Preliminary: please do not quote without permission

Efficient responses to targeted cash transfers

Orazio Attanasio Valérie Lechene

March 18, 2011

Abstract

Money in the hands of women affects consumption choices differently from money in the hands of men. Conditional on household income, there is a positive relationship between food expenditure and income in the hands of women in Mexico, where large cash transfers have been given to women in households living in a randomly selected sample of marginalized communities. Control over resources is also found to affect the allocation of the budget between foods. By comparing the consumption patterns of recipient households to those of similar households living in communities randomly excluded from the program, we can isolate the mechanism through which control over resources affects household choices. Further to control over resources, the relative strength of the family network of household members also shapes household choices. These findings can be rationalised as efficient responses within the collective framework. Tests of collective rationality using $z-conditional$ demands embedded within a Quadratic Almost Ideal Demand System for food components do not reject. These results lend strength to the view that men and women have different roles to play in the development process and should be targeted differently by policies designed to promote growth.

Keywords: Pareto efficiency, collective rationality, social experiment, large conditional cash transfer, power, income elasticities, price elasticities, QUAIDS, food, nutrition.

Acknowledgement 1 *We thank Pierre-André Chiappori, Tom Crossley, Ian Preston, Konrad Smolinski, Richard Spady and Rodrigo Verdu for their comments and suggestions. We also thank participants to the 2009 IFS Workshop on Household behaviour, to the World Bank Workshop on Gender in LAC, to the Cergy Pontoise Workshop on Economics of couples, and to the Empirical Microeconomics seminar at Cambridge University. The authors acknowledge funding from the World Bank Gender in LAC study. Part of this research was funded by Lechene's ESRC Research Fellowship RES-063-27-0002, Orazio Attanasio's ERC grant xxx as well as ESRC grants xxxx .*

Contents

1 Introduction

Since the implementation of *Oportunidades* (started as *PROGRESA* in Mexico in 1998), conditional cash transfer programmes have been put in place in many developing countries and have become extremely popular. They have been shown to result in important reductions in poverty. Beneficiary households in many countries, ranging from Mexico to Brazil, to Colombia and others, have been shown to enjoy higher consumption, increased school enrolment and better children nutritional status (see World Bank, 2009).

The evaluation of many of these programmes has brought to light a remarkable fact: following the injection of cash in the budget of poor households induced by CCTs (in Mexico, about 20% of household income), as total expenditure and consumption increase as expected, the consumption of food increases, proportionally, at least as much, so that the share of food among beneficiaries either increases or stays constant. This contradicts the standard view that, as a necessity, food has an income elasticity less than unity so that when total consumption increases, the share of food should decrease. This fact has been documented in the context of the urban version of the Mexican programme by Angelucci and Attanasio (2007), in rural Mexico by Attanasio and Lechene (2010), in the context of a similar programme in Colombia by Attanasio, Battistin and Mesnard (2008), in the case of a cash transfer programme in Ecuador by Schady and Rosero (2008) and in Nicaragua. A recent World Bank Policy Research Report (2009) documents the same phenomenon in other countries.

In Attanasio and Lechene (2010), we document the fact that the food budget share does not decrease in rural Mexico as total consumption increases and rule out a number of reasons why this could be, such as price increases, changes in the quality of food consumed and homotheticity of preferences as explanations for this puzzle. By estimating a carefully specified Engel curve, we show that food is indeed a necessity, with a strong negative effect of income on the food budget share. In other words, larger levels of income or total expenditure are associated (in a cross section of observations not yet affected by a CCT) with lower levels of the food share. In the case of the CCT, therefore, the decrease in food budget share caused by the large increase in income and total consumption is counterbalanced by some other effect of the programme so that the net effect is nil. We argue that the programme has not changed preferences and that there is no labelling of money. We propose that the key to the puzzle resides in the fact that the transfer is put in the hands of women and that the change in control over household resources is what leads to the observed changes in behaviour.

The focus of Attanasio and Lechene (2010) is to rule out a number of possible mechanisms that could have lead to the observed changes, and by elimination, suggest that the control over resources by women is the only remaining possibility. That work complements our earlier paper (see Attanasio and Lechene, 2002), in which we used PROGRESA as a test of the unitary model and its income pooling implication. In what follows, we go beyond both our previous papers and study a specific model of intrahousehold allocation and show that the restrictions implied by this model are satisfied in the data generated by the evaluation of the same Mexican CCT that we have studied in Attanasio and Lechene (2002, 2010). In other words, the shift in the Engel curves induced by the CCT we observe in the data can be rationalized within the model we study.

In the case of *Oportunidades*/PROGRESA, as with many other CCT programmes, the cash transfers are targeted to women. The mother of the children associated with the programme is typically the one who receives the cash transfers (and participates to the program's activities). The programme, therefore, explicitly and deliberately changes the control of resources within the households, increasing the share of total income controlled by women. Moreover, because of the program, women are involved in new activities that imply that they go out more and have more frequent connections with other women in the locality. This structure makes it possible that the programme changes the balance of power within the household and, as a consequence, the allocation of resources. Implicit in this argument is, of course, that the allocation of resources within the households is a function of who controls them, a clear violation of the so called unitary model.

Standard tests of the unitary model check whether the so-called distribution factors,[1] that is variable that do not affect tastes, relative prices or the total amount of resources, enter the demand functions for different commodities. A

[1] There is a bit of a semantic issue here. In some papers, distribution factors are understood to be any factor that affects the intrahousehold allocation of resources. Here and throughout this paper, we mean by a 'distribution factor' a variable that affects the intrahousehold allocation of resources and does not affect neither the budget constraint nor preferences. Under a unitary model, therefore, a distribution factor should not enter demand equations.

major difficulty in implementing such tests (and therefore in rejecting the unitary model) lies in the difficulty in identifying plausible distribution factors. Many of the variables used in the literature could be claimed to be affecting preferences or budget constraints. The experiment of the introduction of PROGRESA is an ideal testing ground in this sense, because, after controlling for total expenditure (and therefore the additional resources generated by PROGRESA), being in a treatment or control village should not affect the shape of the demand function.[2] Having rejected the unitary model, a legitimate question is whether alternative models are able to explain the data.

One prominent such model has been the so-called collective framework proposed by Chiappori (1988), which only assumes that total resources are allocated efficiently. Efficiency implies some important restrictions in the case in which one has more than one distribution factor. In particular, the distribution factors enter the demand functions for different commodities in a restricted fashion. Such restrictions can therefore be tested. Bourguignon, Browning and Chiappori (2009) have also recently shown that, under some conditions, these same restrictions imply that conditoning on one commodity, no distribution factor enters such conditional demand functions. This approach is particularly attractive because, after showing that PROGRESA constitutes an important distribution factor which affects the shape of several Engel curves in a substantial fashion, testing for the presence of the same factor in the conditional version of the same demand function constitutes a powerful test.

In this paper, we implement these ideas using the same data from the evaluation of PROGRESA we used in Attanasio and Lechene (2002) and in Attanasio and Lechene (2010). The execution of the exercise, however, is not trivial, as it has to solve a number of difficult issues. First, we need to deal with the endogeneity of total expenditure and, in the case of the conditional demand function, of the conditioning good. To control for the endogeneity of total expenditure, we use a control function approach, as discussed in Attanasio, Battistin and Mesnard (2009). Second, as in Attanasio and Lechene (2002), to guarantee that PROGRESA can be considered a genuine distribution factor, in addition to to-

[2] The conditionalities imposed by the program pose another difficult issue. One possible solution, used, for instance, by Attanasio and Lechene (2002), is to condition on the particular type of behaviour induced by the conditionalities (for instance school attendance). Alternatively, one can focus on aspects of the program, such as the primary school grant, which, as they have been shown not to modify behaviour, are effectively unconditional.

tal expenditure we have to control for behaviour induced by the conditionalities of the programme and in particular school enrolment. The demand functions are all conditional on the school enrolment behaviour of children of various ages. We allow enrolment to be simultaneously (and endogenously) determined with the demand functions. Third, we need to identify a plausible distribution factor apart from the (exogenously determined) participation to the programme. One of the innovations of this paper is to use a variable that measure the relative strength of the husband and wife within the household by using data on the network of relatives present in the village and their wealth.

The rest of the paper is organized as follows. In section 2, we present the framework and the theoretical results on which the empirical analysis is based. We show the form taken by the demand functions in the case of two distinct hypothesis on the intra-household negociation process: unitary rationality and collective rationality. We also present the tests of collective rationality based on *z conditional* demands. In section 3, we present the economic context and the data, a sample of poor households from the Mexican population randomly drawn to receive or not large cash transfers. We then document the fact that motivates the analysis: the absence of effect of large cash transfers on the structure of the budget in section 4. In section 5, we discuss the methodological issues pertinent to the estimation of a demand system in the context of a CCT programme. In section 6, we present the empirical results: we estimate a demand system to evaluate the impact of *Oportunidades* on food consumption, and we present tests of efficiency of decisions, using the conditional approach derived in Browning, Bourguignon and Chiappori (2009) within a modified Quaids. Section **??** concludes.

2 Theoretical framework

We consider households with 2 adult decision makers[3] A and B. There are n private consumption goods on which the household can spend, q_i^A and q_i^B, where q_i^j denotes the private consumption of good i by agent j and $i = 1, ..., n$, and

[3]This is not such a restrictive assumption as may appear as we will show below. First, a major part of the sample of poor households we consider are composed of a couple with any number of dependent relatives (children and others). Second, a number of the tests we describe are extended to the case of households with any number of decision makers as we will show. For ease of exposition, we here limit the discussion to the case of nuclear households.

Q denotes the m vector of household consumption of public goods. Household consumption of good i is $q_i = q_i^A + q_i^B$. Vector q^A is the vector of private good consumption of individual A and similarly for B. Household private consumption is $q = q^A + q^B$. Individual preferences are defined on the consumption of private goods and public goods, and they also depend on a set of demographic taste shifter d, called preference factors $v^A(q^A, q^B, Q; d)$ and $v^B(q^A, q^B, Q; d)$. Denoting exogenous total expenditure by x, the budget constraint is

$$p'(q^A + q^B) + P'Q = p'q + P'Q = x \tag{1}$$

where p and P are the price vectors of private and public goods respectively.

