Finding the benefits Evaluating the Impact of the South African Child Support Grant

Marisa Coetzee
University of Stellenbosch

27 - 29 October 2010
Indaba Hotel and Conference Centre
Johannesburg
South Africa

Hosted by

DPRU Development Policy Research Unit

TIPS

Finding the Benefits Evaluating the
Impact of the South African Child Support Grant

Marisa Coetzee
University of Stellenbosch

Draft paper prepared for the DPRU Conference on *The Global Economic Crisis and South Africa: Lessons in Long-Run Economic Growth and Development*, October 2010

Contents

1. Introduction 3
2. Background to the South African Child Support Grant 4
3. An Overview of the Existing Literature 6
   3.1. Impact on Labour Force Participation 7
   3.2. Impact on Education 8
   3.3. Impact on Child Health and Nutrition 9
   3.4. The use of Propensity Scores 10
4. Finding the Benefits - the Choice of Outcome Variables 11
5. The Data 14
6. Delay in Application and Caregiver Motivation (Eagerness) 20

First technique: Matching Results using a Binary Treatment Variable 21

7. Methodology 21
   7.1. The Evaluation Problem 21
   7.2. Exact Matching 23
   7.3. Propensity Score Matching 23
      7.3.1. Balance 24
      7.3.2. Ignorability 24
   7.4. Matching Algorithms 25
8. Matching Results 27
   8.1. Estimating the Propensity Score - Model Specification 27
   8.2. Overlap and Region of Common Support 30
   8.3. Estimating the Average Treatment Effect using Different Matching Algorithms 32

Second technique: The use of Generalised Propensity Scores with a Continuous Treatment Variable 36

9. Methodology 36
   9.1. Theoretical Framework and Underlying Assumptions 37
   9.2. Estimation of the Treatment Effect and Inference 39
10. Results 40
   10.1. Modeling the Conditional Distribution of Treatment 41
   10.2. Estimating the Outcomes Conditioned on the GPS and Treatment 46
   10.3. Estimating the Average Treatment Effect 46
11. Conclusion 50
12. References 51

Appendix A – Method used to identify children eligible for the CSG in NIDS 55

List of Tables

Table 1: Number of Child Support Grants per province – January 2010 5
Table 2: Number of children eligible for and receiving the CSG in 2008 15
Table 3: Descriptive statistics by treatment and eligibility status 17
Table 4: Estimation of the propensity scores (probability of receiving the CSG) 29
Table 5: Estimated average treatment effects (standard errors in parenthesis) 35
Table 6: Maximum likelihood estimates with different samples 41
Table 7: The balancing property: t-statistics for equality of means between treatment terciles . 45
Table 8: Second stage estimates per outcome 46

List of Figures

Figure 1: Main reason CSG was not applied for by eligible caregivers 16
Figure 2: Region of common support for sample 1 31
Figure 3: Region of common support for sample 2 31
Figure 4: Region of common support for sample 3 32
Figure 5: Estimated gains in HAZ 48
Figure 6: Estimated gains in WHZ 48

1. **Introduction**

Anti-poverty programme evaluation has received increasing attention in the recent past as various stakeholders including governments, aid donors and the greater development community 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 a starting point, data on a suitable control group (or counterfactual) must be obtained before any meaningful evaluation can take place. This is an arduous task, specifically when dealing with anti-poverty programmes, since any comparable control group would be so similar to the treated group that there are usually ethical objections against excluding them from treatment. Another stumbling block is the difficulty in identifying the impact of the programme on individual outcomes, i.e. isolating changes in outcomes that are solely as a result of the programme, and not as a result of latent differences between participants and non-participants.

One solution would, of course, be to implement programmes in a randomised fashion and use these experimental data to evaluate the impact thereof (Smith & Todd, 2005: 305). However, anti-poverty programmes are rarely implemented in an experimental or quasi-experimental way. Accordingly, a wide variety of estimation techniques, dealing specifically with solving the evaluation problem in the context of non-experimental data, are available in the literature. Propensity score matching and other techniques utilising propensity scores have repeatedly been used 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),

[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.

which involves an unconditional cash transfer to eligible caregivers of children, who are identified by way of a means test.

This paper assesses the impact which the CSG has on the well-being of recipient children using two techniques in which propensity scores are utilised. For this purpose, education and child health and nutrition have been selected as indicators of the general well-being of children. These outcomes are not only 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 currently administered as a conditional cash transfer programme, conditions regarding proper nutrition, health care and school enrolment of recipient children were included in the initial regulations accompanying the relevant legislation (Leatt & Budlender, 2007).[2] In assessing the success of the CSG, it is therefore important to evaluate the effect of the CSG against these outcomes.

The remainder of this paper is organised as follows. In section 2, some background to the CSG is provided. Section 3 provides a brief overview of the literature on the subject, also focussing specifically on a previous study which made use of propensity scores to evaluate the impact of the CSG. Section 4 explains in further detail the choice of specific outcome variables. Section 5 provides an overview of the data used while section 6 discusses caregiver motivation and application delay. The seventh and eight sections present the methodology and results from assessing the effect of the CSG using the first propensity score technique, while sections 9 and 10 repeat this for the second technique. Section 11 concludes.

