
Finding the Benefits: Estimating the Impact of the South African Child Support Grant

MARISA COETZEE

Stellenbosch Economic Working Papers: 16/11

KEYWORDS: CONDITIONAL CASH TRANSFERS, CHILD HEALTH AND NUTRITION,
CONTINUOUS TREATMENT ESTIMATOR, SOUTH AFRICA
JEL: I38, H53, C21, D13

MARISA COETZEE
DEPARTMENT OF ECONOMICS
UNIVERSITY OF STELLENBOSCH
PRIVATE BAG X1, 7602
MATIELAND, SOUTH AFRICA
E-MAIL: MARISA.COETZEE@GMAIL.COM



UNIVERSITEIT
STELLENBOSCH
UNIVERSITY



**A WORKING PAPER OF THE DEPARTMENT OF ECONOMICS AND THE
BUREAU FOR ECONOMIC RESEARCH AT THE UNIVERSITY OF STELLENBOSCH**

Finding the Benefits: Estimating the Impact of the South African Child Support Grant

MARISA COETZEE*

ABSTRACT

The paper estimates the impact of the South African Child Support Grant (CSG) on child health, nutrition and education. Data from the 2008 South African National Income Dynamics Study (NIDS) are used. Two non-experimental treatment evaluation techniques, both relying on propensity scores, are applied to six different outcome variables. Using propensity score matching with a binary outcome variable, no convincing evidence of improvements on any of the outcome variables is found. A second technique is therefore also applied, using a generalised form of the propensity scores. This follows the approach of Hirano and Imbens (2004) and Agüero *et al.* (2009). The generalised approach estimates a positive treatment effect for children's height-for-age and progress through the school system. Although these estimates do provide some evidence of the positive effect of the Child Support Grant on the lives of children, the estimates are small and do not provide clear evidence that the transfers received by caregivers are spent mainly on improving the well-being of beneficiary children. Some potential and plausible explanations for this result are discussed in the paper. Nevertheless, the findings seem to suggest that some of the cash transferred through the Child Support Grant appears to be spent on improving the well-being of children.

Keywords: Conditional cash transfers, child health and nutrition, continuous treatment estimator, South Africa

JEL codes: I38, H53, C21, D13

1. Introduction

Anti-poverty programme evaluation has received increasing attention in the recent past as various stakeholders have demanded answers on whether these programmes are successful in what they set out to achieve (Ravallion, 2007: 3787). Evaluating these programmes involves some form of treatment evaluation, i.e. a comparison between the outcomes of individuals (be it children, adults, households or communities) who are participating in the programme (also referred to as “the treated”) *versus* those who are not.¹ This type of evaluation is, however, no easy task, as data on a suitable control group (or counterfactual) are not always available, leading to biased estimates of the effect of anti-poverty programmes. It is for this reason that a wide variety of estimation techniques, dealing specifically with solving the evaluation problem in the context of non-experimental data, have been developed. Techniques utilising propensity scores have been used repeatedly in this regard and offer a way in which to create a suitable counterfactual with which to compare outcomes. Under certain conditions, these techniques can provide unbiased estimates of the effect of a programme (Ravallion, 2007: 3805).

In South Africa, as in many other developing countries, anti-poverty programmes are often aimed at improving the conditions of the most vulnerable part of society, i.e. the disabled, elderly and children. One such programme is the South African Child Support Grant (CSG), which involves an unconditional cash transfer to eligible caregivers of children, who are identified by way of a means test.

Introduced in April 1998, the CSG has been lauded as one of the government’s most successful anti-poverty interventions (see, for example, the recent report by UNICEF, 2009). This praise seems to be justified in light of the high take-up of the grant, with 9 351 977 beneficiary children and 5 377 476 recipient caregivers as at the end of January 2010 (South African Social Security Agency (SASSA), 2010: 8, 9). Furthermore, commentators have argued that the post-2000 decline in poverty levels is largely attributable to the introduction of this grant (Van der Berg, Burger and Louw, 2010, Leibbrandt, Woolard, Finn and Argent, 2010: 65). However, the CSG is administered as an unconditional grant, i.e. recipient caregivers have no obligations as far as spending the grant money is concerned. This raises questions as to the effectiveness of the grant, i.e. whether it has any effect in improving of the lives of children.

This paper focuses on answering this question by broadly following the methodology adopted by Agüero, Carter and Woolard (2009) in assessing the impact which the CSG has on the well-being of beneficiary children. Data from the National Income Dynamics Study (NIDS) are

¹ The use of the terms “treatment” and “treated” come from the medical sciences (Cameron & Trivedi, 2005: 860), and makes more sense when used in that context. However, the term is used in this paper in conformity with the bulk of the literature on the subject.

used and two techniques, both relying on the use of propensity scores, are applied to identify the effect of the CSG on three groups of outcome variables. The outcome variables chosen are aimed at measuring children's education, health and nutrition, as three indicators of the general well-being of children. The first technique, used as a baseline against which the final estimates can be compared, uses binary propensity score matching (PSM). The second technique, being the focus of this paper, extends this binary estimate to take into account the fact that children have not all received the CSG for an equal period of their lives (i.e. the "dosage" of CSG received differs from child to child). The second technique, first developed by Hirano and Imbens (2005) and later applied by Agüero *et al.* (2009) is more appropriate in assessing the impact of the CSG, as it assumes that the treatment effect is dependent on the length of grant receipt.

The paper finds no effect on the selected outcomes when applying the first technique, namely binary propensity score matching. However, a positive, albeit small, effect on two of the selected outcomes is found when the effect is estimated with a continuous treatment variable. Although this finding does not provide clear evidence that the CSG transfers received by caregivers are spent mainly on improving the well-being of beneficiary children, it does seem to suggest that some of the cash transferred through the CSG is spent on improving the well-being of children. In this sense the paper contributes to previous findings in the literature indicating a positive impact on the health, nutrition and education of children receiving the CSG.

2. Background to the South African Child Support Grant

The CSG was introduced in 1998 as an alternative to the State Maintenance Grant (SMG), which was phased out as part of the South African Government's reform of its social security system (Kruger, 1998: 3). This reform was aimed mainly at increasing the pool of recipients having access to the SMG. Receipt of the SMG was subject to a means test, and the grant was aimed only at parents who were, for a number of reasons, unable to rely on the assistance of a second parent (Kruger, 1998: 3). The SMG was, however, failing to reach the poorest segment of society, with Africans being almost entirely excluded from receiving any benefit from the SMG (Kruger, 1998: 3). The CSG, on the other hand, was introduced as a means-tested social grant, aimed specifically at the poorest 30% of children in South Africa, irrespective of race (Agüero *et al.*, 2009: 5, Triegaardt, 2005: 249).

The initial roll-out of the CSG involved a cash-transfer of R100 per month to the primary caregivers of all eligible children under the age of 7. The age limit has been raised several times to expand the coverage of the grant.² In response to the alleged success of the CSG,

² The age limit was first increased to 9 years in April 2003, and was again raised to include all children under the age of 11 years in 2004. In April 2005, the age limit was increased to 14 years (McEwen *et al.*, 2009: 2).

the age-limit was again increased to 16 years from January 2010, and will gradually be increased to include all eligible children under 18 years from January 2012 (Government of South Africa, 2009). The amount transferred to recipients has also increased over the period from April 1998, and is currently R250 (at the time of the NIDS survey, the amount was R210).³

The value set for the means test remained unchanged between 1998 and 2008, and was set at a monthly income amount received by the primary caregiver (whether single or married) of R800 in urban areas and R1 100 in rural areas. This was revised in October 2008 in line with inflation, and is currently calculated as 10 times the grant amount. The means test has also been amended to differentiate between single and married caregivers so that for married caregivers, the income threshold is doubled and the primary caregiver's income is added to his/her spouse's income (McEwen, Kannemeyer and Woolard, 2009: 2).

The CSG is currently administered as an unconditional cash transfer programme and requires nothing from the recipients as far as the use of the funds is concerned.⁴ This feature of the CSG distinguishes it from other social grant programmes aimed at the alleviation of poverty affecting children, such as Mexico's conditional cash transfer programme, *Progressa* (now *Oportunidades*). The conditions attached to *Oportunidades* are specifically aimed at improving future human capital; older children are required to attend school, while younger children must be taken to clinics on a regular basis for health and nutritional check-ups (Case, Hosegood and Lund, 2005:468). In addition, the cash transfers made in terms of *Oportunidades* are only made to women, who have been found to apply more resources under their control to the improvement of household nutrition as well as towards the health of the children in the household, compared to men.⁵

The absence of enforced conditions in the CSG programme raises the question whether the funds transferred to recipients have any effect on the welfare (more specifically, the educational, nutritional and health outcomes) of children. The sections that follow are devoted to answering this question in a systematic and rigorous way.

3. An Overview of the Existing Literature

The impact of unconditional cash transfer programmes in South Africa in the form of both the CSG as well as the state old age pension (OAP) has been widely researched. The studies have to a large extent focussed on the poverty-alleviating impacts of these transfers within

³ Eyal and Woolard (2010) provide a detailed account of the age-limit and grant amount from 1998 to 2010.

⁴ It should be noted that a recent amendment of the relevant legislation incorporates a regulation stating that "*the primary care giver must, every six months, submit to the [South African Social Security] Agency the child's proof of school or educational institution attendance*" (Government of South Africa, 2009). This condition holds for all children between the ages of 7 and 18 years. This condition is, however, not currently enforced in awarding the CSG.

⁵ See Thomas (1990 and 1994) for examples of such findings.

the context of three broad categories, more specifically the effect on labour force participation, education, and child health and nutrition (Leibbrandt *et al.*, 2010: 62).

The effect of the CSG and OAP on the labour outcomes of individuals within recipient households has been estimated from various angles. Bertrand, Mullainathan and Miller (2003) indicate that the presence of an age-eligible individual in the household has a significant negative effect on the amount of working hours and the employment rate of both prime-aged males and females. However, this result is sensitive to the inclusion of non-resident household members. Using the same data as Bertrand *et al.* (2003), Posel, Fairburn and Lund (2004) show that African women are significantly more likely to be migrant workers if they originally form part of a household in which the OAP is received, specifically when the pension recipient is female. No such relationship is found for men.

Posel *et al.* (2004) suggest that this positive effect of the OAP on labour participation might be because the CSG allows prime-age females to migrate in search of employment, leaving children in the care of older relatives. This suggestion is supported by Ardington, Case and Hosegood (2009), who find that prime-age adults are significantly more likely to be employed as migrant workers once an individual within the household becomes age-eligible for the OAP.

As for unemployed individuals within such households, Klasen and Woolard (2009) find evidence that, for those individuals who are unable to find employment, social grants are a safety net informing the location decision of the unemployed individual. More specifically, unemployed individuals often choose to remain in rural areas in a household where the OAP is received. This has the adverse effect that it lures job-seeking individuals away from job markets to rural areas where employment opportunities are scarce.

In a very recent paper, Eyal and Woolard (2010) estimate the effect of CSG receipt on participation in the labour market, unemployment conditional on labour market participation and employment status using a sample of African women aged 20 to 45. They find a positive relationship between receipt of the CSG and participation in the labour market, as well as probability of being employed. Women in receipt of the CSG who are participating in the labour market, are also found to be less likely to be unemployed.

As for the effect on children's education, various earlier studies have found a significantly positive effect on the education of children living in households that include OAP recipients (Lund, 1993, May, Atwood, Ewang, Lund, Norton and Wentzel, 1998). More recently, Edmonds (2004) considers the effect of the OAP not only on school enrolment, but also child labour and finds that the presence of an age-eligible pensioner in the household increases school enrolment and decreases child labour.

