Medium and Long Run Effects of Nutrition and Child Care:

Evaluation of a Community Nursery Programme in Rural Colombia[*]

Orazio P. Attanasio[♣] and Marcos Vera-Hernández[♦]

November 2004

Abstract

In this paper we evaluate the effect of a large nutrition programme in rural Colombia on children nutritional status, school achievement and female labour supply. We find that the programme has very large and positive impacts. Dealing with the endogeneity of treatment is crucial, as the poorest children tend to select into the programme. Methods like Propensity Score Matching would even yield negative estimates of the impact of the program. Our results are robust to the use of instruments that do not depend on individual household choices. We also validate our evaluation strategy by considering the effect of the program on pre-intervention variables. Further, we explore the heterogeneity of the impact of the programme. Children from the poorest backgrounds are the ones that benefit the most.

JEL: C21, I12, I38

[*] We are very grateful to Erich Battistin, Jere Behrman, Samuel Berlinski, Pedro Carneiro, Emla Fitzsimons, Luis Carlos Gomez, Alejandro Gaviria, Beatriz Londoño, Costas Meghir, Jairo Nuñez, Jim Smith and seminar audiences at LSE and UCL, Stanford, Universidad de Los Andes, INRA (Paris), the University of Toulouse, the World Bank, the London School of Hygiene and Tropical Medicine for many useful comments and suggestions. We are particularly grateful to the staff of *Instituto Colombiano de Bienestar Familiar (*ICBF), including Maria Francisca Concha, Ana Maria Peñuela and Yaneth Romero, for answering many questions about the details of the programme *Hogares Comunitario* and to the *Departamento Nacional de Planeación* for allowing us the use of the data for the evaluation of the *Familias en Acción* programme.

[♣] UCL, IFS and NBER

[♦] IFS

1. Introduction

The purpose of this study is the evaluation of a large nutrition programme in rural Colombia, called *Hogares Comunitarios de Bienestar Familiar*. This is a large intervention based on community nursery where poor children receive food (purchased by the government) and child care from one of the mothers in the community. Our purpose is to measure the effects of this programme on the nutritional status of young children, its long run effects on school achievement, as well as in female labour supply

Malnutrition is a very serious problem in developing countries. According to Onis *et al.* (2000) about one third of less than five years old children are stunted in growth. There is evidence that inadequate nutrition in childhood affects long term physical development (Martorell and Habicht, 1986, Barker, 1990), as well as the development of cognitive skills (Brown and Pollitt, 1996 and Balazs *et al.* 1986) and educational attainment (Behrman, 1996, Strauss and Thomas 1995). This in turn affects productivity later in life (Dasgupta 1993, Strauss and Thomas 1998, and Schultz 1999).

Because of the importance that obviously malnutrition plays in development and because of the accumulating evidence that early year interventions might be the most important, several different types of nutritional programs have been proposed in the developing world and have received considerable attention in recent years. Given the scarcity of resources and the abundance of different interventions, it is crucial to assess what are, in different situations, the most cost effective.

Nutrition interventions come in many different types. There are interventions that provide food or nutritional supplements to poor households or children, others (such as the programme we are studying) that combine these in-kind transfers with child care, interventions that subsidize prices of some commodities in some areas, interventions that provide cash to poor households with children either unconditionally or, as in some recent programs inspired by the Mexican PROGRESA, in exchange of some forms of behaviour, including the registration of children in growth and development check-ups and vaccinations in health centres.

Some of these interventions have been evaluated. In one of the cleanest evaluations available, the Mexican PROGRESA was shown to have some impact on the height of children aged 12 to 36 months (see Behrman and Hoddinott, 2004). More recently, *Familias en acción*, which is widely perceived as an alternative to the Hogares Comunitarios program, has been evaluated by Attanasio *et al.* (2004a). That evaluation shows that this conditional cash transfer program, similar to PROGRESA, increases the height of children aged 0 to 2, but had limited effects on older children.

Behrman, Cheng and Todd (2003) study a programme in Bolivia called PIDI. This study is particularly relevant for us because PIDI is remarkably similar to *Hogares Comunitarios.* The authors evaluate it using a matching strategy. They show that, given the assumption of selection on observables, the programme has no positive effect on children height. Conditional on participation, however, they find some moderate positive effect of length of exposure. Ruel *et al.* (2000) study a programme very similar, even in name, to *Hogares Comunitarios*, implemented in Guatemala City. Using a 'selection on observables' strategy they find very limited effects of the program.

In this paper we exploit a large and high quality data base recently collected to evaluate a different a new intervention in rural Colombia. In particular, we use information on the children living in the towns where the new programme (which is an alternative to *HC*) did not operate to evaluate the effect of *HC*. The survey on which our study is based contains rich information both on young children, who might be attending a *HC* and on older children. As, for the latter we can reconstruct past attendance to *HC*, we are able to study long run effects of the program. As we do not have a 'control' group, we use an instrumental variable technique. In particular, we will be assuming that, conditional on some observables, the distance of each household from the nearest *HC* is exogenous to the outcome of interest. As always, this type of assumption is debatable: in what follows we discuss it at length and present several arguments and pieces of evidence to justify it in our context.

The results we get are remarkable. We find that the programme has important effects both on the nutritional status of young children and on the academic performance of older children. We also find important effects on female labour supply. Allowing for the endogeneity of treatment is crucial: simple comparison of participants and non-participants find (conditional or non-conditional) find no significant effects of the program. This result is consistent with the evidence from the participation equations that point to the fact that the participating children are the poorest. The programme seems to be able to compensate the difference between these children and those that are slightly better off. In terms of nutritional status these outcomes are large: they amount to two centimetres in height for young children.

We also study the effect of HC on later school attendance and school achievement. There is already some literature that convincingly argues that malnutrition in childhood influence later schooling decisions (Glewwe *et al.* 2001, Alderman *et al.* 2001a). The studies in this literature found that controlling for unobserved variables is very important. Our paper is related to them as we study the effect of HC directly on schooling related outcomes. Clearly, part of this effect could come from improving child malnutrition; but other channels as women empowerment cannot be ruled out. Behrman *et al.* (2003) study the effect of receiving a nutritious supplement when children are between 6 and 24 months on education related outcomes. They find significant and substantial effects of the nutritional supplement on all the outcomes they consider. Their great advantage is that the intervention was randomized at the village level. However, they study a specific intervention that is very different from the one we consider. In particular, the programme they study does not include child care and consequently it is unlikely to have effects on female labour supply.

The rest of the paper is organized as follows. In Section 2, we describe the operation of the programme. In Section 3, we discuss our identification strategy. In section 4, we present the data we use and some descriptive statistics. Section 5 presents our main results. Section 6 discusses heterogonous impacts of the programme.. Section 7 concludes. The definitions of control variables and the full set of results are relegated to the Appendix.

2. The *Hogares Comunitarios* programme

In the late 1970s, the Colombian government legislated a new nutrition intervention targeted towards poor families. The programme, that took the name of *Hogares Comunitarios de Bienestar Familiar,* was legislated in 1979 as the development of previous initiatives where nutrition interventions tried to stimulate community participation and initiatives.

The programme is run by the *Instituto Colombiano de Bienestar Familiar (ICBF).* At the beginning of the programme, which started between 1984 and 1986, the ICBF regional office targeted poor neighbourhoods and localities and encouraged eligible parents with children aged 0 to 6 to form 'parents associations'. Households belonging to the so called SISBEN levels 1 to 3 can participate.[1] After a few meetings with programme officials, the parents association was registered with the programme and elected a *madre comunitaria* (or community mother). This mother had to satisfy some criteria, such as having basic education and a large enough house and would be certified by the regional office of the ICBF. The *madre comunitaria* would then receive in her house the children aged 0 to 6 of the parents belonging to the associations. Each family would pay a small monthly fee (roughly the equivalent of four US dollars), which would be used to pay a small salary to the *madre comunitaria*. Each *madre* would receive up to 15 children. The average number of children is around 12. The parents association would then receive funds from the government to purchase food. The food would be delivered weekly at the house of the *madre comunitaria* who would keep it in her fridge. The menu varies regionally and is established by a nutritionist in the regional office of ICBF. In addition to the food included in the regional menu, the children would also be given a nutritional beverage called *bienestarina*. Children are fed three times: lunch and two snacks. According to ICBF, the food received by the children (including the beverage) would provide them with 70% of the advisable daily amount of calories.

Therefore, in exchange for the small monthly fee, the parents would get child care and some food. The programme objectives included the improvement of the nutritional status of poor children as well as the provision of child care that could stimulate labour force participation of women and the generation of additional income.

The program, whose cost is financed with a 3% tax on the wage bill, expanded very rapidly in Colombia. It is now the largest welfare programme in the country: there are roughly 80,000 HC across the country and more than a million children that attend one. The cost of the programme is approximately 250 million US$, or almost 0.2% of GDP.

As we discuss below, the location of the *HC* plays an important role in our identification strategy. After the start of the programme and its rapid growth, the turnover among the *madre comunitarias* seems to be substantial. According to officials of the ICBF, between 10

[1] In Colombia all welfare programs are targeted through the so-called SISBEN indicator. This indicator is computed using a number of different indicators of economic well-being. SISBEN is constructed on the basis of an index that is the first principal component of a number of variables related to poverty. Depending on the value of the index, each household is assigned to one of six levels. Information on the variables used in the construction of SISBEN is collected periodically. For most welfare programs, only households belonging to level 1 and 2 are deemed eligible. FA households are in SISBEN 1.

and 15% of the existing *HC* are relocated in each year, in that a *madre comunitaria* ceases to be such and a new *madre* starts to operate it. Moreover, if a household moves to a certain neighbourhood, it can normally register its children in an existing *HC*. It seems that over time, the *HC* have evolved into relatively mobile and informal nurseries and have lost some of the tight connection with the original parents association. In rural and very disperse areas, an apparently common problem is the difficulty to set up a new *HC* because the ICBF does not register new *HC's* unless there is a sufficient number of children attending it.

3. Evaluation strategy

The *HC* programme, even though is the largest welfare programme in Colombia, has barely been studied. One of the reasons for the paucity of systematic evidence on the programme, is the lack of a control group, in turn explained by the speed with which the programme was developed when it first started. The *HC* programme now covers all of Colombia, both in rural and urban areas. Besides some early internal studies which considered mainly the operation of the program, the only attempt at measuring the effect of the programme was a study in 1996 (published in 1998) that used a relatively large survey designed for the explicit purpose of evaluating the *HC* programme (Profamilia, 1998). However, that study only measured children observed in *HC's*. No measurements were taken of children not attending the programme. While the study provides a wealth of useful statistics and observations about the children and the *madres comunitarias,* the basic (and implicit) evaluation strategy is to compare the anthropometric measures of *HC* children with those of children of *similar socio-economic background* (observed in other surveys). The most striking observation was the fact that most nutritional indicators, such as height per age, did not systematically differ. The implicit conclusion reached in that study was that the programme fails to improve the nutritional status of poor children substantially. Such a conclusion, obviously, ignores selection problems.