2. **Background to the South African Child Support Grant**

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 recipient children and 5 377 476 beneficiary 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 *et al*, 2010, Leibbrandt *et al,* 2010: 65).

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

In addition, analysis indicates that the grant appears to be reasonably well targeted. For instance, receipt of the CSG seems to be more concentrated in the poorest and most rural provinces, including KwaZulu-Natal, Eastern Cape and Limpopo. A breakdown of the grants per province from SASSA is provided in the table below.

Table 1: Number of Child Support Grants per province – January 2010

Province	**Number of CSG Recipients**
Eastern Cape	1 634 612
Free State	514 559
Gauteng	1 125 183
KwaZulu-Natal	2 378 853
Limpopo	1 425 392
Mpumalanga	735 648
Northern Cape	218 709
North West Province	708 428
Western Cape	610 593

Source: SASSA (2010: 8).

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.[3] In response to the alleged success of the CSG, 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 survey used in this paper, the amount was R210).

As indicated above, the CSG is a means-tested social grant, and its initial aim at introduction was to provide assistance to the poorest 30% of children in South Africa (Agüero *et al*, 2009: 5). 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

[3] 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).

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 *et al*, 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 cash transfer programme, *Progressa* (now *Oportunidades*), which is conditional on the recipient household meeting certain conditions. These conditions 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 *et al,* 2005:468). In addition, the cash transfers 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 in the literature. The studies have to a large extent focussed on the poverty-alleviating impacts of these transfers within 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).

[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 aim of this section is to provide an overview of the existing literature. Although the current paper's focus is on the health, nutritional and education outcomes of recipients of the CSG, the effect of the CSG on labour force participation is also discussed to provide broader context. In addition, although there are obvious differences between the OAP and the CSG[6], both grants have been found to impact poverty in a similar fashion and, accordingly, literature on the impact of the OAP is also discussed.

3.1. **Impact on Labour Force Participation**

Various previous studies have focused on the incentives created by the CSG and OAP with respect to the labour force participation of members in households other than the grant recipients themselves.

Bertrand *et al* (2003) use the 1993 Project for Statistics on Living Standards and Development (PSLSD) data to evaluate the effect of the OAP on the employment patterns of other household members. More specifically, the focus of the study is on prime-aged African males and females (16 to 50 years) living in three-generation households. The results 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. In addition, the additional marginal rand of pension income flowing to a female pensioner reduces the labour supply of prime-aged individuals approximately three times more than the additional marginal rand of pension income flowing to a male pensioner.

Using the same data, Posel *et al* (2004) show that the results by Bertrand *et al* (2003) are sensitive to the inclusion of non-resident household members. The authors include data on these individuals in their sample to analyse the effect of the OAP on the labour supply of migrant household members. They conclude 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.

The authors suggest that this positive effect of the OAP on labour participation might be because the funds received from the grant allows prime-age females to migrate in search of employment, leaving children in the care of older relatives. More concrete evidence

[6] Two apparent differences are the fact that the monetary value of the monthly CSG is much less than the OAP and the fact that the CSG is aimed at the caregivers of children while the OAP is explicitly aimed at the elderly.

that this might be the case is found by Ardington *et al* (2009). The authors use panel data from northern KwaZulu-Natal to show 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, while these individuals are less likely to be migrant workers in cases where the age-eligible individual leaves the household. This is in accordance with the hypothesis that the pension income relieves the credit constraints of poor households and enables older members to look after children, freeing younger individuals up to migrate in search of employment.

However, 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.

Although only indirectly related to the question of labour force participation, Udjo (2009) examines whether payment of the CSG to female caregivers has any effect on age specific fertility rates in South Africa. Although no conclusive causal link can be found based on the data used in the study, Udjo (2009: 297) highlights the relationship between the CSG and the fertility of prime-age African women as a future area of research.

3.2. **Impact on Education**

The effect of social grants on the educational attainment of children has long been the topic of research. A number of studies from the 1990's find evidence that income from the OAP has a significantly positive effect on the education of children in the household (Lund, 1993 and May *et al*, 1998).

More recent studies such as Samson *et al* (2001) confirm these findings for the OAP. Using the 1997 October Household Survey, Samson *et al* (2001) focus on the school enrolment of children of school-going age in three-generation households. A significant positive effect on the probability of school enrolment is found in the bottom quartile of the income distribution for children living with grandparents, while the effect is larger for girls than for boys.

Edmonds (2004) considers the effect of the OAP not only on school enrolment, but also child labour. He finds that the presence of an age-eligible pensioner in the household has a positive effect on school enrolment and also decreases child labour. Interestingly, and contrary to previous evidence on the effect of the OAP on non-recipient household members, the study concludes that child labour decreases and school enrolment increases more if the individual eligible for the OAP is male, rather than female.