Case *et al.* (2005) focus on the effects of the CSG and find further evidence of increased school-enrolment for children in receipt of the CSG compared to their older siblings who are not receiving the CSG. It is suggested that higher enrolment rates could be as a result of an increase in children's health and nutrition, improving school-readiness. In addition, it is plausible that the CSG income is used to pay tuition fees and other school-related expenses.

As for child health and nutrition, Case (2001a & 2001b) find that the presence of a pension-recipient in a household increases the height-for-age of African and Coloured children by approximately one standard deviation (between 3 to 5 centimetres), equivalent to almost 6 month's growth for children aged 0 to 6 in the sample.

Duflo (2003) provides further insight into the intra-household dynamics of this positive effect of the OAP on child health and nutrition. She finds that effects of the cash received from the old-age pension has a significantly positive effect on the weight-for-height z-scores (WHZ) and height-for-age z-scores (HAZ)⁶ for girls (but not boys) if the transfer were made to a female within the household.

This paper follows the methodology of Agüero *et al.* (2009), where the impact of the CSG on the HAZ of children during the first 36 months of life is estimated, using the KwaZulu-Natal Income Dynamics Study (KIDS). The continuous treatment evaluation technique⁷ is used, which provides for the fact that not all children in the sample received the treatment for the entire 36 months under study. The study finds significant improvement in the HAZ for children receiving the CSG in comparison with children in the control group (Agüero *et al.*, 2009: 26). More specifically, the finding is that, while no gain can be detected for children who only received the CSG for 50% of the 36 months or less, a significant increase in HAZ can be found for all children receiving the CSG for a greater part than 50% of their first 36 months, with a maximum gain recorded when treatment is received for three-quarters of the first 36 months.

4. Finding the Benefits -- the Choice of Outcome Variables

The positive effect of adequate health and nutrition on economic development and poverty alleviation has been emphasised in the literature (Barro, 1996).⁸ An array of measures for nutritional outcomes exist, however, this paper will focus on the standardised HAZ and WHZ. Both of these measures calculate the z-score (z) of each child, defined as the deviation of an individual's measure (in this case height and weight) from the median of the reference population, standardised by the standard deviation of the measure within the reference population (World Bank, 2006). In the case of HAZ this is given by $z = \frac{h - \bar{h}}{\sigma_h}$, where h is height

⁶ Described in detail in the next section.

⁷ Described in further detail below.

⁸ For a discussion of the literature relating health and nutrition to the increase in the standard of living, labour supply and income, refer to Strauss and Thomas (1998).

in cm, and \bar{h} and σ_h are respectively the mean height and the standard deviation of height, given age in years.⁹

Height-for-age serves as an accurate measure of the long-run health and nutritional well-being of children, measuring the accumulated investment in children's nutrition and health since birth. As set out in Duflo (2003: 12) and Agüero et al. (2009: 7), nutritional deficiencies causing stunting, especially in the early stages of a child's life, can lead to permanent damage on a person's health and productivity later in life.

On the other hand, the weight-for-height measures the short-run nutritional status of the child (how well-fed the child is, given his/her height) (Duflo, 2003: 3). A lower than average WHZ is often as a result of temporary malnutrition, for example as a result of droughts, and may still be corrected as soon as sufficient nutrient intake is resumed (Duflo, 2003: 8).

Food expenditure is also examined as an outcome and interpreted as a measure of health and nutrition. Food expenditure is measured as the total average monthly expenditure which is spent on food items per household. This outcome provides some indication of the relative importance of food in a household *vis-à-vis* other households exhibiting comparable observable characteristics. However, its many shortcomings have been well documented (see, for example, Deaton, 1997). First, since the measurement thereof is dependent on the *ex post* recollection of expenditure per item by respondents in a survey, it is prone to error and at best results in noisy data. Food expenditure is also not an accurate measure of actual nutritional intake, as it does not provide an indication of the nutritional content and quality of the food purchased or the food actually consumed by household members (rather than wasted or served to guests)¹⁰(Strauss & Thomas, 1998: 793).

In addition, discrimination within the household makes it almost impossible to determine with certainty which proportion of the expenditure is allocated to children (Deaton, 1997: 223). As such, it is at best a measure at household level (rather than a child-specific measure as the rest of the outcomes listed here). Nevertheless, food expenditure provides some indication of the changes in spending behaviour by households containing recipients of the grant, compared to those which do not contain any recipients and as such can be used as a measure of the impact of the CSG.

The paper approaches the question of intra-household distribution of resources from another angle, investigating the dominance of expenditure on items that are bought exclusively for the use of adults in the households (adult goods). Examples include men's or women's clothing, tobacco and alcohol (Deaton, 1997:227). In the current data, such a variable has

⁹ WHZ on the other hand is calculated similarly as $z = \frac{w - \bar{w}}{\sigma_w}$, where w is weight, and \bar{w} and σ_w are respectively the mean weight in kilograms and the standard deviation of weight, given height in cm.

¹⁰ One would, however, expect food wastage to be less in the households eligible for the CSG, and therefore of interest for the current analysis.

been created using expenditure on alcohol and cigarettes (two items which are exclusively consumed by adults). It is argued that, if receipt of the CSG has a positive effect on the spending habits of the household, expenditure should be channelled away from these items towards items which are beneficial to the development of the children in the household such as food, healthcare and education. Similar *caveats* mentioned in relation to food expenditure apply to this variable as well.

Both HAZ and WHZ have the advantage of overcoming most of the issues relating to the expenditure variables. Height, weight and age can be measured more accurately than food expenditure, requiring only that the physical recording of the measure be done without error. To reduce measurement error the fieldwork manual relating to the NIDS survey required all fieldworkers to take the height and weight measure of each child at least twice, and to repeat the process until the two measures did not differ by more than 0.5 cm or 0.5 kg.¹¹ An evaluation of the data shows that this was done effectively, with differences above 0.5cm recorded for only approximately 1.7% of all children sampled, and differences above 0.5kg recorded for approximately 1.8% of all children sampled. Furthermore, birthdates were in most cases confirmed by the age on the child's clinic ("Road to Health") card.

It should, however, be noted that a significant increase of children's HAZ and WHZ as a result of the cash received from the CSG makes two assumptions: first, that there will not be any unwanted behavioural responses to the cash received (i.e. that it will be used by the recipient caregivers to purchase more food); and second, that it will be possible physically for the children receiving the food to absorb the additional nutrients and transform these nutrients into an improved nutritional status (Agüero *et al.*, 2009: 7). If either of these assumptions proves not to hold true for the treated children in the sample, it will be more complicated to gauge the effect of the CSG.

In addition to these measures, children's school attendance is used as a measure of the effect of the CSG on education. However, since near-universal primary school enrolment has been observed for South Africa (UNICEF 2010),¹² the effect of the CSG in this area might be limited. An alternative measurement of school attainment is therefore also used in this paper, namely whether a child has ever repeated a grade or not. This measure captures the educational attainment of a child, rather than just measuring whether he/she attends school. It may also provide an indication of parental support and how conducive the child's living conditions and social environment are to learning.

Not only are these outcomes in line with the literature in this field, but also accord with the general purpose of the policy underlying the CSG. Although the CSG is administered as a conditional cash transfer programme, conditions regarding proper nutrition, health care and

¹¹ See Agüero *et al.* (2009: 8) for a similar discussion of the KIDS data.

¹² In the data used in this paper, the school enrolment rate is approximately 99%.

school enrolment of beneficiary children were included in the initial regulations accompanying the relevant legislation (Leatt & Budlender, 2007).¹³ In assessing the success of the CSG, it is therefore important to evaluate the effect of the CSG on these outcomes.

While there are *caveats* pertaining to each of the above outcome variables individually, the set of outcomes covers an extensive range of the dimensions of child well-being and welfare when considered collectively. This set of outcomes outlined in the paragraphs above has been selected to provide a wide lens through which we can study how the CSG impacts various aspects and features of the lives of poor adults and children.

5. The Data

The data used in this paper come from the National Income Dynamics Study (NIDS), which is the first wave of what will become a panel dataset. The survey, which was completed during 2008, incorporates data on some 7 305 households, containing 31 170 household members as well as data on 9 336 children under the age of 14 years. It also includes data from a variety of questions aimed at increasing the existing information on the receipt of social assistance by South Africans (McEwen *et al.*, 2009: 1).

At the outset, it is useful to examine the data in order to ascertain how effectively the CSG is targeted. This can be done by examining the number of eligible children (identified in terms of both their age and the means test) receiving the grant as well as those not in receipt of the grant. Both of these groups are important to illustrate the errors of inclusion (i.e. the group of individuals who are not eligible, but nevertheless receive the grant)¹⁴ and the errors of exclusion (i.e. the group of individuals who are eligible, but do not receive the grant).

For the purpose of evaluating the effect of receiving the CSG on the health, nutritional and educational outcomes of children, the latter group (i.e. children who have not received the CSG though being eligible to receive it, based on their age and the means test) is important as a potential counterfactual, as one would expect this group to be more similar to the treated group in terms of background characteristics than the group of untreated, non-eligible children.

Using the means test as it was applied during 2008 (the time at which the NIDS survey was conducted)¹⁵, a simple simulation of the number of children receiving the CSG *versus* the number of children eligible for receiving CSG is conducted in order to get an initial indication of the targeting of the CSG. The algorithm used to identify eligible children is based on the

¹³ These conditions were later removed from the legislation and re-introduced only as non-enforceable conditions (Leatt & Budlender, 2007).

¹⁴ One possible reason for the existence of this group is the fact that the means test is only administered at the time when the grant is first applied for (McEwen *et al.*, 2009: 5).

¹⁵ During the time of the survey used in this paper, the age limit was set at 14 years (i.e. all eligible children under the age of 14).

income, age and caregiver data provided in NIDS. A detailed description thereof is included in Appendix A.

Table 1 applies this algorithm of eligibility to the data to highlight the fact that there are both errors of inclusion and errors of exclusion in the implementation of the grant. Although smaller than the group of eligible individuals rightfully receiving the CSG, there seem to be approximately 2 million children who are eligible for the CSG but not in receipt of the CSG in 2008. This figure is in the same region, although slightly less, than what has previously been reported in Woolard, Hatting and Klasen (2010: 15). This discrepancy can, however, be explained by the fact that the figures reported in Woolard *et al.* (2010) are based only on age-eligibility, and do not take into account the means test (presumably since it has not been possible to link caregivers and their income to specific children in previous datasets, something which NIDS has now made possible).


Table 1: Number of children eligible for and receiving the CSG in 2008

Eligible	Receiving CSG		Total
	Yes	No	
Yes	5 688 256	2 070 819	7 759 075
No	1 939 688	3 574 010	5 513 698
Total	7 627 944	5 644 829	13 272 773

Source: NIDS (2008)

Further interrogation of the NIDS data reveals the main reason why so many eligible children are not receiving the CSG. Figure 1 lists the main reasons for non-application by caregivers of eligible children. The administrative burden is listed as the main reason for these caregivers not applying. More worrying, however, is the fact that there seems to be a certain amount of apathy regarding the CSG displayed by the caregivers, with just over 14% of the caregivers indicating that they just “haven’t gotten round to it yet” and slightly less than 10% indicating that they just “couldn’t be bothered”.

Figure 1: Main reason CSG was not applied for by eligible caregivers in the sample



The descriptive statistics per treatment status are listed in Table 2. For comparative purposes, receipt of the CSG is defined both as a binary (i.e. receiving treatment or not) as well as continuous variable (the length of treatment, or “dosage” received). The continuous treatment variable is defined as the percentage of a child’s life during which he/she received the CSG at the time of the survey. This is calculated using the data on the month and year in which children were born as well as the responses from caregivers regarding the initial date of receipt of the CSG. The group of treated children is divided into terciles based on the distribution of the continuous treatment variable, with children divided in accordance with receiving either a low, medium or high dosage (a low dosage defined as receipt of the CSG for 0-34% of the child’s life, medium dosage as 34-67% and high as 67-100%). The group of untreated children is divided into two groups based on eligibility.