Individual preferences are in general not identical so that there must exist some mechanism by which households reach decisions. We consider two such mechanisms. One leads to a standard unitary model and the other to a general collective model. We show how the demand functions differ in these two cases. In what follows, we will denote ζ_i the demand function for good i, irrespective of whether it is a private or public good when we discuss properties which are shared by public and private goods. Browning, Chiappori, Lechene (2006) give a detailed discussion of the distinction between unitary and collective models when there are price variations.

2.1 Demand functions in the unitary model

One way to rationalise a unitary model based on individual preferences is to assume that households maximise a weighted sum of individual preferences where the weights are fixed.

$$Max_{q^A,q^B,Q} \mu v^A(q^A, q^B, Q; d) + (1-\mu) v^B(q^A, q^B, Q; d) \tag{2}$$

subject to the budget constraint (1). With *fixed* weights μ, this is equivalent to assuming the existence of a utility function $U(q^A, q^B, Q; d)$ which, maximised, gives rise to demand functions $\zeta_i(x, p, P, d)$ for $i = 1, ..., n$. The quantity demanded for any good i depends on total expenditure x, prices p and P and taste shifters d. For well behaved individual utility functions, the demand functions must verify adding up, homogeneity, symmetry and the Slustky matrix of compensated price responses must be negative semi definite.

2.2 Demand functions in the collective model

A well known alternative to the unitary model is the so-called collective model (Chiappori 1988, 1992). Unlike in the unitary model, it is not assumed that the weights given to the utility of each individual in the household are fixed, but that they can vary with a variety of factors. The only restriction imposed on the negotiation mechanism in the collective model is that it yields efficient allocations of resources, that is that outcomes are Pareto efficient, given the preferences of the individuals in the household.

Within the collective model, the weights μ in equation (2) can depend on a variety of factors, including prices and factors that affect the budget constraint. We furthermore assume that there exist some observable factors z which play a role in the negotiation but do not affect either the budget constraint or individual preferences. Following the literature these are called distribution factors. Notice that while variables that affect the weights but also enter the budget constraint or affect preferences (such as prices or total income) might be rationalized within the unitary model, distribution factors should not appear in the demand functions associated with such a model. Therefore, variables that can be plausibly be defined as distribution factors, are extremely useful to test the collective model as an alternative to the unitary model. As we argue below, if one can identify more than one distribution factor, one can construct powerful tests of the collective model as well, in that the model imposes strong restrictions on the way these factors enter the demand functions.

The assumptions above define a general collective model, that is a model where decisions are only assumed to be Pareto efficient. When there exist multiple distribution factors z, *Pareto efficiency implies restrictions on the manner in which they affect demand. These restrictions follow from the fact that distribution factors, as they do not affect preferences or budget constraints, enter only through the index that defines the relative weights of the two adults in the Pareto problem.*[4] Household decisions can be represented as resulting from the maximisation of a generalised household welfare function, subject to the house-

[4] In this part of the discussion, we assume that it is possible to find a set of variables which are incontroversially distribution factors. Whether it is possible to find any such characteristics is of course an important question. In the absence of a theory of marriage and of the determination of power, the decision whether a given characteristic is treated as a distribution factor z or as a preference factor d is an (untestable) identifying assumption.

hold budget constraint (1):

$$Max_{q^A,q^B,Q}\mu(x,p,P,d,z)v^A(q^A,q^B,Q;d)+(1-\mu(x,p,P,d,z))v^B(q^A,q^B,Q;d)$$

For any good, private or public, the demand function for good i derived from the programme above is $\xi_i(x,p,P,d,z)$ which depends on total expenditure x, prices, p and P, preference factors d and distribution factors z. Demand functions in the collective model satisfy adding up and homogeneity. However, it is well known that they do not satisfy symmetry, but rather that the Pseudo Slustky matrix of compensated price responses is the sum of a symmetric matrix and a matrix of rank 1 (Browning, Chiappori, 1998)

2.3 Tests of collective rationality

Tests of collective rationality differ depending upon whether the data contains price variation or not, and whether distribution factors are observed. We focus here on tests with distribution factors.

Browning, Bourguignon and Chiappori (2009) show that testing for collective rationality is equivalent to testing any of the following three conditions:

$$\xi_i(x,p,P,d,z)=\Xi_i(x,p,P,d,\mu(x,p,P,d,z)) \qquad \forall i=1,...,n \qquad (3)$$

$$\frac{\partial\xi_i/\partial z_k}{\partial\xi_i/\partial z_l}=\frac{\partial\xi_j/\partial z_k}{\partial\xi_j/\partial z_l} \qquad \forall i,j,k,l \qquad (4)$$

$$\frac{\partial\theta_i^j(x,p,P,d,z_{-1},C_j)}{\partial z_k}=0 \qquad \forall i\neq j, \quad \text{and } k=2,...,K \qquad (5)$$

The first condition states that the functional form of the demand function is restricted so that the distribution factors only affect demands through an index. The second condition is a proportionality restriction which states that the ratio of partial derivatives of the quantities demanded with respect to the distribution factors have to be equal across goods. This restriction follows easily from the first and has been tested for instance in Bourguignon et al. (1993).

To derive the final condition, let us assume that there exists at least one good j and one observable distribution factor z_1 such that $\xi_j(x,p,P,d,z)$ is strictly monotonic in z_1. Then invert demand for j so that $z_1 = \zeta(x, p, P, d, z_{-1}, C_j)$. Replacing z_1 by this expression in the demand for any other good i, one obtains the $z-$*conditional* demand for good i : $C_i = \xi_i(x, p, P, d, z_1, z_{-1}) = \theta_i^j(x, p, P, d, z_{-1}, C_j)$. From this, the third condition is easily derived. It states that, conditional on C_j, the demand for any C_i should be independent of any z_k (other than z_1). Note that because the unobservables of the demand for C_j now appear in the demand for C_i, the former is endogenous in the demand for C_i. One obvious instrument for C_j is the omitted distribution factor z_1. Note also that all these tests require at least two distribution factors and at least two demand functions. It should also be stressed that one of the distribution factors has to be such that one can invert one of the demand functions: one therefore needs a continuous factor and that one demand function is monotonic with respect to that factor.

In this paper we implement a test of collective rationality based on $z-$*conditional* demands. To our knowledge, it is the first such test in the literature.[5] The main difficulty in implementing such a test is in the identification of two variables that can be plausibly labeled as distribution factors, that is variables that affect the allocation of resources within the family but do not affect household preferences or budget constraints. One of the innovative features of this paper is the fact that we work with two such variables.

3 PROGRESA and its evaluation surveys

The data set we use is unique for a variety of reasons. First, it is the survey collected to evaluate the impact of a welfare programme in part motivated by the desire to change the position of women within rural families in Mexico. Second, the evaluation design was based on a rigorous randomized design and involved the collection of a rich and high quality survey. Third, the nature of the data allows us to construct some credible distribution factors. In this section we give some background information on the programme and the data and present some descriptive statistics.

[5] Bobonis (2009) has developed independently some of the same ideas. Unfortunately, the implementation of Bobonis's test is seriously flawed. In additional to a number of technical problems we discuss in Attanasio and Lechene (2011), the two main issues with his approach are: (i) he uses rainfall as a distribution factor, which is clearly problematic; (ii) he uses a PROGRESA dummy as the distribution factor z_1 to invert the demand function, instead of a continuous distribution factor.

3.1 PROGRESA.

After a major crisis in 1994/5,and partly in reacton to it, the Zedillo administration started an innovative programme, PROGRESA, one of the first of a new generation of 'conditional cash transfers' programmes that have since become extremely popular throughout Latin America and eslewhere. PROGRESA, which was later expanded to urban areas and changed its name into *Oportunidades,* was initially targeted to poor and marginalized rural areas and had, as its stated objectives to introduce incentives to the accumulation of human capital while at the same time alleviating short run poverty by providing poor households with cash conditional on certain investments.

Several practical aspects pertaining to the implementation of the programme are relevant for our analysis. PROGRESA/ *Oportunidades* is a conditional cash transfer programme, in the sense that receipt of the grants is conditional on the fulfillment of criterions further to the fact of being identified as poor in the sense of the program. The first set of conditions are related to health seeking behaviour. Women have to take their young children to health centres and they have to attend a number of courses organized by the programme. The second set of conditions is pertinent only for the education component of the grant. Receipt of this component is conditional on school attendance. In practice, nearly all children go to primary school. However, as about 60% of children continue to secondary school, for households with children who have finished primary school, the conditions might be binding. Importantly, the grants are paid to the women, in person, on the basis of fulfillment of the programme conditions during the preceeding period.

PROGRESA is considered a success in many dimensions, and the gold standard of welfare programmes. Replicated in most of Central and South America, and even in poor areas of New York city, the programme has been found to lead to decreases in short term poverty, and to some improvements in health, educational attainment and investment in human capital.[6] It also marks important changes in the design and delivery of interventions and welfare programmes. Price subsidies and transfers in kind are replaced by monetary transfers; eval-

[6] Detailed information can be obtained from the Oportunidades website (http://www.oportunidades.gob.mx/), or IFPRI (http://www.ifpri.org/themes/progresa.htm) the organisation which conducted the initial evaluation of the program, or in a recent World Bank Policy Research Report, 2009.

uation is conducted from the beginning of the programme; possibilities of appropriation of the programme money are removed by using private banks and other institutions to deliver the cash, and finally, the transfers are put in the hands of women. Women's role and involvement in the programme has been heralded as one of the keys of its success. We come back to this aspect below.

At the start in 1997, 300,000 families were PROGRESA beneficiaries. Now, *Oportunidades* covers 5 million households, or 25 million individuals representing 25% of the population. *Oportunidades* has the largest budget of all human development programmes.

The aim of the programme is to increase human capital investment of the poorest households in rural Mexico, through investment in education, health and nutrition. The grants have three components, designed to address these three aims. The amount of the education grant varies with the gender and age of the child, from 65 pesos for a boy in third grade to 240 pesos for a girl in third grade in secondary school (Hoddinott and Skoufias, 2004). At the start of the school year, another component of the education grant is paid to beneficiary households, towards the cost of school supplies. The education grants are capped at 490 pesos per month and per household from January to June 1998 rising to 625 pesos from July to December 1999 (Hoddinott and Skoufias, 2004). The grants are paid to the households every two months. For rural households, the programme constitute an important component of their income. For the average beneficiary, the PROGRESA grant constituted abour 20% of household income.