Case *et al* (2005) confirm that the CSG also has a significantly positive effect on school enrolment. The authors use longitudinal data from the Umkhanyakude district in KwaZulu-Natal collected through the Africa Centre for Health and Population Studies. From the data, the authors are able to conclude that receipt of the CSG increases the probability of children aged 6 and 7 being enrolled at school in the years following receipt of the grant. In addition, the study uses older children who were not eligible for the CSG (as a result of the relatively young cut-off age which applied at the time of the study) as a counterfactual to measure the size of this effect, concluding that these older children are significantly less likely to be enrolled in school when they were aged 6 and 7. The authors suggest 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.

3.3. **Impact on Child Health and Nutrition**

Case (2001a & 2001b), using data on 300 households in the Langeberg health district in the Western Cape, evaluates the effect of receiving the OAP on health outcomes of household members living with the pensioner. The effect of the OAP receipt is estimated by comparing the self-reported health status of adults and the height-for-age of children living with an OAP recipient *versus* those living in a household without any pension recipient. The study evaluates the effect on health status based on whether the income received by the household is pooled or not. It finds that in households where income is pooled, receipt of the OAP has a positive effect on all adult household members; however, where income is not pooled, only the recipient benefits from the OAP in terms of improved health.

As for children, the study shows that the presence of a pension-recipient in the household increases the height-for-age of African and Coloured children by approximately one standard deviation (between 3 to 5 centimeters), 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. Using the 1993 PSLSD data, and focusing on both the HAZ and WHZ of children between 0 and 60 months, she concludes that the cash received from the old-age pension has no significant effect on the health and nutrition of either boys or girls, if the cash transfer were received by a male within the household. However, the effects of the cash received from the old-age pension has a significantly positive effect on the HAZ and WHZ for girls (but not boys) if the transfer were made to a female within the household.

In a more recent study, Agüero *et al* (2006 & 2009) find similar results for children's HAZ. This study is discussed in more detail in the section below.

3.4. **The use of Propensity Scores**

Agüero *et al* (2009) evaluates the impact of the CSG on the weight-for-height z-scores (HAZ)[7] of children during the first 36 months of life using the KwaZulu-Natal Income Dynamics Study (KIDS) (Agüero *et al,* 2009: 3).

The continuous treatment evaluation technique[8] 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 authors specifically include a variable controlling for the eagerness of the caregiver in applying for the CSG. This eagerness variable is calculated as the deviation from average delay in applying to receive the CSG for children in the same age and location cohort (Agüero, 2009: 13), and is included specifically to control for latent caregiver characteristics which are not necessarily sufficiently controlled for by other observable child, caregiver and household characteristics. In other words, conditional independence is ensured, since, conditional on this eagerness variable and the other covariates, the extent of the CSG received is random and depends only on the birthplace and date of the child (which, it is of course assumed, is not influenced by the CSG) (Agüero *et al*, 2009: 20).

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

[7]Described in detail in the next section.
[8] Described in further detail below.

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-fourths of the first 36 months. As an example of how these gains translate into increases in human capital stocks, it is calculated that for children receiving the grant for two-thirds of the 36 month period, an increase of 0.2 in a child's HAZ, translates to a 1.8cm or 1.1% gain in adult height, which again translates into an increase in monthly wages of approximately R67 to R92 (Agüero *et al*, 2009: 28).

The current paper draws from both of these studies and applies two techniques utilising propensity scores to evaluate the CSG impact.

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 emphasized in the literature (Barro, 1996).[9] An array of measures for nutritional outcomes exists, including the standardized weight-for height (WHZ), height-for-age (HAZ) and weight-for-age (WAZ) z-scores, iron deficiency anemia (IDA) and sub-clinical vitamin A deficiency (VAD) (World Bank, 2006: 210). In this paper, as a result of data limitations, only HAZ and WHZ will be used. 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 in cm, and $\bar{h}$ and σ_h are respectively the mean height and the standard deviation of height, given age in years.[10]

These two measures differ in that 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. Nutritional deficiencies causing stunting, especially in the early stages of a child's life, can lead to permanent damage as far as a

[9] 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).
[10] 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.

person's health and productivity later in life is concerned (Duflo, 2003: 12; Agüero *et al*, 2009: 7).

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 measure of health and nutrition. Food expenditure is measured as the *per capita* 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)[11] (Straus & 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 been created using expenditure on alcohol and cigarettes (two items which are

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

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 survey used for this analysis 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 0.5cm or 0.5 kg.[12] 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, birth dates 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 physically be possible 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),[13] 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

[12] See Agüero *et a*, (2009: 8) for a similar discussion of the KIDS data.

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

school. It may also provide an indication of parental support and how conducive the child's living conditions and social environment is to learning.

While there are caveats and problems pertaining to each of these 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 later 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 both 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 is important as a potential counterfactual, as one would expect this group to be quite similar to the treated group in terms of background characteristics.

Using the means test as it was applied during 2008 (being 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 results are set out in Table 2 below (additional information on the algorithm used to identify eligible children is set out in Appendix A).

[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, 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).

Table 2 highlights 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.

Table 2: Number of children eligible for and receiving the CSG in 2008

	Receiving CSG		
Eligible	**Yes**	**No**	**Total**
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)

Most of the eligible children who are not receiving the grant come from the poorest households, with approximately 80% of these children coming from households where monthly per capita income is R500 or less.