There are expected differences between the two untreated groups (the last two columns of Table 2). First, caregivers of untreated eligible children are more likely to be unemployed and married than the caregivers of children who do not qualify for the CSG. In addition, the child-level characteristics vary across these two groups, with eligible untreated children being on average more than a year younger than non-eligible children (as a result of the age-eligibility criteria) and a slightly greater proportion of the children in the eligible group being female.

Table 2: Descriptive statistics by treatment and eligibility status

Variable	Treated				Not Treated	
	Low [0-34%]	Medium [34-67%]	High [67-100%]	All	Eligible	Not Eligible
Outcomes						
HAZ	-0.60 (1.34)	-0.45 (1.30)	-0.32 (1.40)	-0.50 (1.37)	-0.48 (1.33)	-0.19 (1.37)
WHZ	0.23 (1.76)	0.35 (1.59)	0.13 (1.36)	0.18 (1.48)	0.11 (1.49)	0.39 (1.46)
Monthly household expenditure spent on food items [#]	798.36 (487.88)	833.72 (643.67)	757.08 (478.52)	783.871 (506.17)	770.97 (566.63)	1763.50 (1290.08)
Household expenditure on adult goods per adult[1] in the household [#]	22.79 (50.65)	24.89 (36.79)	26.72 (43.60)	22.49 (36.27)	32.67 (49.23)	182.20 (215.86)
Proportion of children enrolled at school ^{##}	0.99 (0.10)	0.99 (0.08)	0.98 (0.13)	0.99 (0.10)	0.99 (0.08)	0.99 (0.11)
Proportion of children ever repeated school year ^{##}	0.25 (0.44)	0.21 (0.41)	0.16 (0.37)	0.24 (0.43)	0.25 (0.43)	0.18 (0.38)
Caregiver Characteristics						
Motivation	-0.24 (0.61)	-0.13 (1.02)	0.49 (1.06)	0.12 (1.05)	-0.42 (0.16)	- -
Delay in application (in days)	1180.10 (863.63)	832.43 (549.06)	298.94 (246.27)	705.72 (638.99)	2127.19 (1385.75)	- -
Proportion employed	0.16 (0.37)	0.15 (0.35)	0.18 (0.39)	0.15 (0.36)	0.07 (0.25)	0.46 (0.50)
Proportion married	0.32 (0.47)	0.33 (0.47)	0.32 (0.47)	0.34 (0.47)	0.36 (0.48)	0.66 (0.47)

Variable	Treated				Not Treated	
	Low [0-34%]	Medium [34-67%]	High [67-100%]	All	Eligible	Not Eligible
Child Characteristics						
Age in years	7.35 (4.52)	6.58 (4.29)	5.44 (3.03)	6.90 (3.91)	6.21 (4.59)	7.44 (4.64)
Proportion male	0.48 (0.50)	0.51 (0.50)	0.53 (0.50)	0.50 (0.50)	0.46 (0.50)	0.51 (0.50)
Household Characteristics						
Proportion access to electricity	0.71 (0.45)	0.76 (0.43)	0.74 (0.44)	0.73 (0.44)	0.74 (0.44)	0.89 (0.31)
Proportion access to piped water	0.61 (0.49)	0.61 (0.49)	0.59 (0.49)	0.56 (0.50)	0.58 (0.49)	0.86 (0.35)
Proportion access to landline	0.09 (0.28)	0.08 (0.27)	0.08 (0.27)	0.08 (0.27)	0.12 (0.32)	0.29 (0.45)
Proportion access to flush toilet	0.33 (0.47)	0.36 (0.48)	0.35 (0.48)	0.31 (0.46)	0.38 (0.49)	0.75 (0.43)
Proportion female headed	0.59 (0.49)	0.52 (0.50)	0.56 (0.50)	0.57 (0.50)	0.58 (0.49)	0.30 (0.46)
Per capita expenditure	411.38 (415.04)	492.10 (613.58)	409.44 (478.78)	393.19 (459.16)	475.02 (838.90)	2238.01 (2880.70)

Notes: Mean values with standard deviation in parenthesis.

[#]Treatment calculated at the household level.

^{##}Excluding children younger than 5 years old.

The last difference between these two groups relates to the household characteristics, with that of the untreated eligible sample on average indicative of worse living conditions than those who are not eligible. Eligible children not receiving the CSG generally belong to households which are more likely to be female headed, less likely to have access to basic amenities and with a much lower per capita expenditure (a difference of approximately R1 763) than children who do not qualify for the grant.

However, when comparing the observable characteristics of the treated group as a whole with that of the eligible untreated group (columns 5 and 6 of Table 2 respectively), it would appear that these two groups are much more similar. Untreated groups, in the form of rejected or tardy applicants, have in the past been used as a counterfactual to estimate the treatment effect (Agüero *et al.*, 2009:11). Since caregivers in the eligible untreated group have not yet applied for the CSG, although they are eligible to apply, they may in a sense be classified as tardy applicants. In this sense, they potentially could be used as a counterfactual in estimating the treatment effect.

Looking at the breakdown of the treated group into terciles based on the treatment dosage (columns 2, 3 and 4 of Table 2), it can be concluded that the differences in the observable characteristics of these three groups are mostly small and negligible.

The fact that there is a difference between the age of children receiving the grant for a small proportion of their lives (7.35 years), those who have received a medium dosage (6.58 years) and those who have received a high dosage (5.44 years) could potentially be explained by the way in which extent of treatment is measured. More specifically, since the extent of treatment is measured as the percentage of a child's life during which he/she has been receiving the CSG, older children who have been receiving the grant the same number of months as younger children will automatically fall into the low treatment group, as the treatment period as a percentage of their lives will be less than for younger children. Since take-up of the grant has increased each year (Case *et al.*, 2003), one would expect a number of older children for whom application for the CSG was only made recently, resulting in the extent of their treatment being relatively low.

As for the various outcomes, the HAZ and WHZ variables were created using, as a reference, the 2000 US CDC Growth Chart (which is premised on the anthropometric measures of a sample of well-nourished US children). This reference is chosen as a standard against which children in both the treated and control samples can be measured so as to normalise the z-scores (Duflo, 2003: 6).

The differences in outcome variables between the three groups receiving the CSG in a low, medium or high dosage provide tentative corroborating evidence that the presence of the CSG does have a positive impact on the lives of the children receiving it and that this impact

increases with the duration of grant receipt. For most of the outcome variables, there are definite improvements when the dosage is increased from low to medium. The only two exceptions are for school enrolment, where no significant change is observed, and expenditure on adult goods, where an increase in spending is observed as the dosage increases.

As for the change from medium to high, further improvements (albeit insignificant) are observed for all of the outcome variables except WHZ, food expenditure and school enrolment. The change in school enrolment does not warrant any further investigation, given the fact that the change from medium to high is merely 1 percentage point (from 99% to 98%).

However, the decrease in WHZ and food expenditure deserves further mention, as these might be as a result of the way in which the treatment variable is measured. Given the difference in the mean ages of the children receiving a medium *versus* high dosage of the CSG, it is possible that the decrease in WHZ and food expenditure is an indication of possible thresholds effects of the CSG. In other words improvements of these outcome variables might only be observable once the CSG has been received for a number of years (which, given the average age of the children in the high treatment group, has not been the case for most children in that group).

In line with the approach taken by Agüero *et al.* (2009), Table 2 also contains a variable labelled caregiver motivation, which was created to take into account unobserved differences between caregivers which could cause certain caregivers to apply for the CSG earlier than others. The creation of this variable is discussed in the section below.

6. Delay in Application and Caregiver Motivation (Eagerness)

The descriptive statistics in the previous section set out observable characteristics which potentially would have an effect on the extent of treatment, in other words how long the child's caregiver delayed before applying for and receiving the CSG. These include caregiver, household and child-specific characteristics.

Agüero *et al.* (2009) rightly point out that there is, however, another underlying force affecting the extent of treatment, which includes unobserved factors influencing the motivation or eagerness of each caregiver to take up the grant. These unobserved factors, if not controlled for, could potentially bias the estimation of the treatment effect.

Some initial evidence of the differences in caregiver motivation is found in the number of days which caregivers delayed before applying for the child support grant. While the application delay for children receiving only a low dosage is more than three years, children in the high dosage group have an application delay of less than two years. Using the date on which the NIDS interview was conducted (instead of the application date) as the cut-off date, the

“application delay” for untreated eligible children is almost 6 years, illustrating the differences in tardiness between the caregivers in the different groups.

As set out in Agüero *et al.* (2009), this unobserved motivation is a function of the effectiveness of the CSG roll-out in the area where the caregiver lives. Although the programme was rolled out simultaneously in all areas, the data reveal that the delay in take-up for older children (who were already born at the introduction of the CSG) is much shorter for urban areas (an average of 886 days) compared to rural areas (an average of 919 days).

In addition, the delay for younger children is also much less than for children who have only recently become eligible,¹⁶ with a delay of more than three years being recorded for children who were already born at the time of the introduction of the CSG *versus* less than a year (approximately 10 months) for children born in 2008 (i.e. the year of the NIDS survey). This confirms previous evidence that the initial roll-out of the CSG was not as effective as in later years (Agüero *et al.*, 2009).

Following the approach in Agüero *et al.* (2009), a variable capturing these differences in caregiver motivation was created. In the first place, the expected delay was estimated as a function of the child’s age and whether the child lived in a rural or urban area. This was done using OLS, and data of children born two years or more prior to the NIDS survey. This approach is in line with Agüero *et al.* (2009: 13) and evidence by Case *et al.* (2005) that the average delay for take-up of the CSG stabilised after two years from birth (in other words the average delay of children under two years will be underestimated, as there are many children in this age cohort who have not yet taken up the CSG).¹⁷ These younger children are therefore excluded in the estimation of the expected delay.

Thereafter, the difference between actual delay and the expected delay was calculated and then standardised, as a measure of the motivation with which a specific caregiver took up the CSG, compared to other caregivers of the same age and location cohort.¹⁸

Table 2 sets out the positive relationship between caregiver motivation and the extent of treatment (increasing from 0.24 standard deviations below the average to almost half a standard deviation above). It is also clear that treated children have more motivated caregivers (0.12 standard deviations above the average) than untreated children (0.42 standard deviations below the average).

¹⁶ It should be noted that only children born two years or more prior to the NIDS survey were taken into account in this calculation, as the average delay of children born close to the survey date is most likely under-estimated (see Agüero *et al.*, 2009: 13).

¹⁷ Of the sample of children under the age of two years, there are approximately 40% who are eligible for the CSG but are not yet receiving the grant. In contrast, the corresponding figure for children older than two years is 22%.

¹⁸ For children already born at the time of introduction of the CSG, the delay was calculated from the date of introduction of the CSG. For children who were eligible but had not yet taken up the CSG at the time of the NIDS survey, the delay was calculated up to the day of the survey interview with the particular child.

7. Description of the Samples used

To incorporate the various outcome variables (and estimate the average treatment effect based on samples for which evaluation of the effect of the CSG on the specific outcome variable would make sense), three sub-samples are specified.

More specifically, the first sample is used to analyse the effect of the CSG on HAZ and WHZ. This includes the full sample of all children eligible for receiving the CSG, i.e. all children aged younger than 14 years and whose caregiver's income is below the threshold specified by the means test. There are 4914 children in this full sample, of which 3550 receive the CSG while 1303 do not. In the baseline estimation using OLS and binary PSM, both treated and untreated eligible children are included in the sample. However, for the purpose of the continuous treatment estimator, the sample is limited to eligible children younger than 14 years who are receiving the CSG.