Like in the evaluation of most social interventions, the fact that a programme is not assigned randomly, can create substantial problems. A comparison of children attending a *HC* to children not attending one, even if we control for observed characteristics, can yield very misleading results as it ignores the endogeneity of the participation decisions: the children whose parents decide to send to a *HC,* are in all likelihood very different from the children that are not sent to a *HC*.

In this section, we first discuss our definition of outcomes and treatment. We then discuss our identification strategy for the effects of the programme. Finally, we propose a simple model which constitutes an attempt to go beyond the simple measurement of the programme and understand the channel through which the programme operates.

3.1 Outcomes and treatment

In our analysis of the impact of *HC* we define several outcomes of interest. First, we look at several anthropometric measures that are available in our data base. In particular, for children aged 0 to 6, we consider height, weight and leg length. Height and weight, standardized by age and gender, are extremely common indicators of nutritional status in the literature. Leg length has recently received some attention because of evidence that reflects well the stock of past nutrition flows and is a good predictor for illnesses in adulthood.

Following the literature, we do not use height and weight directly, but we construct the so-called z-scores for these variables standardizing them by age and sex according to the World Health Organization/Centre for Disease and Control (WHO/CDC) reference population.[2] In particular, the z-score for height per age is obtained from the height of a child, subtracting the median height of WHO/CDC reference population of the same age and gender and dividing by the standard deviation of height of the WHO/CDC reference population of the same age and gender. An analogous procedure is followed for weight per age. As the growth patterns of the WHO/CDC reference population might be different from the ones of our population, we always introduce additional controls for age and age interacted by sex in our regressions.

Probably the most interesting of the three measures is height for age, which is considered to be a good index of the stock of malnutrition. A child is considered as chronically malnourished if his or her z-score is below -2, that is, if his or her height is 2 standard deviations below the median of the reference population for the same age and gender.

A slightly less usual measure we use is leg length, which is obtained, in children aged 2 to 6 as the difference between standing height and sitting height. There is some evidence in the literature that leg length is a good marker of the stock of malnutrition and a very good

[2] This reference population are mostly conformed by the 1975 US children population. At the moment this is the reference population most widely used. The World Health Organization is working in a project to build a true international reference population.

predictor of illnesses in adulthood (see, for instance, Buschang *et al.* (1986), Gunnell *et al.* (1998, 2003) and Davey-Smith *et al.* (2001)). To the best of our knowledge, there are no z-scores for leg-length, so that for this variable we perform the analysis simply controlling for a polynomial in age and sex.

The long run effects of nutritional programs have recently received considerable attention. There is increasing evidence that nutrition at young ages can have long lasting impacts on school achievement and even future earnings. To assess the possibility that the programme has long term effects, we consider some outcomes for children that are no longer attending a *HC* because they are past the age limit, but that might have attended in the past. For these children we consider two measures of academic achievements: whether they are currently attending school and whether they progressed a grade between the baseline and the follow up survey.

In addition to variables directly related to the welfare of children, we also look at the potential effect that the programme has on other outcomes, such as female employment rates and hours of work.

As for treatment, we use several alternative definitions. For children younger than 6, we define treatment on the basis of current attendance to a *HC*. Moreover, for each child we can reconstruct, for each age between 0 and 6, the number of months in which the child has attended a *HC*. Therefore, both for children aged 0 to 6 and those aged between 8 and 17, we use the number of months as a continuous definition of treatment. For children aged 0 to 6 we also normalize the number of months during which he or she attended a *HC* by the child's age in months, therefore defining treatment as the fraction of his or her life spent in a *HC*. For children older than 7, in addition to the number of months we also construct a discrete treatment definition and consider a child as treated if he or she has ever attended a *Hogar Comunitario* between ages 0 and 6. In the case of female labour supply, we define a mother as 'treated' if she has at least one of her children attending a *HC*.

3.2 Identification

Given that the *HC* programme has a very extensive geographic coverage it is difficult to identify a 'control' group that could be used to estimate the impact of the program. Of course, as we will see, it is not difficult to find children that do not attend a *HC*. But the choice to attend is likely to be related to the outcomes of interest. To solve this problem, we decide to adopt an Instrumental Variable approach, that is, to identify at least one variable that is likely to affect the decision to send a child to a *HC* but is unlikely to affect directly the outcomes of interest. We discuss below the choice of the instrumental variable.

Given outcome y_i for child i, we will be estimating the following equation:

(1) $$y_i = \beta' x_i + \gamma\, p_i + u_i$$

where x represents some control variables that we assume to be exogenous (such as mother height or village variables), and p the treatment, defined above to be either current attendance or fraction of the child's life spent into the program. The assumption we make is that the treatment is defined by the equation:

(2) $$p_i = \theta' x_i + \pi\, z_i + v_i,$$

where the variables z_i represent the instrumental variables. The possibility that v_i and u_i are correlated makes the OLS estimation of (1) yield biased estimates of the parameter of interest γ. In order to obtain consistent estimates we assume that both $\pi \neq 0$ and that z_i is uncorrelated with u_i. Given the evolution of the HC programme and in particular the high turnover of mothers in the last few years we believe that both the distance from the household to the nearest HC, and this distance averaged at the town level will be good instruments. We will present evidence of the extent to which both the household distance to the nearest HC and its town average affects participation choices, that is, that $\pi \neq 0$. We also believe that these two distances are unrelated to nutritional outcomes, conditional on the other variables x_i we control for. However, we acknowledge that this assumption is not uncontroversial and should be justified. We do this below.

Two obvious problems can arise if the location of the *HCs* is endogenous or if the location of individual households relative to the *HC* is endogenous. The first problem could be relevant if the government, through the process of formation of parents associations, implicitly targets the programme towards the parents that care the most or have most to gain from the program. The second problem might arise if the parents that care the most or have most to gain actively located themselves closed to a *HC*.

The exogeneity of an instrument cannot be tested. However, we provide several pieces of evidence that can justify our approach. First, conversations with programme officials indicated that, especially in isolated rural areas, which make a substantial proportion of our sample, there might be severe supply restrictions induced by the need of a minimum number of children for ICBF to register a new *HC*. Moreover, after the first few years of the program, the turn over of *madre comunitarias,* induced by a variety of factors, contributed to substantially weaken the link between the original parent association and the location of the *HC*. It seems that many of the current clients of *HC* are households that move to a given neighbourhood and access an existing *HC*. Second, we can provide evidence that households do not move to be closer to a HC. Between the first and second survey, approximately 1,900 households (more than 16% of the sample) changes address of residence. Of these we were able to re-interview 1423.[3] To these households, we asked the reason for changing address. None of them said that they moved to be closer to a HC, even though moving closer to a *HC* was explicitly listed as a possible reason to move.[4] Third, we include a rich set of household level controls, including the distance from the household to the nearest school, and to the nearest health care centre. We believe that these are valuable controls because HC could be located close to clinics or school and access to health care services and nutritional information could be potentially important in determining child nutritional status. Moreover, they would capture the fact that households living in somewhat 'central' locations could be both closer to a *Hogar Comunitario* and systematically different in terms of nutritional status. Fourth, we also use as an instrument the average distance from the households to the nearest HC in the town, that is, the density of *HC* in a town. This is a valuable instrument because it is independent of the location decisions of individual households. Of course, the variability we are exploiting in this case is only across municipalities. For this reason, we also include a rich set of municipality level controls in equation (1). The results we obtain when we use the individual household distance from to the nearest HC or when we use its town level average are very similar, therefore supporting our identification strategy.

[3] It seems that most of the movers we lost were households that moved to large cities.

[4] 301 households moved to find a better equipped house, 284 moved due to labour related motives, 124 moved to be closer to a relative, 54 moved to be closer to a school, 41 moved due to violence, 13 moved to be closer to the village centre, 0 moved to be closer to a HC, and 606 moved due to other reasons.

The final piece of evidence we present in support of our identification strategy appeals to the idea that if any effect we find is driven by a correlation between the error term of the outcome equation and the instrument we use, we would be likely to find effects on variables related to the outcomes we study but on which the programme should not have any effect. For this reason we look at children's birth weights and mother's height: the programme should not affect such variables as they are realized before the exposure to it.

4. The data

The data we use in this paper was collected with the specific purpose of evaluating a different and new welfare program. For this reason, the sample we use is concentrated in a certain type of communities. In this section we first describe the nature of the data set and then present some descriptive evidence on the children who compose our sample.

4.1 The Familias en Acción programme and the evaluation database.

Between 2001 and 2002, the Colombian government started a new intervention in towns with less than 100,000 inhabitants, modelled after the PROGRESA programme in Mexico and financed with a loan from the World Bank and the Inter American Development Bank. This program, called *Familias en Acción* (FA from now on) has an education, a health and a nutrition component and is directed to the poorest families living in the municipalities targeted by the program. As in the case of PROGRESA, the targeting of the programme is first done at the community level and then, within the chosen communities, at the individual level. The targeted communities were chosen on the basis of several criteria. First, they had to be relatively small towns (less than 100,000 inhabitants and no departmental capitals). Moreover, given that *FA* is a conditional cash transfer program, a town could be included only if it had enough education and health infrastructure. Finally, for security reasons in delivering the payments, the presence of a bank in the municipality was also a condition for qualifying.[5] At the individual level, the programme was targeted to households with children aged 0 to 17 belonging to the lowest level of the so called SISBEN index (see footnote 1).

[5] An additional condition (that turned out to be binding in some situations) was that the mayoral office had to process some documents and have a list of potential beneficiaries ready.

The nutrition component of *FA* consists of a cash subsidy that is given to the mother of children aged 0 to 5 living in beneficiary households. The subsidy is about 15 US dollars per month and is conditional on certain behaviours. In particular, the children have to be registered and taken to growth and development check ups and the mother is supposed to attend some courses on hygiene and vaccination. Clearly such a programme is very different from *Hogares Comunitarios* and, indeed, is widely perceived as a substitute for it. While *HC* provides childcare, in-kind transfers and up to certain extent nutritional insurance, *FA* relies on monetary transfers under a conditionality of visits to health care professionals. Moreover, in the targeted municipalities, households entitled to the nutrition component of *FA* have to choose between that programme and *HC,* in that they cannot send their children to an *HC* if they register for *FA*.