Hoddinott and Skoufias (2004), the World Bank CCT Policy Research Report (2009), and IFPRI reports (http://www.ifpri.org/) contain detailed descriptions and analysis of the effects of PROGRESA/*Oportunidades*. The programme's website contains up to date description of the programme and of its impacts: http://www.oportunidades.gob.mx/index.html.

3.2 The PROGRESA evaluation sample.

The means of the evaluation of the effect of the programme was built in from the start, with randomisation of recipients and regular surveys of the populations treated and of control populations. In particular, the evaluation exploited the fact that the expansion of the programme to the poopulation targeted in the first phase would take about two years. The first phase of the programme was targeted to villages identified as poor, but in possession of a certain level of amenities in terms of school and health provision. Of the 10,000 localities included in the first expansion phase, 506 localities are in the evaluation sample and 320 of them were randomly chosen to have an early start of the programme (in June 1998), while the remaining 186 were put 'at the end of the queue' and were excluded from the programme until the last months of 1999. In the 320 'treated' villages the households that in the initial (August 1997 and March 1998) surveys qualified as eligible, started receiving the cash transfers (subject to the appropriate conditionalities) in June 1998, while in the 186 'control' villages, although households were defined as eligible or non-eligible in the same fashion as in the treatment villages, no payment was made until November 1999.

In the evaluation sample, extensive surveys were administered roughly every six months from August 1997 to November 2000. In each of the selected villages, the survey is a census, which is crucial for the measurement of one of the variables we use. We use two survey waves, October 1998 and May 1999. In subsequent survey waves, starting from November 1999, poor households in control villages start being incorporated in the programme and receive part or all of the transfer they are entitled to by the programme.

The Sample Our interest is in the effect of the programme on outcomes for the poor, so that we select a sub sample of households considered as eligible for the programme in 1997, residing either in control or treatment villages.[7] In order to work with a homogenous sample in terms of number of decision makers, we also select households in which there are no more than two adults and any number of children. The sample contains 14,397 households, of which 7,400 observed in October 1998 and 6,997 observed in May 1999. Of these, 62.47% (8994 households) are in treatment villages and 37.53% (5403 households) are in control villages.

[7] In August 1997, on average, just about half the households in the targeted localities turned out to be eligible for PROGRESA. It was subsequently thought that the individual targeting had been too tight and, in March 1998, a new set of households was made eligible, so that, on average, about 78% of the households in the targeted localities turned out to be eligible. However, many of the new eligible households did not receive the transfer, for reasons that are not completely clear, for some time. To avoid dealing with these problems, in what follows we focused on the households that were originally defined as poor and that started receiving the program from its start.

3.3 Descriptive statistics

In Table 1, we report some descriptive statistics from the sample. In the first column, we report the average of each of the relevant variables in the control sample, while in the second, we report the same average in the treatment sample. A formal comparison of the two averages shows that the two samples are balanced, as reported in Behrman and Todd (1999).

Table 1

Descriptive statistics: Means

	C	T		C	T
Educ head	2.19	2.22	Household size	5.99	5.99
Educ spouse	2.15	2.15	Nb young children	2.42	2.44
Head indigenous	0.38	0.38	Nb old children	1.56	1.55
Age of head	39.55	39.39	Children in primary	1.49	1.54
Head male	0.96	0.96	Children in sec.pre.	0.30	0.35
Townsize	405.67	390.35	Distance sec. school	2311	2162
Guerrero	0.07	0.10	Dummy secondary school	0.24	0.26
Hildalgo	0.12	0.19	Distance primary school	0.61	0.23
Michoacan	0.12	0.13	Family network	0.42	0.42
Puebla	0.15	0.16	Relatives eat in	0.07	0.08
Queretaro	0.05	0.04	Household members eat out	0.02	0.02
San Luis Potosi	0.14	0.14			
Veracruz	0.35	0.23	Nb obs	5485	9284

The sample reflects the fact that we are dealing with a very poor population. Education of head and spouse are low[8]. About 60% of the sample has primary education only. The average family size is 6. Just under 40% of households are of indigenous origin. The sample is drawn from 7 different states (Guerrero, Puebla, San Luis Potosi, Michoacan, Queretaro, Veracruz and Hidalgo). About a quarter of the localities have a secondary school in the village. Few households have relatives or other outsiders eating in the house, and similarly few household members declare eating outside the house. We will control for this in the empirical analysis to correct for the direct effect on food expenditure of

[8] Education categories are: incomplete primary, primary, incomplete secondary, secondary and above.

either. We will discuss the construction of the family network variable below, in section 3.5. For now, suffices to say that there does not appear to be a difference between the mean values of the this variable in control and treatment villlages.

3.4 Definition of Commodities and Prices

In what follows, we implement a test of collective rationality on z-conditional demands. To do this, however, we have to consider at least two distribution factors (which we discuss below) and two commodities. We study the demand for the components of total food expenditure, which, in our sample, represents about 80% of non durable expenditure on average. The PROGRESA data contains very detailed information on food: the survey collects information on many narrowly defined commodities and include information both on expenditure and consumption.In computing the shares of the different foods, we include a valuation of in kind consumption.

Obviously it would not be feasible to model the demand for several dozens of food items: we therefore aggregate our data to create consumption and budget shares of 5 different commodities: (i) starches; (ii) pulses; (iii) fruit and vegetables; (iv) meat, fish and dairy; and (v) other foods. For each of the individual commodities that make our five commodities, we compute consumption so to include both what has been bought and quantities obtained from own production, payments in kind and gifts. These quantities are valued in pesos using locality level price information derived from unit values. Particular care is taken to avoid duplication induced by household production.[9]

Unit values are very important for our analysis and are used for two purposes. First, as we mentioned above, we use them to evaluate consumption in kind. Second, we use them to compute price indexes for each of the composite commodities. Unit values can be computed for each household that purchases a given commodity, as they report both quantity and value of the pourchase. 'Prices' for individual commodities at the locality level are obtained taking the median unit value of the households that purchased that product in a given locality.

[9] If a household has consumed some tortilla that were produced in the house, we include the value of the tortillas (valued at average prices in the town) but do not include the value of the flour that was purchased to make the tortillas.

Locality level prices for individual commodities are then used to compute price indexes for each of the composite commodities, averaging individual level prices and using as weights locality level budget shares in each of the individual commodities. Details on the computation of the unit values and their use to compute price indexes can be found in the Appendix and in Attanasio et al. (2010).

Spatial and temporal differences in prices of foods mean it is important to condition demands on prices. It is worth noting that the prices of foods decreased considerably between October 1998 and May 1999. As mentioned above, prices do not seem to have moved differentially between treatment and control communities. Having said that, however, it is clear that the data present a considerable amount of price heterogeneity across communities. To estimate demand functions, therefore, it will be necessary to take into account price variability even if we were considering a single cross section. The necessity to take into account variation in prices is compounded by the fact that we use two separate waves of the survey, October 1998 and May 1999.

3.5 Distribution factors

As we mentioned above, the main difficulty in testing the restrictions implied by the collective model in the way distribution factors affect the demand for commodities consists in identifying variables that can be plausibly defined as distribution factors. That is, for many variables that have been used in the literature, it is often possible to think of reasons why such a variable could affect preferences and/or budget constraints. One of the best examples, in the case of couples, is the share of income earned by the wife. While it is plausible that such a variable affects the distribution of resources within the family, if preferences are non separable between female leisure and consumption, one might find that the share of women's income, which is obviously related to female leisure, appears in the demand system even if the unitary model holds.

The fact that PROGRESA targets its cash transfers explicitly to women and that, within the evaluation sample, it was allocated randomly to a subset of villages, makes the assignment of the program an ideal distribution factor. Conditioning on total resources (including those provided by the program) and on behaviour (including those required by the program, such school enrolment), there is no reason why the demand of a household within PROGRESA should be different from that of a household outside. Moreover, the fact that the program is randomly allocated across villages allows us to observe variation in the assignment which is exogenous by construction. The amounts paid to households within the programme depend on the number, gender and age of the children, a dimension we exploit as well in the empirical implementation.

The identification of a second plausible distribution factor is more problematic, but the *Oportunidades* data offers several alternatives worth exploring. In our empirical approach we consider two alternatives.

Decision making questions. The first alternative stems from a set of questions directly addressing decision-making in the households, and perceived decision power. Women are asked to state who is usually in charge of making a range of decisions, from taking a child to the doctor, to the spending of hypothetical additional income they would receive. The possible answers are: the wife, the husband and both together. In previous work (Attanasio and Lechene (2002)), we have used these data extensively. In Table 2, we report some descriptive statistics on how these questions are answered. In particular, in the table, we report, for each question, the percentage of each of the three possible answers in both treatment and control villages.

While the majority of households report that many decisions are taken together, there is a substantial amount of variation in our sample. Especially on questions like the decision on how to spend 'the wife's additional money' and on investment in human capital (as proxied by children schooling and health investments), a considerable fraction of women report that they are mainly responsible for the decision. Moreover, as we shall see in Table 3, less than half the households answer 'together' to all eight questions about decisions.

Given these considerations, one can think of using these questions to construct an index of bargaining strength within the household. In Attanasio and Lechene (2002), we drew the conclusion that reported decision power is linked to control over monetary resources. Since it is also generally established that control over resources affects economic outcomes, reported decision power is a candidate for a distribution factor. For the sample we consider here, we confirm that answers to the the questions about who makes decision are affected by monetary resources. Table2 compares the distribution of answers to each of the eight decision questions in treatment and control villages. We find that for four of the eight questions, we reject at the 5% level the hypothesis that the distribution of answers to the question of whether it is the husband, the wife or both together who decide is the same in control and treatment villages. For another two the hypothesis is rejected at the 10% level.