The second sample is used for the expenditure outcome variables. Since the expenditure variables are evaluated at the household level, all eligible children who are not in receipt of the CSG, although living in a household where another child receives the CSG, are dropped from the sample. This essentially limits the analysis to a household level (i.e. children who are in a household where they are the CSG recipients and children who are in households where no-one is receiving the CSG). This is done to avoid including eligible children who are in a household where they are not a recipient of the CSG, but another child in the household is (if one assumes a pooling of resources in the household, income from the CSG would be spent both on recipient and non-beneficiary children within that household). It should be noted that, although this has a substantial effect in limiting the sample for the first (baseline) technique, it does not limit the sample where treatment is treated as a continuous variable.

The third sample limits the original data by excluding all eligible children under the age of 5 years. This is done to obtain the effect on schooling for children who are of school-going age.

8. Baseline Estimates using OLS and Binary Propensity Score Matching

As a baseline, the treatment effect was estimated using OLS. The results of these estimates are set out in Table 3. For the sake of comparability, the covariates used in the regressions are the same as those used in the estimation of the average treatment effect.

The coefficients from the OLS regressions are as expected, apart from the coefficient on caregiver motivation, which is significantly negative in the regression for food expenditure. The negative sign of this coefficient could be explained by the fact that food expenditure decreases slightly in the sample for the highest level of caregiver motivation. The results from the OLS regressions do not provide any evidence of a significant treatment effect for any of the outcome variables, apart from expenditure on adult goods, which is negative and

significant at the 10% level, indicating a decrease in expenditure on adult goods for households that are receiving the CSG *versus* those that are not.

Table 3: OLS estimation of treatment effect

Variable	Sample 1		Sample 2		Sample 3	
	HAZ	WHZ	Food Expenditure	Adult Expenditure	School Enrolment	School Progress
Child receives CSG	-0.08	-0.26	44.25	-8.91*	0.00	-0.02
Caregiver Characteristics						
Motivation	-0.02	0.19***	-20.06*	0.92	0.00	-0.01
Employed	-0.09	-0.06	20.83	-7.24*	-0.01	-0.05
Married	0.12	-0.27	-0.18	14.13***	0.00	-0.09***
Child Characteristics						
Age	-0.07	-0.98***	-23.39*	0.5	0.03**	0.17***
Age squared	0	0.08***	1.48	-0.04	-0.00**	-0.01***
Boy	-0.15**	-0.08	-53.49**	-1.11	-0.00*	0.16***
Household Characteristics						
Electricity	-0.15*	-0.25	88.20***	0.56	0.00	-0.01
Piped water	-0.09	-0.24	117.86***	2.47	0.00	0.00
Landline	0.09	0.01	269.86***	3.84	0.00	0.03
Flush toilet	0.30***	0.12	45.81	14.75***	-0.01**	-0.08**
Female head	0	0.09	-77.62***	2.98	0.00	-0.05**
Log pc expenditure	0.08	0.03	-	-	0.00	-0.01
Constant	-0.07	3.08**	694.54***	14.79**	0.87***	-0.58***
Observations	1872	456	3530	680	2033	2054

*Significant at 10% level.

**Significant at 5% level.

***Significant at 1% level.

The baseline estimates are, however, not an accurate estimate of the treatment effect since the treated and control units may differ substantially, giving rise to potential selection bias. More specifically, treatment evaluation requires the estimation of the average treatment effect, written as

$$\begin{aligned}\delta &= E[y_{1i} | D_i = 1] - E[y_{0i} | D_i = 0] \\ &= E[y_{1i} - y_{0i} | D_i = 1] + \{E[y_{0i} | D_i = 1] - E[y_{0i} | D_i = 0]\}\end{aligned}$$

where y_1 and y_0 are the treatment and control outcomes respectively, and D_i is the binary treatment variable for individual i , with $D_i = 1$ indicating the presence (absence) of treatment (Cameron & Trivedi, 2005: 872). In the equation above, the first term in the second line represents the average treatment effect on the treated, while the second term in the second line (i.e. the term in parenthesis) represents a bias term. Clearly the estimate of the treatment effect potentially is confounded in a situation where this bias term is not equal to zero, which will be the case where treatment was not assigned randomly.

One proposed solution to reducing the bias is PSM, i.e. the conditional probability measure of participation in treatment, given the underlying characteristics contained in X (Cameron & Trivedi, 2005: 864). Under certain conditions, PSM provides a way in which to control for selection bias. These conditions ensure that, after conditioning on the propensity scores of individuals (i.e. the propensity of being selected into treatment of both the treated and untreated), none of the selection bias remains.

PSM was applied to all three samples, for all six outcome variables. However, no significant treatment effect was estimated. One exception is the adult expenditure variable, for which a significant (and fairly robust) treatment effect was estimated. However, this result is not convincing, as the overlap in the propensity score between treated and untreated units is very small. A more detailed discussion of the implementation of and results from PSM is contained in Appendix B.

The failure of PSM to provide any significant (or convincing) results raises questions regarding the validity of the method in these specific circumstances. One criticism which may be raised against the use of PSM to estimate the treatment effect is that the CSG is not implemented in a binary fashion. In addition, as already mentioned, the effect of receipt of the CSG on the outcomes may only be observable subsequent to a certain monetary threshold being reached (after receiving the grants for a certain period). CSG would fail to measure this effect.

9. The use of Generalised Propensity Scores with a Continuous Treatment Variable

The bulk of the treatment evaluation literature employing propensity scores so far have made use of PSM in order to measure the average effect of treatment. However, as indicated above, one of the underlying assumptions of PSM is the existence of a binary treatment variable. More specifically, it assumes that all individuals may be classified as falling into one of two categories -- those who received treatment and those who received no treatment at all. Implicit in this there is an assumption that those that were treated are comparable and received equal dosages of treatment. However, as discussed in Agüero *et al.* (2009), this assumption clearly does not hold in the case of the CSG. There are large discrepancies in the length of exposure to the CSG, depending on the rollout and expansion schedule of the grant

as well as individual characteristics relating to motivation and the take up of the grant. Conventional applications of PSM fail to take this variation into account and it has been argued that this may partly explain inconclusive results regarding the treatment effect of CSG (Agüero *et al.*, 2009). It may thus in this instance be more appropriate to think of treatment as a continuous variable.

Hirano and Imbens (2004) were the first to extend the use of propensity scores to develop a continuous treatment estimator. They use data collected from winners of the Megabucks lottery in Massachusetts during the mid-1980s to estimate the average labour response of winners, more specifically their earnings six years after winning. The data were generated as a random process, and should, accordingly, satisfy the general unconfoundedness condition. However, there were significant non-responses to the survey which may have introduced bias, necessitating the use of a technique which reduces bias and provides more credible estimates than those obtained using simple regressions (Hirano and Imbens, 2004: 73).

Hirano and Imbens (2004) show how, following the PSM literature discussed earlier, it is possible to remove the bias relating to treatment status by introducing an unconfoundedness assumption similar to the one set out above. They obtain a generalised version of the propensity score in the binary treatment case (henceforth referred to as the “generalised propensity score” (GPS)), producing unbiased estimates of the effect of treatment. This paper now provides a brief overview of the theoretical basis on which the technique is founded before applying it to the NIDS data.

9.1. Theoretical Framework and Underlying Assumptions

To formalise the above, suppose the analysis draws from a random sample of N units indexed $i = 1, \dots, N$. For each unit, the set of potential outcomes is given by the random variable $Y_i(d)$, called the “unit-level dose-response function” by Hirano and Imbens (2004: 74). In other words, $Y_i(d)$ is the potential response of unit i to receiving a dosage of treatment d , where $d \in D$. For current purposes, the dosage d is the percentage of each child’s life during which he/she has received the CSG as at the time of the survey, and the set of outcomes $Y_i(d)$ include all relevant health, nutritional and educational benefits accruing as a result of the grant.

In the binary treatment case discussed above, $D = \{0,1\}$. However, where treatment is a continuous variable, as is the case with receipt of the CSG, $D \in [d_0, d_1]$, with d_0 and d_1 representing the upper and lower bounds of the treatment interval, set at 0%-100% for current purposes.

It is assumed that, for each of the units within the sample, it is possible to observe the vector of covariates x_i ; the level of treatment, D_i ; and the outcome corresponding to the treatment received $Y_i = Y_i(D_i)$ ¹⁹.

The object of interest is referred to by Hirano and Imbens (2004) as the “average dose-response function”, $\mu(d) = E[Y_i(d)]$. Following the authors, the subscript i is dropped until the next section to simplify notation.

The fundamental assumption of Hirano and Imbens (2004) extends the weak unconfoundedness assumption of Rosenbaum and Rubin (1983) to continuous treatments, and states that

$$Y(d) \perp D \mid x, \quad \forall d \in D.$$

Put differently, outcome and treatment should be independent once the covariates have been controlled for. This general weak unconfoundedness assumption must merely hold for each level of treatment, and is not required to hold jointly for all potential outcomes (Hirano & Imbens, 2004: 74).

In the current context, this assumption requires that, conditional on the covariates, there are no additional factors which influence both take-up of the CSG (extent of treatment) and the outcome variables. Controlling for a variety of household and individual-level characteristics, as well as caregiver motivation, it is argued that extent of treatment and health and educational outcomes are unconfounded, unless, as Agüero *et al.* (2009: 20) note, more dedicated caregivers postponed having children until after the introduction of the CSG. However, since the CSG effectively became available immediately upon announcement to all eligible children, this is an unlikely scenario.

Next, the GPS is defined by Hirano and Imbens (2004:74). If $r(d, x)$ is the conditional density of the treatment, given the covariates, as in

$$r(d, x) = f_{D|x}(d|x).$$

then the GPS is defined as $R = r(D, x)$

In addition, just as in the binary case, the GPS has a balancing property in that, within strata of the same value of $r(D, x)$, the extent of treatment is not dependent on the value of the covariates (Hirano & Imbens, 2004: 74). This can be stated as

$$x \perp 1(D = d) \mid r(d, x)$$

¹⁹ It is also necessary to make assumptions regarding the probability space of these variables, which are listed in Hirano and Imbens (2004:2) and Bia and Mattei (2008: 355).

If both the balancing property as well as the general weak unconfoundedness assumption holds, Hirnao and Imbens (2004: 75) show that assignment to treatment is unconfounded, conditional on the GPS. Then, for each level of treatment d ,

$$f_D\{d|r(d, x), Y(d)\} = f_D\{d|r(d, x)\}.$$

It is important to note that the conditional density of treatment, given the covariates, is evaluated at each level of treatment. Accordingly, there are as many propensity scores as there are levels of treatment (Hirano & Imbens, 2004: 75).

Last, Hirano and Imbens (2004: 76) illustrate how the GPS can be used to produce an unbiased estimate of the dose-response function at each level of treatment. First, the conditional expectation of the outcome variable, given the treatment level D and the estimated GPS, R (both scalar variables) is estimated, as in $\lambda(d, r) = E[Y|D = d, R = r]$ (Hirano & Imbens, 2004: 76). Then, this conditional expectation is averaged over the GPS at the relevant level of treatment $\hat{\mu}(d) = E[\lambda(d, r(d, x))] = E[Y(d)]$. It is emphasised that the average is not taken over the entire GPS R , but rather over the GPS at the specific treatment level, i.e. $r(d, x)$ (Hirano & Imbens, 2004: 76).

9.2. Estimation of the Treatment Effect and Inference

To facilitate the practical implementation of the above, the paper follows a three-step approach in line with Hirano and Imbens' (2004) original methodology, as well as subsequent work by Bia and Mattei (2008) and Agüero *et al.* (2009).