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 10% of the caregivers indicated that they just “haven’t gotten round to it yet” and slightly less than 9% indicating that they just “couldn’t be bothered”.

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

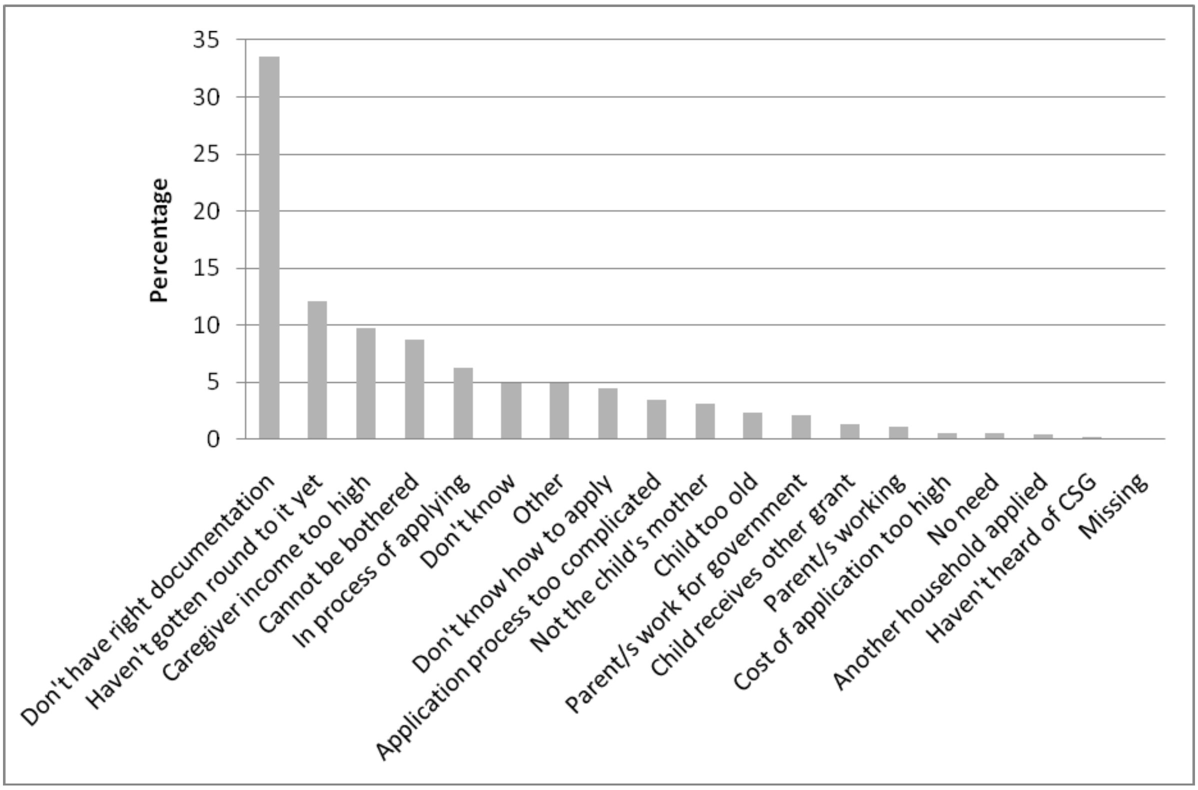

Source: NIDS (2008)

The descriptive statistics per treatment status are listed in Table 3 below. For comparative purposes, receipt of the CSG is defined both as a binary (i.e. receiving treatment or not) as well as continuous (the length of treatment, or "dosage" received) variable. 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 0-35% of the child's life, medium dosage as 35-68% and high as 68-100%). 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 3). First, untreated eligible children typically have caregivers with lower levels of education than children who do not qualify for the grant (6.88 *versus* 10.03 years on average). In addition, the caregivers of untreated eligible children are significantly more likely to be unemployed than children who do not qualify for the CSG.

