Food and Cash Transfers: Evidence from Colombia[♠]

Orazio Attanasio, Erich Battistin, Alice Mesnard

June 2011

Abstract

We study food Engel curves among the poor population targeted by a conditional cash transfer programme in Colombia. After controlling for the endogeneity of total consumption and for the price variability across villages, our estimates imply that an increase in consumption by 10% would lead to a decrease of 1% in the share of food. However, quasi-experimental estimates of the impact of the programme show that the share of food increases. This result is not inconsistent with the hypothesis that the programme could increase the bargaining power of women, inducing a more than proportional increase in food consumption.

Total words: 98

JEL classification : C52, D12, I38.

The description of demand patterns is one of the oldest endeavours in applied economics. And yet, many unresolved problems still make the estimation of a demand system a difficult exercise. When considering, for instance, how consumption shares vary with total consumption and prices, there is no consensus on the specific functional form to be used for the relationships to be estimated, how to address the endogeneity of total consumption, how to model the effect of

---

[♠] Orazio Attanasio: UCL Department of Economics, Gower Street, London WC1E6BT, UK, email: o.attanasio@ucl.ac.uk. Erich Battistin: Department of Statistics, Via Cesare Battisti 241, 35123, Padova, Italy, email: erich.battistin@unipd.it. Alice Mesnard: Deaprtment of Economics, City University, London, email: alice.mesnard.1@city.ac.uk. This research was financed with ESRC grant RES-000-22-0583 on "Household and Individual Consumption Decisions under Extreme Poverty". We received useful comments from Richard Blundell and Arthur Lewbel, as well as from the editor and two anonymous referees.

prices when they are not observed, what is the estimation approach that is more effective. All these issues are key for a correct estimation of demand systems and the relationship between consumption shares and total consumption. These relationships are not only of academic interest but have important implications for the design of policies.

In this paper, we study consumption patterns among poor households in rural Colombia. This study has three main goals. First and foremost, we want to characterize the demand for food for our population. This is an interesting exercise *per se* because of the very nature of the population. The extreme poverty of the households in our sample makes some facts taken for granted among other populations, in particular that the food income elasticity is less than one, questionable. Such elasticity is relevant for the design of policies aimed at improving the nutritional status of children and other poor and vulnerable individuals.

Second, as the data were collected for the evaluation of a conditional cash transfer programme in Colombia, one can assess the extent to which the structural equation defined by the demand for food can predict the changes in consumption patterns implied by the quasi-experimental variation in our sample. Since the latter can be estimated with some confidence given the way in which the evaluation was designed, we can use these results and our demand estimates to validate the latter. The identification of specific inadequacies of the demand system we estimate in predicting how the structure of consumption changes with the policy intervention might be suggestive of the channels through which the policy operates and of richer behavioural models that should be fitted to the data.

Finally, by addressing the various methodological problems we will be dealing with and by exploring alternative modelling choices, the paper gives a methodological contribution to the study of demand patterns. In particular, we will be addressing the issue of the functional form for the demand system, the appropriate instrumenting of total consumption and how to control for price differences when prices are not observed.

The main results of the paper can be summarised as follows. After investigating different econometric techniques we conclude that estimates of the structural parameters of the Engel curve obtained using a control function approach seem to be the most reliable. It is clear that taking into account the endogeneity of total consumption is important and affects in an important fashion the shape of the curve. OLS estimates seem to indicate that food is a luxury at very low levels of total consumption and become a necessity only at sufficiently high levels. This evidence, however, disappears when we instrument for total consumption. We find that food consumption is indeed a necessity for almost every household in our sample. This inference is important given that our sample is made of very poor individuals. We also show that it is important to take into account the variability of relative prices across villages (which we do not observe perfectly). However, by far the most important aspect turns out to be to control for the endogeneity of total consumption.

An issue that we discuss and about which there is no consensus in the literature is what type of instrument one should use for total consumption. The data set we use is particularly useful in this respect as it contains an interesting variable (that we will refer to as "expected income" in what follows) that seems particularly appropriate to instrument total consumption in the context of Engel curve estimation. Because of the way the survey questions were formulated, the instrument is likely to be valid even in the case of non separability between consumption and leisure choices. To the best of our knowledge, such a variable has not been used before in other studies.

Having obtained a preferred specification for the Engel curve, we use it (together with quasi-experimental estimates of the increase in total consumption) to estimate the impact of the programme on food shares. We then compare these estimates with quasi experimental estimates for the same outcome derived from the programme evaluation. By maintaining the assumption that the latter estimates are purged from any source of selection bias, we find that the two sets of results are statistically different. We argue that a possible explanation is the fact that the grant is targeted to women and therefore is likely to change the balance of power within the household and, in general, change the way choices are made. Implicit in this argument is a mis-specification of the Engel curves. We discuss possible alternatives in the conclusions.

The remainder of the paper is organized as follows. In Section 2 we introduce the conceptual framework within which we will be discussing the various estimation problems. In Section 3 we present the data we will be used in the analysis. In this section we also describe the welfare programme for whose evaluation the data were collected. In Section 4, we present the results of our empirical analysis of demand patterns in Colombia using different approaches. After having established which of the alternative approaches considered yields our preferred specification, in Section 5 we relate the impacts estimates to the estimates of the Engel curves, as discussed above. Section 6 concludes.

2. Estimating Engel curves

In this paper we study Engel curves for food, that is, the relationship between the share of total consumption devoted to food and total consumption. Such a relationship can be derived within a standard demand system. If one assumes that individual households (conceived as a single decision unit) maximize utility subject to a budget constraint, one can obtain demand curves where consumption (shares) on individual commodities depend on total consumption, prices and preference shifters that might include demographic and other variables. The tension in an exercise of this type is between equations that are flexible enough to fit the data and yet are consistent with the restrictions implied by the theory.

Deaton and Muellbauer's (1980) Almost Ideal Demand System (AIDS) has been widely used and, for a given level of prices, implies a linear relationship between consumption shares and the log of total consumption. Banks, Blundell and Lewbel (1997) (BBL from now on) have proposed a quadratic generalization of such a system (the Q-AIDS). It could be argued that the

AIDS and its quadratic generalization constitute one of the most flexible theory consistent functional forms available in the literature. Therefore, in our discussion, we use the BBL specification as a starting point.

2.1 Functional forms and price effects

As detailed in BBL, a Q-AIDS demand system can be derived from the following indirect utility function V:

(1) $$\ln V = \left\{ \left[\frac{\ln m - \ln(a(p))}{b(p)} \right]^{-1} + \lambda(p) \right\}^{-1}.$$

where m is total consumption, $\mathbf{p}$ the vector of prices and the functions $a(p)$, $b(p)$ and $\lambda(p)$ are defined as follows

By applying Shephard's lemma to (1) one can get the following share equations:

$$w_i = \alpha_i + \sum_{j=1}^{n} \gamma_{ij} \ln p_j + \beta_i (\ln m - \ln a(p)) + \frac{\lambda_i}{b(p)} \left[\ln \frac{m}{a(p)} \right]^2.$$

To these equations one can add demographics (either as affecting the intercept α_i or the price coefficients or even the coefficients on total consumption). We re-write this equation so to include a residual term u_i to reflect unobserved taste shocks and measurement error:

(2)

As discussed in BBL, the demand system in (2) combines functional form flexibility and consistency with theory, in that it is integrable. The last term in (2) makes the demand system of rank 3, the highest admissible rank for a theory-consistent demand system that is exactly aggregable, in that it is linear in function of total consumption.

BBL discuss extensively the importance of a quadratic term in the demand system, as such a term allows some commodities to be necessities at certain levels of total consumption and luxuries at others. This aspect is potentially very important in our context. We will be particularly interested in the consumption elasticity of food. We will therefore want to avoid imposing ex-ante the linearity implied by a standard AIDS system and allow for the additional flexibility afforded by the quadratic term in (2).

BBL show that any theory-consistent system is of rank 3 (and therefore allows some commodities to have quadratic terms and some not) only if the coefficient on the quadratic term is a function of prices, as is the case in equation (2). This issue is of particular relevance for us because, although the data we use to estimate versions of equation (2) are from a single cross section, they come from more than 100 small villages that exhibit a substantial amount of variation in relative prices. Moreover, as in our data price information is limited to some food items and we do not have any price information on other commodities, we cannot compute the relative price of 'food'. Therefore, we will have to work under the assumption that prices are unobservable.

One possibility, of course, is to assume the problem away. If one uses data from a single cross section and is willing to assume that prices faced by the consumers in that cross section are uniform (within and across towns), one does not need to worry about the issue of unobservable prices. It should be remembered, however, that in such a situation, the size of the coefficient on the quadratic term cannot be extrapolated to different contexts, as it would depend on the level of prices prevalent in the cross section used for estimation and would vary in different situations. Moreover, this assumption seems very strong.[1]

After rearranging terms in equation (2), we see that prices enter in three places, that is as an intercept shifts and as shifters of the coefficients of both the linear and quadratic terms:

(3)

Notice, however, that if one is willing to assume that the coefficients on the quadratic term λ_i are zero, then prices enter as a simple intercept shift.

In the absence of detailed information on (relative) prices, we consider two alternative strategies other than the simple strategy that assumes no variation in prices in the cross section. A first and flexible approach is to control for prices by village dummies. While this approach is robust, testing for the presence of quadratic terms in (2) in this context becomes problematic as the coefficient on the quadratic term becomes village specific and varies with prices. Notice that if the quadratic term has no effect on consumption shares, the estimation procedure is greatly simplified as village dummies enter only as intercept shifts.[2]

Alternatively, one can try to capture differences in relative prices across villages by means of village level variables reflecting the economic environment that are relevant for the determination of relative prices. These variables might include the size and population density of the villages, the number of shops, the altitude and the level of some representative prices on which information is available. Of course such an approach implies that all the systematic variability in relative prices across villages is captured by these variables.

[1]Attanasio and Frayne (2005) show that, in the same data we are using, there is a substantial amount of variability of unit values for individual food items both within and across villages.

[2] Price interacts with the linear term in log consumption only through the cross product in equation (3). If $\lambda_i = 0$ there are no interactions between prices and log consumption.

In what follows we will be looking at these different approaches. Obviously, if one does not reject the hypothesis that the quadratic terms are absent, the analysis, even in the presence of unobservable prices, is greatly simplified.

2.2 Endogeneity of consumption

There are several reasons why the terms in log total consumption might be correlated with the residuals of the demand system. The usual interpretation of a static system is as the second step of a two stage budgeting, where the first step determines the allocation of total consumption across time periods, and the second determines the allocation within the period. If heterogeneity in intertemporal preferences are correlated with (unobserved) taste shifters in the demand system, one would obtain that the residuals of the latter are correlated, across individuals, with the allocation of resources over time and therefore with log consumption. It is possible, for instance, that individuals that have a relatively stronger preference for food are also relatively impatient and therefore have a higher level of current consumption as well as a high share of food consumption.

Another reason for the possible correlation between residuals and log consumption is the presence of measurement error. A useful source of exogenous variation in this context, therefore, may come from a variable that explains the cross sectional variability of log consumption but is unlikely to be correlated with taste variables and/or with measurement error.