As a first step, the conditional distribution of treatment given the covariates is modelled, assuming a normal distribution, as in

$$D_i|x_i \sim N(\beta_0 + \beta_1'x_i, \sigma^2)$$

where β_0 , β_1 and σ^2 are estimated using maximum likelihood, and the GPS as

$$\lambda(D_i, \hat{R}_i) = E[Y_i|D_i, \hat{R}_i] = \alpha_0 + \alpha_1 D_i + \alpha_2 D_i^2 + \alpha_3 \hat{R}_i + \alpha_4 D_i \hat{R}_i$$

As part of this step, Bia and Mattei (2008: 358) also suggest testing whether the balancing property holds by dividing the sample into strata and calculating the mean difference of each covariate between units in different strata, conditional on the GPS.

The second step involves the estimation of the function $\lambda(d, r)$. This estimation can be done using flexible polynomial functions; Bia and Mattei (2008: 358) suggest using a quadratic approximation with an interaction between the treatment level and the GPS, as in

$$\lambda(D_i, \hat{R}_i) = E[Y_i|D_i, \hat{R}_i] = \alpha_0 + \alpha_1 D_i + \alpha_2 D_i^2 + \alpha_3 \hat{R}_i + \alpha_4 D_i \hat{R}_i$$

These coefficients may be estimated using ordinary least squares or maximum likelihood (depending on the type of outcome variable). Hirano and Imbens (2004: 76) emphasise that the estimated coefficients do not have any causal interpretation.

Finally, the parameters estimated in the previous step are used to derive the dose-response at a particular treatment level, by computing the average potential outcome at treatment level d

$$E[\hat{Y}(d)] = \frac{1}{N} \sum_{i=1}^N \left(\hat{\alpha}_0 + \hat{\alpha}_1 d_i + \hat{\alpha}_2 d_i^2 + \hat{\alpha}_3 \hat{r}(d, x_i) + \hat{\alpha}_4(d_i) \hat{r}(d, x_i) \right)$$

To derive the entire dose-response function, this estimation is repeated for each treatment level of interest (Hirano & Imbens, 2004: 77).

Following Bia and Mattei (2008: 364), this can be taken a step further. After deriving the dose-response functions at each level of treatment, the treatment-effect function can also be estimated. This is done by comparing the dose-response at a specific level of treatment (or dosage, d) with the outcome at a lower reference level ($d - \rho$), where ρ represents a treatment gap of size ρ

$$\mu(d) - \mu(d - \rho) = E[Y_i(d)] - E[Y_i(d - \rho)] \forall d \in D.$$

In the current paper, the treatment effect of interest is the improvement of child health, nutrition and education at each level of receipt of the CSG (d), compared with children who have received the CSG for a smaller percentage of their lives ($d - \rho$). The results from applying this technique to the NIDS data are set out below.

10. Results

The analysis of the outcome variables is done using the same three sub-samples of the data used in the application of the PSM in the previous section. The results in this section follow the three-step approach described above.

10.1. Modelling the Conditional Distribution of Treatment

As a first step, a maximum likelihood estimation was conducted for all three samples. The model is specified to contain covariates which influence both take-up of the CSG and the outcome variables, in line with the literature (Caliendo & Kopeinig, 2008: 38). In addition, it must be specified such that the conditional independence assumption listed above is satisfied. Here the model is specified which controls for observable caregiver characteristics as well as observable child characteristics that may influence child well-being as well as take-up of the grant. The model also includes a range of covariates which are most likely to influence both take-up and the outcome variables, including access to basic amenities (such as piped water, flush toilets, electricity and a

landline telephone) and controls for household poverty by including log of monthly per capita expenditure and a binary variable for female-headed households (which are on average worse off than male-headed households)²⁰.

One of the assumptions to satisfy when implementing the GPS technique is that, conditional on the covariates specified in the model, extent of treatment and the outcome variable are conditionally mean independent. This condition is satisfied by controlling for a range of child-, household- and caregiver-specific covariates which together predict take-up of the CSG as well as the extent of treatment. However, in addition to these covariates, as indicated by Agüero *et al.* (2009), the caregiver-motivation variable is also added to the model in order to control for unobserved caregiver characteristics. Since the use of the generalised propensity score per implication only includes treated units, the three samples discussed previously are limited to only included children who have been receiving the CSG for some period of their lives.

The results from the maximum likelihood estimation are set out in Table 4. As expected, the coefficient of the caregiver-motivation variable is large and significant, indicating the expected positive relationship between more motivated caregivers and larger dosages of treatment. In addition, the fact that a child's caregiver is employed has a significant effect on the length of the child's treatment in the first two samples.²¹

The child-age variables are jointly significant²² and confirm previous evidence that older children are more likely to be receiving the CSG for a smaller percentage of their lives, possibly as a result of the fact that the roll-out of the grant became more effective with time. As for household characteristics, having a landline decreases the probability of the child having received the CSG for a large portion of his/her life. Having electricity in the household has the opposite effect, which is correlated with the fact that more households in urban areas have access to electricity and were therefore first to know of and apply for the CSG.

²⁰ The data reveal that the mean monthly per capita expenditure of female-headed households is approximately R700, while that for male-headed households is approximately R1300.

²¹ The similarities in the results of these two samples are to be expected, since the same sample is essentially used (as discussed above).

²² At the 1% significance level for the first two samples, not for the third.

Table 4: Maximum likelihood estimates with different samples

Variable	Sample 1		Sample 2		Sample 3	
	Coefficient	Robust Standard Error	Coefficient	Robust Standard Error	Coefficient	Robust Standard Error
Caregiver Characteristics						
Motivation	10.01***	0.49	10.01***	0.49	7.26***	0.55
Employed	2.51*	1.50	2.46*	1.48	1.83	1.86
Married	1.36	1.20	1.35	1.20	1.25	1.57
Child Characteristics						
Age	5.14***	0.49	5.14***	0.49	-1.43	1.97
Age squared	-0.52***	0.04	-0.52***	0.04	-0.16	0.11
Boy	0.12	1.02	0.12	1.02	0.11	1.37
Household Characteristics						
Electricity	4.24***	1.25	4.20***	1.24	4.88***	1.67
Piped water	-3.48**	1.35	-3.51***	1.34	-2.92	1.80
Landline	-0.87	1.99	-0.91	1.99	-1.94	2.57
Flush toilet	0.52	1.43	0.49	1.42	-0.46	1.91
Female head	-0.87	1.10	-0.87	1.10	-0.84	1.50
Log pc expenditure	-0.18	0.76	-	-	-0.60	0.99
Constant	49.90***	4.36	48.97***	1.79	80.05***	10.16
Observations	2375		2375		1391	

*Significant at 10% level.

**Significant at 5% level.

***Significant at 1% level.

As a next step, the balancing property for all three samples is explored in Table 5. The first part of the table deals with the data before adjusting for the propensity score. To test whether the mean of each of the covariates within one of the treatment terciles is different from the combined mean of the other two groups, the t-statistic is used.

The second part of the table deals with the balancing property of the data after adjusting for the GPS, following the algorithm suggested by Hirano and Imbens (2004). In the first place, for each of the three treatment terciles, the probability that each observation within the tercile would have received the median dosage of treatment is estimated. More specifically, if the median treatment level in the first terciles is given by d_t^{low} , then in each tercile, $r(d_t^{low}, x_i)$ is estimated.

Next, all observations are pooled again and then divided into five blocks or strata based on the estimation of $r(d_t^{low}, x_i)$. Within each of these five blocks, the differences between the mean covariates of those observations falling within the lowest treatment tercile and those which do not, are calculated and tested for statistical significance. This process is then repeated for the medium- and high-treatment terciles.

From the table it can be seen that the balancing property is not satisfied for the unadjusted data. The GPS improves the balancing of the data, reducing the highest t-statistics (especially for the age and caregiver motivation) to make the data more balanced. However, it should be pointed out that, even after conditioning on the GPS, the child-age and caregiver-motivation covariates remain unbalanced between the treatment terciles. In the case of the caregiver-motivation covariates, this imbalance is a result of the way in which the variable was generated (i.e. to capture the differences in caregiver motivation between caregivers in these three terciles). In the case of the child-age covariates, the imbalance again indicates the differences in the ages of children in different treatment terciles as a result of the slow initial roll-out of the CSG.

It is argued that it is for this reason that the age variables should be included in the model when estimating the GPS (as set out in Table 4), since it is important when estimating the GPS and the average treatment effect to also control for the age-group in which the child falls (so as to ensure that the effect on the outcome variable is measured for children of the same age group, who have been exposed to the same length of treatment).

Table 5: The balancing property: t-statistics for equality of means between treatment terciles

				Sample 1			Sample 2			Sample 3		
Unadjusted for GPS				Adjusted for GPS			Adjusted for GPS			Adjusted for GPS		
Treatment Intervals				Treatment Intervals			Treatment Intervals			Treatment Intervals		
Variable	Low	Medium	High	Low	Medium	High	Low	Medium	High	Low	Medium	High
	[0,34)	[34,67)	[67,100]	[0,34)	[34,67)	[67,100]	[0,34)	[34,67)	[67,100]	[0,34)	[34,67)	[67,100]
Caregiver Characteristics												
Motivation	11.64	6.21	-16.93	6.39	2.88	-4.19	6.41	2.79	-4.07	4.66	1.76	-1.51
Employed	0.75	0.53	-1.17	-0.81	-0.22	1.99	-0.81	-0.24	1.94	-0.20	0.89	-0.67
Married	0.68	-0.15	-0.48	-0.34	-1.14	0.94	-0.34	-1.13	0.99	0.64	0.16	0.28
Child Characteristics												
Age	-7.06	-3.40	9.64	-0.23	-9.61	5.17	-0.37	-9.74	5.36	1.78	-7.42	4.92
Age squared	-9.83	-4.95	13.80	0.57	-11.41	6.08	0.42	-11.56	6.27	1.50	-7.61	5.81
Boy	0.60	0.15	-0.68	-0.57	0.09	-0.64	-0.49	0.10	-0.50	0.29	-1.13	0.64
Household Characteristics												
Electricity	2.43	-1.00	-1.31	-0.68	-1.52	3.02	-0.69	-1.46	2.97	0.97	-1.26	0.41
Piped water	-0.85	0.24	0.56	0.26	0.63	-1.42	0.16	0.68	-1.37	0.25	-0.66	-0.26
Landline	-0.89	0.50	0.35	0.13	1.13	-0.96	0.09	1.11	-0.94	-0.07	0.24	-0.69
Flush toilet	-0.24	1.19	-0.87	-1.04	1.16	-0.20	-1.06	1.23	-0.05	-0.25	0.16	0.14
Female head	-1.22	0.15	0.98	-0.02	0.04	-0.58	-0.04	0.07	-0.50	-0.42	-0.29	-0.43
Log pc expenditure	0.38	-1.27	0.81	0.68	-1.16	0.62	-	-	-	0.62	-0.67	-0.41

10.2. Estimating the Outcomes Conditioned on the GPS and Treatment

The results of the estimates in the second stage are presented in Table 6. For the continuous-outcome variables, ordinary least squares estimation was used on a linear model, while a logit model was specified for the binary school outcome variables.

As indicated above, the coefficients from these estimations have no direct meaning, however, if the coefficients in this estimation are statistically different to zero this can be interpreted as evidence that they introduce no bias (Bia & Mattei, 2008: 359).

The results of this test are reported at the bottom of Table 6. The null-hypothesis can be rejected for three of the models, i.e. HAZ, progress through the school system and expenditure of adult goods. The analysis of the remaining three outcome variables is therefore taken to the next stage in order to estimate the treatment effect.