When the *FA* programme was started, a large scale evaluation of its impact was also started. In particular, a large data collection project was undertaken in 122 municipalities, 57 of which were targeted by the programme. The remaining 65 were chosen as 'controls'. While the assignment of the programme to municipalities was not random, the control towns were chosen so be as similar as possible to the random sample of 57 'treatment' municipality. In practice, most of the control towns satisfy most of the conditions imposed by the programme with the exception of the bank presence.[6]

The *FA* evaluation survey is a longitudinal dataset whose collection started with the baseline survey in the summer of 2002. A total of 11,502 households in the 122 survey towns were administered a detailed questionnaires including detailed information on a large number of individual and household level variables.

The households included had to satisfy the eligibility rules of *Familias en Acción,* that is they had to be registered as SISBEN 1 as of December 1999 and have children aged 0 to 17. This implies that our sample is representative of the poorest households in small towns. In addition to a very large number of questions covering consumption, income, school attendance, labour supply and a variety of other variables, every child aged 0 to 6 was

[6] The municipalities were classified in 25 strata according to geographical region, population size living in the urban part of the municipality, the value of synthetic index for quality of life (QLI) as well as education and health infrastructure. Two treatment municipalities were randomly selected within each stratum among the municipalities participating in *Familias en Acción*. For each treatment municipality, a control municipality was chosen as the most similar to the treatment municipality in terms of population size, population living in the urban part of the municipality, and QLI among the set of municipalities not participating in *Familias en Acción* but belonging to the same stratum than the treatment municipality.

weighted and measured. In particular, his or her standing and sitting height was measured with a precision instrument. The questionnaire included a number of questions about current and past attendance of each child to a *HC*. In particular, for each child, we know whether he or she is currently attending a *HC*, and, for each year of the child's life, how many months he or she had attended a *HC*. Finally and important for our identification strategy, for each household, regardless of whether it has children attending a *HC*, we know the distance to the nearest *HC*.

In the summer of 2003, when *FA* was operating in all treatment towns, the same households in the baseline survey were re-contacted in a follow-up survey, during which a questionnaire very similar to the baseline was administered. Two noticeable additions to the questionnaire were a question for children aged 7 to 10 about past attendance to a *HC* and additional questions about the distance of the household residence from a variety of public structures, such as the main hospital and health centre, the city hall and so on.

As we are interested in evaluating the impact of the *HC* programme and we want to avoid contaminations by the new programme (*FA*), in what follows we focus on the towns where *Familias en Acción* was not implemented. That is data from the 65 'control' municipalities. In the baseline, in these municipalities 4,689 households were interviewed, including 4,147 children aged 0 to 6. As these households are all SISBEN 1 then these children are eligible to participate in *HC*. In the first follow up, we re-interviewed 4,426 of these households. As in what follows we use some municipality level averages, we drop from our sample municipalities where we observe less than 30 households. This leaves us with 54 of the 65 control municipalities.

As some of the impacts we will be measuring are likely to take some time to build up, we focus on cross sectional 'stock' outcomes (such as height) rather than longitudinal (growth) outcomes. We obtain most of the results below by pooling the baseline and follow up data. As all standard errors are computed taking into account cluster effects at the municipality level, we also allow for correlation among children interviewed twice. Most of our results do not change if we use only follow-up or only baseline data.

4.2 Descriptive statistics

Our sample is made of very poor households, mostly living in very difficult conditions. Most towns have been widely affected by the civil war in Colombia. A striking indication of the effect that these conditions have had on our sample is its age structure that we report in Figure 1 (taken from Attanasio *et al.* (2003)). The number of individuals aged 18 to 36 one can find in our sample (and by extension in the towns we are studying) is considerably less than what would be implied by a standard age structure of a population with relatively high fertility rates. This reduced number of individuals in these age groups is now also being reflected in a much reduced number of very young children.

Tables 4.1a to 4.1c report the main descriptive household level statistics for our population. Table 4.1a reports village level statistics, many of which are taken from the census and other sources. Table 4.1b reports household level statistics. The households considered in this table are those where the children included in the analysis live and are computed using our survey. In table 4.1c we report children level statistics.

Figure 1.
Age structure

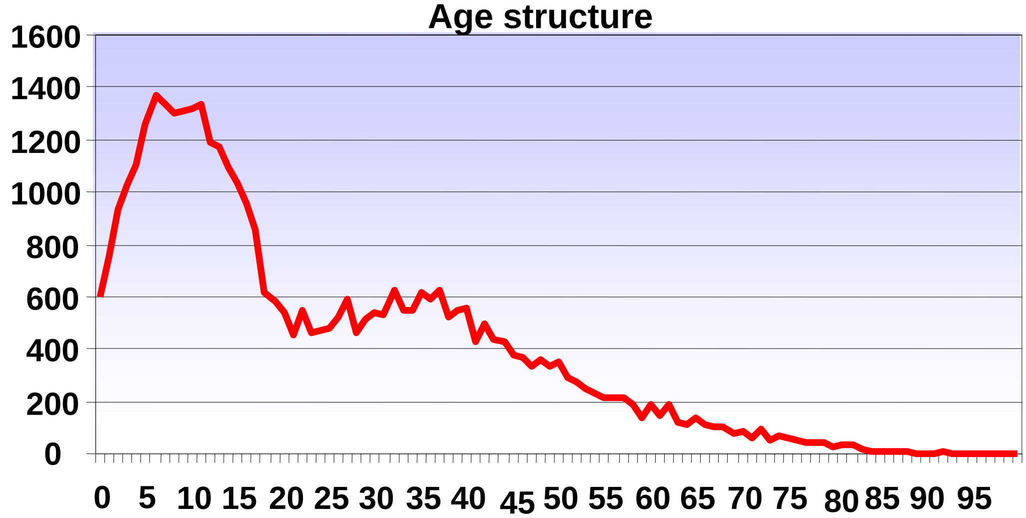

The towns in our sample are reasonably small: the average (median) population in 2001 was 25k (20k) and even the town at the 75% percentile had less than 30k. However, the area over which these municipalities extend is at time substantial: the average size in square kilometres is 674. Typically, there is a substantial fraction of the population that lives in the so called 'cabecera municipal' (the main town) (the average is 14k) while the rest is dispersed in the country side. The variation in altitude reflects the geographic complexity and diversity of Colombia.

Table 4.1a:
Municipality statistics

(54 towns)				
	Mean	Standard deviation	Median	P75
Altitude	760	876	416	1350
Area in Km2	674	1686	192	615
Population (2001)	25154	24659	20591	29743
Population in urban areas	13950	17115	8608	17736
Population sisben 1 (% of total pop.)	4606 (0.50)	4917 (0.32)	3066 (0.44)	7678 (0.81)
NBI	0.55	0.17	0.58	0.68
% of household with sewage access	0.31	0.26	0.26	0.46
% of household with piped water	0.62	0.27	0.63	0.86
Students per teacher	22	5.5	22	27
Class Square meters per student	2.6	1.8	2.2	3.1
# of hospitals	0.68	0.46	1.00	1.00
# of health centres	0.74	1.01	1.00	1.00
Average distance[a] to city hall	59	58	48	65
Average min. distance[a] to closest health cntr.	46	38	37	55
Average min. distance[a] to closest school	15	7	13	19
Average min. distance[a] to *Hogar Comunitario*	29	23	19	38
Average fee to attend *HC*	4490	4038	3717	5250
Average age of *madres comunitarias*	38	3.8	37	40
Average experience of *madres comunitarias*	7.7	2.9	7.6	9.9

a. Distance is measured in minutes.

A substantial fraction of the inhabitants of the towns in our sample are potential beneficiaries: in the average town 50% of the population is registered in SISBEN 1. The importance of poverty is also reflected in the NBI index, which is supposed to measure the percentage of households with 'unsatisfied basic needs': this index averages 0.55 in our towns. On average, only 31% of households in our towns have access to sewage and 62% of households have access to piped water. These percentages are substantially lower for our sample.

The next group of variables in the Table provides information on the infrastructure present in the municipalities. Schools do not seem to be particularly crowded and class sizes are not very large. Likewise, there is on average 1 health centre and one hospital per town (although not necessarily both). The average distance to the health centre is 46 minutes.

Another indication of the reasonably supply of schools in this towns is the fact that, even though the population appears to be quite disperse (the average travel time to the city hall is, in the average town, an hour) the average distance to schools is only 15 minutes and even for the town on the 75th percentile, it is less than twenty minutes.

The last group of variables in Table 4.1a gives information, at the municipality level, on the *HC*. The average distance to a *HC* is just short of half hour. The average age of a *madre comunitaria* is 38, and she has an average experience of 7 years. Interestingly, the average monthly fee is, in our data, only 4500 pesos (less than 2 US$). According to the ICBF, the monthly fees should be between 7,000 and 14,000 pesos: however, as also confirmed by our field workers, in our towns there are many *HC* where the *madre comunitaria* charges very little, if at all.

In Table 4.1b we report information at the household level *for the children used in the analysis.* As poorer families have more children than less poor families, the statistics in Table 4.1b over-estimate the level of poverty as measured, for instance, by the share of food consumption in total. The population we are dealing with is obviously very poor. The average family size is 7. Average consumption is about 114 US dollars per month, which includes our estimates for consumption of food produced or acquired as remuneration of work.[7,8] The average share of food consumption in total consumption is 73%. 85% of our households report consumption 'in-kind'. On average, this accounts for 25% of food consumption. The education level of household heads and spouses is very low: in our sample, 20% of the children have a mother with no education.

---

[7] The data base contains information on the quantities of 98 types of food consumed and on prices of each of these commodities at the town level.

[8] According to the 2003 Quality Life Survey, the average consumption in Colombia is $432, excluding auto-consumption.

Table 4.1b
Household level statistics

	Mean	Standard deviation	Median	P75
Monthly Consumption	114	67.2	100.5	142.1
Food share	0.727	0.146	0.753	0.835
Family size	6.7	2.5	6	8
% of mother with no education.	20.1	0.40	-	-

The numbers in this table refer to averages of household variables for the children in our samples. Consumption is reported in US$ using an exchange rate of 2,600 to the peso.

In table 4.1c, we report mean, median and standard deviation of the z-score for height per age, weight per age, weight per height and leg length. Z-scores are computed as the variable of interest (say height) minus the median value for the same variable for children of the WHO/CDC reference population of the same age (or height in the case of weight per height) and gender, divided by the standard deviation of the same group of children of the WHO/CDC reference population. This is the normalization most commonly used in the literature

If we take at face value the figures in Table 4.1c, we observe that our population presents substantial deficits in height per age, which are much reduced for weight per age and almost non-existent in weight per height.

In the case of leg length, which will be one of our outcomes, we do not have a normalization or standard Z-scores. In Table 4.1c we report the mean, median and standard deviation of this variable in our sample. This variable has been shown to be an important marker of the stock of malnutrition as well as a very good predictor of illnesses in adult age.