In the literature, there is no strong consensus on the appropriate variable to use for instrumenting total consumption. Many studies use total household income. However, if labour supply enters the utility function in a non-separable manner, income might be correlated with taste shifters in the same way as total consumption is. Moreover, in the presence of large transitory shocks, current income can constitute an inefficient instrument for total consumption even if it is uncorrelated with taste shifters.

A possible alternative to the use of income is the use of wages, which may be considered as a price that the individual household takes as given An even more conservative stance would be to use village level wages as instruments for total consumption of an individual household and is unlikely to be correlated with measurement error or taste shifters. Such an approach has been tried, for instance, by Attanasio and Lechene (2002). Wages, however, are an invalid instrument if leisure and consumption are not separable. Moreover, as our analysis allows consumption patterns to vary depending on village prices, we cannot use village wages as an instrument for total consumption.

Our survey contains a variable that can proxy for household expected income, using information relating to the variation in households' future income stream. This variable is constructed using two questions on the lowest and highest income a household is expecting to receive in the next month.[3] This variable has many advantages: given the way these questions are asked, these bounds are exogenous to labour supply choices. Moreover, as the variable we try to construct captures *expected* income, it should be uncorrelated with transitory shocks. In Section 3, we will show some descriptive statistics on this somewhat unusual variable and show how it co-varies with income and, most importantly, with total consumption. When estimating the food Engel curves, we will discuss how such a variable performs in the first stage and some potential problems that may arise when using it as an instrument for total consumption.

The presence of quadratic terms in equation (2) introduces additional problems to the instrumenting approach. Once the instrument has been established, one can use powers in this instrument to take into account the presence of non linear terms. Such an approach, however, often yields very imprecise estimates. To overcome this problem we adopt an alternative strategy based on a control function (CF) approach as proposed in this context, for instance, by

[3] The two questions are: (1) Suppose that next month the members of your family who want to work, get a good job (alternatively: imagine the harvest is good). How much money do you think would be earned/would come into the household in that month? (2) Now imagine the opposite: that they have very little work next month (alternatively: imagine the harvest is bad) and that they have just this to live on, as well as what people give them (which is very little). How much money do you think the household would receive in that month?

Blundell, Duncan and Pendakur (1998). Accordingly, one uses the residuals of the first step regression for total consumption to control for the endogeneity of this variable in equation (2) by introducing a polynomial in the residuals as additional regressors.

2.3 Unitary or non unitary models?

The structure described above is derived under the assumption that decisions are taken by a single decision unity that maximizes a well defined utility function. In what follows, we will suggest that such a representation might not be accurate, as the consumption patterns of poor Colombian families might be the result of the interactions of more than one decision maker. A model that has been proposed to deal with these issues is the so-called collective model of Chiappori (1988) which imposes the restriction that decisions are made in an efficient fashion. Browning and Chiappori (1998) have studied some of the features of household demand systems that emerge from such a framework and used a QAIDS model very similar to the one we use to exemplify their results.

In the collective model, efficiency implies that the household maximizes a weighted average of the utility functions of the household members with weights reflecting the relative power that the different members have within the family. We will argue below that a conditional cash transfer targeted to women might shift the weights in favour of women and therefore change the nature of the demand system. Browning and Chiappori (1998) show that under certain circumstances such a shift can take the form of a change in the intercept and possibly the slope of the Engel curves.[4] If one does not allow for the effects of what Browning and Chiappori (1998) call 'distribution factors', of which a conditional cash transfer can be one, these would be reflected in changes in the unobserved component u in equation (2).

[4] The coefficient on prices, which we do not estimate explicitly here, also changes.

3. Consumption in rural Colombia.

The survey we use for the estimation was collected to evaluate a welfare programme sponsored by the World Bank and the Inter American Development Bank. The programme, modelled after the Mexican PROGRESA, consists of conditional cash transfers targeted to poor households living in small towns with certain features. To evaluate the effects of the programme, two waves of a large data set were collected over a period of one year. In the first part of this section we describe the programme. We then move to describe in detail the dataset that we use in the estimation of Engel curves. We finally present the effects of the programme on household consumption.

3.1 The Familias en Acción programme

In 2001 the World Bank and the Inter American Development Bank decided to sponsor a large welfare intervention in Colombia, inspired by the Mexican conditional cash transfer programme PROGRESA. As in PROGRESA, the *Familias en Acción* (*FeA* from now on) programme consists of three components: health, nutrition and education. The 'nutrition' component is a cash transfer (40,000 pesos per month or 15 US$) eligible households receive if they have children under the age of 6 and participate into the health component. The latter consists of a number of growth and development check ups for young children, a vaccination programme and some courses for mothers on various health issues. The education component consists of grants for school age children that are received by each child who attends regularly school. The grant is 14,000 monthly pesos (5.5 US$) for primary school children and 28,000 (11 US$) for secondary school children. As in PROGRESA, the money is received by mothers.

The programme was first targeted geographically. Of the 900 odd municipalities in Colombia, 627 were chosen as targets. The targeted municipalities had to have less than 100,000 inhabitants, they could not be department capitals, had to have enough education and health infrastructures, had to have up-to-date lists of welfare recipients and had to have a bank in town. Within each municipality eligibility was established using the so-called SISBEN indicator. SISBEN is an indicator of economic well-being that is used throughout Colombia for targeting welfare programmes as well as for the pricing of utilities. In theory, each Colombian household is classified periodically in one of six levels, on the basis of an indicator determined by the value of several variables periodically measured. In the case of *FeA*, only households in the first level of SISBEN as of December 1999 were eligible. Eligible households, which in what follows will be referred to as SISBEN 1, constitute approximately the bottom twenty percent of Colombian households living in rural areas (see Vélez, Castaño and Deutsch, 1998).

The programme started, with a few exceptions, in the second half of 2002 and the take up among eligible households was over 90%. By 2003 about 340,000 households throughout Colombia were covered by the programme. *FeA* was subsequently expanded to larger towns and as of the end of 2008, more than 1.5 million households were involved in it. The programme is now very visible and probably constitutes the largest social intervention in Colombia. For the households in our sample, the grant received constitutes typically about 20% of household monthly consumption for participant households, and is thus likely to have an important impact on their consumption.

3.2 The data set

As the *FeA* programme was started, the Colombian government decided to launch a large scale evaluation of its effects. The evaluation work started with the collection of a large scale data base in 2002. The evaluation is based on the comparison of SISBEN 1 households in municipalities targeted by the programme (hereafter 'treatment areas') to SISBEN 1 households living in 'control' municipalities. As the random allocation of the programme was not feasible, the evaluation survey was constructed by first choosing a stratified random sample of targeted communities. The stratification was done on the basis of geographic areas and the level of health and education infrastructures, for a total of 25 strata. Within each of these strata, the evaluation team chose 'control' municipalities that were as similar as possible to the municipalities included in the 'treatment' sample in terms of size, population, an index of quality of life as well as health and education infrastructures. In each of the municipalities in the sample, 10 geographic clusters were randomly drawn, with weights proportional to the population, of which three clusters are in the urban centre (*cabecera municipal*) and seven are more rural. Finally, in each of the clusters, about 20 households were drawn from the SISBEN 1 lists. Given non-response rates and mobility about 10 households per cluster entered the final evaluation sample, which was, in the end, made of about 11,500 households living in 122 municipalities, of which 57 were 'treated' and 65 used as 'controls'.

Political pressure resulted in the programme starting in 26 out of 57 treatment municipalities before the baseline was collected. For this reason, at baseline, we have two types of treatment municipalities: the 'early treatment', where the programme was already operating and the 'late' treatment, where it was not.

The 11,500 households in the baseline survey were interviewed in 2002. A year later, in 2003, after the programme had started in all treatment municipalities, the same households were followed up and re-interviewed. The presence of a baseline and a follow up survey allows the evaluation to be based on a difference in difference approach, which can be combined with matching methods to take into account observable differences. The attrition rate between the baseline and follow up survey was only 6%, partly because of the low mobility of our population and partly because of the effort made in following up movers..

For the estimation of the Engel curves, we use data from a *baseline* survey -collected *before* the programme started- in order to investigate household consumption patterns that are not affected by the programme. From the baseline sample we exclude households living in 'early treatment' municipalities, as they were already affected by the program. We do make use of data from the baseline and follow-up surveys (collected *before* and *after* the implementation of the programme on the *same* households) in both treatment and control municipalities to compare the estimated effect of *FeA* on consumption to the implications of Engel curves on the same outcome variable.

The data set includes a large set of variables. In addition to information on family composition and the like, we have detailed information on consumption and, in particular, food consumption. This includes information on about 100 different commodities. Respondents are asked to report expenditure and consumption on each of these commodities during the week preceding the interview. Information on consumption is collected to include consumption in kind (either produced or received in payment or gift). Consumption on non food items is recorded at different intervals, to avoid a large number of zero due to the infrequency of purchases. In what follows, all information on consumption is converted to the equivalent amount they represent in monthly pesos and is reported in Colombian pesos. At the time of the survey, one dollar was worth about 2,600 pesos.

We kept in our working sample only households for which we do not have missing information for expected income, which, as we mentioned above, we use as an instrument for total consumption. Our chosen instrument is missing for about 20% of the sample, thus reducing the size of our final sample to 5,218 households.[5] However, even if data are missing in some systematic fashion, as long as the instrument is valid, this would not result in any bias to our estimates (although it may raise some issues related to the support, which we address below). Moreover, the results obtained by estimating Engel Curves via OLS before and after this selection step turned out very similar. Finally, by comparing changes in consumption from before to after the implementation of the programme across treatment and control areas using a difference in differences approach (see Section 5), the results we find are similar to those in Attanasio and Mesnard (2006). This we took as informal evidence to rule out selection issues in the data due to the missing instrument.

[5]As a consequence of this selection, tests showed that variables such as the number of children aged 7-11, number of adults above 60, head having less than a primary level, total consumption and total food are significantly different between the sample we use and the full sample.

3.3 Descriptive statistics

Tables 1a and 1b summarize area and household level variables for our final sample. From now on we will consider geographical clusters defined by *villages* instead of municipalities, as some neighbouring municipalities that are very small and adjacent were grouped together to form the same cluster for the analysis. This geographical definition of areas led us to 75 villages in our working sample, 25 of which were treated. In Table 1a we notice that about 55% of our sample lives in the '*cabecera municipal*' (the urban centre of the municipality). We will be defining these as 'urban' households, although the villages in our sample are relatively rural and small, the average village having less than 30,000 inhabitants. The location of the villages in our sample is spread all over Colombia, with a relatively smaller proportion of villages (12%) in the Pacific Region. The average altitude of our villages is about 650 meters above sea level. However, there is a large dispersion around this mean (750 meters) reflecting the large geographic diversity of Colombia. A large proportion of households in the villages in our sample are not connected to the sewage system (44% on average) and 13% of them do not have piped water. These numbers are indicative of the relatively high poverty levels in these villages. It is worth noting that, perhaps not surprising given the way control municipalities were chosen for the evaluation, the distribution of the variables reported in Table 1a turned out the same in treatment and control villages: a binary regression of the indicator for treatment villages on the full set of village characteristics led to a p-value of the F statistic for their joint significance of about 32%.

