household members of each gender in ages brackets 0-5, 6-9, 10-12, 13-15, 16-18, and 19 or older; and controls for the village-level marginalization index and migrant share.

Table 8: Z-Conditional Cooperative Test Results by Household Tradition

	(1)	(2)	(3)	(4)
Panel A: All Households				
Cooperative test statistic	52.2	59.4	55.3	66.4
	[0.000]	[0.000]	[0.000]	[0.000]
Panel B: Households with Low Tradition Scores				
Cooperative test statistic	63.2	58.3	68.3	45.0
	[0.000]	[0.000]	[0.000]	[0.000]
Panel C: Households with High Tradition Scores				
Cooperative test statistic	19.1	23.2	19.4	23.2
	[0.085]	[0.026]	[0.080]	[0.026]
Panel D: Comparing Subsample-Specific Curves				
Test statistic for equal slopes over subsamples	36.8	24.0	37.2	17.7
	[0.000]	[0.020]	[0.000]	[0.126]
Linear in log expenditure	×		×	
Quadratic in log expenditure		×		×
Expenditure interacted with state-by-survey wave fixed effects			×	×

Notes: Test results in columns 1 and 3 are from estimating models 5 and 6 respectively. Test results in columns 2 and 4 include interactions between the log consumption terms and state-by-survey wave fixed effects. Panel D reports the results of testing the hypothesis that the coefficients on the included distribution factor, Progresa eligibility, are equal for more and less traditional households across all budget shares. The tradition index is defined as the simple average of three standardized variables: the age gap between spouses, the husband's education and the husband's age. Standard errors are clustered at the village level. p-values are shown in brackets. The system of equations, for more and less traditional households, is estimated simultaneously in a seeming unrelated regression model. All equations include state-by-survey wave fixed effects; controls for household head age, education literacy, and indigenous descent; controls for the number of household members of each gender in ages brackets 0-5, 6-9, 10-12, 13-15, 16-18, and 19 or older; and controls for the village-level marginalization index and migrant share.

Table 9: Z-Conditional Cooperative Test Results by Household Head Age in 2007

	(1)	(2)	(3)	(4)
Panel A: All Households				
Cooperative test statistic	69.1	53.9	68.5	53.7
	[0.000]	[0.000]	[0.000]	[0.000]
Panel B: Households with Younger Household Heads				
Cooperative test statistic	39.2	30.1	38.9	25.1
	[0.000]	[0.003]	[0.002]	[0.015]
Panel C: Households with Older Household Heads				
Cooperative test statistic	31.0	21.3	31.1	22.5
	[0.002]	[0.047]	[0.002]	[0.032]
Panel D: Comparing Subsample-Specific Curves				
Test statistic for equal slopes over subsamples	18.6	14.6	19.3	15.2
	[0.100]	[0.264]	[0.082]	[0.230]
Linear in log expenditure	×		×	
Quadratic in log expenditure		×		×
Expenditure interacted with state-by-survey wave fixed effects			×	×

Notes: Test results in columns 1 and 3 are from estimating models 5 and 6 respectively. Test results in columns 2 and 4 include interactions between the log consumption terms and state-by-survey wave fixed effects. Panel D reports the results of testing the hypothesis that the coefficients on the included distribution factor, wife's income share, are equal for older and younger households across all budget shares. Standard errors are clustered at the village level. p-values are shown in brackets. The system of equations, for young and old households, is estimated simultaneously in a seeming unrelated regression model. All equations include state-by-survey wave fixed effects; controls for household head age, education, literacy, and indigenous descent; controls for the number of household members of each gender in ages brackets 0-5, 6-9, 10-12, 13-15, 16-18, and 19 or older; and controls for the village-level marginalization index and migrant share.