Table 3: Descriptive statistics by treatment and eligibility status

Variable	Treated				Not Treated	
	Low [0-35%)	Medium [35-68%)	High [68-100%]	All	Eligible	Not Eligible
Outcomes						
HAZ	-0.65	-0.46	-0.32	-0.50	-0.48	-0.19
	(1.35)	(1.33)	(1.37)	(1.37)	(1.33)	(1.37)
WHZ	-0.05	0.56	0.14	0.20	0.13	0.39
	(1.39)	(1.69)	(1.39)	(1.47)	(1.46)	(1.46)
Per capita monthly household	141.27	152.09	139.43	147.70	182.27	425.39
expenditure spent on food items#	(119.31)	(147.65)	(89.13)	(99.50)	(159.66)	(412.45)
Household expenditure on adult	25.82	25.99	24.53	24.28	32.67	182.23
goods per adult[16] in the household#	(46.95)	(38.53)	(39.80))	(27.15)	(49.23)	(216.63)
Proportion of children enrolled at	0.99	0.99	0.98	0.99	0.99	0.98
school##	(0.11)	(0.11)	(0.13)	(0.11)	(0.08)	(0.13)
Proportion of children ever	0.25	0.23	0.13	0.23	0.24	0.17
repeated school year##	(0.43)	(0.42)	(0.34)	(0.42)	(0.43)	(0.38)
Caregiver Characteristics						
Motivation	-1.08	0.17	0.77	0.56	-1.65	N/A
	(1.59)	(1.59)	(0.91)	(1.94)	(0.62)	N/A
Years of education	6.91	7.99	8.48	7.43	6.88	10.03
	(3.91)	(3.85)	(3.77)	(4.04)	(4.16)	(3.45)
Proportion employed	0.18	0.14	0.18	0.15	0.07	0.46
	(0.38)	(0.35)	(0.38)	(0.36)	(0.25)	(0.50)
Proportion female	0.97	0.97	0.95	0.96	0.97	0.89
	(0.17)	(0.18)	(0.22)	(0.20)	(0.18)	(0.32)
Proportion married	0.32	0.34	0.32	0.34	0.36	0.66
	(0.47)	(0.47)	(0.47)	(0.47)	(0.48)	(0.47)
Child Characteristics						
Age in years	9.10	6.50	5.15	6.91	6.20	7.44
	(3.52)	(4.25)	(3.36)	(3.91)	(4.60)	(4.64)
Proportion male	0.47	0.51	0.53	0.50	0.46	0.51
	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)
Household Characteristics						
Proportion access to electricity	0.68	0.74	0.75	0.73	0.74	0.89
	(0.47)	(0.44)	(0.43)	(0.44)	(0.44)	(0.31)
Proportion access to piped water	0.59	0.61	0.60	0.56	0.57	0.86
	(0.49)	(0.49)	(0.49)	(0.50)	(0.50)	(0.35)
Proportion access to landline	0.08	0.09	0.08	0.07	0.12	0.29
	(0.27)	(0.28)	(0.27)	(0.27)	(0.32)	(0.45)
Proportion access to flush toilet	0.32	0.33	0.36	0.31	0.38	0.75
	(0.47)	(0.47)	(0.48)	(0.46)	(0.49)	(0.43)
Proportion female headed	0.57	0.55	0.55	0.57	0.58	0.30
	(0.49)	(0.50)	(0.50)	(0.50)	(0.49)	(0.46)
Per capita expenditure	419.92	433.64	437.94	393.34	475.06	2239.25
	(427.44)	(551.65)	(509.85)	(459.96)	(840.33)	(2881.32)

Notes: Mean values with standard deviation in parenthesis.

[16] Defined as individuals aged 15 years and older.

#Treatment calculated at the household level.
##Excluding children younger than 5 years old.

In addition, the child-level characteristics vary across these two groups, with eligible untreated children being on average almost a year younger that non-eligible children and a slightly greater proportion of the children in the eligible group being female.

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 R2 041) 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 3 respectively), it would appear that these two groups are much more similar, with differences between children, their caregivers and living conditions of these two groups being small if not insignificant. Based on observable characteristics, it would therefore appear that this untreated eligible group could potentially form a counterfactual which compares well with children who are receiving the CSG.

Looking at the breakdown of the treated group into terciles based on the treatment dosage (columns 2, 3 and 4 of Table 3), it can be concluded that the differences in the observable characteristics of these three groups are mostly small and negligible. Two characteristics however deserve further mention. First, the relatively large gap between the average caregiver's years of education in each of these three groups is in accordance with the belief that more educated caregivers would be more efficient in applying for the CSG as early as possible (possibly because they are more literate but presumably more educated caregivers would also be more motivated to ensure receipt of the CSG as early in their children's' lives as possible).

Second, the fact that there is such a significant difference between the age of children receiving the grant for a small proportion of their lives (9.10 years) and those who have received a medium dosage (6.50 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 with the *zanthro* command in *Stata* (see Vidmar *et al,* 2004).[17] As a reference, the 2000 US CDC Growth Chart, which is premised on the anthropometric measures of a sample of well-nourished US children, is used. 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 listed above, there are definite improvements when the dosage is increased from low to medium. The only two exceptions are for school enrolment and expenditure on adult goods, where no significant change is observed.

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).

[17] These commands calculate the z-score (z) of each child.

Last, Table 3 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 below.

6. **Delay in Application and Caregiver Motivation (Eagerness)**

The descriptive statistics in the previous section set out observable characteristics which would potentially 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.

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. For reasons discussed more fully below, these unobserved factors, if not controlled for, could potentially bias the estimation of the treatment effect.

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,[18] with a delay of more than 3 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 at the time 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

[18] 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).

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).[19] 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.[20]

Table 3 sets out the positive relationship between caregiver motivation and the extent of treatment (increasing from one standard deviation below the average to almost one standard deviation above). It is also clear that treated children have more motivated caregivers (0.56 standard deviations above the average) than untreated children (1.65 standard deviations below the average). Given the discussion of the data, the first technique employed in the paper is set out in the next section.