Mean, standard deviation and different percentiles of the distribution of household characteristics are reported in Table 1b. On average, the households in our sample have 6 members. The large majority of them do not have a good health insurance, as only 5% of them benefit from an unsubsidised health insurance, which is typically associated to having a good job in the formal sector. Moreover, only 4% of their heads or spouses have more than a secondary education level and just above 20% of them have less than a primary education. Only few household characteristics were found to significantly explain the probability of living in treatment villages, the latter being characterized by a larger proportion of households with spouses having less than a primary education level and children under 7 years old.

Finally, in the bottom part of Table 1b we also report the mean, 25^{th}, 50^{th} and 75^{th} percentiles of the distribution, at baseline, of total consumption, food consumption and the share of food. Food is a very important component of consumption of the very poor households in our sample, with an average share of about 70%. Note that these figures do not include only consumption, as we imputed the value of commodities consumed but not purchased (i.e. produced or received as a gift) using information on food prices in the municipality. This measure of 'food-in-kind' represents around 18% of food consumption for the 75% of households in the sample who report consumption-in-kind. All these household characteristics point to the fact that the households in our sample are amongst the poorest households in rural Colombia.

As expected future income constructed from subjective expectations is somewhat unusual, it is important to check how this variable co-varies with more standard indicators such as income and consumption. The correlation patterns amongst these variables are investigated in Table 2 where we report the results of several regressions. Columns (1) and (4) show that expected future income co-varies significantly with both income and consumption. Moreover, the same pattern is confirmed in Columns (2) and (5), after netting off the effect of a large number of controls, such as demographics and indicators for the education attainment of the household heads. Finally, in Column (3) we report the results of a regression of total consumption on both income and expected income, including the same set of controls used above. As in what follows, we use expected future income as one of our instruments for total consumption, this equation represents our first stage regression. While we do not report the complete set of results for the sake of brevity, it is worth noting that the regression has an R-squared of 0.26 and the t-statistics for income and expected income are 16 and 14, respectively. The significance of these coefficients is obviously important for identification.

4. Estimating Engel curves for food

In this section, we present the estimation results for the food Engel curve. The purpose of this exercise is twofold: on the one hand we want to characterize the behaviour of poor, rural Colombian households. On the other, we want to check whether we can predict the impact of the programme on food shares using our (quasi-experimental) estimates of the impact of the programme on total consumption and a well specified structural model. As we want to use our structural estimates to predict the effect of the conditional cash transfer programme on food shares, we estimate our Engel curves using a sample of households who do *not* receive the cash transfers. In this sense, our 'predictions' are truly out-of-sample predictions.

It is easy to predict that a conditional cash transfer increases household total consumption, consumption although it is important to quantify this effect precisely. To predict the effect that this increase in total consumption has on the structure of consumption, knowledge of the shape of the Engel curves is crucial. We will therefore devote much effort to establish whether, for the case at hand, food Engel curves are better approximated by a linear or quadratic specification. In other words, we want to establish whether food might be a luxury at some levels of total consumption.

To answer these questions in a rigorous fashion we need to deal with a number of methodological issues. In particular, we want to deal with the possibility that total consumption consumption is endogenous and we want to control for the possibility that (unobserved) relative prices differ across localities. Our approach to these issues is mainly empirical, in that we will test on our data which of different alternatives fits the data best. We will be comparing linear and quadratic specifications and we will be comparing estimates obtained under the assumption that total consumption is exogenous to those we obtain allowing for endogeneity. We will also report tests of exogeneity of total consumption. To deal with the possibility of heterogeneity in relative prices, we will be reporting estimates obtained with three different approaches to the problem of unobservable prices. This gives us a total of twelve different specifications.

We summarize our results in Table 3a, which focuses on the coefficient(s) on total log consumption and (in the case of the quadratic specification which implies heterogeneity) on the distribution of derivatives of the food share with respect to log consumption. In Table 3b, we present a sensitivity analysis of our results to check their robustness with respect to the estimation method employed and the sample selection criteria adopted.

4.1 Basic specification: no price effects and exogenous consumption

As we have mentioned above, we use a Q-AIDS model for the share of food. In particular, we will be interested in whether for a sizeable fraction of the population food is actually a luxury, so that its share increases with total spending. A finding of this nature could potentially explain why, on average, the share of food stays constant or increases slightly with increases in total consumption.

We start by re-writing equation (3) under a very strong assumption, namely that consumers in our cross section face the same relative prices, which we normalize to one, regardless of the town in which they live. This implies the following specification for the food share:

(4) $$w_f = \alpha_f + \beta_f \ln m + \lambda_f \ln m^2 + \theta_f ' z + u_f ,$$

where the *f* index stands for food and we omit the individual index for simplicity. The vector z includes controls, such as demographic variables, that enter the demand system as determinant of the intercept of equation (3).[6]

If we are willing to assume that total consumption is uncorrelated with the residual term u_f , we can estimate equation (4) by OLS. We report the results of such an exercise in Column (1) in the bottom panel of Table 3a. In the presence of a (significant) quadratic term, equation (4) implies that the slope of the Engel curve changes with total consumption. For this reason, in

[6] We also experimented with the possibility that the demographic variables enter the 'constant' of the price index a(p) of the same equation (3). In this case, as with the price effects, the demographics were interacted with the log of total consumption and its square. The results we obtained with this richer specification, which we do not report for the sake of brevity, were very similar to those presented here.

the bottom panel of Table 3a, we report some percentiles of the distribution of the first derivatives of the Engel curve implied by our estimates and the distribution of log consumption. Finally, in Figure 1, we plot how the share of food varies with (logged) total consumption according to the estimates reported in Table 3a. Here and in what follows, in plotting the Engel curve, we set the intercept at an arbitrary point, so that the only relevant information is the shape of the curve.

A remarkable feature of the Engel curve estimated by OLS is that both coefficients on the linear and quadratic terms are strongly significant. The presence of a significant and sizeable quadratic term and the distribution of total consumption imply a substantial amount of variability in the first derivatives of the Engel curve: the interquartile range of the slopes in column (1) is 0.034. Although the confidence interval that would result from the precision with which the two coefficients are estimated is quite wide, the estimates do imply that the share of food increases at low levels of total consumption and starts decreasing at levels that are close to the 10th percentile of consumption in our sample (between 5 and 5.5 in Figure 1). Effectively, as reported in the last row of column (1) in Table 3a, these estimates imply that food is a luxury for almost 13% of our sample and a necessity for the remaining households.

As a reference, the first column in the top panel of Table 3a reports OLS estimates of the coefficient on log total consumption when fitting a linear specification. Ignoring the quadratic term yields a negative coefficient on total consumption. The coefficient is estimated at -0.024 and is statistically different from zero.

This result appears to be a feature already discussed in other studies that use data from developing countries. For example, Kedir and Girma (2008) using data from the Ethiopian Urban Household Survey find that for a non-negligible proportion of households the share of food increases at low levels of the total consumption distribution. There are several possible explanations for this finding. First, it could be that for these very poor households food is indeed a luxury: the necessities might be constituted by housing consumptions (rent, utilities) and the rest of what they consume goes into food. Increases in total consumption, therefore, are translated into increases in the share of food, as these households increase the amount of food consumed and possibly start switching from diets based almost entirely on basic staples (rice, potatoes and so on) to increase the frequency with which they consume foods rich in proteins (chicken, beef and so on). A second possible explanation could be measurement error for the households reporting very low levels of total consumption. In their study, Kedir and Girma (2008) use an approach suggested by Lewbel (1996)[7] and find that the curvature of the Engel curve is robust to the presence of measurement error in the data. By adapting the same approach to our context with and without price heterogeneity we found no significant effect of measurement error on the results presented in this and the next section, thus concluding that the non-linearity of food Engel curves estimated by OLS is robust to the presence of such non classical type of measurement errors (results are available on request). A third explanation may be that the shape is induced by the fact that we are ignoring endogeneity of total consumption and possible price effects. It is to these issues that we now turn, starting with the possibility that ignoring heterogeneous prices can introduce significant biases.

4.2 Heterogeneous prices across villages

To account for the possibility that relative prices are different in different towns, we use two different approaches. First, we proxy relative prices with village level dummies; second we parametrize relative prices with a number of village level, in particular the prices of some common goods for which we have price information.

4.2.1 Proxing prices by village dummies

As it is clear from equation (3), in the case in which the coefficient on the quadratic log consumption term is non-zero, prices interact both with the linear and quadratic terms. We therefore estimate the following version of equation (3):

[7] Errors in the measurement of food consumptions induce a complicated form of non-classical measurement error in the estimation of Engel curves, as both the right hand side and left hand side variables are error ridden in a complicated fashion. IV does not eliminate the biases induced by this type of error.

(5) $$w_f = \alpha_f^v + \beta_f^v \ln m + \lambda_f^v \ln m^2 + \theta_f\,'z + u_f$$

where the v superscript stands for village. The village specific coefficients in equation (5) are estimated adding to the regression village dummies and their interactions with the linear and quadratic consumption terms. We use the same vector of controls as in the previous specification and, as before, we ignore the possible endogeneity of $\ln m$ and estimate equation (5) by OLS. We parametrized the village dummies so to interpret the coefficients on the linear and quadratic terms as the average coefficients in the sample. We report them in Column (3) of Table 3a and plot the profile of the Engel curve in Figure 1. As with column (1), the top panel reports the average coefficient for a linear specification, while the central panel reports the average coefficients for the quadratic specification and the percentiles of the distribution of first derivatives. As with column (1), the slope of the Engel curve is different for different consumers and, for that reason, in the bottom panel of the Table we report some percentiles of the distribution of slopes. We should notice that, unlike in column (1), the slopes are heterogeneous not only because of the variation in log total consumption but also because of variability in the coefficients across villages. As before, we compute the individual level slope and report the moments of its distribution as well as the percentage of households with a negative slope.

There is substantial variation in the estimated coefficients across villages. To document this heterogeneity, in Figure 2, we plot the deviation of the village coefficients (that is, the coefficients of village dummies and their interaction with the linear and quadratic terms in log consumption) from their average. Many of these coefficients are statistically different from zero. Remarkably, however, in both the linear and the quadratic specification, the average coefficients in column (3) are not too dissimilar from those in column (1). As a consequence, the shape of the implied Engel curve implied by the average coefficients is not very different from that implied by equation (4) – See Figure 1.

While Figure 2 indicates a substantial amount of heterogeneity, it is not very informative about the variation in the shape of the village level Engel curves, which is ultimately what we are interested in. Rather than plotting the 75 Engel curves implied by these coefficients, we focus on the distribution of the first derivative of the Engel curve implied by the estimated parameters. Relative to Column (1), the slope of the Engel curve varies not only because total consumption varies across households but also because the coefficients of the Engel curve vary across villages. In our sample, both the mean and the median first derivative of the Engel curve across households are negative (at -0.030 and -0.033 respectively). The variability of the first derivatives in Column (3) is substantial and is substantial larger than that in Column (1): the interquartile range is now 0.085 (compared to 0.034 in Column (1)). A substantial fraction of the households (29%) has a positively sloped Engel curve for food. This is considerably more than the percentage in Column (1).