Table 6: Second stage estimates per outcome

Variable	Sample 1		Sample 2		Sample 3	
	HAZ	WHZ	Food Expenditure	Adult Expenditure	School Enrolment	School Progress
Parameter	Coefficient (Robust Standard Error)					
α_0	-0.39 (0.15)	-0.37 (0.21)	131.45 (9.01)	28.34 (9.50)	4.15 (1.27)	-2.08 (0.32)
α_1	0.01 (0.01)	0.00 (0.01)	0.67 (0.39)	-1.19 (0.41)	0.05 (0.05)	0.03 (0.01)
α_2	0.00 (0.00)	0.00 (0.00)	-0.01 (0.00)	0.01 (0.00)	0.00 (0.00)	0.00 (0.00)
α_3	-48.11 (18.93)	113.11 (50.89)	-64.06 (1179.49)	3280.65 (1303.19)	60.00 (170.24)	122.70 (35.66)
α_4	0.77 (0.34)	-1.74 (0.71)	-7.21 (20.82)	-42.77 (22.05)	-2.83 (2.85)	-2.95 (0.68)
Observations	1695	980	2375	490	1265	1175
Prob>F	0.01	0.10	0.23	0.01		
Prob>Chi2					0.37	0.00

10.3. Estimating the Average Treatment Effect

The average treatment effect can now be estimated as the difference between the dose-response functions at different levels of treatment. As described above, the treatment effect is estimated as the difference between the outcome at a specific treatment level

compared to the outcome of a lower reference level ($d - \rho$), where ρ represents a treatment gap of size ρ . For the current analysis, the value of ρ is set to 10, i.e. the difference is calibrated to 10%. The results are presented as Figures 2 and 3, with the treatment level indicated on the y-axis and the change in the outcome variable at each level of treatment indicated on the x-axis.

The 95% confidence intervals (calculated using bootstrapped standard errors with 1000 repetitions) are also included for each of the outcome variables in order to provide a sense of the significance of the estimated treatment effect. The specific region for which a significant effect was estimated, is marked with vertical lines on each figure.

Since no significant treatment effect was estimated for the adult expenditure outcome variable (based on the confidence intervals calculated), this figure is not included.

Figure 2 reflects positive gains from the CSG on HAZ. The estimated average treatment effect is, however, not very large. Positive and statistically significant gains in HAZ from receiving the CSG are recorded for children receiving the grant for 40-60% of their lives. The gains in HAZ reach a maximum where children have been receiving the CSG for 50% of their lives, compared to a children who have only been receiving it for 40%, where the increase in the z-score is approximately 0.06, translating into 6% of the standard deviation. This estimated gain in HAZ occurs at a lower level of treatment and is somewhat less than the effect estimated in the previous study by Agüero *et al.* (2009), where an increase of 0.2 in the z-score for HAZ was estimated.

The treatment effect of the CSG on HAZ is estimated to be negative at higher dosages of treatment. However, this portion of the curve may be ignored since the confidence interval band lies on both sides of zero, indicating that there is no statistically significant evidence of any effect for such high treatment dosages.

Looking at Figure 3, there appears to be some positive impact on the probability of repeating a school year. A significant decrease in this probability is estimated for children receiving the CSG between 30% and 60% of their lives. The maximum impact is a decrease in the probability of approximately 0.04 percentage points for receiving the CSG between 40% and 50% of the child's life. Again, the portions where a positive effect on the probability of repeating a school year is estimated can be ignored, as these fall outside the confidence interval.

Although these estimates suggest that the CSG does have some positive impact on the lives of beneficiary children, these effects are small and do not provide conclusive evidence that the transfers received by caregivers are spent mainly on improving the well-being of children. This might be explained by the fact that the transfers are unconditional and may accordingly be channelled towards the purchase of other goods

Figure 2: Estimated gains in HAZ





Figure 3: Estimated change in probability of repeating a school year (children older than 5 years)



which are not only to the benefit of children. In addition, given the fact that the grant amount is relatively small (compared to, for example, the OAP), it might also be that when the cash is spread across the entire household the observable effect on children is small.

Nevertheless, the results set out above do provide some evidence of a positive impact of the CSG and seems to indicate that some of the cash from the grant does, at least, filter through to beneficiary children in the form of better health and nutrition, as well as faster progression through the school system.

11. Conclusion

This paper set out to evaluate the impact of the CSG on child health, nutrition and education following the approach by Agüero *et al.* (2009) in order to ascertain if the cash transferred has any effect on the well-being of beneficiary children. The evaluation is of particular importance as the CSG is currently administered as an unconditional cash transfer programme.

Since the CSG was made available to all eligible children at introduction of the grant and not implemented in a randomised fashion, a simple estimation of the effect of treatment between individuals receiving the grants and those who remain untreated is subject to selection bias. Accordingly, this paper employs two techniques using propensity scores to estimate the average treatment effect of the CSG on children and households. Under the assumptions set out in the paper, these techniques can provide unbiased estimates of the treatment effect of the CSG.

Applying PSM to the NIDS data, the paper finds no convincing evidence of improvements in child health, nutrition or education. These results may be driven by the fact that PSM assumes a binary treatment variable, which is not the case with the CSG.

Applying a second technique using a generalised form of the propensity scores, results in positive treatment effects for children's HAZ and whether the child has repeated a school year. The treatment effect estimated for HAZ is in line with, although slightly less than, a previous estimation by Agüero *et al.* (2009). Although these estimates do provide some evidence of the positive effect of the CSG on the lives of children, the estimates are small and do not provide clear evidence that the transfers received by caregivers are spent mainly on improving the well-being of children. Some potential and plausible explanations of this have been offered.

Nevertheless, the findings in this paper seem to suggest that some of the cash transferred through the CSG appears to be spent on improving the well-being of children and in this sense contributes to previous findings in the literature indicating a positive impact on the health, nutrition and education of children receiving the CSG.

12. References

- Ardington, C., Case, C. & Hosegood, V. 2009. Labor Supply Responses to Large Social Transfers: Longitudinal Evidence from South Africa. *American Economic Journal: Applied Economics*, 1(1): 22–48.
- Agüero, J.M., Carter, M. & Woolard, I. 2006. From Flows to Stocks: The Impact of Unconditional Cash Transfers on Human Capital. *Paper presented at the CSAE Conference on “Reducing Poverty and Inequality: How can Africa be included”*. 20-21 March 2006, Oxford.
- Agüero, J.M., Carter, M. & Woolard, I. 2009. The Impact of Unconditional Cash Transfers on Nutrition: The South African Child Support Grant. *Working Paper, International Policy Centre for Inclusive Growth*.
- Barro, R. Determinants of Economic Growth: A Cross-Country Empirical Study. *Journal of Comparative Economics*, 26(4): 822-824.
- Becker, S.O. & Ichino, A. 2002. Estimation of average treatment effects based on propensity scores. *The Stata Journal*, 2(4): 358-377.
- Bertrand, M.S., Mullainathan, D., Miller, D. 2003. Public Policy and Extended Families: Evidence from Pensions in South Africa. *World Bank Economic Review*, 17(1): 27-50.
- Bia, M. & Mattei, A. 2008. A Stata package for the estimation of the dose-response function through adjustment for the generalized propensity score. *The Stata Journal*, 8(3): 354-373.
- Budlender, D. Rosa, S. & Hall, K. 2005. *At all Costs? Applying the Means Test for the Child Support Grant*. Cape Town: Children’s Institute and Centre for Actuarial Research, University of Cape Town.
- Caliendo, M. & Kopeinig, S. 2008. Some Practical Guidelines for the Implementation of Propensity Score Matching. *Journal of Economic Surveys*, 22(1): 31-72.
- Cameron, A.C & Trivedi, P.K. 2005. *Microeconometrics: Methods and Applications*. Cambridge University Press, New York: NY.
- Case, A. & Deaton, A. 1998. Large Cash Transfers for the Elderly in South Africa. *Economic Journal*, 108: 1330-1361.
- Case, A. 2001a. Health, Income and Economic Development. *The Annual World Bank Conference on Development Economics 2001/2001*: 221-241.
- Case, 2001b. Does Money Protect Health Status? Evidence from South African Pensions. *National Bureau of Economic Research Working Paper No. 8495*.

- Case, A., Housegood, V. & Lund, F. 2005. The reach and impact of Child Support Grants: evidence from KwaZulu-Natal. *Development Southern Africa*, 22(4):467-482.
- Deaton, A. 1997. *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*, Johns Hopkins University Press, Baltimore: MD.
- Dehejia, R & Wahba, S. 2002. Propensity Score-Matching methods for Nonexperimental Causal Studies. *The Review of Economics and Statistics*, 84(1): 151-161.
- Duflo, E. 2003. Grandmothers and Granddaughters: Old-Age Pensions and Intrahousehold Allocation in South Africa. *The World Bank Economic Review*, 17(1): 1-25.
- Edmonds, E. 2004. Does Illiquidity alter Child Labor and Schooling Decisions? Evidence from Household Responses to Anticipated Cash Transfers in South Africa. *National Bureau of Economic Research Working Paper No. 10265*.
- Eyal, K & Woolard, I. 2010. Female Labour Force Participation and the Child Support Grant in South Africa, *Paper presented at the ERSA Public Economics Workgroup, University of Pretoria 9-10 September 2010*.
- Government of South Africa. 2009. *Amendment to the Regulations relating to the Application for and Payment of Social Assistance and the Requirements or Conditions in respect of Eligibility for Social Assistance*, Government Notice No R. 1116 of 27 November 2009.
- Heckman, J., Ichimura, H. & Todd, P. 1997. Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Program. *Review of Economic Studies*, 64(4): 605-654.
- Hirano, K. & Imbens, G.W. 2004. The Propensity Score with Continuous Treatments, in A. Gelman and X. Meng (eds.), *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, Chichester: Wiley.
- Klasen, S. & Woolard, I. 2008. Surviving Unemployment Without State Support: Unemployment and Household Formation in South Africa. *Journal of African Economies*, 18(1): 1-51.
- Kruger, J. 2008. From single parents to poor children: Refocusing South Africa's transfers to poor households with children, *ISSA's 2nd International Research Conference on Social Security, Jerusalem, 25-28 January*.
- Lalonde, R. 1986. Evaluating the Econometric Evaluations of Training Programs. *American Economic Review*, 76(4): 604:620.
- Leibbrandt, M., Woolard, I., Finn, A. and Argent, J. 2010. Trends in South African Income Distribution and Poverty since the fall of Apartheid. *OECD Social, Employment and Migration Working Papers*, number 101.