First technique: Matching Results using a Binary Treatment Variable

7. Methodology

7.1. The Evaluation Problem

As set out above, in order to assess the impact of the CSG on the nutrition and health of the children receiving the grant, some form of treatment evaluation is required. In other words, the challenge is to identify the change in the outcome variables (child health, nutrition and education in this instance), resulting from the “treatment” (the cash transfer received from government in this instance), while holding all other relevant factors constant (Cameron & Trivedi, 2005: 861, Ravallion 2007: 3791).

This challenge is often best met in a situation where experimental data are available, and treatment is assigned randomly (Cameron & Trivedi, 2005: 862; Lalonde, 1986: 609). This ensures that the counterfactual is identified so as to represent the treated group as closely as possible, while the actual treatment status is not influenced by the underlying characteristics of the individuals. However, experimental data are not always available, and it is often necessary to use observational data to evaluate the effect of programmes on outcome variables. In a non-experimental context, treatment is often assigned based

[19] 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%.

[20] 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 end of the period in which the survey was completed.

on some observable characteristic, in the current instance the CSG is awarded based on a means test. This effectively removes the possibility of identifying the counterfactual (Cameron & Trivedi, 2005: 871), creating the need for alternative, non-experimental methods.[21]

In essence, treatment evaluation requires the estimation of the average treatment effect, written as

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

where y_{1i} and y_{0i} are the treatment and control outcomes respectively,[22] and D_i is the binary treatment variable for individual i, with $D_i = 1\ (0)$ 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 is potentially 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 in the absence of experimental data or randomization that has received considerable attention during the past few years, is matching. As a substitute for randomization data are obtained from a potential comparison group (which may or may not be drawn from the same population as the treated group) with pre-treatment observable characteristics or covariates (contained in the vector x) matching those of the treated units (Cameron & Trivedi, 2005: 871). Matching therefore uses the newly created counterfactual to obtain an estimate of the outcome variable of the treated group in the absence of treatment.

[21] It should be noted that this is an additional difference between the South African CSG and *Oportunidades,* which involved a programme roll-out closely resembling an experimental study, with rural areas being randomly selected into programme participation, allowing for the counterfactual to be formed from the areas not selected (Agüero *et al,* 2009: 2).

[22] Some authors prefer the simplified notation Y^T and Y^C for treated and control units respectively (see Ravallion, 2007: 3801), however, this paper will follow the notation of Cameron and Trivedi (2005), as set out above.

7.2. **Exact Matching**

One option available is to match units directly on observable characteristics. This assumes that, conditional on the observable characteristics of the individuals, the outcome is independent of treatment. In other words,

$$y_0, y_1 \perp D \mid x$$

(Cameron & Trivedi, 2005: 863). This assumption has been referred to in the literature as the "conditional mean assumption", "selection on observables" or "unconfoundedness" (Caliendo & Kopeinig, 2008: 32) and these terms are used interchangeably in this paper.

In addition, matching also assumes that for all x, there is a positive probability of participating or not participating (Smith & Todd, 2005: 312), i.e.

$$0 < Pr(D = 1|x) < 1$$

As discussed in more detail by Dehejia & Wahba (2002: 151), this method may prove to be problematic if the dimensionality of the observable characteristics is high. More specifically, if there are *n* covariates of interest, and each covariate is divided into two levels (e.g. in the case of binary variables), matching on x would have to be done on 2^n levels (Dehejia & Wahba, 2002:153).

7.3. **Propensity Score Matching**

An alternative to the above is propensity score matching (PSM), i.e. the conditional probability measure of participation in treatment, given the underlying characteristics contained in x (Cameron & Trivedi, 2005: 864). The term "propensity score" was first used in the seminal paper by Rosenbaum and Rubin (1983), which contains a detailed discussion of the technique. PSM has since been used as a solution to the dimensionality problem set out above.

The propensity score is expressed as $p(x)$, or more specifically

$$p(x) = \Pr[D = 1|X = x],$$

i.e the probability of receiving treatment, conditional on observable characteristics (Cameron & Trivedi, 2005: 864). Since the propensity score is essentially the predicted probability of selection into treatment, it can be estimated by using either a logit or probit model. Instead of matching directly on the covariates of interest, matching now occurs on the propensity score (i.e. the probability of receiving treatment).

The above statement makes it clear that, once the propensity score has been calculated, 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". If sufficient overlap between the propensity scores of the treated and untreated units does not exist, results can be severely biased (Caliendo & Kopeinig, 2008: 45). It is accordingly important to ascertain whether this condition holds.

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

7.3.1. **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

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

7.3.2. **Ignorability**

The second condition, referred to by Rosenbaum and Rubin (1983: 42) as "*strongly ignorable treatment assignment*" is essentially an application of the conditional independence assumption on propensity scores. 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 (which of course is a function of the covariates in x), 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. However, unlike experimental data, it is uncertain whether all of these factors are contained as covariates in the vector x (Rosenbaum & Rubin, 1983: 43). Therefore, in order to obtain unbiased estimates of the treatment effect, there must be no omitted variable bias once x is controlled for. If this is not the case, assignment to treatment will be confounded given the propensity score and results will be biased.