4.2.2 Parametrizing relative prices

Our second approach to take into account heterogeneity in prices across villages is to assume that the variability in the vector of relative prices across villages can be completely spanned by a linear combination of a vector of nominal prices of certain important commodities and some village level variables. In particular, we assume that the variability across villages of the price indexes in the demand system can be controlled for by the log price of potatoes, rice, coffee as well as the average level of men wages and some other village level variables. We chose potatoes, rice and coffee as we have good quality data on their prices, they are widely consumed by most households and their prices are not too correlated. Village level variables include population size, altitude and its square, an index of quality of life in 1993. We therefore estimate the following equation:

(6) $$w_f = \alpha_f(\xi_v) + \beta_f(\xi_v)\ln m + \lambda_f(\xi_v)\ln m^2 + \theta_f{}'z + u_f ,$$

where ξ_v is a vector of village level variables, including the representative prices mentioned above. As in equation (5), we allow the demographics to enter only the intercept of the demand for food and not the price indexes, although allowing for these effects does not change the results substantially. We express the village level variables ξ_v in terms of deviation from the mean, so that we can interpret the coefficients on log total consumption and its square as the average coefficient across villages. We report estimates of some of the parameters in equation (6) obtained by OLS in Column (5) of Table 3a. As in Columns (1) and (3), we consider both a linear and quadratic specification and the coefficients on log total consumption are averages across the villages.

Once again, there is evidence that price heterogeneity plays a statistically significant role in the specification we estimate. The coefficient on the village level variables, including the prices of several important commodities, are statistically significant. However, the average coefficients in Column (5) are even more similar to the coefficients in Column (1) than those in Column (3). The average Engel curve estimated by OLS assuming that price heterogeneity is approximated by a vector of observable variables are virtually identical to the case in which price heterogeneity is ignored: for this reason we do not plot the average profiles in a new figure. The degree of heterogeneity in the slopes of the Engel curve is also more similar to that observed in Column (1) than to the figures in Column (3): the interquartile range stands at 0.034 (compared to 0.034 and 0.085 in Columns (1) and (3), respectively) The only noticeable difference relative to Column (1) is that the percentage of households with a positive slope of the Engel curve is slightly higher (14.1% versus 12.6%).

The conclusion of this section is that while there is evidence of heterogeneity in prices that manifests itself in significant coefficients on either village dummies or in village level variables that include prices, this does not change the evidence on the shape of the Engel curves and on the fact that for a sizeable fraction of households food seems to be a luxury. Between the two specifications with heterogeneous prices we tend to favour the second, where prices differences are assumed to be captured by a number of village level variables. We suspect that the much of the variability observed in the specification with village dummies reflects low precisions induced by our attempt to estimate a large number of coefficients.

4.3 Allowing for endogenous total consumption

So far we have assumed that the log of total consumption is uncorrelated with the residuals of the share equation. We now allow for the presence of correlation, which can be caused by any of the reasons discussed in Section 2.2. To obtain consistent estimates of the coefficients of interest we used a CF approach, although we also considered the possibility of using Instrumental Variables (IV). In a linear model the two methods deliver the same estimates, in a non linear case the same is not true. The results we obtain with the two methods are not very different, the results from CF being slightly more precise.

Both IV and CF imply the choice of specific instruments for the first stage model for the log of total consumption. Different models justifies the use of different instruments: for instance, if one assumes that labour supply and the various commodities modelled in the demand system are not separable in the utility function, one cannot use wages as an instrument for total consumption. One also needs to consider the variability of the instruments and their ability to span the relevant support of the endogenous variables. In our exercise we experimented with three instruments: average town wages, as in Attanasio and Lechene (2002), expected future income, which is available in our data set, and total household income, which is commonly used in the literature.

From a theoretical point of view, our favourite instrument is the one based on expected future income. Because of the way the questions about max and min future income are formulated (as discussed in Section 2.2), the variable we construct is independent of labour supply behaviour and, therefore, would not be affected by non separability between leisure and consumption in the utility function. Moreover, expected future income should not be affected by temporary shocks to income that might introduce other problems to the use of such a variable as an instrument for total consumption. However, when we used it as an instrument, we faced an important issue.[8] As we are interested mainly in the shape of the Engel curve at low levels of total consumption, where the OLS estimates indicate the possibility that food is a luxury for the poorest households, we want to make sure that the results we obtain are not driven by the failure to span the values of total consumptions in the lower tail of the distribution. For this reason, for each set of instruments we use, after running the first stage regressions we plot the density function of actual and fitted log of total consumption (i.e. the endogenous variable). Although in what follows we report the results obtained from CF, we think this exercise is informative about the ability of an instrument set to span the variable being instrumented.

In the top panel of Figure 3, where the fitted value is obtained using only expected income, we notice that the support of this variable is much narrower than that of actual consumption: there are virtually no observations over the range between 4.5 and 5, which is where we observe a positive slope of the Engel curve according to the OLS estimates. For this reason, we also consider using total household current income as an instrument for total consumption. We repeat the same exercise performed to construct the graph in the top panel of Figure 3 using both variables as instruments and obtain the bottom panel of Figure 3. Now, although the density of the fitted value of total consumption is more concentrated in the middle of the support, the tails of the distribution cover approximately the same support of actual total consumption, including low values of total consumption.

Variation in the support that depends on the instruments set is not necessarily a problem for the internal validity of our estimates, but rather for their external validity. Given that our interest lies in the shape of the distribution at low levels of consumption, we will report results obtained

[8] We are grateful to the editor for drawing out attention to this issue.

using both instruments. We will comment on how these results change when we use only expected income (see Table 3b).

To implement the CF approach we first run a regression of log total consumption on the same controls included in the share equation and the instrument. We then add to equations (4), (5) and (6) a third order polynomial in the estimated residuals of the first stage regression. The results do not change if we add higher powers of the residuals. Columns (2), (4) and (6) of Table 3a report the CF estimates corresponding to the specifications estimated by OLS in columns (1), (3) and (5) respectively. Also reported is a test of endogeneity, derived by considering the joint significance of the coefficients associated to the polynomial in the residuals.

Starting with column (2), we notice that, while the quadratic term is still significantly different from zero, over the relevant range the shape of the Engel curve is not too different from a linear and decreasing Engel curve. This is evident from Figure 4, which plots the three Engel curves estimated by CF. Moreover, the estimates in column (2) imply that for no households food is a luxury. As with the OLS case, we also estimate a linear version of the Engel curve, whose coefficients are reported in the top panel of Table 3a. In the case of column (2) we see that the coefficient is considerably larger in absolute value than the corresponding OLS estimate.

Similar evidence emerges from the specifications, in column (4) and (6), that control for price heterogeneity using the same approach used above. Especially in the case of column (6) where relative prices are parametrised as a function of some village level variables, the results are virtually identical to those in column (2). Moreover, 98% of households have a negative slope of the Engel curve. The picture is less clear in column (4) where the percentage is a bit lower at 83.1%. We suspect however, that this is induced by greater noise and heterogeneity induced by the village level dummies when we estimate a quadratic specification. Notice that when we estimate a linear specification with village level dummies, then the estimates of the Engel curve are not very different from those reported in columns (2) and (6).

In Table 3b, we analyse two issues: to what extent our result vary when we limit our estimation sample to a subset of the support of log total consumption and with different choices of instruments. The first issue is particularly relevant because we are interested in the shape of the Engel curve and we would not want that to be affected unduly by few observations in some relatively unusual values of log total consumption. The second issue is interesting because our instrument is somewhat unusual. For the sake of brevity, we report the analysis only for the specification in which prices are parametrized using village level characteristics.

We start from the results obtained via OLS and CF run on the subsample where log expenditure is between about 5.3 and 6.2, that is the middle range of the distributions we have plotted in figure 3, where there is a lot of support in the first stage predictions. These results are in columns (1) and (2) of Table 3b. Given the reduced number of observations and the limited variability of log total consumption, obviously our estimates become much less precise. This is particularly true for the quadratic specification. We notice, however, that the slope of the Engel curve estimated by CF is, if anything, more negative than in the corresponding specification in Table 3a. Our main conclusion, therefore, is not affected by limiting the support of log total consumption.

We then consider in the next two columns the robustness of CF estimates to alternative sets of instruments. In column (3), rather than considering the average future income, we consider separately expected income in the best case and worst case scenarios. In column (4), instead, we consider as the only excluded instrument expected income. We see that the results do not change significantly relative to what we have reported in Table 3a.

To test formally that the shape of the quadratic Engel curves estimated by CF (rather than OLS) is not too different from that of linear Engel curves, in Table 4 we compare the shapes predicted from a linear and a quadratic model for the three specifications in columns (2), (4) and (6). In particular, for each of several percentiles of the consumption distribution, we compute the difference between the slope of the linear and quadratic Engel curve and report the 95%

bootstrap confidence interval for this difference.[9] This statistics essentially tests whether the shape that results from a quadratic model is appreciably different from the slope coefficient in a linear specification, and does so at different points of the expenditure support.

In column (1) and (3) of Table 4, which correspond to columns (2) and (6) of Table 3a, we observe that the slope of the quadratic Engel curve is significantly higher (meaning less negative) at low percentiles of the consumption distribution and lower (more negative) at higher percentiles). However, such differences are not very large, so that a linear approximation of the Engel curve would not make much violence to the data. In the case of column (2) of Table 4 (corresponding to column (4) of Table 3a, where price heterogeneity is proxied by village level dummies) we do not observe any significant difference in the slopes of the linear and quadratic Engel curves over the relevant range of consumptions. This confirms our intuition that the estimates of the quadratic Engel curves with village specific slopes are quite noisy.

We conclude this section by noticing that a linear specification seems the appropriate one for all three cases considered.

4.4 Preferred specification.

Establishing whether the Engel curve we have been studying is linear or quadratic is important for several reasons. First and foremost for our context, a linear (or monotonically decreasing)

[9] Our procedure essentially tests the "general" model:

against the "restricted" model:

which in fact amounts to testing the hypothesis . In particular, the procedure employed contrasts the shape of the Engel curved obtained under the alternative hypothesis, that is:

with the shape obtained under the null hypothesis, that is . The test statistic is constructed as:

and is computed at different percentiles of the distribution of logged consumption. Reported in Table 4 are confidence bands that result for the last quantity across percentiles.

Engel curve for food could not explain the evidence we discussed in the Introduction (and that we present for our context below) that an increase in total consumption induced by a Conditional Cash Transfer does not change or even increases the share of food in total consumption. A second advantage of a linear specification is that, even in the presence of heterogeneous prices across villages, we do not need to interact their proxies with the total consumption terms. Intercept shifts will be sufficient.

The evidence we have presented in Tables 3 and 4 seems to indicate that a linear specification is not too strongly at variance with the data. Moreover, we also know that when we use a set of instruments that might be more credible in certain contexts, than we do not have any evidence of an even partially upward sloping Engel curve. We therefore conclude this section by saying that a linear specification provides an adequate specification of the Engel curves for our sample.

5. Changes in the structure of consumption: quasi-experimental evidence and structural predictions

The aim of this section is twofold. First, we estimate the *causal* effect of the transfers made by *FeA* on the food share exploiting the quasi experimental design of the evaluation data. We then investigate if the *structural behavioural* model we have estimated in Section 4, that is the food Engel curve, is able to predict this effect.