- Leuven, E. & Sianesi, B. 2003. Psmatch2: Stata module to perform full Mahalanobis and propensity score matching, common support graphing, and covariate imbalance testing, version 3.1.5. Available: <http://ideas.repec.org/c/boc/bocode/s432001.html> (accessed 21 June 2010).
- Leatt, A. & Budlender, D. 2007. Under what conditions? Social security for children in South Africa. In Minujin, A. & Delamonica, E. (eds) *Social Protection Initiatives for Children, Women and Families*, New York: The New School & UNICEF.
- Lund, F. 1993. State Social Benefits in South Africa. *International Social Security Review*, 46(1):5–25.
- May, J., Atwood, H., Ewang, P., Lund, F., Norton, A., and Wentzel, W. 1998. *Experience and Perceptions of Poverty in South Africa*, Durban: Praxis Publishing.
- McEwen, H., Kannemeyer, C. & Woolard, I. 2009. Social Assistance Grants: Analysis of the NIDS Wave 1 Dataset. *Discussion paper on the National Income Dynamics Study*, Available: <http://www.nids.uct.ac.za/home/index.php/?Nids-Documentation/discussion-papers.html> (accessed: 23 November 2009).
- Posel, D., Fairburn, J & Lund, F. 2004. Labour Migration and Households: A Reconstruction of the Effects of the Social Pension on Labour Supply in South Africa. Ninth Annual Conference on Econometric Modelling for Africa, 30 June – 2 July, Cape Town.
- Ravallion, M. 2001. The Mystery of the Vanishing Benefits: An Introduction to Impact Evaluation. *World Bank Economic Review*, 15(1): 115-140.
- Ravallion, M. 2007. Evaluating Anti-Poverty Programmes. In Evenson, R.E. & Schultz, T.P. (eds.), *Handbook of Development Economics Volume 4*. Elsevier: Amsterdam, North Holland: 3787-3846.
- Rosenbaum, P.R. & Rubin, D. 1983. The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika*, 70(1): 41-55.
- Smith, J.A. & Todd, P.E. 2005. Does matching overcome LaLonde's critique of nonexperimental estimators? *Journal of Econometrics*, 125: 305-353.
- South African Social Security Agency (SASSA). 2010. Statistical Report on Social Grants, Report number 26 (30 January 2010), Available <http://www.sassa.gov.za/STATISTICAL-REPORTS/STATISTICAL-REPORTS-654.aspx> (accessed: 5 July 2010).
- Strauss, J. & Thomas, D. 1998. Health, Nutrition and Economic Development. *Journal of Economic Literature*, 36(2): 766-817.

- Thomas, D. 1990. Intra-household Resource Allocation: an Inferential Approach. *Journal of Human Resources*, 25(4): 635-664.
- Thomas, D. 1994 Like Father, like Son; Like Mother, like Daughter: Parental Resources and Child Height. *The Journal of Human Resources*, 29(4): 950-988.
- Triegaardt, J. 2005. The Child Support Grant in South Africa: a social policy for poverty alleviation? *International Journal of Social Welfare*, 14: 249-255.
- Udjo, E.O. 2009. Trends in the age pattern of fertility, 1995-2005, in the context of the child support grant in South Africa, *Development Southern Africa*, 26(2): 289-299.
- UNICEF. 2010. Education indicators for South Africa. Available: http://www.unicef.org/infobycountry/southafrica_statistics.html (accessed 11 July 2010).
- UNICEF. 2009. *Situation Analysis of Children in South Africa*. Analysis commissioned by the Presidency of the Republic of South Africa. UNICEF: Johannesburg
- Van der Berg, S., Burger, R and Louw, M. 2010. Post-Apartheid South Africa: Poverty and Distribution. In Nissanke, M. and Thorbecke, E. 2010. *Trends in an Era of Globalization: The Poor under Globalization in Africa, Asia, and Latin America*. Oxford: Oxford University Press.
- Vidmar, S., Carlin, J., Hesketh, K. & Cole, T. 2004. Standardizing Anthropometric measures in children and adolescents with new functions for egen. *The Stata Journal*, 4(1): 50-55.
- Woolard, I., Hatting, K. & Klasen, S. *The evolution and impact of social security in South Africa. Paper prepared for the Conference on "Promoting Resilience through Social Protection in Sub-Saharan Africa", organised by the European Report of Development in Dakar, Senegal, 28-30 June, 2010.*
- World Bank. 2006. *Repositioning Nutrition as Central to Development: A Strategy for Large-Scale Action*. Directions in Development, The World Bank.

Appendix A – Algorithm used to identify children eligible for the CSG in NIDS

Receipt of the CSG was captured by two sets of questions in the NIDS survey. The first set of questions required respondents to indicate whether “*the household receives a social grant*”, followed (if the first question was answered in the affirmative) by a question asking respondents to indicate “*what type of grant*”. In addition, a question was included later in the survey asking respondents to indicate whether “*a child support grant was received on behalf of the child*”. In constructing a binary variable indicating receipt of the CSG, a positive answer for either of the two sets of questions was taken as an indication that the CSG was indeed received for that specific child.

As for the algorithm applied in identifying eligible children, all children under the age of 14 were taken as being age-eligible. In addition, monthly income data for all adults were generated from the imputed income variables included in the data set, and included all regular as well as temporary income apart from grant income and UIF payments (which is in accordance with the government regulations regarding the CSG) (Budlender *et al*, 2005: 13).

Caregivers were matched to specific children using the question included in the child data set requesting children to indicate the person code of their primary caregiver (if the caregiver resided in the household). In addition, married caregivers were identified and matched to their spouses. Monthly income data for married caregivers were calculated as using income for both adults (i.e. caregiver and spouse).

Since data on income (even after imputations) were missing for many caregivers, all cases where caregivers indicated that both they and their spouses were unemployed and where income data were missing were automatically coded as being eligible for the CSG.

Lastly, all children living in rural areas as well as in informal dwellings in urban areas were coded as falling into the rural classification specified by the means test (again, this is in accordance with the relevant government regulation) (Budlender *et al*, 2005: 9).

Appendix B – Baseline Estimates of OLS and Binary Propensity Score Matching

1. The Evaluation Problem

PSM will only reduce the bias in the OLS estimates if certain conditions hold. In the first place, a meaningful estimation of the treatment effect can only be made if there is a region where the propensity scores of the treated and control units overlap (Ravallion, 2007: 3797). This has been referred to in the literature as the “region of common support”. It has been shown that the results from PSM can be severely biased if sufficient overlap between the propensity scores of the treated and untreated units does not exist, (Caliendo & Kopeinig, 2008: 45).

In addition, following Rosenbaum and Rubin (1983: 42), there are two qualifying conditions for the use of PSM, which are listed below (notational adjustments have been made for the sake of consistency).

1.1. Balance

As a first condition, it is held that the propensity score must be such that the conditional distribution of the covariates in x , given the propensity score $p(x)$, must be the same for treated ($D = 1$) and control ($D = 0$) units. In other words, for units with the same propensity score, selection into treatment is random, and these units should be identical in terms of the observable characteristics in x (Cameron & Trivedi, 2005: 865). This is a testable hypothesis (as will be illustrated below), and can be stated as

$$D \perp x \mid p(x).$$

1.2. Ignorability

The second condition, referred to by Rosenbaum and Rubin (1983: 42) as “*strongly ignorable treatment assignment*” is also referred to as the conditional independence assumption. Rosenbaum and Rubin (1983: 44) show that the conditional independence condition also holds with the use of propensity scores, as in

$$y_0, y_1 \perp D \mid x \rightarrow y_0, y_1 \perp D \mid p(x).$$

Therefore, conditional on the propensity score, the outcome variable is independent of treatment. This second condition takes into account the fact that selection into treatment is often confounded where data are not obtained from non-experimental sources. In other words, the same factors influencing selection into treatment could also potentially influence the outcome variable. This is, of course, not a testable assumption, which emphasises the importance of the choice of covariates.

2. Matching Algorithms

Once a set of potential control units have been identified (based on the propensity score) for each treated unit, this set is weighted and matched to gauge the average difference in outcome so as to calculate the average treatment effect. The way in which the weight is specified, will of course influence how the treatment effect is estimated (Cameron & Trivedi, 2005: 874). Three methods are used in this paper, namely

- Nearest neighbour matching chooses, for each treated unit i , the set of control units where the difference in propensity score is minimised (i.e. the “nearest neighbour” to the treated unit (Cameron & Trivedi, 2005: 875).
- Kernel matching match treated units with a weighted average of all the control units, with the weight being inversely proportional to the distance of the propensity score of each of the control units to the propensity score of the treated unit (Becker & Ichino, 2002: 361).
- Radius (or calliper) matching, which matches each treated unit to the control units found within a defined radius or calliper.

3. Matching Results

3.1. Estimating the Propensity Score - Model Specification

The propensity score (i.e. the probability of each child in the sample receiving the CSG) is estimated using a linear probability model. The sample is restricted to only include eligible children, an approach which takes into account the fact that the similarities in the covariates of these two groups would potentially assist a greater overlap in the propensity scores of treated and control units, as opposed to units from the greater untreated, non-eligible population.

As set out above, without the inclusion of the caregiver-motivation variable, the model may not comply with the conditional unconfoundedness assumption discussed above (Agüero *et al* 2006:22). Accordingly, the variable was initially included in the model. However, as a result of the way in which the variable was created, it is a very strong predictor of treatment status, which causes the common support assumption to be violated (and introduces bias from that perspective). In addition, the balancing property does not hold as the caregivers of treated and untreated children have vastly different levels of motivation. The propensity into treatment estimated conditional on caregiver motivation is accordingly very different for treated and untreated children, which translates into the balancing condition not being complied with. In other words, caregiver motivation will never be similar for units with the same propensity score. Accordingly the caregiver-motivation variable was left out of the final model specification. It should, however, be

noted that the inclusion of the caregiver-motivation variable does not have a substantial effect on the results reported below.

Children whose caregivers are married are less likely to receive the CSG, potentially because married caregivers have a better support system and are therefore less likely to apply for the CSG. As expected, the age of the child has a significant positive effect on whether treatment is received, which enters the model non-linearly. As for the variables controlling for household characteristics, it is only access to a landline telephone and having a flush toilet which are statistically significant (both have a negative effect on the probability of receiving the CSG, as expected). In addition, the effect of monthly household *per capita* expenditure has a negative effect on the probability of receiving treatment.

Table B1 sets out the results of the propensity score estimations using the three samples.

Table B1: Linear Probability Model estimation of the propensity scores (probability of receiving the CSG)

Variable	Sample 1		Sample 2		Sample 3	
	Coefficient	Robust Standard Error	Coefficient	Robust Standard Error	Coefficient	Robust Standard Error
Caregiver Characteristics						
Employed	0.01	0.03	0.02	0.02	0.01	0.03
Married	-0.05***	0.02	-0.06***	0.02	-0.05***	0.02
Child Characteristics						
Age	0.09***	0.01	0.04***	0.01	0.09***	0.01
Age squared	-0.01***	0.00	0.00***	0.00	-0.01***	0.00
Boy	0.03*	0.02	0.02	0.01	0.03*	0.02
Household Characteristics						
Electricity	0.00	0.02	0.01	0.01	0.00	0.02
Piped water	0.03	0.02	0.01	0.02	0.03	0.02
Landline	-0.08**	0.03	-0.04	0.03	-0.08**	0.03
Flush toilet	-0.11***	0.03	-0.13***	0.02	-0.11***	0.03
Female head	0.01	0.02	0.02	0.02	0.01	0.02
Log pc expenditure	-0.03**	0.01	-	-	-0.03**	0.01
Constant	0.74***	0.08	0.8***	0.03	0.74***	0.08
Observations	4853		4914		4853	

*Significant at 10% level.

**Significant at 5% level.

***Significant at 1% level.


3.2. Overlap and Region of Common Support

Before continuing with the estimation of the average treatment effect, it is important to confirm the existence of a region of common support. As indicated above, sufficient overlap between the estimated propensity scores of the treated and control groups is required to ensure that the results are not biased, since the average treatment effect is estimated from the observations included in this region of common support.²³

Although various methods of confirming compliance with this condition have been used in the PSM literature, Caliendo and Kopeinig (2008: 45) suggest a visual inspection of the density distribution of the estimated propensity scores of both the treated and control groups. Figures B1 to B3 set out the histograms of the propensity scores for both the treated and control groups in all three samples.

A visual inspection of these three figures reveals that the distribution of the propensity scores seem to be sufficiently similar between the treated and control groups for the first and third sample. However, the overlap in the second sample appears to be small and limited to the region around the top-end of the distribution, which is a possible cause for concern. To limit the possibility for bias, the treatment effect is evaluated along the common support for this sample.

Figure B1: Region of common support for sample 1



²³It is for this reason that the caregiver-motivation variable has been left out of the model, as discussed in the previous section.