Table 10: Z-Conditional Cooperative Test Results by Spousal Age Gap in 2007

	(1)	(2)	(3)	(4)
Panel A: All Households				
Cooperative test statistic	69.1	53.9	68.5	53.7
	[0.000]	[0.000]	[0.000]	[0.000]
Panel B: Households with Lower Spousal Age Gaps				
Cooperative test statistic	66.6	58.7	64.6	57.6
	[0.000]	[0.000]	[0.000]	[0.000]
Panel C: Households with Higher Spousal Age Gaps				
Cooperative test statistic	37.3	19.0	36.8	18.4
	[0.000]	[0.088]	[0.000]	[0.105]
Panel D: Comparing Subsample-Specific Curves				
Test statistic for equal slopes over subsamples	39.4	22.4	39.0	23.5
	[0.000]	[0.033]	[0.000]	[0.024]
Linear in log expenditure	×		×	
Quadratic in log expenditure		×		×
Expenditure interacted with state-by-survey wave fixed effects			×	×

Notes: Test results in columns 1 and 3 are from estimating models 5 and 6 respectively. Test results in columns 2 and 4 include interactions between the log consumption terms and state-by-survey wave fixed effects. Panel D reports the results of testing the hypothesis that the coefficients on the included distribution factor, wife's income share, are equal for older and younger households across all budget shares. Standard errors are clustered at the village level. p-values are shown in brackets. The system of equations, for young and old households, is estimated simultaneously in a seeming unrelated regression model. All equations include state-by-survey wave fixed effects; controls for household head age, education, literacy, and indigenous descent; controls for the number of household members of each gender in ages brackets 0-5, 6-9, 10-12, 13-15, 16-18, and 19 or older; and controls for the village-level marginalization index and migrant share.

Table 11: Investment in education: share of children (a) who completed primary school (age 11-16) and secondary school (age 16+) and (b) aged 6-11 or 12-16 and currently enrolled in school.

	Completed primary school (age [11-16])			Completed secondary school (age $\geq$ 16)		
	(1)	(2)	(3)	(4)	(5)	(6)
Difference from baseline	0.019	0.034	0.034	0.020	0.018	0.025
(old-young; $\alpha_3 + \alpha_6 \bar{P}$)	[0.011]*	[0.011]***	[0.009]***	[0.005]***	[0.005]***	[0.005]***
Treatment effect	-0.023	-0.022	-0.019	0.002	0.003	0.009
(young; α_4)	[0.020]	[0.020]	[0.017]	[0.013]	[0.013]	[0.011]
Difference in treatment	0.030	0.027	0.033	-0.008	-0.009	-0.008
effect (old-young; α_7)	[0.023]	[0.023]	[0.018]*	[0.017]	[0.017]	[0.015]
Baseline mean	0.56	0.56	0.51	0.07	0.07	0.16
(st. dev.)	(0.38)	(0.38)	(0.42)	(0.12)	(0.12)	(0.26)
Drop households with potentially eligible kids	Yes	Yes	No	Yes	Yes	No
Controls for wealth and education	No	Yes	No	No	Yes	No
N	12,642	12,587	19,834	12,738	12,682	20,472
	Enrolled in school (age [6-11])			Enrolled in school (age [12-16])		
	(7)	(8)	(9)	(10)	(11)	(12)
Difference from baseline	-0.020	-0.018	-0.012	-0.046	-0.016	0.004
(old-young; $\alpha_3 + \alpha_6 \bar{P}$)	[0.005]***	[0.006]***	[0.006]**	[0.013]***	[0.013]	[0.011]
Treatment effect	-0.008	-0.008	-0.003	0.047	0.046	0.053
(young; α_4)	[0.010]	[0.010]	[0.010]	[0.023]**	[0.023]**	[0.019]***
Difference in treatment	0.037	0.038	0.021	0.031	0.037	0.007
effect (old-young; α_7)	[0.014]***	[0.014]***	[0.012]*	[0.026]	[0.026]	[0.021]
Baseline mean	0.95	0.95	0.92	0.55	0.55	0.58
(st. dev.)	(0.18)	(0.18)	(0.24)	(0.42)	(0.42)	(0.43)
Drop households with potentially eligible kids	Yes	Yes	No	Yes	Yes	No
Controls for wealth and education	No	Yes	No	No	Yes	No
N	12,623	12,574	28,018	15,306	15,253	22,934