As the programme led to an increase in total consumption and the curve, in our preferred specification, slopes down monotonically (once endogeneity has been taken into account), one would expect the food share to decrease for beneficiary households. We instead show that this *ex ante* expectations is at odds with the evidence on the causal effects of the programme on the food share.

In its bare essential, this comparison can be interpreted as an over-identification test of the structural model. Throughout our discussion, we maintain the assumption that the estimates of the effect of the programme on the food share and total consumption can be given a causal interpretation, and that the quasi-experimental methodology we use overcomes the problem that arises from non-random allocation of the programme. We also maintain a linear specification for the structural model represented by the Engel curve. As we discussed in the last section, this was our favourite specification. Finding discrepancies between *ex-ante* and *causal* results indicates that the estimates obtained in Section 4 might not be structural, in that the model is not able to predict the impact of the programme. However, we want to go beyond that and present evidence on how the programme affects the parameters of the structural model we estimate. We argue that this evidence is suggestive of the channels through which the effects of *FeA* took place.

5.1 Causal effects of the programme on household consumption

A first step in our analysis is to quantify the effect of the programme on total consumption and food shares. As the programme was not allocated randomly between treatment and control municipalities, we need to control for differences among them. As we have information on both sets of municipalities before and after the programme we combine matching with a difference in difference approach as in Attanasio and Mesnard (2006), who look explicitly at the effect of the programme on consumption.

Let $w_j(t)$ and $m_j(t)$ be the potential outcomes for the share of food and total consumption that are observed whether or not the programme is in operation (see Rubin, 1974), where t=1,2 denotes baseline and follow up periods, respectively, and j=0,1 "no programme" and "programme" regimes, respectively. There are *four* groups of observations defined by the evaluation design resulting from the combination of control and treatment villages denoted, respectively, by *d*=0 and *d*=1, in the baseline and follow-up periods.

Using this notation, in what follows we will estimate the following causal parameters:

$$\Delta_m = E[\ln m_1(2) - \ln m_0(2) \mid d = 1],$$
$$\Delta_w = E[w_1(2) - w_0(2) \mid d = 1],$$

which we interpret as the causal effect of the programme on (logged) total consumption and on the budget share for food, respectively. In obtaining these parameters, we allow for the effect of a number of control variables X on the outcomes of interest and for possible impact heterogeneity across different X's. The analysis yields correct causal impacts if the following conditions:

(7)

hold conditional on X. This amounts to assuming that, conditional on the control variables X, there would have been no differential trends in the outcomes between control and treatment villages in the absence of the programme. The validity of the conditions in (7), which is discussed at length in Attanasio and Mesnard (2006), will be assumed in the remainder of this section.

As in Attanasio and Mesnard (2006), we estimate the impact of the programme by using the following parametric specifications:

(8)

where $\Delta \ln m$ and Δw are changes in the outcome variables between the baseline and the follow up periods, and X is the set of control variables at both individual and village level described in Table 1 (measured at baseline). The programme impact are given by the coefficients τ_1 and .

To check the sensitivity of results, we also estimate the equations (8) using different techniques. First, we consider a simple OLS regression and implicitly assume homogeneous impacts of the programme. Second, we allow for heterogeneous programme effects by adding interactions of D with the X's, and calculate the effect of interest by taking the average of the X's across treatment observations. We refer to this specification as "fully interacted OLS". Finally, we check the robustness of our results to possible support problems in the distribution of the X's in treatment and control areas. To this end, we compare the average of $\Delta \ln m$ and Δw for households in treatment villages to the weighted average of $\Delta \ln m$ and Δw for "similar" households in control villages, the similarity being defined with respect to the *propensity score* P[d=1|X]. Weights are defined using a Gaussian kernel truncated at a 1% distance, resulting in higher weights for households that are most similar with respect to the propensity score. As a result of this matching procedure only about 1% of the households in treatment villages were discarded from the analysis.

Table 5 reports estimates of the impact of the programme on logged total and food consumption, as well as on the food share, obtained using the three different methods just mentioned. The impact of the programme on total consumption is estimated by OLS at 13.3%, while that on food consumption is estimated at 15.9%. Both estimates are statistically different from zero and are reasonably similar to those obtained with the other two methods, which stand at 14.8% (for both methods) for total consumption and at 17.6% and 17.0% for food consumption.[10] The effect on the share of food is estimated at around 1%, but is not statistically

[10] In our data 13.3% (14.8%) of total consumption is about 60,000 (67,000) pesos, which compares with an average monthly grant (conditional on being paid) of about 100,000 pesos. This implies either that a part of the grant is saved or that there is a reduction in other sources of income. To rule out the latter, we estimated with the same methods the impacts of the programme on logged total household income excluding the programme subsidy. We obtain a point estimate of a reduction of around 10,000 pesos, which is not statistically significant different from zero. These results, available upon request, indicate that part of the grant might be saved. The point estimates we obtain are slightly below those obtained for other conditional cash transfer programmes, such as PROGRESA, where the effect on consumption corresponded to roughly 80% of the grant. It should be stressed, however, that even in PROGRESA the impact on consumption was initially quite low and grew only as the program reached its stability and maturity. Something similar might be happening here. The saving of part of the PROGRESA grant in Mexico is studied extensively in Gertler, Martinez and Rubio-Codina (2009).

different from zero. It should be noticed that the latter result is somewhat surprising: food is usually considered a necessity, so that its share should decrease with total consumption.[11]

5.2 Ex ante effects of the programme on household consumption

In this section, we develop the notation needed to derive the causal parameter from the Engel curves estimated in Section 4 (and from). As for the Engel curve, we work throughout with a linear specification. Omitting the additional variables considered in our specification (such as demographics) to ease notation, we define the following *potential* Engel curves (PECs):

(9)

Equation (9) represents the <u>unfeasible</u> regression of w_0 (or w_1) on $\ln m_0$ (or $\ln m_1$) and household preferences v_0 (v_1), the expectation being taken across households. The quantities (w_0, $\ln m_0$) are revealed in the pre-programme period in both treatment and control villages.

The potential variables are unobserved, and may represent heterogeneous preferences. The possibility that this heterogeneity is correlated with the term is what requires the use of instrumental variables or a control function to estimate the parameters and .[12]

We define the model in equation (9) as *structurally stable* if its parameters do not vary over time and with the programme implementation status:

,

,

and if:

[11] Although we used a selected sample relative to the one used in Attanasio and Mesnard (2006) (as explained in Section 3.2), the results above are largely consistent with those reported in that study and with results for other Conditional Cash Transfers in different contexts (see Attanasio and Lechene, 2009, for rural Mexico, Angelucci and Attanasio, 2009, for urban Mexico, Schady and Rosero, 2007, for Ecuador, and Macours et al., 2008, for Nicaragua).

[12] As we discussed this issue extensively in Section 4, to keep the notation at a minimum we do not consider this issue now. However, our framework will continue to hold if one wants to allow for the endogeneity of total consumption. Note also that, in this specification, the shape of the curves is not affected by the location of villages.

(10) .

Under these conditions and the definition of the PCEs, it must be that:

.

Under structural stability this equation establishes the relationship between the change in the budget share and the change in total expenditure through the parameter of the Engel curve. The assumption of structural stability implies that equation (9) captures the behaviour of households both in the policy on and in the policy off regimes, and in both time periods. Knowledge of the effect of the programme on total consumption (and the relevant demographics in the Engel curve equation) is all we need to estimate the impact on the food share.

The stability of the 's and of the 's can be tested exploiting the evaluation design, against data. In particular, we can test:

(a) whether and do not vary across treatment and control villages in the pre-programme period;
(b) the validity of and exploiting longitudinal variation from before to after the roll out of the programme in control villages;
(c) whether and differ from the 'policy-off' parameters in the previous two bullet points.

Results of this test are reported in Table 6. The top part of the table presents estimates of the parameters for a linear specification of the Engel curve using the three estimation methods discussed in Section 4, and across the four cells defined by the treatment status of the village and the time period. The bottom part of the table presents the p-values for the tests in (a), (b) and (c). Point estimates of the α 's and the β 's are remarkably similar for three of the four cells.

The one cell that stands out is the 'treatment' cell at follow up: the only one where the programme was operating. It is clear that the stability of the coefficients across cells is dubious, and rejected across the three specifications at the conventional significance levels, when we compare the coefficients in the 'treatment on' cell against each of the other three cells (although the relatively low precision of the estimator employed makes this inference not extremely powerful at times). On the other hand, results of the tests do not reject the restrictions: [13]

(11)

(12)

implying that the lack of stability of these parameters is mainly driven by a change in the programme implementation status.

If stability of the 's and the 's is rejected but (10) is maintained, it is easy to work through the definition of the PEC's and the causal parameters defined in the last section to write:

(13)

The causal parameter on the left hand side of this expression can be computed from the parameters reported in Table 6, and the counterfactual term , which is also identified under the restrictions in (7). Equation (13) can therefore be tested: Table 7 contains estimates and standard errors of the quantity:[14]

(14) ,

which confronts ex-ante and causal effects of the programme. We impose the restrictions (11) and (12) (which are clearly not rejected) but not that the intercept and slope of the Engel curve is the same before and after the implementation of the programme in the treatment areas and

---

[13] Note that the parameters and are well defined, but in fact are not identified by our evaluation design. They define the curvature of the Engel curve in a pre-programme period (t=1) under the policy on regime (j=1) – see the notation defined at the beginning of Section 5.1. Because of this they are not included in equations (11) and (12).

[14] Note that the following quantity is actually over-identified in our setting, through the relationships in (11) and (12).

derive the bootstrap distribution of this quantity using 1000 pseudo-samples. Working again through the definition of the PEC's it is easy to see that values away from zero of this difference imply:

and therefore imply a rejection of restriction (10). Intuitively, we want to know whether the effect of the programme on our structural model can be summarized as a shift in intercept and slope or whether the programme has more complex effects on preference shifters and therefore results in a correlation between the 's and d. Results are overall not in favor of the hypothesis that (14) is centered at zero. The probability mass of the bootstrap distribution is concentrated over values which are well above zero, although values in the lower tail somehow depend on the estimation method being used.[15]

5.3 Structural vis-à-vis non-experimental estimates of programme effects

Table 5 indicates that log consumption increased between 13% and 15% depending on the estimation strategy adopted. The results on the estimation of the Engel curve reported in Table 3a imply that, food being a necessity, the share of food should *decrease* as a consequence of a positive shift in Δ_m induced by the programme. The analysis presented in Table 7 combines this expectation and the possible lack of stability of the parameters of the Engel curve documented in Table 6. The main result in this table is that the consumption increase in Table 5 cannot be reconciled with the ex-ante predictions obtained from our preferred specification of the Engel curve.

[15] As a sensitivity check, we derived the counterparts of Table 6 and Table 7 when structural estimates are obtained estimating a quadratic Engel curve using OLS. Although we discussed in Section 4.4 that this is not our preferred specification, it at least prima facie has some potential to fit the programme results given that there is some upward sloping part to the Engel curve. The results of this analysis are reported in the Appendix, where Table A1 and Table A2 should be read as the counterparts of Table 6 and Table 7, respectively. We find that the results presented in the main text are confirmed, and - if anything - a quadratic specification would amplify in magnitude the discrepancies between non-experimental estimates and predictions obtained from the structural model.