This is, of course, not a testable assumption, which emphasises the importance of the choice of covariates. In addition, and more importantly, this condition emphasises the quality of the data which are used, so as to facilitate the choice of covariates. This point has been emphasised repeatedly in the literature (see Ravallion, 2007: 3804; Smith & Todd, 2005: 319).

An additional point to make at this stage is that PSM, although often sufficient to solve the issue of confoundedness between selection into treatment and the outcome variables, is only useful when selection into the treated group is done only on the basis of observable characteristics (Cameron & Trivedi, 2005: 871). In other words, should there be selection into treatment based on unobservable child, caregiver or household characteristics, conditional independence will not hold and some form of selection bias will remain (Ravallion, 2001: 125; 2007: 3828), leading to an under-or over-estimation of the effect of treatment.

7.4. **Matching Algorithms**

Once a set of potential control units have been identified (based on the propensity score) for each treated unit, this set can matched to each treated unit to gauge the average difference in outcome so as to calculate the average treatment effect. The set of control units which could potentially be matched with the treated unit i with covariates x_i can be defined as $A_j(p(x)) = \{p_j | p(x_j) \in c(p(x_i))\}$, where $c(p(x_i))$ denotes the neighbourhood of x_i. If N_c denotes the number of control units in this set, and $\omega(i,j)$ the weight allocated to the j^{th} unit (control) being compared to the i^{th} unit (treated), then the average treatment effect (ATT) can be calculated as

$$ATT = \frac{1}{N_1} \sum_{i \in \{D=1\}} (y_{1,i} - \Sigma_j \omega(i,j) y_{0,j}).$$

The way in which the weight is specified, will of course influence how the treatment effect is estimated (Cameron & Trivedi, 2005: 874).

a. 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). This set is defined as

$$A_j(p(x)) = \left\{p_j \middle| \min_j |p_i - p_j|\right\}.$$

This method essentially ignores the weighting function discussed above. Although the nearest neighbour matching technique has the advantage that all treated units will be matched (i.e. none of the treated units will be ignored because suitable control units could not be found), this means that treated units may be matched with unsuitable control units, which could potentially lead to imprecise results (Becker & Ichino, 2002: 361).

b. As a solution, Kernel matching specifies the weight function as

$$\omega(i,j) = \frac{K\big(p_i(x) - p_j(x)\big)}{\sum_{j=1}^{N_{0,j}} K\left(p_i(x) - p_j(x)\right)},$$

where K denotes the Gaussian kernel

$$K = \frac{1}{\sigma\sqrt{2\pi}} \varepsilon^{-\frac{p(x)^2}{2\sigma^2}}.$$

Therefore, Kernel matching has the advantage that treated units are matched 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).

Both of these estimation techniques have the advantage that no functional form is assumed when estimating the outcome equation (Cameron & Trivedi, 2005: 875).

c. An additional method of estimating the average treatment effect includes radius (or calliper) matching, which matches each treated unit to the control units found within a defined radius or calliper. The set defined above now becomes (Cameron & Trivedi, 2005: 876)

$$A_j(p(x)) = \left\{p_j \middle| \min_j |p_i - p_j| < r\right\}.$$

Two additional issues should be considered before implementing PSM (Dehejia & Wahba, 2002: 153).

- First, the question arises whether matching should take place with or without replacement, i.e. whether the control units should be matched multiple times (i.e. with replacement) or whether control units should be "ignored" once they have been matched once (i.e. without replacement). Matching with replacement has the advantage of ensuring that the distance between the propensity scores of the treated and control units remain as small as possible, reducing bias in the estimates. However, in the case where there are only a few control units which are similar to the treated units, matching without replacement, although possibly increasing bias, has the advantage that it reduces the variance and improves the precision of the estimates. The choice therefore appears to be a trade-off between bias and precision (Dehejia & Wahba, 2005: 153).
- The second question is the number of comparison units which should be matched to each treated unit. Again, it is a question of trade-off between bias and precision, since a greater number of control units will lead to more precise estimates, but at the cost of potentially inducing bias in the estimates (Dehejia & Wahba, 2005: 153).

8. Matching Results

Matching was done in accordance with the methodology developed by Leuven and Sianesi (2003). The specific methodology applied and the results obtained are discussed in detail in the sub-sections below.

8.1. Estimating the Propensity Score - Model Specification

As a first step, the propensity score (i.e. the probability of each child in the sample receiving the CSG) should be estimated. Previous studies have used either a logit or probit specification to estimate the propensity score. Both models yield comparable results and accordingly either could be used for his purpose. For the current study, a logit model is used, following the approach of Dehejia & Wahba (2005).

The sample was 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 indicated in Table 3 containing the descriptive statistics of the data.

The logit model should contain covariates which influence both take-up of the CSG and the outcome variables (Caliendo & Kopeinig, 2008: 38). In addition, it must be specified such that the conditional independence assumption listed above is satisfied. Here the basic model specified Agüero *et al* (2006: 19) is used as a starting point, controlling for observable caregiver characteristics (gender, employment status, education and marital status) as well as observable child characteristics (gender and age). 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)[23].

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 originally 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). 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.