Table 2
Decision Making Questions

Who decides:	The husband		The wife		Together		$\chi^2(2)$ $(P > \chi^2)$
	C	T	C	T	C	T	
whether to send a sick child to the doctor	10.52	9.43	11.45	12.05	77.10	77.48	5.18 (7.50%)
whether a child has to go to school	11.50	10.42	10.59	11.18	76.77	77.16	4.80 (9.10%)
whether a child can go outside	15.55	14.85	8.08	8.52	75.35	75.58	1.95 (37.80%)
how to spend wife's extra money	2.94	2.07	33.29	35.64	63.61	61.97	17.45 (0.00%)
about household's important exp.	32.67	31.67	4.16	4.84	62.32	62.49	4.58 (10.10%)
about exp. on electrical goods	29.77	28.23	3.87	4.72	65.29	65.86	8.66 (1.30%)
about exp. on children clothing	23.76	21.18	5.32	6.79	69.94	70.94	22.45 (0.00%)
about exp on food	17.79	16.12	17.85	20.53	63.72	62.57	18.93 (0.00%)

In order to use the answers to perceived decision power, the information contained in these variables needs to be summarised into a decision power index. There is no obvious manner in which to transform multidimensional qualitative information into such an index. We experimented with a number of alternatives to summarise this information, and upon finding that the results are robust to the way in which we aggregate the information, we decided to present those obtained with a very simple index, which is constructed by assigning the value 1 if the husband decides alone, 2 if they decide together and 3 if the wife decides alone, and adding the values of the answers to the decision-making questions in table 2. The index takes 24 values from 1 indicating the least perceived decision-making power, to 24 for 2.85% of wives, who report that they decide alone for everything. The average value of the decision-making power index is 15.36 in control villages, and 15.54 in treatment villages. A chi-square test strongly rejects the hypothesis that the distribution of the index is the same in treatment and control villages.

Table 3
Distribution of decision power index
in control and treatment villages

	Control	Treatment		Control	Treatment
1	0.02	0.01	13	5.63	5.35
2	0.04	0.07	14	6.89	6.69
3	0.22	0.11	15	8.42	7.64
4	0.09	0.02	16	42.06	41.78
5	0.02	0.06	17	15.75	17.01
6	0.02	0.08	18	2.48	3.03
7	0.15	0.09	19	0.88	1.18
8	0.57	0.57	20	0.93	1.07
9	2.33	2.16	21	0.40	0.53
10	3.81	3.30	22	0.46	0.41
11	1.79	1.46	23	0.40	0.28
12	4.14	3.69	24	2.37	3.14

$\chi^2(24)$	45.10
P under H_0	0.60%

The index we construct is not without problems. First, while the index is obvioulsy a proxy that reacts to changes in the control of resources (as indicated by the fact that it is affected by the PROGRESA grant), it is not obvious what are the distribution factors that it represents. Second, it is obviously an imperfect measure of bargaining power with important focal points that limit its variability and usefulness. The fact that for 42% of observations takes the value 16 is an indication of this problem. Third, for the exercise we want to perform, we are looking for a distribution factor other than that represented by the granting of PROGRESA to women. When using this index as the second distribution factor we consider, we implicitly assume that the indexis affected by distribution factors other than those changed by PROGRESA

Family networks. Another distribution factor can be constructed using an idea developed in an innovative paper by Angelucci, Di Giorgi, Rangel and Rasul (2007) (ADRR08, from now on). ADDR08 use the fact that the PROGRESA evaluation survey is a census within each locality and the convention of Spanish last names to map the network of siblings and cousins within each community. In Spanish speaking countries, individuals get two surnames. The first is the (first) surname of his/her father, while the second is the (first) surname of his/her mother. As in one of the waves of PROGRESA/Oportunidades, both surnames of all individuals are available, one can identify the family network for a large fraction of the sample households. We construct an algorithm which is very similar to that used by ADDR08 (and which is described in detail in the Appendix) and construct, for each individual in the evaluation sample, the number of siblings and cousins that are present in the same locality.

We then use these data to explore the possibility that power within the familiy is linked to the relative size of an individual's extended family and to the resources controlled by this extended family. We restrict the influence of family to being from siblings residing in the same village.

We measure the relative size of the wife's family network in two different ways. First, we simply consider the ratio of the number of her siblings residing in the village over the total number of siblings she and her husband have in the village, $s_2/(s_1 + s_2)$. Second, we also try to take into account the relative importance of the siblings and not only their number. More specifically, a second index is constructed as the ratio of the food consumption of the wife's siblings over the food consumption of all siblings (husband and wife). Out of approximately 11000 households, in 7562 households either the wife or the husband or both have siblings present in the village. For these households, the family network variable is straightforward. For 3080 households, neither wife nor husband have siblings living in the village. When both spouses have an identical positive number of siblings in the village, the relative strength of the family network F takes the value 1/2. We therefore code 1/2 for households with no siblings in the village, highlighting the fact that what matters is to have an equal number of siblings. Finally, for 2013 households, it is not known whether husband or wife have siblings in the village.

In Table 4, we report some descriptive statistics for the two measures of relative family networks strength we have considered in the analysis. The first column contains information about the relative number of siblings, and the second column contains information about the relative wealth of the family network. This table shows that both variables exhibits a considerable amount of variation and, therefore, have the potential of capturing variation in the barganing strength of women in different households. The correlation coefficient between the two variables is very high, at 0.9904.

Table 4
Family networks

	Siblings	Wealth
Minimum	0	0
Maximum	1	1
25%	0	0
Median	0.5	0.5
75%	0.5	0.58
Mean	0.42	0.42
Std Dev.	0.35	0.36
Correlation	0.9904	
Nb Obs	11052	

Of course, the number of siblings (of both spouses) can have a direct effect on the demand for food. To control for these possible direct effect that siblings can have on consumption, we consider the number of relatives who share meals with the household in the demand system. The survey contains explicit information on this variable.

4 Effect of the transfers on budget structure

Given the availability of the experimental setup, we can estimate the impact of the programme on total expenditure, on the share of food and on the share of the five commodities in food in a very simple fashion and with a minimal set of assumptions. The strongest of these assumptions is probably that there is no effect (maybe through anticipation) on the control localities. [10]

[10] Notice that this is different from the absence of spillover effects on individuals not receiving the transfer. As the program was randomized across communities, we can allow for spillover effects of the kind documented in these data by Angelucci and DiGiorgi (2009).

As the program was randomly allocated across localities and treatment and control samples have been proved to be well balanced in terms of baseline characteristics, the impact of the program on any given variable can be simply obtained by comparing averages in treatment and control localities. In this section, we show the effect of the program on total consumption, the consumption of food and the share food. Some of these impacts can be used as inputs in subsequent tests of the theoretical structure. Given a demand system in which, say, the demand for food depends on total consumption, one could take the impact of the programme on total consumption, feed it in an estimated relationship and test whether the model is able to predict the change in food consumption.

Table 5 shows averages for total non durable consumption, total food consumption and the budget share of food in treatment and control villages, in October 1998 and in May 1999. Not surprisingly, the consumption of non durable is considerably higher on average in treatment villages than in control villages. In May 1999, the average difference between non durable consumption in control and treatment villages is 17%, which, when converted in pesos, is still less than the amount of the grant, which accounted for about 20-25% of total consumption on average. This difference is estimated with considerable precision (the standard error is 0.012) and is therefore significantly different from zero. The increase in consumption in treatment villages in October 1998, when the programme had only just started, is considerably smaller, but still sizeable at 7% and statistically different from zero. Such a modest difference might be justified by the fact that the programme was not necessarily perceived as permanent at its inception and by administrative delays in the first few payments. The evidence on total consumption is consistent with what has been reported in the literature. The fact that the increase in total consumption is below the amount of the grant has been noted an interpreted by Gertler, Martinez and Rubio-Codina (2009).

The log of expenditure on food is 7% higher in treatment villages than in control villages in 1998. The difference between treatment and control villages increases to 15% in 1999. These average impacts of the programme, again strongly significant, are remarkably similar to the increases in total non-durable consumption, implying that the share of food does not change much. Indeed, we cannot reject the hypothesis that food shares are the same in treatment and control villages in 1998. In 1999, we actually have a small but statistically significant decrease, of less than one percentage point.

It is therefore the case that in Mexico, as in other countries where similar programmes have been operating, the share of food does not decrease after the transfer and after an increase in total consumption. This is a somewhat surprising result: if food is a necessity, one would expect its share to decrease with total expenditure.

Table 5

Comparison of total (log) consumption and the food share between control and treated villages in October 1998 and May 1999

	October 1998			May 1999		
	Cont.	Treat.	Diff.	Cont.	Treat.	Diff.
ln(cons. exp.)	6.72	6.79	0.07	6.68	6.85	0.17
	(0.47)	(0.45)	(0.011)	(0.48)	(0.48)	(0.012)
ln(food exp.)	6.52	6.59	0.07	6.45	6.60	0.15
	(0.46)	(0.46)	(0.011)	(0.47)	(0.48)	(0.012)
Share of Food	83.44	82.64	−0.48	80.22	79.64	−0.70
	(10.94)	(11.48)	(0.30)	(12.16)	(12.20)	(0.30)
Nb of obs	2659	4339		2421	4122	

Budget shares are multiplied by 100; Nb in parenthesis are standard errors for differences; standard deviations elsewhere.

In Attanasio and Lechene (2010), we have argued that the reason the share of food does not decline significantly as total consumption increases might be because the cash transfer was targeted to women and that has changed the balance of power within the household. Here, we want to check whether the restrictions implied by a specific non-unitary model of intrahousehold resource allocation, the collective model, hold in the data and can explain this evidence.

As discussed in Section 2, to perform this test, we need at least two distribution factors and at least two independent demand functions. The latter and adding up of expenditure share imply considering three commodities. Given that food accounts for such a large fraction of these families' budget and the fact that the quality of the non food items is not as high as that on food consumption, in what follows we focus on the demand for food components. This choice is also motivated by the fact that the information we have on unit values seem to indicate a large level of heterogeneity in prices across villages. To test the predictions of the collective model on a demand system, it will therefore be important to control for prices and we do not have that information for non-food components of consumption. Finally, as we document below, even when food consumption increases, the programme seems to induce relatively small changes in the composition of food consumption. It is therefore particularly interesting to check whether the demand system we estimate is able to generate this type of patterns.

In table 6, we consider the effect of the programme on the composition of food consumption. We consider consumption of five food groups: starches, wheat and rice; pulses; fruit and vegetables; meat, fish and dairy products; and finally other foods. Our figures include a valuation of in-kind consumption.

Starches account for 40% of food consumption and, therefore, about 30% of total consumption. The size of this share is another reminder of the level of poverty of these households. By contrast, expenditure on meat, fish and dairy products, which are important sources of proteins, account for only 18% of total food, while fruit and vegetables, account for 12%. Notice that almost 8% of households report zero consumption of meat, fish and dairy products in the previous week.