How does one interpret and reconcile this evidence? The difference in the two predictions can be explained by a positive value of Δ_u that offsets the value of the right hand side term in (13). The u terms can be thought as "preferences", which might have changed because the programme is likely to change the balance of power within the household and, more generally, the way choices are made. One possible reason for the discrepancy between the prediction of the quasi-experimental impacts and those based on the model is thus a failure of the unitary model behind the derivation of the Engel curve. It is possible, instead, that household decisions are reached taking into consideration the utility function of more than one agent. This interpretation might be particularly plausible in the case of the Familias en Acción because the programme, like many other Conditional Cash Transfers, does not only increase household resources but also targets them to women. Attanasio and Lechene (2002, 2008) propose this explanation in the case of the Mexican PROGRESA.

Our results, which are also consistent with the evidence reported from many Conditional Cash Transfers programmes about their impacts on consumption and its composition (see for instance, Attanasio and Lechene (2002) for rural Mexico, Schady and Rosero (2006) for Ecuador; Angelucci and Attanasio (2009) for urban Mexico, Macours et al. (2008) for Nicaragua and Attanasio and Mesnard (2006), for Colombia), might suggest that such programmes change, somehow, the decision process. One possibility, mentioned in several of these papers, is that the programme changes the relative weight given to women in the decision process. Another possibility is that information changes. Without additional evidence and a more structural analysis it is difficult to establish how.

To conclude, we speculate that the misspecification of the Engel curve that we detect because of their inability to predict the change in food shares induced by the program can stem from a misspecification of the model that fails to take into account distributional factors that shift the power within the family. Conditional cash transfers targeted to women could be equivalent to shifts in the unobserved component of the Engel curve captured in equations (2) and (3). The small increase (and insignificantly different from zero) in the share of food estimated in Table 5 combined with a sizeable increase in total consumption is suggestive of shift in preferences. This might be linked to a shift in the balance of power within the household, as suggested in Attanasio and Lechene (2002).

6. Conclusions

In this paper, we have studied the shape of food Engel curves in a data set collected to evaluate the impact of a large welfare programme in Colombia. *Familias en Acción* is a conditional cash transfer (CCT) that has become one of the main social programmes in Colombia. We use the evaluation data set to study consumption patterns among the poor population targeted by the programme. In particular, we estimate food Engel curves for this population. Such an exercise, and the availability of quasi-experimental estimates of the impact of the programme on total consumption can be used to predict the effect of the programme on the share of food. This prediction and the quasi-experimental estimates can then be used to validate the specification of the Engel curves.

The first aim of the paper was, therefore, methodological. We wanted to establish the best specification for food Engel curves in our population and the best technique to estimate their parameters. In this respect we established that it is crucial to control for the endogeneity of total consumption and to control for the (unobserved) variability of prices across towns (at the same point in time). We conclude our analysis of food Engel curve by saying that, in our data set, the best fit seems to be provided by a log-linear specification (estimated by a control function method) with a coefficient of -0.1 on (log) total consumption. This implies that, *coeteris paribus*, an increase in total consumption by 10% would lead to a decrease of 1% in the share of food.

The introduction of the conditional cash transfer programme *Familias en Acción* is a useful testing ground for our specification of Engel curves. The introduction of the programme led to an increase in total consumption between 13% and 15% depending on the estimation strategy adopted. This number cannot be replicated ex-ante using predictions from the preferred specification of the Engel curve. This evidence on the effect of CCTs on the share of food is consistent with that of other CCTs in different countries, such as Mexico, Nicaragua and Ecuador. We speculate that the mis-specification of the food Engel curve might be explained by the fact that these are targeted to women. Attanasio and Lechene (2002, 2009) suggest in the context of the Mexican programme PROGRESA, that a failure of the unitary model could explain this type of observations. A shift in power towards the women would lead an increase in total consumption to induce a more than proportional increase in food consumption because in addition to the income effect, the CCT would imply a modification of weights towards mothers' preferences. The evidence we present here is not inconsistent with that hypothesis.

Further work is surely needed. In particular, it would be interesting to repeat our exercise for subcomponents of food consumption. For these, prices are observable and, under the assumption of separability between food and non-food, one could estimate a demand system that controls for prices explicitly. One could then compare the predictions of the Engel curves derived from a unitary model to the quasi-experimental evidence. The allocation of resources across commodities (and within the household) is important not only from an academic point of view but also from a policy perspective. CCT have explicitly targeted women because of the perceived need to improve the standing of women within households. Moreover, these programmes have a strong emphasis on nutrition and provide, in addition to cash, advice on best health and nutrition practices. It is therefore important to check whether these programmes are having the desired effects. Understanding the mechanisms at play behind the effects is important to the design and re-design of policy interventions. This paper is a first attempt towards an understanding of these mechanisms.

Affiliations:

Orazio Attanasio: IFS, UCL, NBER and CEPR.

Erich Battistin: Università di Padova and IRVAPP, Trento.

Alice Mesnard: City University, IFS and CEPR.

References

Angelucci, M. and Attanasio, O.P. (2009). 'The effects of Conditional Cash transfers on consumption and its components: evidence from urban areas and implications for household behaviour', Mimeo, University of Arizona.

Attanasio O.P. and Frayne, C. (2005). 'Do the poor pay more?', Mimeo, Institute for Fiscal Study, London.

Attanasio, O.P. and Lechene, V. (2002). 'Tests of Income Pooling in Household Decisions', *Review of Economic Dynamics*, Vol. 5, No. 4, pp. 720-748.

Attanasio, O.P. and Lechene, V. (2009). 'Efficient Responses to Targeted Cash Transfers', Mimeo, Institute for Fiscal Studies, London.

Attanasio, O.P. and Mesnard, A. (2006). 'The Impact of a Conditional Cash Transfer Programme on Consumption in Colombia', *Fiscal Studies*, 27 (4), pp. 421–442.

Bhalotra, S. and Attfield, C. (1998). 'Intrahousehold resource allocation in rural Pakistan: a semiparametric analysis', *Journal of Applied Econometrics*, Vol.13 (5), pp. 463-480.

Banks, J., Blundell, R. and Lewbel, A. (1997). 'Quadratic Engel curves, welfare measurement and consumer demand', *Review of Economics and Statistics*, Vol. 79, No. 4, pp. 527-539.

Blundell, R., Duncan, A. and Pendakur, K. (1998). 'Semiparametric Estimation and Consumer Demand', *Journal of Applied Econometrics*, Vol. 13, No 5, pp.435-461.

Browning , M. and Chiappori, P.A. (1998). 'Efficient Intra-Household Allocations: A General Characterization and Empirical Tests', *Econometrica*, Vol. 66(6), pp. 1241-1278.

Chiappori, P.A. (1988). 'Rational Household Labor Supply', *Econometrica*, Vol. 56(1), pp. 63-90.

Deaton, A. and Muellbauer, J. (1980). *Economics and Consumer Behaviour*, Cambridge University Press.

Gertler, P., Martinez, S. and Rubio-Codina, M. (2006). 'Investing Cash Transfers to Raise Long Term Living Standards', World Bank Policy Research Working Paper No. 3994.

Imbens, G.W. and Angrist, J.D. (1994). 'Identification and Estimation of Local Average Treatment Effects', *Econometrica*, Vol. 62, No. 2, pp.467-475.

Kedir A. and Girma, S. (2007). 'Quadratic Engel Curves with Measurement Error: Evidence from a Budget Survey', *Oxford Bulletin of Economics and Statistics*, Vol. 69, No. 1, pp. 123-138.

Lewbel A. (1996). 'Demand Estimation with Expenditure Measurement Errors on the Left and Right Hand Side', *The Review of Economics and Statistics*, Vol. 78, No. 4, pp. 718-725.

Macours, K., Schady, N. and Vakis, R. (2008). 'Can Conditional Cash Transfer Programs Compensate for Delays in Early Childhood Development?', Mimeo, Johns Hopkins University.

Rubin, D.B. (1974). 'Estimating Causal Effects of Treatments in Randomised and Non-randomised Studies', *Journal of Educational Psychology*, 66, pp. 688-701.

Schady, N. and Rosero, J. (2008). 'Are Cash Transfers Made to Women Spent Like Other Sources of Income?', *Economics Letters,* 101, pp. 246-48.

Velez, C. E., Castano, E. and Deutsch, R. (1998). 'An Economic Interpretation of Colombia's SISBEN: A Composite Welfare Index Derived from the Optimal Scaling Algorithm', Unpublished manuscript, Inter American Development Bank.

Table 1a: *Village characteristics*

	Mean	Std. Dev.
Proportion of households in town centre	0.55	0.28
Atlantic Region	0.36	0.48
Oriental Region	0.25	0.44
Central Region	0.28	0.45
Pacific Region	0.12	0.33
Altitude of the village (in meters)	646.18	753.90
Total population of the village	28,066.16	23,472.64
Proportion of households with piped water	0.87	0.14
Proportion of households connected to sewage system	0.56	0.36
Number of villages	75	

Table 1b: *Household characteristics*

	Mean	Std. Dev.	Percentiles 25th	50th	75th
Number of household members	6.09	2.42	4.00	6.00	7.00
Number of children under 7	1.18	1.16	0.00	1.00	2.00
Number of children aged 7–11	1.08	0.97	0.00	1.00	2.00
Number of children aged 12–17	1.04	1.03	0.00	1.00	2.00
Number of adults above 60	0.29	0.59	0.00	0.00	0.00
Number of female adults	1.37	0.73	1.00	1.00	2.00
Single headed household	0.19	0.39	0.00	0.00	0.00
Affiliated to a good social security	0.05	0.22	0.00	0.00	0.00
Age of head	44.64	13.06	35.00	42.00	53.00
Head : less than a primary education	0.24	0.42	0.00	0.00	0.00
Head : more than a secondary educ.	0.04	0.20	0.00	0.00	0.00
Spouse has less than a primary educ.	0.28	0.45	0.00	0.00	1.00
Spouse: more than a secondary educ.	0.05	0.21	0.00	0.00	0.00
Total consumption	420778	249825	258306	369517	521452
Total food consumed	296017	174311	183359	261926	365971
Share of cons. devoted to food	0.72	0.14	0.64	0.74	0.82
Log of total consumption	12.80	0.57	12.46	12.82	13.17
Log of expected income	12.47	0.68	12.06	12.49	12.90
Log food consumption	12.44	0.59	12.12	12.48	12.81
Log household income	12.30	0.95	11.91	12.43	12.85
Number of households		5,218			

Note:

Sample selection criteria: we excluded all households living in "early treatment" areas or for which we were not able to compute the "expected income" variable due to missing values in the data. Only baseline data are considered. The exchange rate between the US dollar and the Colombian peso was about 2,600 at the date of the survey. The value of consumption has been converted to monthly amounts.