Table 4.1c
Individual level statistics

	Mean	Standard deviation	Median	P25
Height per age	-1.24	1.11	-1.11	-1.95
Weight per age	-0.80	1.03	-0.86	-1.49
Weight per height	-0.03	0.90	-0.05	-0.60
Leg length	45.18	6.31	45.40	40.4

Chronic malnourishment in children is typically defined in terms of height per age. A child is defined as 'chronically malnourished' if is her or his Z-score for height per age is less than -2 standard deviations from the median of the WHO/CDC reference population of the children of the same age and gender. A child whose Z-score is between -2 and -1 is defined 'at risk' of malnutrition. The Z-scores for weight per age and weight per height are used to define different forms of malnutrition. Height is thought to be the better suited to capture long run trends and the stock of malnutrition.

In Table 4.2a, we report the percentage of children that are defined as chronically malnourished, by age and gender. In this Table we use the measurements taken at the time of the first follow up. The first noticeable feature is the prevalence of malnutrition in our sample. For all age groups except the youngest the fraction of 'chronically malnourished' children is above 20%. There are no large differences in the prevalence of malnourishment between boys and girls: the only significant differences are at age 1 and 6. Interestingly, when we re-do the exercise for the children observed in the baseline, we obtain a table very much similar to Table 4.2a. That is, the only significant difference between genders (and of roughly the same order of magnitude) is for children aged 1. Notice that children aged 2 in follow up are (to a large extent) the same as those aged 1 in the baseline. Boys seem to have caught up with girls by age 2.

		Table 4.2a Proportion of chronically malnourished children (z score of height per age < -2)	
Age	Girls	Boys	Difference
0	0.121	0.154	-0.034
	(0.030)	(0.034)	(0.038)
1	0.216	0.367	-0.152
	(0.038)	(0.42)	(0.053)
2	0.207	0.243	-0.036
	(0.039)	(0.038)	(0.045)
3	0.241	0.236	0.005
	(0.025)	(0.031)	(0.036)
4	0.242	0.250	-0.008
	(0.035)	(0.034)	(0.047)
5	0.215	0.260	-0.045
	(0.037)	(0.026)	(0.039)
6	0.209	0.276	-0.067
	(0.028)	(0.032)	(0.032)

Standard errors in parentheses computed taking into account clustering at the town level
Source: Familias en Acción follow up survey.

.

		Table 4.2b Proportion of globally malnourished children (z-score of weight per age <-2)		
	Boys		Girls	
Age	% z-score <-2	% -2<z-score<-1	% z-score <-2	% -2<z-score<-1
00.067	0.067	0.196	0.075	0.119
1	0.151	0.337	0.115	0.265
2	0.104	0.273	0.106	0.245
3	0.082	0.352	0.113	0.291
4	0.068	0.375	0.129	0.302
5	0.112	0.331	0.075	0.300
6	0.090	0.384	0.076	0.235

Standard errors in parentheses computed taking into account clustering at the town level
Source: Familias en Acción follow up survey.

Table 4.2c				
Weight per height				
	Girls		Boys	
Age	% z-score <-2	% -2<z-score<-1	% z-score <-2	% -2<z-score<-1
0	0	0.055	0	0.081
1	0.024	0.161	0.025	0.130
2	0.015	0.112	0.016	0.131
3	0.00	0.085	0.014	0.099
4	0.00	0.107	0.014	0.122
5	0.011	0.098	0.003	0.090
6	0.006	0.087	0.015	0.100

Standard errors in parentheses computed taking into account clustering at the town level
Source: Familias en Acción follow up survey.

In Table 4.2b and 4.2c, we report the percentage of children with a Z-score less than -2 and between -2 and -1 for weight per age and weight per height by age and gender. These tables confirm that the deficit in terms of these measures is less pronounced than in terms of height per age. This will partly justify our focus on height per age when we will be looking at our results.

In Table 4.3 we report the percentage of children who attend a *HC*. Two features are worth stressing. First, attendance rates have an inverted U shape, being highest at age 3. They are particularly low for very young children. Second, the programme does not seem to be extremely popular. Even for age 3 children, attendance rates do not achieve 50%. These numbers stand in stark contrast with the anecdotic evidence that claims that most children from *SISBEN 1 and 2* households attend *HC*.

For each child that does not attend an *HC* we ask the main reason for not attending. In Table 4.4, we report the percentages reporting a specific reason, for different age groups. The most popular reason for not attending is the availability of child care at home. As to be expected, this is particularly relevant for the youngest children. For the oldest children, the importance of the 'other' reasons is explained by the fact that a significant proportion of these children are in school. Interestingly for our analysis, the distance from the nearest *HC* appears as an important reason for not attending a *HC*.

Table 4.3

Percentage of children attending *Hogares Comunitarios*

Age	Girls	Boys
0	3	3
1	19	16
2	38	44
3	46	48
4	36	34
5	20	24
6	11	11

Table 4.4

Reasons for not attending a *HC*

	Age: 0-1	Age:2-4	Age:5-6
Available caregiver at home	71%	38%	21%
No *Hogar* or too far away	16%	27%	19%
Cannot afford fee	7%	10%	5%
Does not like food	2%	6%	5%
Other	4%	18%	50%

In Table 4.5, we report the some of the statistics of the distribution of travel distances to the nearest *HC*. As can be seen there is a substantial amount of variation in distances, especially in rural areas. In these areas, the 75th percentile is 12 times the 25th percentile! In urban areas, however, the variation is much more limited.

Table 4.5

Distribution of travel distances in minutes

	All	Urban	Rural
25th perc.	5	4	5
Median	10	5	20
Mean	26	11	43
75th perc.	30	10	60

3. The impact of *Hogares Comunitarios*

We start our analysis of the impact of *HC* from the first stage regressions that model the take up of the programme. We then present, for several outcomes of interest, the results we obtain on the overall impact of the programme. The results on children height, weight and school achievement are then complemented with various pieces of evidence that support our identification strategy. Our estimates of the impact are, for most specifications, obtained using Instrumental Variables, where the matrix of instruments is formed by non-linear prediction of the treatment variable. For instance, in the case of exposure, as defined in section 3.1, we run a Tobit where the fraction of life spent in a *HC* by a child is explained by a number of control variables and a polynomial in the distance variables. We then predict exposure using this model and use it and its square as an instrument.[9] The tests of over-identifying restrictions we report in the tables are Sargan statistics robust to the presence of heteroscedasticity of unknown form and cluster effects. Analogously, for attendance we use predictions from a Probit model, while for the number of months we use predictions from a Negative Binomial specification. All the results we report in this section are obtained from linear specification in the outcome equation and do not consider the possibility of heterogeneity of programme's effects. These issues, and others, are taken up in the next section. In the last part of this section, we discuss possible mechanisms through which the

[9] Standard errors are adjusted to take into account the 2SLS type of technique used (and, obviously, cluster effects at the village level).

programme might be operating, such as female labour supply, and discuss the robustness of our identification strategy.

5.1 First stage regressions

As we mentioned above, we use three different definitions of 'treatment'. First we look at whether a child is currently attending a *HC*. We then consider exposure, that is, the number of months a child has spent in a *HC* divided by the child's age in months. Finally, we also consider the number of months spent in a *HC*. When considering long run outcomes, such as the school outcomes of children aged 8 to 17, we consider the number of months and a binary indicator that tells us if the child has ever attended.

The results on the first stage are reported in Tables 5.1 and 5.2. While for instrumenting purposes we use predictions of the non-linear models in this tables, here we report also the results obtained with linear models. In the case of attendance we use a Probit, in the case of exposure we use a Tobit model and for months of attendance, we use a Negative Binomial regression.

As mentioned above we use two sets of instruments. First, for each single child, regardless of whether he or she attends a *HC*, we consider the distance to the nearest *HC* from the child's household residence in minutes and its square. Second, in each town, we compute the average distance for all the households in the town to the nearest *HC*. Again we consider both the level and the square of this variable. The first stage regressions, both in their linear and non linear incarnations, included, in additional to these identifying instruments all the controls used in the outcome equations. The fact that the distance from the household to the nearest *HC* affects participation in the programme is not surprising. This is evident even from the self-reported reasons not to participate in HC that can be found in Table 4.4. The average distance from the households to the HC in the municipality measures the density of the programme in the town. This can also influence participation because if the closest HC is full, the density of the programme on the town becomes relevant for the participation.

Table 5.1

First stage regressions

Children 0-6

	Exposure		Currently attending		Number of months (children 0-6)	
	OLS	Tobit	OLS	Probit	OLS	Negative binomial
Distance from the nearest *HC*	-0.17 (0.03)	-0.54 (0.08)	-0.36 (0.06)	-1.88 (0.36)	-4.85 (1.96)	-1.69 (0.500)
*(*Distance from the nearest *HC)^2*	0.05 (0.01)	0.08 (0.05)	0.10 (0.02)	0.37 (0.21)	0.871 (0.495)	0.345 (0.174)
Average min. dist. from an *HC* in town.	-0.65 (0.16)	-1.34 (0.21)	-1.16 (0.26)	-4.27 (1.06)	-8.48 (2.11)	-1.62 (0.430)
(Average min. dist. from an *HC* in town.) ^2	0.54 (0.16)	1.18 (0.26)	1.08 (0.26)	3.95 (1.25)	2.92 (0.960)	0.357 (0.164)
N	2445	2445	2554	2554	2028	2028
(Joint significance for the four instruments)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)

Standard errors in parentheses are clustered at the town level.
Distance is measured in minutes and divided by 100.
The mean of the Negative Binomial model is linear exponential
F is a test for the joint significance for the four instruments, for which we report the p-value

In Tables 5.1 and 5.2, we report only the coefficients on the distance variables for a selected set of first stage regression. The reported regressions refer to the equations for children height. Those for other outcomes are slightly different because of a small number of children that might have a missing value for one of the outcomes. The regressions for female labour supply are very similar and are available upon request. In addition to the point estimates and their standard errors, we report the significance of a test of the joint significance of our four instruments. The conclusion to be drawn from these tables is that the instruments are indeed important determinants of the participation choice.