The table also shows the impact of the program on the shares of the five food components we are considering. The structure of the budget is not very different between control and treatment villages both in October 1998 and in May 1999. In October, the statistically significant differences are for pulses, whose share is 0.80 percentage point lower in treatment villages, for meat, fish and dairy, whose share is 1.16.percentage point higher, and other foods, whose share is 0.53 percentage point lower. In May 1999, again, statistically significant differences are not large: the largest differences recorded are for starches and meat, fish and dairy, respectively -2.30 percentage points and 2.54 percentage points different between treatment and control villages. As we see in section 6, estimating the demand for food components on control villages identifies income elasticities much different from one for several commodities. Starches, for instance, are identified as a necessity and meat a luxury. This implies that the size of the effects in Table 6 are surprisingly small.

Table 6
Composition of the food basket
in control and treated villages in October 1998 and May 1999

	October 1998			May 1999			Average %
	Cont.	Treat.	Diff.	Cont.	Treat.	Diff.	of zeros
Cash transfer	0	268		0	291		
Starches	40.26	40.04	-0.21	43.34	41.04	-2.30	0.12
	(14.72)	(13.73)	(0.34)	(15.37)	(14.79)	(0.37)	
Pulses	12.82	12.03	-0.80	11.42	10.63	-0.79	3.64
	(8.24)	(7.72)	(0.19)	(7.26)	(7.47)	(0.18)	
Fruit and vegetables	13.26	13.65	0.38	10.55	11.53	0.97	2.32
	(8.61)	(7.86)	(0.20)	(7.41)	(7.38)	(0.18)	
Meat, fish and dairy	16.03	17.20	1.16	15.99	18.53	2.54	7.91
	(12.47)	(12.36)	(0.30)	(12.75)	(12.78)	(0.31)	
Other	17.62	17.09	-0.53	18.69	18.27	-0.42	0.94
	(9.93)	(8.99)	(0.23)	(9.93)	(10.17)	(0.25)	
Nb of obs	2874	4798		2611	4486		

Budget shares are multiplied by 100; Nb in parenthesis are standard errors for differences; standard deviations elsewhere.

The evidence we have shown in this section confirms that the share of food, suprisingly, has not declined in correpondence of the increased consumption, as one would expect if food is a necessity. Furthermore, the composition of the food basket changes very little even when total food (and total consumption) change substantially. This is the evidence that a structural demand system needs to match.

5 The Demand System: Methodological issues

The estimation of a demand system on data such as those in the PROGRESA evaluation survey raises a number of methodological issues, most of which have been addressed in a number of previous papers, among which Attanasio and

Lechene (2002 and 2010) and Attanasio, Di Maro, Lechene and Phillips (2009). We briefly review these methodological issues here, starting with the functional form of the demand system, followed by the endogeneity of total expenditure, and finally by the endogeneity of schooling.

5.1 Functional form of the demand system

In what follows, we model the demand for the components of food consumption ignoring non-food consumption. In particular, we assume that the shares of the various components of food consumption are a function of total food consumption, demographics, relative prices and, possibly, distribution factors. This modeling assumption can be justified by two stage budgeting and separability of food and the rest of consumption. That is, we assume that households first decide how much to spend on food and then, conditional on that total, decide how much to allocate to each food component. Separability between food and other components implies that this subsystem does not depend on other commodities or on relative prices between food and other items. This is a strong assumption, but we feel that in this context, in which food represents such a big fraction of total consumption, it can be justified.

The first step in the analysis is the specification of a demand system. In the literature on demand, many researchers have estimated a version of the so called Quadratic Almost Ideal Demand System (QUAIDS) proposed by Banks, Blundell and Lewbell (1997) as an extension of Deaton and Muellbauer (1980)'s AIDS model. Banks, Blundell and Lewbell (1997) show how demand functions where expenditure shares are a function of relative prices, log total expenditure and its square can be consistent with utility maximization and how such a system is the one with highest rank theory consistent system in the class of perfectly aggregable systems. The extension proposed by Banks, Blundell and Lewbell (1997) to allow for quadratic effects is particularly important in our context where we want to predict changes in expenditure shares related to a relatively large change in total (food) consumption. In Attanasio, Di Maro, Lechene and Phillips (2009), we find that it is important to allow for income responses to vary with the level of income as permitted by a QUAIDS when estimating a demand system on the PROGRESA data.

The QUAIDS system can be derived from the maximization of a unitary utility function, in which case, the coefficients on the vector of prices have to satisfy a number of restrictions (so that, for instance, the resulting Slutsky matrix is symmetric and negative definite). In the context of a collective model with public goods within the family (that is goods that give utility to both maximizing agents) the shape of the demand functions that would arise is not obvious even when both agents have a utility function that would give rise to a QUAIDS system in the unitary case. Browning and Chiappori (1998) show that symmetry does not hold in the collective model, but that the Pseudo-Slutsky matrix of price responses is the sum of a symmetric matrix and a rank one matrix.and, in their empirical application, have used the QUAIDS specification as a useful parametrization of the household demand function.

In our application, following Browning and Chiappori (1998) and Attanasio, Di Maro, Lechene and Phillips (2009), we specify a QUAIDS, in which expenditure shares are allowed to depend on log total (food) consumption and its square, on prices and on demographics, as they would in a standard QUAIDS. We do not impose symmetry of the Slutzky matrix, but only homogeneity and adding-up.[11]

We also allow the effect of the two distribution factors we consider. As we saw in section 2, the assumption of efficiency of decisions imposes restrictions on household demands. For demand choices to be consistent with collective rationality, distribution factors must affect demands only through an index. This imposes restrictions which can be satisfied in a variety of ways. Suppose, for instance, that in a AIDS or QUAIDS setting, distribution factors enter the consumption shares equations additively. The proportionality restriction is then a restriction on ratio of parameters on pairs of distribution factors across demands. However, nothing prevents the distribution factors to affect demands in more complicated manners. For instance, it could be that the power index enters demands multiplicatively on total expenditure. In that case the restrictions to be tested are much more complicated. In what follows, for the sake of simplicity, we only explore specifications where the distribution factors enter additively, that is as intercept shifts.

In a Quaids, the budget shares take the following form:

[11] Browning and Chiappori (1998)'s restrictions on the Slutzky matrix from the demand systemt of a collective model can also be tested. We leave that exercise, however, for future work.

$$w_i = \alpha_i + \sum_{j=1}^{n} \gamma_{ij} \ln(p_j) + \beta_i \ln\left(\frac{x}{a(p)}\right) + \frac{\lambda_i}{b(p)} \left(\ln\left(\frac{x}{a(p)}\right)\right)^2 + u \qquad (6)$$

where w_i is the share of commodity i in total expenditure on goods, $i = 1, ...n$, x is total expenditure on goods and $a(p), b(p)$ and are price indexes defined by the following equations:

$$\begin{aligned} \ln a(p) &= \alpha_o + \sum_k \alpha_k \ln(p_k) + \frac{1}{2} \sum_k \sum_l \gamma_{kl} \ln(p_k) \ln(p_l) \\ b(p) &= \prod_{i=1}^{n} p_i^{\beta_i} \end{aligned}$$

We assume that the 'constant' α_i is a function of various demographic variables that represent shock to tastes. These include household size and its square, the number of young children, indicators for different regions, and a time trend. We also consider controls for the education of the head of household and his spouse, for the age of the head of household, for whether the head of household is indigenous and the size of the town. Finally, we also control for household members eating out and for relatives eating in, as these can have an influence on the structure of the food budget and on total food expenditure, if it is a frequent occurrence. The variable u_i represents unobserved taste heterogeneity.

Distribution factors also enter the intercept term. What distinguishes distribution factors from demographics is the fact that there are additional restrictions in the manner in which they enter into the demand functions. These restrictions are equivalently the proportionality restrictions or the restrictions on conditional demands, as we saw in section 2.3. To be specific, the equation for consumption shares of the various food components we consider is given by the following expression:

$$w_i = \theta_i' z + \phi_i' d + \sum_{j=1}^{n} \gamma_{ij} \ln(p_j) + \beta_i \ln\left(\frac{x}{a(p)}\right) + \frac{\lambda_i}{b(p)} \left(\ln\left(\frac{x}{a(p)}\right)\right)^2 + u \qquad (7)$$

where z is a vector of demographic variables and d a vector of distribution factors. θ_i and ϕ_i are vectors of parameters. Under the unitary model, the two distribution factor should not enter the demand system, so that the evidence we present also constitutes a test of the unitary model. The collective model imposes cross-equation restrictions on ϕ_i that we will test. Before doing that, however, we need to tackle a number of econometric problems. In particular, we need to deal with the endogeneity of total (food) consumption and with the endogeneity of other conditioning variables, such as school enrolment, which are important to take into account as the PROGRESA grant is conditional on these variables.

5.2 Endogeneity of total expenditure on food

In what follows, we model the five components of food as a function of total food consumption, under the assumption of two stage budgeting. Households first decide how much to allocate to food and then, conditional on total food expenditure, how much to allocate to each food component. The residuals of our equations can be interpreted as unobservable components of tastes that affect budget shares. If taste shocks to the system that determines total food consumption are correlated to the unobserved shocks to food components, then total food will be endogenous in our system. Measurement error in total expenditure is also a likely cause of endogeneity.

An instrument for total expenditure often used in the literature is household income, which implicetly assumes that the measurement error in total expenditure is uncorrelated with measured income. Under the assumption that heterogeneity in tastes is the source of endogeneity of total expenditure, income is a valid instrument if labour supply is separable from consumption.

Even if income is a valid instrument in that it is uncorrelated with the errors in the demand equation, it can be a weak instrument in a context where large transitory shocks may weaken the relationship between income and total expenditure. A possible alternative is to use a locality level measure of the price of leisure, such as average local wages, which are exogenous for the demand system if leisure is separable from commodities. It can be argued that village level agricultural wage is uncorrelated with households unobserved heterogeneity in tastes. Using local wages, however, one looses the dimension of individual variability.

We find that once we introduce distribution factors in the model, only results obtained with the village average agricultural wage are robust across different dimensions. We therefore present these results.