Table 2: *The relations between consumption, income and expected income*

	Log of Total Consumption			Log of Total Income	
	(1)	(2)	(3)	(4)	(5)
Log of Expected Income	0.303*** (0.023)	0.222*** (0.011)	0.172*** (0.012)	0.498*** (0.034)	0.320*** (0.019)
Log of Total Income			0.141*** (0.009)		
Additional controls	NO	YES	YES	NO	YES
Number of observations (villages)	5,218 (75)	5,218 (75)	4,598 (75)	4,598 (75)	4,598 (75)

Note:

*Sample selection criteria: see Table 1. Standard errors (clustered at the village level) in parentheses. The smaller number of observations in columns (3) (4) and (5) is due to missing values in household income. The specification in column (3) is used to define the first stage regression in Section 4. Additional controls include: number of household members, number of elderly adults, number of children less than 6, number of children between 7 and 11, number of children between 12 and 17, number of female adults, education dummies for head and spouse, age of the household head and its square, dummy for single household head, dummies for affiliation to social security. ***denotes statistical significance at the 1 percent level or less.*

Table 3a: *Engel Curves*

Linear specifications

	No price controls		**Village dummies**		**Village prices**	
	OLS	CF	OLS	CF	OLS	CF
	(1)	(2)	(3)	(4)	(5)	(6)
Beta	-0.024***	-0.129***	-0.021***	-0.095***	-0.023***	-0.175***
	(0.007)	(0.016)	(0.006)	(0.013)	(0.007)	(0.012)
Endogeneity test		19.82		15.10		24.27
(p-value)		(0.000)		(0.000)		(0.000)

Quadratic specifications

	No price controls		**Village dummies**		**Village prices**	
	OLS	CF	OLS	CF	OLS	CF
	(1)	(2)	(3)	(4)	(5)	(6)
Beta	0.254***	0.119	0.169***	0.024	0.262***	0.096
	(0.062)	(0.104)	(0.009)	(0.015)	(0.056)	(0.088)
Lambda	-0.024***	-0.021**	-0.017***	-0.008***	-0.025***	-0.016**
	(0.005)	(0.009)	(0.001)	(0.001)	(0.005)	(0.008)
Endogeneity test		16.94		224.35		5.79
(p-value)		(0.000)		(0.000)		(0.000)
First derivatives						
10th	-0.063***	-0.161***	-0.117***	-0.197***	-0.067***	-0.134***
	(0.011)	(0.021)	(0.008)	(0.016)	(0.010)	(0.017)
25th	-0.048***	-0.147***	-0.074***	-0.144***	-0.050***	-0.116***
	(0.009)	(0.018)	(0.007)	(0.016)	(0.008)	(0.015)
50th	-0.031***	-0.132***	-0.033***	-0.082***	-0.034***	-0.099***
	(0.008)	(0.016)	(0.008)	(0.013)	(0.007)	(0.015)
75th	-0.014	-0.117***	0.011	-0.024*	-0.016**	-0.079***
	(0.008)	(0.015)	(0.010)	(0.014)	(0.007)	(0.016)
90th	0.004	-0.101***	0.057***	0.030	0.010	-0.058***
	(0.010)	(0.018)	(0.010)	(0.019)	(0.009)	(0.017)
% of negative first derivative	87.4	100	70.9	83.1	85.9	98.0

Note:
*Estimation results obtained using baseline data and the sample selection criteria described in Table 1. Standard errors are clustered at the village level and given in parentheses. ***denotes statistical significance at the 1 percent level or less, ** at the 5 percent level or less. Coefficients reported for specifications in columns (3), (4), (5) and (6) are village averages. Covariates include number of household members, elderly adults, children less than 6, children between 7 and 11, children between 12 and 17, adult females; education dummies for head and spouse; household head age and its square; dummies for single household head and for affiliation to social security. In all CF specifications, the excluded instruments used in the first stage are expected income and household income. Additional covariates included in column (2): a third order polynomial in the residuals of the first stage regression. Additional covariates included in column (3): village dummies and their interactions with log. total expenditure and its square. Additional covariates included in column (4): a third order polynomial in the residuals of the first stage regression and their interaction with village dummies. Additional covariates included in columns (5) and (6): log prices of coffee, potatoes, rice, sugar, male wages, altitude and its square, index of quality of life. In column (6) these variables are also interacted with a third order polynomial in the residuals of the first stage regression. In the control function specifications in columns (2), (4) and (6), the endogeneity test is an F-test of the joint significance of the coefficients of the polynomial in the residuals of the first stage equation.*

Table 3b: *Robustness analysis - Parametrizing village prices*

Linear specifications

	Restricted Sample		**Alternative Instruments**	
	OLS	CF	CF	CF
	(1)	(2)	(3)	(4)
Beta	-0.043***	-0.429***	-0.117***	-0.126***
	(0.012)	(0.080)	(0.015)	(0.018)
Endogeneity test (p-value)		10.91 (0.000)	26.67 (0.000)	25.85 (0.000)

Quadratic specifications

	Restricted Sample		**Alternative Instruments**	
	OLS	CF	CF	CF
	(1)	(2)	(3)	(4)
Beta	1.017***	-0.376	0.031	0.026
	(0.392)	(1.285)	(0.097)	(0.099)
Lambda	-0.092***	-0.001	-0.012	-0.013*
	(0.034)	(0.109)	(0.008)	(0.008)
Endogeneity test (p-value)		2.37 (0.001)	8.79 (0.000)	6.85 (0.000)
First derivatives				
10th	-0.179***	-0.592***	-0.158***	-0.164***
	(0.024)	(0.114)	(0.021)	(0.023)
25th	-0.100***	-0.486***	-0.136***	-0.143***
	(0.014)	(0.090)	(0.018)	(0.020)
50th	-0.049***	-0.373***	-0.116***	-0.125***
	(0.012)	(0.083)	(0.017)	(0.019)
75th	0.005	-0.276***	-0.094***	-0.103***
	(0.016)	(0.089)	(0.017)	(0.020)
90th	0.079***	-0.195*	-0.072***	-0.081***
	(0.028)	(0.100)	(0.018)	(0.020)
% of negative first derivative	73.2	98.4	99.1	99.34

Note:
*Estimation results are obtained only for the specification with village prices (see Table 3a). Standard errors are clustered at the village level and given in parentheses. ***denotes statistical significance at the 1 percent level or less, ** at the 5 percent level or less. Columns (1) and (2) correspond to columns (5) and (6) of Table 3a when the sample is restricted to values of logged total consumption between 5.3 and 6.2 (see Figure 3). In column (3) the excluded instruments are expected income in the best case scenario and expected income in the worst case scenario. In column (4) the only excluded instrument is expected income.*

Table 4: *Testing linear versus quadratic specifications*

Consumption percentile	(1) no price controls	(2) village dummies	(3) village level variables
1st	(0.017,0.104)	(-0.044,0.690)	(0.016,0.101)
5th	(0.011,0.066)	(-0.045,0.412)	(0.010,0.066)
10th	(0.008,0.047)	(-0.039,0.268)	(0.008,0.049)
25th	(0.003,0.019)	(-0.045,0.088)	(0.000,0.023)
50th	(-0.010,-0.002)	(-0.090,0.025)	(-0.012,0.005)
75th	(-0.034,-0.007)	(-0.336,0.019)	(-0.034,-0.003)
90th	(-0.057,-0.011)	(-0.487,0.023)	(-0.055,-0.008)
95th	(-0.072,-0.013)	(-0.582,0.029)	(-0.069,-0.011)
99th	(-0.101,-0.020)	(-1.328,0.028)	(-0.096,-0.019)

Note:
The first column compares the shape of the Engel curve for the linear and quadratic specifications, without price controls and allowing for endogenous total consumption (see column 2 of Table 3a). The second column compares the shape of the Engel curve for the linear and quadratic specifications, controlling for price heterogeneity by village dummies and allowing for endogenous total consumption (as in column 4 of Table 3a). The third column compares the shape of the Engel curve for the linear and quadratic specifications, controlling for price heterogeneity by village level variables and allowing for endogenous total consumption (as in column 6 of Table 3a). See footnote 9 for further explanations of the test. Numbers reported are 95% confidence intervals for the difference in the slope of the Engel curve calculated at the 1st, 5th, 10th, 25th, 50th, 75th, 90th, 95th and 99th percentile of the distribution of total consumption using a linear vis-à-vis a quadratic specification. These were derived using a bootstrap procedure based on 1000 replications for the difference of the slope under the two specifications on each pseudo sample. The confidence intervals reported are based on the 2.5th and 97.5th percentiles of the bootstrap distribution from the 1000 replications.

Table 5: *Difference in differences estimates of the effect on total consumption and food consumption*

Estimate (standard error)	Log Total Consumption	Log Food Consumption	Share of Consumption Devoted to Food
OLS	0.133*** (0.043)	0.159*** (0.045)	0.010 (0.010)
Fully interacted OLS	0.148*** (0.048)	0.170*** (0.050)	0.009 (0.011)
Matching	0.148*** (0.053)	0.176*** (0.055)	0.009 (0.013)

Note:
Estimation results obtained from equation (4) using baseline and follow-up information and the sample selection criteria described in Table 1. Standard errors clustered at the village level are given in parentheses and obtained from 1000 bootstrap replications. Only treatment and control households on the "common support" are considered (5,163 out of 5,218), the latter being defined from the regression of a dummy for living in treatment areas on the controls described in Table 1.
****denotes statistical significance at the 1 percent level or less.*

Table 6: *Control Function estimates for the linear specifications*

	No price controls		**Village dummies**		**Village prices**	
	α_f	β_f	α_f^v	β_f^v	$\alpha_f(\xi_v)$	$\beta_f(\xi_v)$
(A): Control areas at baseline	1.444***	-0.123***	1.193***	-0.081***	1.226***	-0.086***
N = 2,711	(0.115)	(0.019)	(0.084)	(0.013)	(0.099)	(0.016)
(B): Treatment areas at baseline	1.452***	-0.122***	1.124***	-0.066**	1.194***	-0.078***
N = 1,887	(0.223)	(0.037)	(0.149)	(0.025)	(0.125)	(0.021)
(C): Control areas at follow up	1.327***	-0.122***	1.160***	-0.094***	1.051***	-0.075***
N = 2,615	(0.132)	(0.022)	(0.121)	(0.020)	(0.119)	(0.019)
(D): Treatment areas at follow up	0.991***	-0.057**	0.844***	-0.033	0.795***	-0.025
N = 1,858	(0.143)	(0.023)	(0.158)	(0.026)	(0.143)	(0.023)
P-values for the equality of coefficients:						
(A) = (B)	0.9727	0.9845	0.6871	0.6221	0.8413	0.7733
(A) = (C)	0.4169	0.9581	0.8175	0.5868	0.2676	0.6859
(A) = (D)	0.0105	0.0238	0.0535	0.1055	0.0140	0.0340
(B) = (D)	0.0268	0.0629	0.1696	0.3160	0.0250	0.0699
(C) = (D)	0.0910	0.0468	0.1224	0.0705	0.1751	0.1049
(A) = (B) = (C)	0.7046	0.9986	0.9144	0.6965	0.5220	0.9127

Note:

Reported are estimates of the parameters of a linear Engel curve suing the three specifications in Section 4, separately for the four groups of observations defined by the policy-off and policy-on regimes across time periods. Standard errors clustered at the village level are given in parentheses and obtained from 1000 bootstrap replications.