Table 5.2
First stage regressions
For long run effects

	(children 8-12)			(children 13-17)		
	Number of months		Ever attended	Number of months		Ever attended
	OLS	Negative binomial	Probit	OLS	Negative binomial	Probit
Distance from the nearest *HC*	-11.40 (2.45)	-1.84. (0.39)	-1.19 (0.24)	-13.86 (2.90)	-2.31 (0.54)	-1.03 (0.24)
(Distance from the nearest *HC*)^2	2.58 (0.70)	0.50 (0.17)	0.30 (0.06)	3.24 (1.08)	0.64 (0.23)	0.23 (0.08)
Average min. dist. from an *HC* in town.	-88.89 (26.32)	-4.69 (1.87)	-4.63 (1.60)	-70.58 (20.81)	-2.82 (2.17)	-4.18 (1.56)
(Average min. dist. from an *HC* in town.) ^2	71.47 (30.74)	0.66 (2.17)	2.78 (1.98)	56.25 (24.48)	-1.46 (3.09)	2.18 (2.03)
N (F p-value)	3410 (0.000)	3410 (0.000)	3419 (0.000)	2975 (0.000)	2975 (0.000)	2991 (0.004)

Standard errors in parentheses are clustered at the town level.
Distance is measured in minutes and divided by 100.
The mean of the Negative Binomial model is linear exponential
F is a test for the joint significance for the four instruments, for which we report the p-value.

In addition to the reported coefficient, the first step regressions include a large number of control variables both at the municipality and the household level. The full specification of the regression can be found in the Appendix. The most interesting element that comes out of the analysis of these tables is the fact that the poorest children are those that seem to be more likely to attend the HC. Variables such the education levels of both parents confirm this.

5.2 Programme Impacts

In this section we present our estimates of the programme impacts. We will start with the impacts on anthropometric measures: height, weight and leg length. We then move to the impacts on long run effects on the school achievements of older children.

5.2.1 Height per age

We start our analysis with one of the most interesting outcomes: height per age. In Table 5.3 we report our estimates of the impact of *Hogares Comunitarios*. The left hand side of our equation is the Z-score for height per age. While in the Table we report only the estimates of the coefficients most relevant for our discussion, all the specifications include a variety of controls for the outcome of interest, including a polynomial in age, the height of the child's mother, her education and several other household and community level variables. The complete set of coefficients is reported in the Appendix. In the last rows of the table, we report a test of the over-identifying restrictions that can be constructed as we use both the individual distance and the average distance in the municipalities. In section 5.3, we will explore the issue of the validity of our instruments.

Table 5.3
Effect of HC on height per age

	Attendance		Number of months		Exposure	
	OLS	IV	OLS	IV	OLS	IV.
Impact	-0.059 (0.050)	0.486 (0.156)	-0.002 (0.002)	0.013 (0.007)	-0.042 (0 .099)	0.780 (0.340)
N	4557	4557	4384	4384	4384	4384
(p-value)		0.552	-	0.825	-	0.314

Standard errors in parentheses are clustered at the town level. p-value refers to the over-identifying restriction test)

In the first two columns, we report the coefficient we get on current attendance, that is, a dummy that is zero for children who are not currently attending a *HC* and one for those who are. In columns 3 and 4, we report the coefficient on the number of months. Finally, the last two columns contain the coefficients on exposure, that is, the number of months a child attended HC over her age in months. In Columns 1, 3 and 5, we report the estimates we obtain by OLS, that is, without controlling for the endogeneity of participation into a *HC,* while in columns 2, 4 and 6 we report the IV estimates.

The OLS estimates are small and negative numbers. However, when we instrument participation, months or exposure using distance from a *HC* and average distance in a town the estimated coefficient is positive and significantly different from zero. We never reject the over-identifying restrictions.

The effects we report are large. Current attendance is estimated to have an effect of 0.4486 standard deviations on the Z-score. This effect corresponds to 2.36 centimetres for a boy (2.39 for a girl) aged 72 months. Even more interestingly, if we look at exposure we obtain that the effect of having attended a HC during the first six years of life is 3.78 centimetres for a boy (3.83 for a girl) aged 72 months. The fact that the OLS estimates are negatively biased is an interesting result in its own right. This piece of evidence is consistent with the evidence from the 1998 study we mentioned in the introduction, which did not find significant differences between children attending *HC* and children of 'similar socio-economic background'. It is also consistent with the fact that the participation equations seem to indicate that the poorest households are those that are sending children to *HC*. This indicates that the programme is remarkably well targeted, in that the households most in need seem to self-select as *HC* customers. It seems that the programme allows the poorest children to keep up with their better off peers.

As we mentioned above, all specifications include a large set of controls. The complete set of estimates is reported in the Appendix. It is worth noting that the specifications we have estimated included a large number of town and areas specific variables. The reason for our un-parsimonious specification in this respect is our worry that our instruments could capture some unobserved feature of the environment where the households live and have a direct effect on the outcome of interest. While such an identification assumption is clearly un-testable, we will provide indirect evidence of the validity of our approach in section 5.3. The remaining coefficients have, reassuringly, the expected sign: tall mothers have tall children, as do better educated mothers and so on. Notice that the polynomial in the age of the child (interacted with gender) is strongly significant, indicating that difference between our population and the reference population is not invariant to age.[10] Several of the environmental variables, including those referring to the availability of health infrastructure, seem to be important determinants of children height (and other outcomes).

5.2.2. Leg length and weight per age

The second outcome we look at is leg length. As we mentioned repeatedly, there are no common standardizations for leg length, so that controlling for a flexible and gender specific function of age is crucial. As with Table 5.3, In Table 5.4 we only report the estimates of the impact. The complete set of results can be found in the Appendix.

Table 5.4
Effect of HC on leg length

	Attendance		Number of months		Exposure	
	OLS	IV	OLS	IV	OLS	IV.
Impact	-0.383 (0.118)	0.288 (0.416)	-0.008 (0.005)	0.037 (0.024)	-0.467 (0.299)	1.324 (0.944)
N	3813	3813	3650	3650	3650	3650
(p-value Overidentifying Restriction test)	-	(0.693)	-	0.755	-	0.651

Standard errors in parentheses are clustered at the town level.

As for height, we report the effects of attendance, number of months and exposure. The results on leg length are roughly consistent with those on height per age. This is to be expected, as both variables are good indicators of the stock of past nutrition. Three aspects are worth noting, however. First, some of the OLS estimates are *significantly* negative. Second,

[10] This is a general finding in nutritional data from developing countries, see Shrimpton (2001)

the positive effects obtained by IV are not as precisely estimated as in the case of height per age. This might be a consequence of the smaller number of observations for which the leg measurement is available. Finally, notice that the effect of current attendance is remarkably smaller than that of the number of months or what we define as 'exposure'. This is consistent with the fact that the variable is more reactive to long run than short run nutrition.

Finally, on Table 5.5, we report the estimates of the programme impact on weight per age. In this case, we do not find any significant effect of the program. As we mentioned above, weight per age is more reactive to short run nutrition, than long run interventions. The result, therefore, is not inconsistent with those mentioned earlier.

Table 5.5
Effect of HC on weight per age

	Attendance		Number of months		Exposure	
	OLS	IV	OLS	IV	OLS	IV.
Impact	0.006 (0.046)	0.274 (0.170)	-0.002 (0.002)	0.001 (0.006)	-0.003 (0.097)	0.132 (0.340)
N	4557	4557	4384	4384	4384	4384
(p-value Overidentifying Restriction test)		0.207	-	0.432	-	0.699

Standard errors in parentheses are clustered at the town level.

5.2.3. Long run effects

We finally move to the analysis of the effect of the programme on long run outcomes. In particular, we consider as an outcome two measures of school achievement by children aged 8 to 17. The first measure is current attendance in school. The second measure is the probability of having advanced a grade between 2002 and 2003. Attendance, obviously, is only a necessary condition for advancing a grade. As mentioned above, we consider as treatment whether the child has ever attended in his/her life a *HC* and the number of months spent in a *HC*.

We instrument the 'treatment' variable, as we did for the other Tables, following the IV procedure discussed in Section 3. In particular, to model the 'ever attended' variable we use a Probit, while for the number of months we used a negative binomial regression. The first step results were reported in Table 5.2. As with the other specifications, in addition to the treatment variable, we also consider a variety of controls that include parental education variables, age, and many village infrastructure variables, such as the number of children per teacher in the village. Finally, rather than considering all children together, we split the sample and consider separately children aged 8 to 12 and those aged 13 to 17. Attendance rates among the first group are very high in our sample, while they start dropping quite dramatically at age 13. Once again, in table 5.6, we only report the estimates of the long run effects of the programme. The coefficients of the complete specification are relegated to the Appendix.

The results we obtain for older children (in the lower panel of Table 5.6) are quite remarkable. While for the younger group we are unable to identify any significant effect,[11] we find substantive effects for the older group. For instance, having attended an *HC* when less than 6, increases the probability of being in school of children aged 13 to 17 by 0.198. Given that the average attendance rate for children aged 13-17 in our sample is 0.63, this effect is very strong. Similarly, attending a *HC* increases the probability of progressing a grade by 0.165 for the same group, though the effect is not statistically different from zero at the 95%. More information can be obtained by looking at the exposure to the HC program. The last two columns of Table 5.6 show the effect of the number of months in school attendance and achievement. The average number of months attending a HC is 29 for older children that ever attended a HC. According to the estimates in Table 5.6, having attended a HC during 29 months increases the probability of school attendance in 0.208 percentage points.

Notice that, consistently with the evidence on nutritional outcomes, the OLS estimates seem to be affected by a negative bias, indicating that the poorest and most disadvantaged children have been the customers of *HC* for some time now.

---

[11] Attendance rates are very high for young children. In this sense, higher ages, when children are starting to drop out of school, are more interesting.