5.3 Schooling

Conditional cash transfer programmes impose minimum schooling requirements for children of the recipient households to receive the largest component of the grant. The grant amounts are devised with the aim to cover the opportunity cost of schooling for the household, which is why they vary with the age and gender of the child. The conditionality might affect consumption behaviour, if sending children to school imposes related costs, such as for uniforms, shoes or books. Conversely, children might be fed in school, which would also have an impact on the budget share of food and its components. It is thus necessary to control for schooling of children, over and above controlling for household composition. However, it could be that unobserved taste for school is correlated with unobserved taste for certain foods, so that schooling could be endogenous in the demand system. To allow for this possibility, we instrument schooling with an indicator for the existence of a secondary school and distance from secondary school if it is not in the village (and zero if it is in the village). The average distance to a secondary school is 2.2 kilometers, with a maximum distance of 14 kilometers. In only about 25% of villages is there a secondary school. Primary school is not considered as endogenous in the demand system. Indeed, primary school attendance is high and not affected by the programme. We follow previous literature on PROGRESA in that we condition on the number of children attending primary school in the demand system, for the reasons mentioned above.

6 Empirical results

We are now ready to present our empirical results. We divide this section in three parts. First we present results for the first stage regressions. We then discuss our estimates of the demand systems and informally compare the predictions implied by the model for the impact of the grant on consumption shares. Finally, we present the formal tests of the restrictions of the collective model.

6.1 First stage regressions

As we discussed in the previous section, there are two potentially endogenous variables in the system: total expenditure on food and number of children in secondary and "preparatoria" school. Table 7 shows the first stage regressions for the log of total expenditure and the number of children in secondary school and "preparatoria" school. In the first two columns of the table, the instrument for total expenditure is the village average agricultural wage, and household income is the instrument in the last two columns. The other instruments we use are an indicator of the presence of a secondary school and the distance to the secondary school.

The results reported in Table 7 refer to the entire sample, which includes both treatment and control localities. Results on the control sample, for which we report estimates of the demand system in the next section, are substantially similar and are available upon request.

The instruments have the expected effects in the first stage regressions. In the equation for log total expenditure, we find that both village median agricultural wage and income have power in explaining total food expenditure and that both the linear and the quadratic terms are important. As for the equation for secondary school enrolment, we find that distance to secondary school influences the number of children attending secondary school in the expected direction. Education of both parents play a role in determining the number of children in secondary education.

The first stage results are also conditional on the distribution factors we consider: namely a treatment indicator, a grant level and the family network size ratio. The first stage results obtained with the alternative measure of family network, based on consumption, as well as those obtained with the perceived power index, are not different from those presented in Table 7.

Table 7

First stage regression for total expenditure and schooling

Instrument	Village wage		Income	
Variable	Tot.Exp.	Ch.High sch.	Tot.Exp.	Ch.High sch.
ln(instr)	-0.60	0.24	-0.18	0.05
	(0.31)	(0.12)	(0.11)	(0.07)
ln(instr)^2	0.29	-0.11	0.02	-0.01
	(0.12)	(0.05)	(0.01)	(0.01)
Grant	0.21	1.91	0.23	1.91
	(0.05)	(0.09)	(0.06)	(0.10)
Treatment	0.07	-0.30	0.06	-0.30
	(0.02)	(0.02)	(0.02)	(0.02)
Family network size	-0.02	-0.02	-0.02	-0.02
	(0.01)	(0.02)	(0.01)	(0.02)
Distance high school.	0.01	-0.01	0.01	-0.02
	(0.005)	(0.005)	(0.01)	(0.01)
Indicator high school.	-0.002	0.01	0.01	-0.02
	(0.03)	(0.03)	(0.03)	(0.03)
Children in primary school	0.02	-0.15	0.01	-0.15
	(0.01)	(0.01)	(0.01)	(0.01)
Townsize	-0.09	-0.09	-0.01	-0.02
	(0.04)	(0.04)	(0.004)	(0.05)
Hhld size	0.12	0.30	0.11	0.30
	(0.01)	(0.02)	(0.01)	(0.02)
Hhld size squared	-0.004	-0.01	-0.004	-0.01
	(0.001)	(0.001)	(0.001)	(0.001)
Nb. young. children.	-0.03	-0.18	-0.03	-0.19
	(0.004)	(0.01)	(0.01)	(0.01)
Educ. spouse	-0.01	0.03	-0.01	0.03
	(0.01)	(0.01)	(0.01)	(0.01)
Head indigenous	-0.07	0.01	-0.05	0.01
	(0.02)	(0.02)	(0.02)	(0.02)
Head age	-0.001	0.03	0.001	0.02
	(0.005)	(0.006)	(0.005)	(0.008)
Educ. head	0.01	0.03	-0.003	0.03
	(0.01)	(0.01)	(0.01)	(0.01)
Relatives eat in	0.01	-0.003	0.01	-0.002
	(0.01)	(0.005)	(0.01)	(0.005)
Hhld member eat out	-0.02	0.01	-0.05	0.02
	(0.02)	(0.02)	(0.03)	(0.03)
R2	0.21	0.41	0.19	0.41
N	11671	11671	9018	9018
Test of instrument for total expenditure				
F(2,490)	11.13	2.62	9.03	0.31
$p-value$	0.00	0.07	0.00	73.09
Test of indicator for secondary, distance to secondary				
F(2,490)	2.38	6.52	2.99	6.21
$p-value$	0.09	0.00	0.05	0.00

Region dummies, prices of foods, a time trend, and a constant are also included.

6.2 Demand System

As mentioned above, the demand system is specified as a QUAIDS, which allows both linear and quadratic terms in the log of total consumption of food. This variable and schooling are treated as endogenous and we use a control function approach to deal with this issue. In particular, we add to the equations we estimate a second degree polynomial in the residuals of the first stage regressions. As indicated in Table 7, as instruments we use an indicator for the presence of a secondary school and the distance to secondary school and either village average agricultural wages or household income. By looking at the significance of first stage regression residuals in the demand system, we reject exogeneity of both total expenditure on food and secondary school in the structure of the food budget. In what follows, we present the results obtained with the agricultural wage as an instrument for total consumption of food. The results obtained with income as instrument are in an appendix. Whilst qualitatively similar to those obtained with the wage, they are less precise.

6.2.1 Demand system without distribution factors

We start our analysis of the relationship between the structure of the budget and total food expenditure, by estimating the demand system using only households living in control localities. The demand system, which corresponds to equation (6) above, includes a wide set of demographic controls, including family size and its square, household head age, the number of young chidlren, dummies for indigenous head, for household head and spouse education and for regions. It also control for the number of children in secondary and primary school (with the former being considered as endogenous). The first two columns of Table 8 report the estimates of the coefficients on total consumption of food and its square for the five food components we consider in the demand system A full set of coefficients is reported in the Appendix. We plot the Engel curves implied by the coefficients in the first two columns of Table 8 as the continuous lines in the five panels in Figure 1.

In our sample, starch is a necessity over the whole range of total expenditure, while meat is a luxury everywhere. The relationship between pulses and total food consumption is not precisely estimated. Fruit are luxuries for most of the range of total expenditure, and become necessities at high levels of total expenditure. Finally the category "other foods" appears to be a necessity at low levels of total expenditure and a luxury at high levels of total food consumption.

The next exercise consists in re-estimating the demand system pooling control and treatment households. The coefficients on log total food consumption (and its square) are shown in the next two columns of table 8; once again the complete set of results is relegated to the Appendix..The Engel curves corresponding to these coefficients are plotted again in Figure 1 as the dotted lines. The treatment consists in the injection of large amounts of cash in the budget of treated households. As the treatment was allocated randomly, if the structural model we estimated was well specified, including the treated households in the estimation sample should not make much difference to the point estimates of the income effects, but should increase the precision. What we find however, is entirely different. The estimated coefficients do change substantially. However, fruit remains a luxury and then a necessity and meat remains a luxury over the enitre range of total expenditure. For starch and pulses, the shape of the Engel curves change completely: starch is now a luxury at low levels of total expenditure and then a necessity, while pulses is a luxury over the whole range of total expenditure.

Table 8

Income effects, demand system with no distribution factors,

	(1) Control loc.		(2) Control and treatment loc.	
	Tot. Food	Tot.Food^2	Tot. Food	Tot.Food^2
Starch	-1.81	0.18	2.16	-0.20
	(1.48)	(0.15)	(1.25)	(0.12)
Pulses	0.002	0.0001	0.95	-0.09
	(0.75)	(0.08)	(0.64)	(0.06)
Fruit	2.53	-0.26	2.18	-0.23
	(0.80)	(0.08)	(0.68)	(0.07)
Meat	0.58	-0.05	2.33	-0.22
	(1.18)	(0.12)	(1.03)	(0.10)
Other foods	-1.30	0.12	-7.62	0.74
	(0.94)	(0.10)	(0.80)	(0.08)
Nb obs	5361		14170	

The demand system is as in equation 6 and includes the following demographics: household size and its square, Nb of children, education dummies,indigenous head dummy, regional Agricultural wage as instrument for total expenditure

Clearly, the system estimated in Table 8 is not stable between treatment and control localities. Another way to see this, is to use the estimates corresponding to the first two columns of Table 8 and the estimated experimental impact on total food consumption (and schooling) to predict the impact of the PROGRESA grant on the shares of the five commodities that we are considering. We do this in Tables 9 and 10, respectively for October 1998 and May 1999.

In the first column of the tables, we report the average impact of the program as estimated comparing treatment and control communities. The budget share of starch is 2.26 percentage point lower in treated households than it is in control households in May 1999. This is consistent with the fact that starch has been found to be a necessity. The share of pulse and other foods also decreases significantly at both dates, albeit not dramatically. The shares of fruits and especially of meat increase.

In the second column of Tables 9 and 10, we use the estimates of the demand system in the first two columns of Table 8 to predict the impacts. While the demand system performs reasonably well for meat and fruits, the decline in pulses is underestimated at both dates, and in starches in May 1999, while that in other foods is over-estimated. The most spectacular failure is in the prediction of the share of pulses, where the system predicts no change, while we observe a significant decline.

In the third column of Tables 9 and 10, we use the estimates of the demand system obtained from both treatment and control villages. In this case, in October 1998, two of the five predictions have the wrong sign, and in May 1999, three of the five predictions have the wrong sign. In the case of other foods, the predicted decline is more than ten times as large as the one observed. The only decent prediction is the one for the share of meat.