Table 7: *Structural vis-à-vis diff-in-diff estimates of programme effects on the share of food*

Difference in Differences Estimate		0.010 (0.010)	
	No Price Controls	**Village Dummies**	**Village Prices**
Structural Estimate see equation (14)	-0.8469** (0.3072)	-0.4751 (0.3282)	-0.8124** (0.3169)
Difference in Differences Estimate – Structural Estimate	0.8572** (0.3054)	0.4854 (0.3270)	0.8227** (0.3161)

Note: standard errors in parentheses.
The difference in difference estimate is that reported in the last column of Table 5. Structural Estimates of programme effects are derived using equation (14) in Section 5 under the three specifications considered (No Price Controls, Village Dummies, Village Prices). The standard error of the difference is obtained via bootstrap using 1000 replications.

Fig. 1: *Engel curves estimated by OLS*

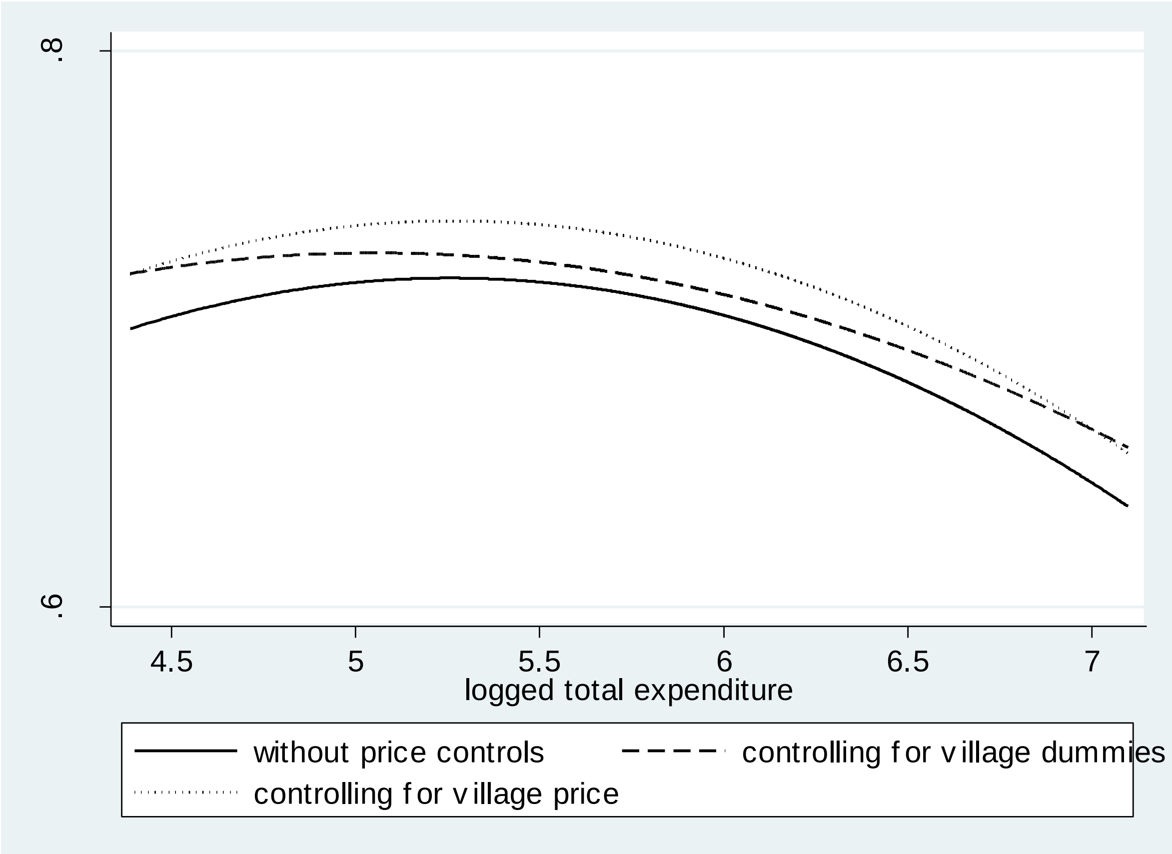

Note. Predicted values from OLS estimation results obtained ruling out endogeneity of total consumption. Estimates "without price controls" are obtained from equation (4) in the main text (see Section 4.1). Estimates "controlling for village dummies" are obtained from equation (5) in the main text, modelling price effects through village dummies (see Section 4.2). Estimates "controlling for village prices" are obtained from equation (6) in the main text, modelling price effects through village characteristics (see Section 4.2). Total consumption is divided by 1000.

Fig. 2: *Distribution of estimated coefficients across villages*

a OLS

b allowing for endogenous total consumption

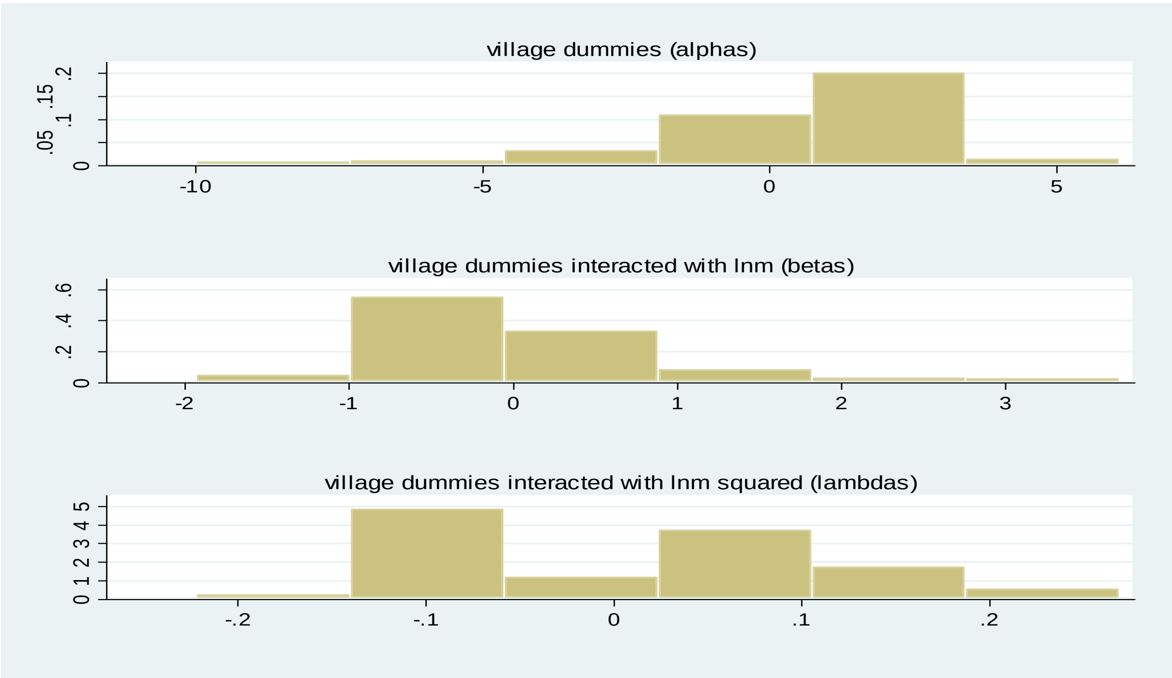

Note. Distribution of the estimated coefficients α_f^v, β_f^v and λ_f^v of the equation (5). Distributions displayed in panel a are obtained ruling out endogeneity of total consumption, whereas distributions displayed in panel b are obtained through control function approach.

Fig. 3 : *Distribution of log expenditure and its fitted value*

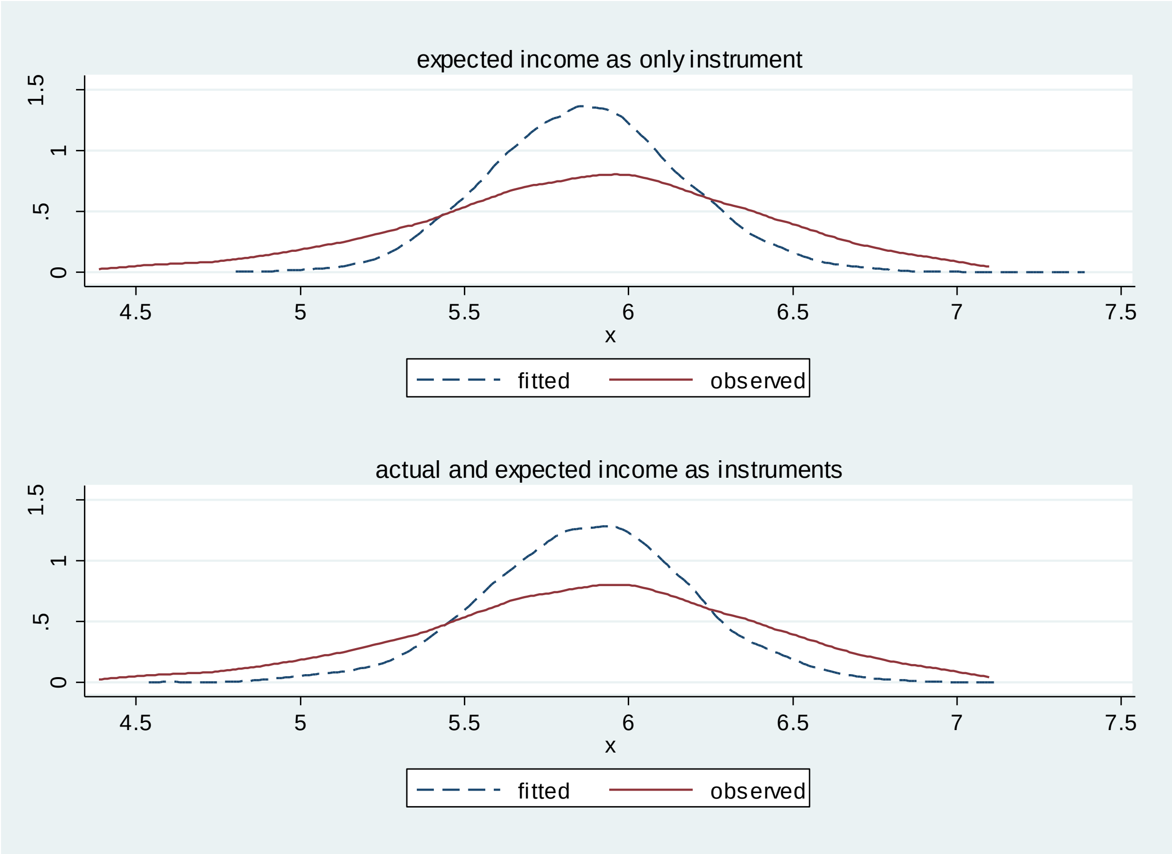

Note. Non-parametric estimates of the probability distributions for logged total consumption (observed) and predicted values of logged total consumption obtained from the first stage regression (fitted). See Section 4.3 for details. Total consumption is divided by 1000.

Fig. 4: *Engel curves allowing for endogenous total consumption*

Note. Predicted values from CF estimation results obtained allowing for endogeneity of total consumption. Estimates "without price controls" are obtained from equation (4) in the main text (see Section 4.1). Estimates "controlling for village dummies" are obtained from equation (5) in the main text, modelling price effects through village dummies (see Section 4.2). Estimates "controlling for village prices" are obtained from equation (6) in the main text, modelling price effects through village characteristics (see Section 4.2). Total consumption is divided by 1000.