Table 5.6
Effect of HC on school achievement

	Children 8-12			
	Ever attended		Number of months/72	
	OLS	IV	OLS	IV.
Effect on probability of attending school	0.044 (0.011)	0.045 (0.070)	0.089 (0.021)	0.100 (0.128)
Probability of progressing a grade between 2002 and 2003	0.008 (0.015)	0.02 (0.067)	0.025 (0.028)	0.022 (0.137)
	Children 13-17			
	Ever attended		Number of months/72	
	OLS	IV	OLS	IV
Effect on probability of attending school	0.026 (0.020)	0.198 (0.093)	0.040 (0.040)	0.517 (0.170)
Probability of progressing a grade between 2002 and 2003	0.01 (0.019)	0.165 (0.100)	0.055 (0.036)	0.411 (0.181)

Standard errors in parentheses are clustered at the town level

We next consider the effects of HC separately boys and girls aged 13 to 17. The results, shown in Table 5.7, indicate that the programme affects the probability of passing a grade for girls but not for boys. The probability of school attendance increases by a striking 0.28 percentage points for girls, and 0.19 for boys. In general, it seems that HC has a larger long-term impact on girls than boys.[12] If it were the case that HC increased school attendance by improving children's health, this finding would indicate that girls school attendance is

[12] When we estimate the effect of the programme on the nutritional status of boys and girls separately, we do not find strong gender differences. Indeed, if anything, the effect generally goes in the opposite direction, with slightly stronger observed effects on height per age and leg-length for boys than for girls.

relatively more sensitive to their health. This interpretation is consistent with evidence from Pakistan reported by Alderman *et al.* (2001a), who find that an improvement in nutritional status increases the school attendance of girls but not of boys.[13]

Table 5.7
Effect of HC on school achievement

	Boys 13-17			
	Ever attended HC		Number of months/72	
	OLS	IV	OLS	IV.
Effect on probability of attending school	0.015 (0.026)	0.193 (0.105)	0.0043 (0.050)	0.506 (0.188)
Probability of progressing a grade between 2002 and 2003	-0.014 (0.025)	0.045 (0.098)	0.008 (0.048)	0.148 (0.182)
	Girls 13-17			
	Ever attended HC		Number of months/72	
	OLS	IV	OLS	IV
Effect on probability of attending school	0.035 (0.024)	0.284 (0.115)	0.026 (0.043)	0.393 (0.206)
Probability of progressing a grade between 2002 and 2003	0.036 (0.029)	0.35 (0.13)	0.104 (0.046)	0.667 (0.227)

Standard errors in parentheses are clustered at the town level

[13] There is also evidence of other characteristics having differential effects across boys and girls. For instance, we find that distance to school has a relatively large effect on girls' rather than boys' school attendance. This is also found in data from Pakistan (Alderman *et al.* 2001b). Of course, many other explanations are also possible. It could be that HC did not improve school attendance through improving children nutritional status, but the benefit comes indirectly through other siblings. It could be that older girls living closer to a HC benefit from the programme because they do not have to take care of their small siblings who attend the HC instead. However, we detect an important impact of HC even on older girls that live in households without small siblings.

5.2.3 Female labour supply

In this section, we look at the effect of the programme on female labour supply, both in terms of employment rates and number of hours worked. This might be important, as the childcare aspect of the program could allow mothers to work and earn additional resources that might benefit the child indirectly.

Table 5.8
Effect of HC on female labour supply

	Employment rates		Hours	
	Simple probit	Bi-variate probit	OLS	IV
Programme effect	0.121 (0.027)	0.377 (0.061)	16.639 (3.914)	75.302 (13.504)
N	2936	2936	2920	2920

Standard errors are corrected for clusters at the municipality level. The coefficient on the probits are marginal effects

We report our estimates of the effects of *HC* on female labour supply in Table 5.8. As outcomes, we consider both employment rates and number of hours; as treatment we define a binary variable that is one if the mother has at least one child currently attending *HC*. In the case of employment rate, we estimate a bi-variate Probit for employment and programme participation. Consistently with our IV strategy, the distance variables enter the participation decision, but not the employment rate decision. The assistance variable enters the index for the employment decision.

The effects of the programme on female employment rates and participation are remarkable. Once again, taking into account the endogeneity of programme participation increases substantially the estimated effect of the program. In the case of employment rates, the probability of employment increases from 0.12 to a staggering 0.37. We see similar evidence about the effect on hours worked. The programme increases the number of hours worked by 75 hours per month. In our sample, 37% of the mothers are working. They work, on average, 39 hours per month.

5.3 Is the identification strategy credible?

In this section we present some evidence that justifies our identifying strategy. First, we supplement the evidence provided by the tests of over-identifying restrictions and identify the parameters of interest with different sets of instruments. Second, as we mentioned above, we consider whether, using our identification strategy, the programme would be shown to have an effect on variables on which it should not.

5.3.1 Impact results using household and municipality instrument separately

The impacts presented above were estimated using two instruments: distance of the household to the nearest HC and average minimum distance in the municipality between the household and the nearest HC (and their square). Here we show the estimates obtained using only one instrument at a time. Some researchers would trust more the average minimum distance as an instrument because the household does not determine it. Consequently, it is more likely to be independent of household unobserved characteristics.

Table 5.9

Effect of HC on anthropometrics. Estimates using household distance and municipality average minimum distance separately as instruments

	Attendance		Number of months		Exposure	
	Household	Municipality	Household	Municipality	Household	Municipality
Height for age z-score	0.438 (0.161)	0.602 (0.205)	0.013 (0.007)	0.012 (0.009)	0.644 (0.342)	0.923 (0.475)
Leg length (cms)	0.3 (0.419)	0.185 (0.608)	0.047 (0.023)	0.03 (0.031)	1.516 (0.966)	1.056 (1.430)
Weight for age z-score	0.278 (0.168)	0.33 (0.216)	0.003 (0.006)	-0.001 (0.008)	0.121 (0.354)	0.107 (0.454)

Standard errors in parentheses are clustered at the town level.

Table 5.9 shows the impact of HC on anthropometrics using either household minimum distance or average minimum distance in the municipality as an instrument. Consistently with the non-rejection of the over-identifying restrictions test reported in the previous three tables, the estimates are quite similar to the ones that we find when we use the both instruments at the time. Obviously, these estimates are less accurate than those obtained with both instruments used simultaneously. The impact estimates using household and municipality instruments do not differ by more than one standard deviation. More interestingly, the estimate of the impact on height using the average minimum distance in the municipality as an instrument is further away from the OLS results than that obtained using the household distance. This is important because the municipality level distance would be more credible as an instrument for some researchers because the household itself does not determine it. We believe that the similarity between the results obtained with the household and municipality level instrument provides a strong support for our identification strategy.

5.3.2 Evidence using municipality fixed effects

Table 5.10

Impact of HC on Height for Age. Estimates using municipality fixed effects (FE)

Attendance		Number of months		Exposure	
Without FE	With FE	Without FE	With FE	Without FE	With FE
0.438 (0.161)	0.415 (0.182)	0.013 (0.007)	0.013 (0.007)	0.644 (0.342)	0.594 (0.40)

Standard errors in parentheses are clustered at the town level.
Without FE estimates are copied from Table 5.3
Instrument: household distance to the HC

The programme administrators could follow a compensatory rule and keep more HC opened in municipalities that have worse municipality infrastructure. If that was the case, our identifying assumption would be invalid unless we are adequately controlling for municipality infrastructure. In order to examine this issue, we estimate the impact of the HC controlling for municipality fixed effects, and using only distance from the household to the HC as instrument. The results, shown in Table 5.10, show that the introduction of fixed effects only increases the estimated standard errors but hardly changes the parameter estimates. Consequently, our results are robust to the correlation between household distance and unobserved municipality variables.

5.3.3 Evidence on birth weight and mother height

Table 5.11
Effect of HC on birth weight (kilograms)

	Attendance		Number of months		Exposure	
	OLS	IV	OLS	IV	OLS	IV.
Impact	0.057 (0.081)	-0.041 (0.205)	-0.0037 (0.0024)	-0.0096 (0.0074)	-0.119 (0.135)	-0.468 (0.434)
N	1140	1140	1096	1096	1096	1096
(p-value Overidentifying Restriction test)		0.673		0.268		0.506

Standard errors in parentheses are clustered at the town level.

As an additional check on the validity of our instrumental variable approach, we estimate, using the same approach as in section 5.2, the impact that *HC* has on two variables on which it should not have an effect: children's birth weight[14] and mother height. If we were to find that current attendance or exposure to *HC* had an effect on birth weight or mother height, one would suspect that the result is driven by the wrong choice of instrument. It could be that children from households that live closer to a *HC* are healthier (and therefore heavier when they are born and taller at later ages) for some reason other than the exposure or participation into *HC*. The evidence in this section, therefore, constitutes an important specification test of our IV approach. In table 5.11 we report the results we obtain by OLS and IV of the effect of *HC* on birth weight. Both the OLS and the IV estimates tend to be

[14] We refer to birth weight of the same children for whom we have estimated the impact of HC on anthropometrics

negative, small in size and no statistically significant. Consequently, the IV results suggest no significant effect of HC on birth weight. This evidence, therefore, supports our claim that the instrument we have been using is a valid one.[15]

Although birth weight is a pre-intervention variable, it is possible that a household might be shifting resources from a child already attending a HC towards a pregnant mother. In this case, attendance by another child could be increasing birth weight of a child still to be born. Also, due to increases in labour supply, the pregnant mother could be better nourished. However, notice that if that criticism would be true, we would be expecting to find positive impact of the programme on birth weight, but our estimates are negative and not statistically significant.

Table 5.12
Effect of HC on mother's height (cms)

	Attendance		Number of months		Exposure	
	OLS	IV	OLS	IV	OLS	IV.
Impact	-0.111	1.709	-0.003	-0.028	-0.137	0.07
	0.360	1.791	0.014	0.050	0.810	3.551
N	1585	1585	1549	1549	1549	1549
(p-value Overidentifying Restriction test)		0.727		0.197		0.480

Standard errors in parentheses are clustered at the town level.

We perform a further robustness check. Mother's height is an important determinant of children's height. However, it can also be considered a pre-intervention variable in the sense that a child attendance to HC cannot affect his or her mother's height. Still, we could find an effect if our instrument was correlated with certain household unobservable characteristics. Assume that a mother inherits part of the nutritional knowledge or awareness

[15] Notice that the number of observations reported in Table 5.11 is smaller than in the previous tables. This is partly because of missing values in birth weight but mostly because; unlike for height, we obviously just have one observation on birth weight for each kid.

that she received when she was young. Assume further that (i) mothers that, when they were young, lived far away from the places where today there is a HC, still live far away from a HC, and (ii) mothers living far away from a HC have worse nutritional knowledge. In that instance, our instrument will be invalid and we would find a positive effect of HC on mother's height. However, Table 5.12 shows that there is no impact of HC on mother's height, validating further our evaluation strategy.

In addition to these results, we also considered the possibility that our instrument is correlated with household and municipality observed characteristics, conditional on the rest of the variables. Though clearly this does not test for correlation with unobservable characteristics, we would find our identification assumption untenable if the instrument was correlated with many observable characteristics. However, we can clearly see from Table 5.13 that this is not the case.

Starting with the individual variables, which are probably the most important for us, we find that the distance to the HC is not partially correlated with mother's height, or with her education or head of the household's education. This is very important as we could think that some unobservable characteristics such as awareness of the importance of nutrition will be correlated with education. Regarding municipality variables, we find that distance to the HC is correlated with municipality wages, but it is not with municipality infrastructure that influences child nutritional status such as coverage of piped water network and hospital.[16] We do find that the distance to the HC is strongly correlated with the distance to the nearest school and clinic, but not with distance to the town hall. This suggests that the HC are not all located close to the town hall, but scattered around the municipality close to where schools and health clinics are. Though we do not show it in Table 5.13, the distance of each household to the *HC* is much more correlated with the distance to the school than with the distance to the health clinics.