Table 9

Actual and predicted effects of the program on the budget structure

No Distribution factors - October 1998

	Average impact	Impact pred w/ coef of (1)	Impact pred w/ coef of (2)
Starch	-0.32	-0.66	1.96
	(0.34)	(0.43)	(0.64)
Pulses	-0.80	0.02	0.82
	(0.18)	(0.22)	(0.33)
Fruit	0.46	0.58	0.24
	(0.18)	(0.23)	(0.35)
Meat	1.21	1.00	1.79
	(0.30)	(0.34)	(0.53)
Other foods	-0.54	-0.94	-4.81
	(0.23)	()	()

Standard errors in parentheses.

Predictions are obtained from 1998 data.

Table 10
Actual and predicted effects of the program on the budget structure
No Distribution factors - May 1999

	Average impact	Impact pred w/ coef of (1)	Impact pred w/ coef of (2)
Starch	-2.26	-0.68	3.28
	(0.36)	(0.75)	(0.96)
Pulses	-0.82	0.05	1.35
	(0.19)	(0.38)	(0.49)
Fruit	0.96	0.24	-0.35
	(0.19)	(0.41)	(0.52)
Meat	2.62	1.87	2.86
	(0.31)	(0.60)	(0.78)
Other foods	-0.51	-1.48	-7.15
	(0.24)	()	()

Standard errors in parentheses.
Predictions are obtained from 1999 data.

These estimates are obtained taking the estimated demand system (without distribution factors) as a structural relationship and inputing into it the impact that the programme has on total food consumption and schooling. Other features of the programme might be relevant, such as the fact that the transfer is put in the hands of women. We will now investigate whether this is the case by introducing distribution factors in the demand system.

6.2.2 Demand system with distribution factors

We now re-estimate the demand system allowing the expenditure shares to be affected by two distribution factors: the presence of the PROGRESA programme and the variable measuring the relative size of wife and husband family networks that we discussed above. The consumption shares we estimate correspond to equation (7) and include the same demographic variables used when estimating the demand system reported in Table 8. We enter the first distribution factor as a indicator which equals to one if the household lives in a treatment locality. We have also considered the possibility that the potential grant the household is entitled to as an alternative way to enter PROGRESA as a distribution factor, but discovered that it was not significant.

The second distribution factor is crucial for the test of the collective model that we report below. We therefore investigate the possibility that it enters non-linearly the consumption shares equations. Allowing for the presence of quadratic (or higher order) terms is important for two reasons. First, the theory is silent about the specific form in which this (or any other) distribution factors enters the share equations. Second, the test of the collective model we propose requires that there is at least one commodity for which one of the distribution factors enters monotonically, so that the relationship can be inverted.

In Table 11, we report the estimates of the coefficients on the PROGRESA dummy and on the family network size and its square that we obtain on the whole sample using the agricultural average wage as an instrument for total food consumption. This set of results is representative and robust across different specifications. The main finding is that, while the coefficient on the PROGRESA dummy is significantly different from zero in four out of five share equations, the quadratic terms of the family size network never attracts a coefficient that is significantly different from zero. The results were virtually identical when we used the index based on the relative wealth of husband and wife's networks.

Table 11

Effect of distribution factors on the consumption shares

	Starch	Pulses	Fruit	Meat	Other foods
Treatment	0.02	-0.018	-0.007	-0.02	0.03
	(0.01)	(0.004)	(0.005)	(0.01)	(0.01)
Family network size	-0.08	0.07	0.13	0.03	-0.15
	(0.11)	(0.06)	(0.06)	(0.09)	(0.07)
Family network size^2	0.34	-0.43	-0.81	0.54	0.35
	(1.13)	(0.58)	(0.61)	(0.92)	(0.71)
Nb obs	11975				

Instrument for total food is average:wage. Standard errors in parentheses

Given these findings, we decided to use a specification of equation (7) that is linear in relative network sizes In Table 12, we report the estimates we obtain for the coefficients on the two distribution factors and for the linear and quadratic log total food consumption. Several comments are in order. First, both distribution factors are strongly significant. The PROGRESA dummy is significantly different from zero at standard levels in four of the five share equations. The relative size of family networks is also significantly different from zero for four of the five shares equations. Notice that under the unitary model, neither of these variables should enter the demand system. The linearity of the relationship between shares and the the relative size of the family networks implies that we can use any of these relationships to perform the inversion described in Section 2.3 and construct the z-conditional demands.

Table 12					
Effect of distribution factors on the budget, inst:wage					
	Starch	Pulses	Fruit	Meat	Other foods
Treatment	0.02	-0.018	-0.007	-0.02	0.03
	(0.01)	(0.004)	(0.005)	(0.01)	(0.01)
Family network size	-0.05	0.03	0.05	0.08	-0.12
	(0.04)	(0.02)	(0.02)	(0.03)	(0.03)
ln Tot. Exp.Food	-2.97	0.53	2.58	3.05	-3.20
	(0.99)	(0.51)	(0.53)	(0.80)	(0.62)
ln Tot.Exp.Food^2	0.29	-0.05	-0.26	-0.29	0.31
	(0.10)	(0.05)	(0.05)	(0.08)	(0.06)
Nb obs	11975				
Estimates of some of the coefficients of Equation (7).					

As for the coefficients on total log food consumption, we notice that for all commodities, except for pulses, the quadratic effects are strongly significant. For pulses, indeed, even the linear term is not significantly different from zero. However, the two terms are jointly significant at the 3% level. We plot the Engel curves implied by these coefficients in Figure 2. Starches, as before, are a necessity over the whole relevant range of food consumption, pulses and meat are luxuries and fruit is a luxury at low levels of total expenditure, then becoming a necessity.

In the next sub-section, we present the formal test of the restrictions implied by the collective model. One more informal but informative way to check whether our specification fits the data generated by the PROGRESA experiment, which explicitly changed the control of resources within the family in a controlled way, is to check whether the specification of the demand system in equation (7) is able to predict the changes in consumption shares reported in Table 6. We therefore, re-do the exercise reported in Tables 9 and 10, but using the coefficients of the demand system that includes the distribution factors we have considered. The results of this exercise are reported in Table 13 and 14.

The results we obtain now are much different from those in Tables 9 and 10. All predicted impacts are of the same sign as the observed changes in consumption shares, apart from other foods in October 1998. Apart from starches and fruit in May 1999, the predicted impacts are estimated with sufficient precision so that they are significantly different from zero. Moreover and more importantly, they are considerably closer to the actual experimental impacts.

Table 13:
Actual and predicted impacts of the program on
the commodity shares with distribution factors October1998

	Average impact	Predicted impacts
Starch	-0.32	-0.11
	(0.34)	(0.32)
Pulses	-0.80	-0.92
	(0.18)	(0.17)
Fruit	0.46	0.29
	(0.18)	(0.17)
Meat	1.21	0.51
	(0.30)	(0.26)
Other foods	-0.54	0.23
	(0.23)	()

Predicted impacts computed using the model in Table 12.

Table 14:
Actual and predicted impacts of the program on
the commodity shares with distribution factors May 1999

	Average impact	Predicted impacts
Starch	-2.23	-0.59
	(0.36)	(0.32)
Pulses	-0.82	-0.47
	(0.19)	(0.17)
Fruit	0.96	0.12
	(0.19)	(0.17)
Meat	2.62	1.99
	(0.31)	(0.26)
Other foods	-0.51	-1.04
	(0.24)	()

Predicted impacts computed using the model in Table 12.

We interpret this evidence as indicating that there is scope, in the context of the PROGRESA programme, for investigating the role played by features of the programme which cannot be rationalised within the standard framework of unitary household choices, but need to be accounted for. In the next subsection, we turn to the formal test of the collective model.

We will not comment here on other aspects of the estimation of the demand system. Interested readers are referred to Attanasio, Di Maro, Lechene and

Phillips (2009 and 2010) for in depth analysis of income and price responses and welfare analysis in this context.

6.3 Test of Efficiency

The test for collective rationality requires that we can observe at least two goods and two distribution factors. One aspect which is crucial in the analysis is that the distribution factor and the good which is used for the test have to have a statistically significant link, otherwise the test has no power. This is not a problem here, since the relative size of husband and wife's networks is significant in the demands for four goods out of five. Similar considerations apply to the results that we obtain with the reported power and the alternative measure of relative networks.

The good we use for the test of conditional rationality is animal protein, or meat, fish and dairy. Both distribution factors significantly influence the demand for animal protein, and the relationship between demand and the family network is monotonic.

The table below gives the results of test of collective rationality using z-conditional demand with animal proteins as the conditioning good and relative family network size as the distribution factor used to invert the demand for meat. To deal with the endogeneity of the conditioning good, we use, once again a control function approach, where the identifying instrument, consistently with the collective model, is the distribution factor that is used to invert the demand for the conditioning good. For each good, in the table, we also report the results for the unconditional estimation. We only report the coefficient on the PROGRESA dummy and, in the case of the conditional demands, the coefficient on meat and the coefficient on the residual for the first stage regression for meat, denoted with u_{meat}. As for the other tables, the full set of results is available in the Appendix.

The results are striking: in the unconditional demand system, treatment is significant in the 4 goods. In the z conditional demand system it is nowhere significant. We cannot reject the collective model.

Table 13								
Collective rationality test with family network distribution factor,								
	Starches		Pulses		Fruit		Other foods	
	Uncond.	z cond.	Uncond.	z cond.	Uncond.	z cond.	Uncond.	z cond.
Treat	0.02	0.01	-0.018	-0.01	-0.007	0.01	0.03	-0.01
	(0.01)	(0.01)	(0.004)	(0.01)	(0.005)	(0.01)	(0.01)	(0.01)
Meat		-0.61		0.40		0.65		-1.44
		(0.47)		(0.28)		(0.30)		(0.35)
u_{meat}		-0.08		-0.52		-0.69		1.29
		(0.47)		(0.28)		(0.30)		(0.35)
R^2		0.38		0.17		0.09		0.20
Nb	11975							
Instrument for total food consumption: average agricultural wage								

7 Conclusions

In this paper we offer an explanation for a phenomenon that has been observed in the context of a number of conditional cash transfer programs, namely the fact that in the face of the large change in total consumption that follows the injection of cash implied by these programs, the structure of consumption, that is how total consumption is allocated to different commodities changes in ways that are hard to reconcile with perceived wisdom or even with the estimates of state of the art demand systems. We suggest that this might be due to a violation of the unitary model and to the fact that the cash tranfsers delivered by these programmes are targeted to women.