Figure B2: Region of common support for sample 2





Figure B3: Region of common support for sample 3



3.3. Balance

Caliendo and Kopeinig (2008: 47) discuss a variety of ways in which compliance with the balancing condition can be evaluated. One of these is stratification tests, a method used by Dehejia and Wahba (2002). In terms of this method, all observations are divided into strata based on their estimated propensity scores. Within each stratum, the difference in the mean of the covariates of treated and untreated units is assessed using t-tests.

In the first and third sample, the observations were divided into 8 strata, while the second sample was divided into 15 strata. Table B2 provides an indication of the propensity score balance within each stratum by providing the mean propensity score within the stratum, the number of treated and control units as well as the difference in the propensity score which, in most cases, is zero). Table B3 provides the t-statistics of the difference between the covariates of treated and untreated children within each stratum of each sample.

4. Estimating the Average Treatment Effect using Different Matching Algorithms

Last, the matching algorithms discussed above are now used to estimate the average treatment effect for each of the outcome variables (using the three samples set out above). The results are reported in Table B4.

The average treatment effect is first estimated by using the nearest neighbour technique. Although there is substantial overlap between the propensity scores of the treated and control units, the sample of control units is relatively small, and accordingly this technique is applied with replacement. Nearest neighbour matching is conducted so that each treated unit is matched to only one control unit (i.e. its “nearest neighbour”) as well as the nearest 10 units. The common support option is also selected.

The average treatment effect is next estimated using the Kernel matching technique. Matching is again done with replacement. However, as set out above, this technique involves a weighted average of all control units to be matched to the treated unit.

Caliper matching is done while specifying three different caliper sizes (starting with the largest specification and then decreasing the size of the caliper). The advantage of this matching method compared to the nearest neighbor method is that it allows control units matched to each treated unit to be quite similar to the treated unit.²⁴ Again, the common support condition is imposed.

²⁴There is a trade-off between variance and bias, since caliper matching could potentially provide more unbiased results, however at the cost of increased variance since fewer matches are performed (Caliendo & Kopeinig, 2008).

With the exception of the adult expenditure outcome variable, none of the PSM techniques yield a significant average treatment effect for any of the outcome variables under consideration (both analytical and bootstrapped standard errors are estimated).

The average treatment effect of the CSG on expenditure on adult goods, however, seems to be significant and robust. All of the PSM techniques applied estimate a negative effect of receipt of the CSG on household expenditure of adult goods per adult. There are, however, reasons why this result should not be taken as conclusive evidence of the effect of the CSG.

In the first place, all of the reservations mentioned previously regarding the measurement of these expenditure variables remain relevant. In addition, as indicated in the previous section, the region of common support is not very large for this sample, and observations seem to be concentrated around the top-end of the distribution. Accordingly, the number of observations used to estimate the average treatment effect for this outcome variable is very small. The concern is therefore that PSM produces biased estimates of the treatment effect.

Table B2: Propensity score balance

Strata	Mean p(x)	Number of control units	Number of treated units	Difference in propensity score	Standard Error
Sample 1					
1	-	-	-	-	-
2	0.37	15	11	0.00	0.01
3	0.54	306	301	-0.01	0.00
4	0.63	165	243	-0.00	0.00
5	0.68	176	378	-0.00	0.00
6	0.73	169	490	0.00	0.00
7	0.78	181	729	-0.00	0.00
8	-0.00	284	1398	-0.00	0.00
Sample 2					
1	-	-	-	-	-
2	-	-	-	-	-
3	0.57	14	20	-0.02	0.01
4	0.66	61	128	0.00	0.00
5	0.71	37	58	0.00	0.00
6	0.74	25	115	-0.00	0.00
7	0.76	30	41	-0.00	0.00
8	0.77	26	117	0.00	0.00
9	0.79	36	167	-0.00	0.00
10	0.83	140	683	0.00	0.00
11	0.88	179	1276	-0.00	0.00
12	0.91	87	776	-0.00	0.00
13	0.93	39	454	-0.00	0.00
14	0.94	14	337	-0.00	0.00
15	0.95	8	45	0.00	0.00
Sample 3					
1	-	-	-	-	-
2	0.37	15	11	0.00	0.01
3	0.54	306	301	-0.01	0.00
4	0.63	165	243	-0.00	0.00
5	0.68	176	378	-0.00	0.01
6	0.73	169	490	0.00	0.00
7	0.78	181	729	-0.00	0.00
8	0.84	284	1398	-0.00	0.00

Table B3: Covariate balance

		T-Statistic per Block														
Strata		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
		Sample 1														
Caregiver Employed	-	0.85	1.04	0.55	0.99	0.61	-0.21	0.09								
Caregiver Married	-	-0.27	-0.83	-2.02	2.44	1.67	0.79	0.28								
Child's Age	-	-0.43	-1.10	-2.19	2.30	0.92	-0.96	1.83								
Child's Age Squared	-	-0.55	-0.88	-1.91	1.90	0.86	-1.02	2.18								
Boy	-	-0.61	-0.24	2.33	1.80	0.15	0.18	-0.29								
Electricity	-	0.00	-0.41	-1.00	1.07	0.49	-0.58	-1.49								
Piped water	-	-0.85	-1.11	-0.55	1.30	-0.28	2.40	0.97								
Landline	-	-0.18	0.51	-0.13	-0.12	1.07	1.11	-0.67								
Flush toilet	-	1.18	-1.05	-1.02	2.31	-0.98	1.28	1.10								
Female head	-	-0.27	0.46	1.23	-1.29	0.02	-0.64	0.74								
Log pc expenditure	-	-0.14	0.71	-0.74	2.14	-0.49	-0.07	0.05								
		Sample 2														
Caregiver Employed	-	-	0.00	0.08	0.66	-0.94	0.40	0.11	-0.05	0.48	2.37	0.62	-0.83	-1.50	0.77	
Caregiver Married	-	-	0.28	-0.27	0.11	1.15	0.12	-0.71	-0.75	-0.42	-1.19	0.62	1.27	0.00	0.00	
Child's Age	-	-	0.92	-0.83	1.56	0.20	0.37	2.05	0.27	-0.73	-0.01	0.70	-0.91	-0.34	0.57	
Child's Age Squared	-	-	0.2	-0.88	1.44	-0.20	0.12	1.82	-0.00	-0.19	0.14	0.88	-0.74	-0.45	0.64	
Boy	-	-	-0.68	1.31	1.47	1.40	1.34	1.10	0.81	-0.91	0.13	1.10	0.45	0.77	-1.06	
Electricity	-	-	-0.92	-0.30	0.61	-0.67	0.29	1.20	1.40	-0.12	0.20	-2.17	-1.34	-0.41	0.00	
Piped water	-	-	0.28	-0.32	1.10	1.53	0.30	1.96	1.12	-1.63	0.82	1.57	0.89	-1.12	0.74	
Landline	-	-	1.51	0.88	1.72	0.42	-1.08	-1.63	-0.51	0.61	-0.89	-1.25	-1.03	-0.41	0.00	
Flush toilet	-	-	0.00	0.00	1.39	1.53	1.33	2.26	1.64	-1.95	1.81	0.00	0.00	0.00	0.00	
Female head	-	-	0.92	-0.24	0.71	0.32	-1.42	0.89	0.04	-0.18	0.57	-1.67	0.88	0.40	0.60	
Log pc expenditure	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	

Sample 3								
Caregiver Employed	-	0.85	1.04	0.55	0.99	0.61	-0.21	0.09
Caregiver Married	-	-0.27	-0.83	-2.02	2.44	1.67	0.79	0.28
Child's Age	-	-0.43	-1.08	-2.20	2.30	0.92	-0.96	1.83
Child's Age Squared	-	-0.55	-0.88	-1.91	1.90	0.86	-1.02	2.18
Boy	-	-0.61	-0.24	2.33	1.80	0.15	0.18	-0.29
Electricity	-	0.00	-0.41	-1.01	1.07	0.49	-0.58	-1.49
Piped water	-	-0.85	-1.11	-0.55	1.30	-0.28	2.40	0.97
Landline	-	-0.18	0.51	-0.13	-0.12	1.07	1.11	-0.67
Flush toilet	-	1.18	-1.05	-1.02	2.31	-0.98	1.28	1.08
Female head	-	-0.28	0.46	1.22	-1.29	0.02	-0.64	0.74
Log pc expenditure	-	-0.14	0.71	-0.74	2.14	-0.49	-0.07	0.05

Table B4: Estimated average treatment effects (standard errors in parenthesis)

Matching Algorithm	Sample 1		Sample 2		Sample 3	
	HAZ	WHZ	Food Expenditure	Adult Expenditure	School Enrolment	School Progress
Nearest 1 Neighbour [#]	0.07 (0.07) <i>N=3401</i> <i>T=2621</i>	-0.02 (0.12) <i>N=1738</i> <i>T=1400</i>	23.00 (32.11) <i>N=4910</i> <i>T=4213</i>	-8.99 (3.97) <i>N=922</i> <i>T=765</i>	0.00 (0.01) <i>N=2784</i> <i>T=2117</i>	-0.00 (0.03) <i>N=2621</i> <i>T=1978</i>
Nearest 10 Neighbours [#]	0.05 (0.07) <i>N=3401</i> <i>T=2621</i>	-0.01 (0.12) <i>N=1738</i> <i>T=1400</i>	42.46 (32.11) <i>N=4910</i> <i>T=4213</i>	-9.54 (3.97) <i>N=922</i> <i>T=765</i>	0.00 (0.01) <i>N=2741</i> <i>T=2117</i>	0.00 (0.03) <i>N=2621</i> <i>T=1978</i>
Kernel	0.03 (0.07) <i>N=3403</i> <i>T=2623</i>	-0.01 (0.12) <i>N=1743</i> <i>T=1405</i>	24.38 (32.11) <i>N=4914</i> <i>T=4217</i>	-9.37 (3.97) <i>N=924</i> <i>T=767</i>	0.01 (0.01) <i>N=2786</i> <i>T=2119</i>	-0.00 (0.03) <i>N=2623</i> <i>T=1980</i>
Radius (caliper=0.01) [#]	0.04 (0.07) <i>N=3399</i> <i>T=2619</i>	0 (0.12) <i>N=1741</i> <i>T=1403</i>	25.97 (32.11) <i>N=4914</i> <i>T=4217</i>	-9.48 (3.97) <i>N=922</i> <i>T=765</i>	0.01 (0.01) <i>N=2785</i> <i>T=2118</i>	0.00 (0.03) <i>N=2622</i> <i>T=1979</i>
Radius (caliper=0.001) [#]	0.04 (0.07) <i>N=3322</i> <i>T=2542</i>	-0.01 (0.12) <i>N=1565</i> <i>T=1227</i>	27.26 (32.11) <i>N=4740</i> <i>T=4043</i>	-7.86 (3.97) <i>N=547</i> <i>T=704</i>	0.01 (0.01) <i>N=2723</i> <i>T=2056</i>	-0.00 (0.03) <i>N=2559</i> <i>T=1916</i>
Radius (caliper=0.0001) [#]	0.07 (0.07) <i>N=1813</i> <i>T=1033</i>	-0.05 (0.12) <i>N=644</i> <i>T=306</i>	28.73 (32.11) <i>N=3985</i> <i>T=3288</i>	-11.71 (3.97) <i>N=398</i> <i>T=241</i>	0.00 (0.01) <i>N=1369</i> <i>T=702</i>	-0.03 (0.03) <i>N=1270</i> <i>T=627</i>

Notes: [#]With replacement.

All estimations done on the region of common support.

Bootstrapped standard errors (200 repetitions).

N=number of observations, T= number of treated observations used in matching