To sum up, we find that the distance to the HC is partially correlated to other distances (school and health clinics), but there is little evidence that it is correlated with other very important household characteristics. Regarding municipality characteristics, we find that the

[16] Attanasio *et al.* (2004b) estimate the influence of municipality infrastructure in child's nutritional status with this dataset. They find that the presence of a hospital and the coverage of the piped water network are important determinants of children's nutritional status.

distance to the HC is correlated only with municipality wages. The lack of correlation between our instrument and other important observable municipality level characteristics is also reassuring. In any case, given that our main results are robust to the introduction of municipality fixed effects, the correlation of our instrument with some municipality level variables is not particularly worrying.

Table 5.13

P-Values of joint significance of these variables in a regression of distance to the nearest HC from the household

Children age, gender and birth order	0.23
Distance to town hall	0.40
Distance to nearest school and clinic	0.00
Mother's height	0.31
Mother's and head of household's education	0.45
Mother's and head of household's age	0.54
Head of household's without spouse	0.20
Presence of hospital, coverage of sewage and piped water network, index of quality of life, number of students per teacher	0.25
Municipality wages	0.02
Altitude	0.15

6. Heterogeneous impacts

The recent literature on policy evaluation has highlighted the role of heterogeneous impacts on the identification and estimation of treatment effects. The literature has clarified that different methods will, in general, identify different parameters when the benefits of the programme depend on individual characteristics (see Bludell, Dearden and Sianesi 2004 for an up to date review).

The benefits of participation in HC could be heterogeneous across the population. The programme could benefit more children from poor backgrounds, or children who have a larger unobserved tendency to illness and malnutrition. If families take into account these idiosyncratic benefits when taking decisions about children participation into the program, the IV estimates presented before would not identify the average treatment effect, but a Local Average Treatment Effect (Imbens and Angrist, 1994). Control function methods have been advocated as a method to estimate the average treatment effect under heterogeneous impacts but requiring stronger assumptions than IV.

From a policy perspective, when impacts are heterogeneous, it is not only important to estimate the Average Treatment Effect, but also to estimate what types of people benefit differently from the program. There are now several papers that try to estimate differential impacts by individual characteristics (Aakvid *et al.* 2002; Carneiro *et al.* 2003; Blundell, Dearden and Sianesi 2004, Moffitt 2004). In what follows, we explore this issue in two different ways: through interactions between observable characteristics and treatment, and through the use of quantiles methods. To take into account the endogenous selection into the programme we use methods based on control functions.

Before presenting our results obtained with control functions, however, we report the results that can be obtained by propensity score matching. There are several reasons to look at these results. First, propensity score matching has become an extremely popular method to use in the evaluation of public policies when a randomised control group does not exist. Second, PSM does allow for heterogeneous programme effects in a simple and flexible way. However, the assumptions needed for PSM to work are very stringent. We will be arguing that they are not satisfied in our application.

For the sake of brevity, we will focus on the most significant results of Section 5, those on height per age.

6.1 Propensity score matching

As is well known, matching methods assume that conditional on observable variables, the assignment of a programme is random. Given this assumption, Propensity Score Matching (Rosenbaum and Rubin, 1983) uses the propensity score to re-weight the untreated sample. Propensity Score Matching (PSM) improves over OLS with a dummy variable for treatment because (i) PSM is a semi-parametric method, (ii) it does not require homogenous effects to identify the parameter of interest, and (iii) it does not extrapolate outside the common support. However, as OLS, it requires that selection take place on observables only.

Reporting PSM results allows us to check for the lack of common support. The parametric models previously used will extrapolate outside the common support. Our previous results could be seriously biased if the common support was narrow (Heckman, Ichimura, Smith, and Todd 1998).

Table 6.1
Effect of Attendance to HC on height per age
Estimates by Propensity score matching

	Standard Specification	Standard Specification plus house condition variables	Standard Specification plus house condition variables and assets
Impact	-0.098	-0.105	-0.118
	(0.057)	(0.060)	(0.063)
% On the common support	99.25%	98.91%	98.96%

Standard errors in parentheses are clustered at the town level.

Table 6.1 reports the effect of current attendance to HC on height per age using PSM.[17] The second column shows the PSM estimate of the Average Treatment Effect using the same set of covariates for the propensity score than we use for the outcome equation in the IV and OLS models reported above. The PSM estimate of –0.098 is not significantly different from zero and, given the precision of the estimates, substantially similar to the –0.059 we reported in Table 5.3 as our OLS estimate. This similarity is due to two reasons: first, almost everyone in the sample is within the common support, and consequently the OLS model is hardly extrapolating outside the common support. This is clear from looking at Figure 1, which plots the density of the propensity score for participants and non-participants. Second, for most observations the probability of participation is below 0.5, and hence the implicit weights implied by OLS are roughly proportional to the PSM weights (Angrist, 1998).

[17] We use Kernel matching (Heckman, Ichimura and Todd 1997; Heckman, Ichimura and Todd 1998) to estimate the Average Treatment Effect over the common support. We use the programs developed by Leuven and Sianesi (2003).

The third and fourth columns of Table 6.1 add more variables to the specification of the propensity score. The third column shows the results when we include variables about house conditions.[18] The fourth column shows the results obtained when we include dummies for the ownerships of different assets.[19] To include these new variables one needs a very strong assumption: namely one needs to assume that attendance to the HC cannot influence any of the new variables. This assumption is unlikely to hold given that the programme might have (as we saw above) very important impacts on female labour supply and therefore on the amount of resources available to an individual household. However, for a moment, we abstract from that to study the sensibility of our results to the set of covariates. As it is clear, the negative sign seems very robust.

All our results based on methods that require selection on observables, OLS and PSM, yield a negative point estimate of the impact of attending a *HC* on children anthropometrics, although these estimates are small in absolute value and insignificantly different from zero. In fact, Behrman, *et al.* (2003) also found non statistically significant and often negative impacts on anthropometrics of length of exposure to the PIDI programme in Bolivia using a sophisticated matching method that is based on the assumption that selection into different exposures to the programme is based on observables (although it allows participation to be endogenous). As in our setting, children participating in the PIDI programme in Bolivia tend to come from families with lower income and lower parental education. Their evidence together with ours, sheds light on the success of the self-targeting mechanisms of these types of programs. Obviously, this casts certain doubts about the success of methods based on observables to evaluate these types of programs, at least in absence of a baseline. Given these doubt we don't explore the issue of impact heterogeneity using PSM.

This results, of course, do not imply that PSM techniques are never reliable: in our situation, however, where participation into the programme implies complex individual decisions, it seems that unobserved heterogeneity in the selection process can cause important biases.

---

[18] In particular we use variables about the materials of the floor, roof, and walls of the house, whether or not the household gets water and gas by pipe, whether or not the household has sewage system, source of water and energy to cook, whether or not toilet is connected to the sewage system or septic well, and whether or not the household enjoys a service of waste collection.

[19] We use whether or not any household member owns a fridge, sewing machine, colour TV, black and white TV, stereo, motorbike, fan, blender, kerosene lamp, animals, and dummies variables for ownership of the house with property title, without title, or renting.

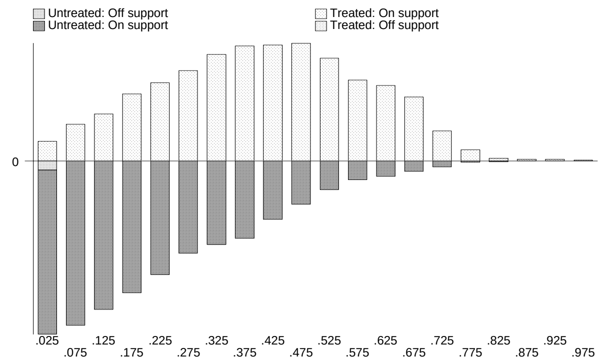

Figure 1. Distribution of the estimated propensity score

6.2 Control functions

The basis of the control function approach is to recover the average treatment effect by controlling directly for the correlation between the error term in the outcome equation with the treatment variable. To do that, one requires an explicit model of treatment choice. The control function approach treats the selection problem as an omitted variable problem and augments the outcome equation by a term to control for this variable omission. The traditional example is the Heckman (1979) sample selection model that augments the outcome equation by an estimate of the Mills Ratio. Using a control function approach we will estimate the impact of HC on both the mean and different quantiles of the conditional distribution of the error term. And when looking at the mean we will allow the effect to be heterogeneous in observable dimensions.

6.2.1 Impact of the programme over the mean outcome

In columns (1) and (2) of Table 6.2, we report the estimates of ATE using the control function approach. In order to compute these estimates, the outcome equation was augmented with a polynomial on the residual of the first stage. [20] The control function estimates the coefficients on exposure to be 1.04 and that for attendance 0.57, which are larger than the IV estimates (0.78, and 0.48), albeit not significantly so. This could be an indication that treatment effect for the average individual is larger than for those that would change treatment status due to changes in distance to the HC.

Table 6.2
Effect of HC on height per age. Control function approach

	Without Interactions		With Interactions	
	(1) Exposure	(2) Attendance	(3) Exposure	(4) Attendance
ATE	1.048 (0.525)	0.573 (0.180)	1.101 (0.505)	0.398 (0.194)
Mother without education, interacted with treatment	-	-	0.542 (0.227)	0.250 (0.114)
Age in months, interacted with treatment	-	-	-0.009 (0.0046).	0.011 (0.013)
Age squared/100, interacted with treatment	-	-	-	-0.018 (0.013)

Standard errors in parentheses are clustered at the town level.

To explore the issue of impact heterogeneity further, we study how the effect of *HC* changes with observables characteristics. We follow Blundell, Dearden, and Sianesi (2004) and we

[20] For exposure, we used the difference between actual exposures and the exposure predicted by the Tobit model as the basis for the polynomial. For attendance, we used a polynomial in the Mills Ratio. For exposure, we used a fifth order polynomial, though the results changed very little with the order of the polynomial. For attendance we used a third order polynomial. The estimate of attendance changed from 0.38 to 0.57 when we went from a second to a third polynomial order. It changed very little when used a fourth or fifth order polynomial, and moreover these polynomial terms were not statistically significant from zero.

interact the treatment variable with the observable characteristic while the first stage residual is entering only additively in the augmented outcome equation. This fiercely exploits the additive structure of the error term. Moffitt (2004) proposes an alternative method that allows the interactions to enter in a more flexible way, but also needs to exploit both the additive structure of the error term and the participation model.[21] As Blundel, Dearden and Sianesi (2004) emphasize, the IV assumptions would require instrumenting the main treatment variable as well as each of the interactions. However, this is unlikely to give good results in practice as the interaction of the observable characteristics and the instrument will normally have little explicative power in the first stage regression.