Of course, the unitary model has been rejected a number of times. In this paper, we go beyond that result and test some of the implications of one of the main alternatives to the unitary model, the so-called collective model that postulates that, however intrahousehold allocations are achieved, they are such that there is no waste of resources and are therefore efficient.

We implement a test of the collective model that has been recently proposed by Bourguignon et al (2009) which requires the analysis of the demand for at least two commodities and at least two distribution factors. The idea is relatively simple: an important implication of the collective model and efficiency is that distribution factors only affect demand through the Pareto weights that defines the efficient allocation. If this is the case, then two distribution or more distribution factors have to enter proportionally or, equivalently, if the relationship between the demand for one good and a distribution factor is monotonic, one can condition on that commodity and 'explain away' all other distribution factors.

We apply this test to the context of rural Mexico and on the data set collected to evaluate teh conditional cash transfer program PROGRESA. This data set is ideal for a variety of reasons. First, the programme is targeted to women with the explicit purpose of changing the balance of power within the households that receive it. The programme itself, whose allocation is randomized across localities within the evaluation sample, is an ideal distribution factor. Second, the fact that the evaluation factor is a census within each village gives researchers the opportunity to map out the family network and allows us to construct an additional distribution factor, the relative size of husband and wife family networks. This measure is continuous and turns out to be an important determinant of the demand for food.

We use the PROGRESA data to estimate a state of the art demand system both with and without the distribution factors we considered. We first confirm that the demand system we estimate without distribution factors is not stable and is unable to predict the impact of the programme on consumption shares. The distribution factors we consider are not only significant, but seem to enter in a fashion which is not inconsistent with the implications of the collective model. Moreover, we are able, with these distribution factors, to predict the impacts of the programme much better than the standard unitary model.

Our results are important because they constitute the first test of this nature of the collective model in a context where the intrahouseholds allocation of resources is certainly salient, as witnessed by the attempt on the part of the government, to change it. Moreover, the fact that we do not reject the implications of the collective model is important because it points to a specific model of intrahousehold resources that can be used to study the household behaviour and establish the consequences of different policies.

In future work, we plan to test addtional restrictions of the collective model, and in particular those on the demand elasticities.

8 Bibliography

References

[1] Angelucci, M., Attanasio, O., (2007), "Oportunidades: programme effects on consumption, low participation, and methodological issues", Economic Development and Cultural Change, forthcoming.

[2] Attanasio, O.P. and V. Lechene (2002), "Tests of Income Pooling in Household Decisions", Review of Economic Dynamics, Vol. 5, pp. 720-748.

[3] Attanasio, O.P. and V. Lechene (2010), "Conditional Cash Transfers, Women and the Demand for Food", IFS Working Paper

[4] Attanasio, O.P., Di Maro V., Lechene V and D. Phillips, "The welfare consequence of increases in food prices in rural Mexico and Colombia"(2009)

[5] Bourguignon, F., Browning, M. and P.A. Chiappori (2009) "Efficient Intra-Household Allocations and Distribution Factors: Implication and Identification ", *Review of Economic Studies, 76,* pp 503-28.

[6] Bourguignon, F., Browning, M., Chiappori P.A. and V. Lechene, (1993), "Intra-Household Allocation of Consumption: A Model and Some Evidence from French Data", Annales d'Economie et de Statistiques, n29, Janvier-Mars 1993, p137-157.

[7] Browning, M., Chiappori, P-A., Lechene, V., (2006), "Collective and unitary models: a clarification", Review of Economic of the Household.

[8] Chiappori, P.A. (1988)"Rational Household Labor Supply", Econometrica, 56, 1988, 63-89.

[9] Chiappori, P.A. (1992)"Collective labor supply and welfare", Journal of Political Economy, 100, 1992, 437-67

Figure 1: Engel curves without distribution factors

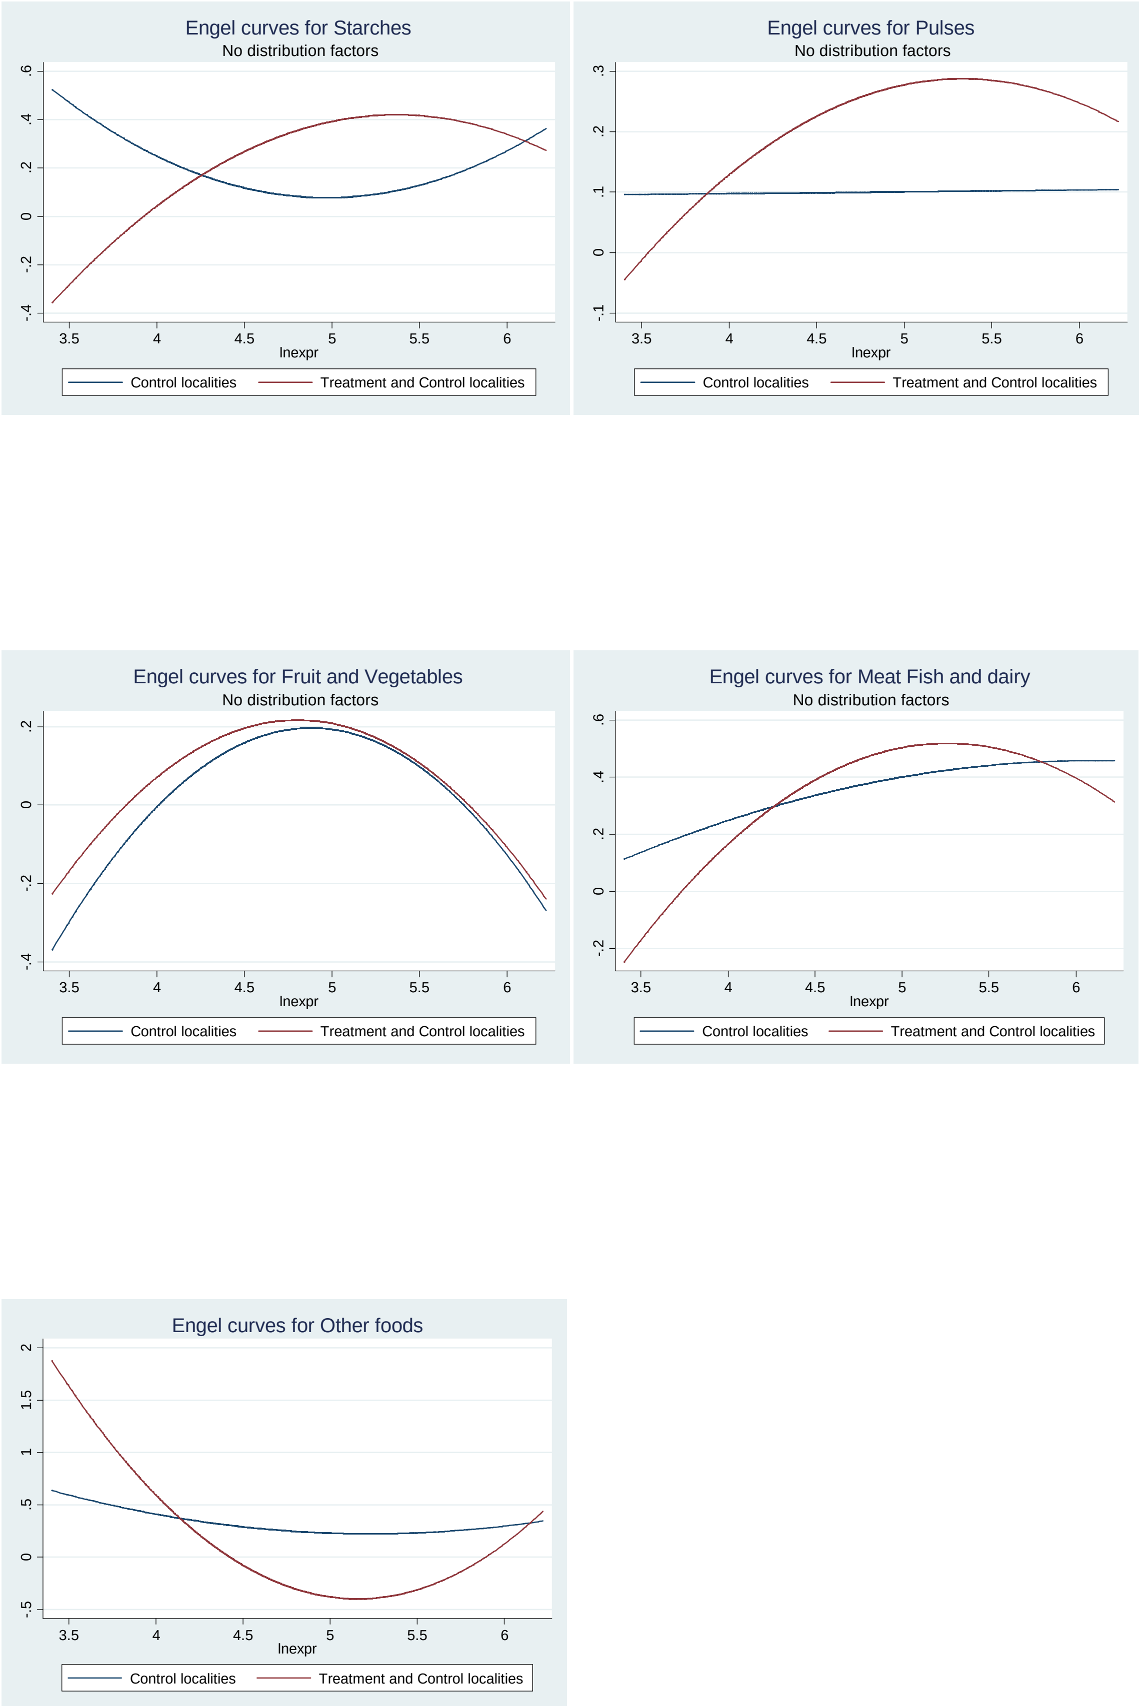

Figure 2: Engel curves with distribution factors