Columns (3) and (4) of Table 6.2 report the results we obtain when we interact the programme with children age and mother education. We also tried interactions in gender, and a dummy for female head of household but they were not individually or jointly significant. For education, we also tried several dummies, for different levels of education. It turned out that the group with mothers without any formal education (about 20% of our sample) is the only systematically different in this dimension. There is clear evidence that children whose mother does not have any education will tend to benefit more from the treatment. There is also some evidence of differences in impact by children age. In the case of attendance the effect is quadratic in age, with a maximum at two years and a half. The P-Value of the joint hypothesis that the age interactions are both zero is 0.005. As for exposure, we only observe a significant linear term. The quadratic coefficient had a t-statistic of 0.63. The effectiveness of the programme (as measured by the coefficient on exposure) seems to be decreasing in age. In interpreting these coefficients, it should be remembered that there are very few children aged less than 2 who attend the program. In general, these estimates support the view that heterogeneity is present in the benefits of this programme.

6.2.2 Quantiles

To further explore impact heterogeneity , we also study the possibility of differentials effects of HC according to where in the conditional distribution of the error term an individual child is. That is, we ask the question whether HC is especially beneficial for those children that would be otherwise in the bottom part of the conditional distribution of height. Here

[21] Moffitt (2004) uses splines to obtain more flexibility in how the treatment status and observable characteristics interact. However, it turns hard to estimate the parameter of the splines in his empirical application.

we follow Lee (2004) and estimate quantile regressions that are augmented by the residuals of the first stage regression. The standard errors are estimated by bootstrap. We use a second degree polynomial in the estimated residuals, as the third term was not statistically significant different from zero.

Table 6.3
Effect of Attendance to HC on height per age percentiles
Estimates by quantile regression using control functions

	10th	25th	50th	75th	90%
Exposure	1.469	1.022	1.101	0.571	0.095
	(0.839)	(0.663)	(0.585)	(0.569)	(0.778)
Attendance	0.470	0.475	0.476	0.321	-0.118
	(0.244)	(0.190)	(0.192)	(0.189)	(0.233)

Standard errors in parentheses are clustered at the town level.

The results, although are not always very precise, are intriguing. Our reading of Table 6.3 is that *HC* has smaller effect for those children that, conditional on the observables variables, would tend to at the top of the height distribution. This, once again, is consistent with the evidence that are the poorest children those who enrol in HC more frequently as they are the ones that benefit more from it.

7. Conclusions

In this paper we have studied what is probably the largest welfare programme in Colombia. *Hogares Comunitarios*, a community nursery programme absorbs about 250 million US$ per year and, although it has been running since the id 1980s, had never been systematically evaluated. In the first part of the paper, we propose an Instrumental Variable approach based on comparing participant and non-participant children and using the distance of the house of residence to the nearest *HC* as an instrument for participation. The remarkable result we find is that, unlike the OLS estimates, the IV results show important effects of the program. In particular, a 72 months old boy (girl) that has attended a HC during all his life will be 3.78 (3.83) centimetres taller than if he had not attended a HC at all. The negative bias that characterizes OLS (and PSM) estimates is a consequence of the fact that the poorest children are those that select into the program. The program, therefore, seems to be well targeted and to help maintaining the nutritional status of the poorest children.

We show what believe to be credible evidence on the reliability of our instrument: when we perform the same exercise on birth weight and mother height we do not obtain any effect of the program. We find similar results when we instrument using only the average distance to a HC in the municipality. This is important because this instrument is independent of individual household location decisions.

The programme also has a long run effect: we identify some positive effects on children aged 13-17 who in the past have attended a *HC:* these children are more likely to be in school and to have advanced a grade. We estimate positive effects of the programme on mothers employment rates and hours of work. Finally, in the last section, we explore the issue of impact heterogeneity and find, once more that the most disadvantaged children seem to be those that most benefit from the programme.

This evidence is important as the Colombian government is currently considering the possibility of scrapping the *HC* programme in favour of conditional cash subsidies. What we have shown should give an element of caution in implementing these plans.

There is much work that still needs to be done. First, it should be stressed that our results cannot be extrapolated to different contexts. The large majority of *HC* in Colombia operates in large cities and urban areas: our results are from a sample of small towns with an important rural component. Second, as we have evidence on the effectiveness of the alternative programme being considered, it would be interesting to compare them explicitly and model impact heterogeneity, especially in terms of children ages. It might be that the two programs, rather than being substitutes, are complements. It might be that the youngest children are best targeted by the conditional cash subsidies, while children aged 2 to 4 might benefit more from the *HC* programme.

References

Angrist, J. (1998) "Estimating the labour market impact of voluntary military service using social security data on military applicatnts", *Econometrica*, 66, 249-88.

Angrist, J. and Imbens, G. (1995) "Two-stage least squares estimation of average causal effects in models with variable treatment instensity", *Journal of the American Statistical Association*, 90, 431-42.

Attanasio, O., Meghir, C. and M. Vera-Hernández (2004a) "Do Conditonal Cash Transfers Improve Child Health and Nutrition. Evidence from Colombia" Mimeo

Attanasio, O., Gómez, L., Gómez, A., and M. Vera-Hernández (2004b) "Child Health in Rural Colombia: Determinants and Policy Interventions" Economics and Human Biology, *forthcoming.*

Attanasio *et al.* (2003): "Baseline Report on the Evaluation of *Familias en Acción*". EDePo-IFS. Available at: http://www.ifs.org.uk/edepo/wps/familias_accion.pdf .

Aakvik, A, Heckman, J. and Vytlacil, E. (2002) "Estimating Treatment Effects for Discrete Outcomes when Responses to Treatment Vary: An Application to Norwegian Vocational Rehabilitation Programs". Mimeo

Alderman, H., Behrman, J., Lavy, V., and R. Menon (2001a) "Child Health and School Enrollment. A Longitudinal Analysis", *The Journal of Human Resources*, Vol. 36(1), pp. 185-205.

Alderman, J., Orazem, P., and E. Paterno (2001b) "School Quality, School Cost, ad the Public/Private School Choices of Low Income Households in Pakistan", *The Journal of Human Resources,* Vol. 36(2), pp. 304-326.

Behrman, J.R., Cheng, Y. and P.E. Todd (2004): "Evaluating Preschool Programs when Length of Exposure to the Programme Varies: A Nonparametric Approach", *The Review of Economics and Statistics*, February, 86(1): 108-132.

Behrman, J.R. and J. Hoddinott (2004): "An Evaluation of the Impact of PROGRESA on Pre-School Child Height" , IFPRI, mimeo.

Behrman, J., Hoddinott, J., Maluccio, J., Quisumbing, A., Martorell, R., and Stein, A. (2003) "The Impact of Experimental Nutritional Interventions on Education into Adulthood in Rural Guatemala: Preliminary Longitudinal Analysis" Mimeo.

Blundell, R., Dearden, L. and Sianesi, B. (2004): "Evaluating the Impact of Education on Earnings in the UK: Models, Methods and Results from the NCDS" *forthcoming* on *The Journal of the Royal Statistical Society:* Series A.

Buschang PH, Malina RM, Little BB (1986) Linear growth of Zapotec schoolchildren: growth status and yearly velocity for leg length and sitting height. *Ann Hum Biol* 13:225-233.

Carneiro, P, Heckman, J., and Vytlacil, E. (2003) "Understanding What Instrumental Variables Estimate: Estimating Marginal and Average Returns to Education" Mimeo

Davey Smith G, Greenwood R, Gunnell D, Sweetnam P, Yarnell Jelwwod P (2001) "Leg length, insulin resistance, and coronary heart disease risk: The caerphilly Study." *Journal of Epidemiology and Community Health,* 55:1-6.

Gertler, P., Rivera, J., Levy, S. and Sepulveda J. (2004): "Mexico's PROGRESA: Using a Poverty Alleviation Program as an Incentive for Poor Families to Invest in Child Health", Mimeo.

Glewwe, P., Jacoby, H., King, E. (2001) "Early Childhood nutrition and academic achievement: a longitudinal analysis" *Journal of Public Economics* 81, pp. 345-368.

Gunnell, D., Smith, G.D., Frankel, S., Nanchahal, K. (1998)Childhood leg length and adult mortality: follow up of the Carnegie (Boyd Orr) Survey of Diet and Health in Pre-war Britain *Journal of Epidemology and Community Health* 52: 142-152.

Gunnell, D., Whitlye, E., Upton, M., McConnachie, A., Smith, G.D., Watt, G. (2003) "Associations of height, leg length, and lung function with cardiovascular risk factors in the Midspan Family Study", *Journal of Epidemology and Community Health* 57: 141-146.

Heckman, J.J. (1979), "Sample selection bias as a specification error", *Econometrica*, 47, 153-61.

Heckman, J.J., Ichimura, H. and Todd, P. (1997), "Matching as an econometric evaluation estimator: evidence from evaluating a job training programme", *Review of Economic Studies*, 64, 605-54.

Heckman, J.J., Ichimura, H. and Todd, P. (1998), "Matching as an econometric evaluation estimator", *Review of Economic Studies*, 65, 261-94.

Heckman, J.J., Ichimura, H., Smith, J. and Todd, P. (1998), "Characterizing selection bias using experimental data", *Econometrica,* 66: 1017-98.

Imbens, G., and Angrist, J. (1994) "Identification and Estimation of Local Average Treatment Effects", *Econometrica*, 62: 467-475.

Lee, S. (2004), "Endogeneity in Quantile Regression Models: A Control Function Approach", Cemmap working paper 08/04. http://cemmap.ifs.org.uk/docs/cwp0804.pdf

Leuven, E. and Sianesi, B (2003), "PSMATCH2: Stata module to perform full Mahalanobis and propensity score matching, common support graphing, and covariate imbalance testing", htpp://ideas.repec.org/c/boc/bocode/s432001.html.

Moffitt, R. (2004) "The Effect of Transfer Programs on Labor Supply in the Presence of Prefence Heterogeneity and Variable Takeup" Mimeo.

Ruel, M.T.; Brière (de la), B.; Hallman, K.; Quisumbing, A.; Coj, N. (2002): "Does subsidized childcare help poor working women in urban areas?: evaluation of a government-sponsored programme in Guatemala City**",** IFPRI, Mimeo**.**

Rosenbaum, P.R. and Rubin, D.B. (1983), "The Central Role of the Propensity Score in Observational Studies for Causal Effects", *Biometrika* 70, 1, 41-55.

Shrimpton, R., Victora, C., de Onis, M., Costa Lima, R., Blössner, M., and Graeme Clugston. (2001), Worldwide Timing of Growth Faltering: Implications for Nutritional Interventions. *Pediatrics*, May; 107: 75.

Wadsworth MEJ, Hardy RJ, Paul AA, Marshall SF & Cole TJ. "Leg and trunk length at 43 years in relation to childhood health, diet and family circumstances; evidence from the 1946 national birth cohort". *International Journal of Epidemiology*