In addition, 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 receiving the CSG, i.e. not only children who were receiving the CSG at the time of the survey, but also those who were eligible for the CSG at the time of the survey, but not receiving the CSG (i.e. who have effectively been receiving the CSG for 0% of their lives).

[23] 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.

The second sample is used for the expenditure outcome variables. Here 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 (where previously they were included as having received the CSG for 0% of their lives). 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-recipient children within that household).

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

Table 4 sets out the results of the propensity score estimations using the three samples.

Table 4: Estimation of the propensity scores (probability of receiving the CSG)

	Sample 1		**Sample 2**		**Sample 3**	
Variable	**Coefficient**	**Robust Standard Error**	**Coefficient**	**Robust Standard Error**	**Coefficient**	**Robust Standard Error**
Caregiver Characteristics						
Education	-0.01	0.04	-0.01	0.05	0.04	0.05
Education squared	0.00	0.00	0.01	0.00	-0.00	0.00
Employed	0.07	0.19	0.18	0.23	0.07	0.22
Female	-0.26	0.27	-0.26	0.30	-0.02	0.29
Married	-0.24**	0.11	-0.46***	0.14	-0.23*	0.14
Child Characteristics						
Age	0.45***	0.04	0.29***	0.05	0.42**	0.19
Age squared	-0.02***	0.00	-0.02***	0.00	-0.03***	0.01
Boy	0.18*	0.09	0.17	0.12	0.15	0.12
Household Characteristics						
Electricity	-0.02	0.11	0.06	0.13	0.04	0.15
Piped water	0.14	0.12	0.18	0.16	0.09	0.16
Landline	-0.39**	0.16	-0.28	0.18	-0.56***	0.18
Flush toilet	-0.56***	0.14	-1.11***	0.18	-0.56***	0.18
Female head	0.08	0.10	0.13	0.13	0.10	0.14
Log pc expenditure	-0.24***	0.07	-	-	-0.20**	0.10
Constant	1.43***	0.50	1.58***	0.36	1.02	1.03-
Observations	4843		4904		3021	

*Significant at 10% level.
**Significant at 5% level.
***Significant at 1% level.

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.

8.2. **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.[24] This condition is especially relevant for cases where Kernel matching is chosen as the relevant matching algorithm, since this algorithm uses all observations to ascertain the average treatment effect. For nearest neighbour and calliper matching, this is less of an issue as matching only takes place between units falling within the defined calliper (or observations closest to the treated unit for nearest neighbour matching) (Caliendo & Kopeinig, 2008: 46).

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 2 to 4 below 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. The average treatment effect can therefore be estimated from the region of common support between the treated and control groups, without introducing evaluation bias into the results. However, the overlap in the second sample

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

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 2: Region of common support for sample 1

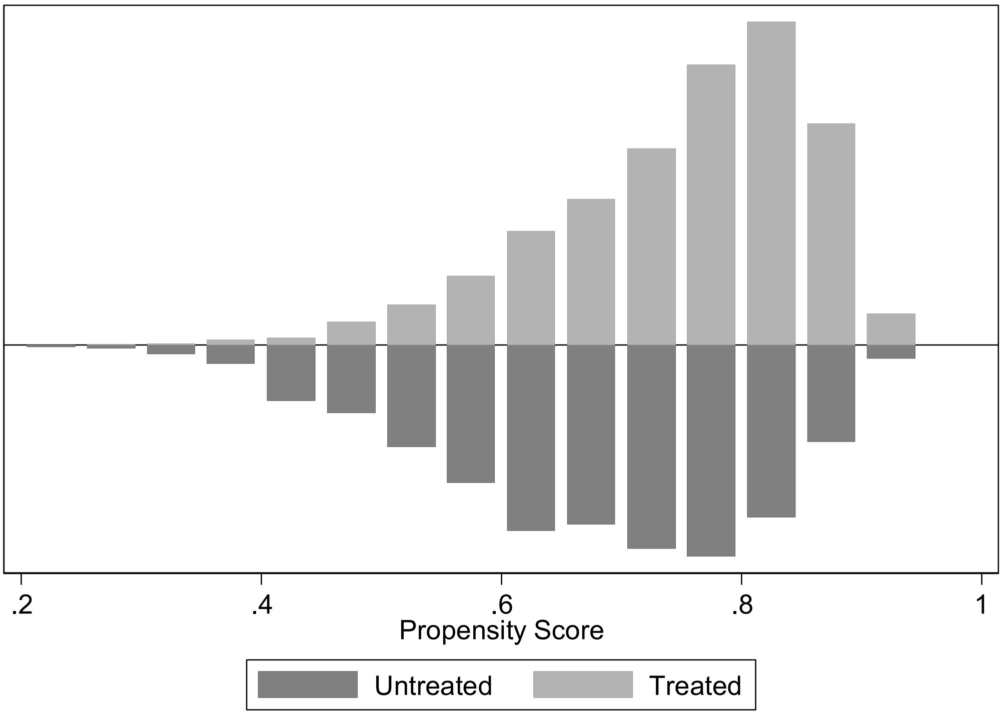

Figure 3: Region of common support for sample 2

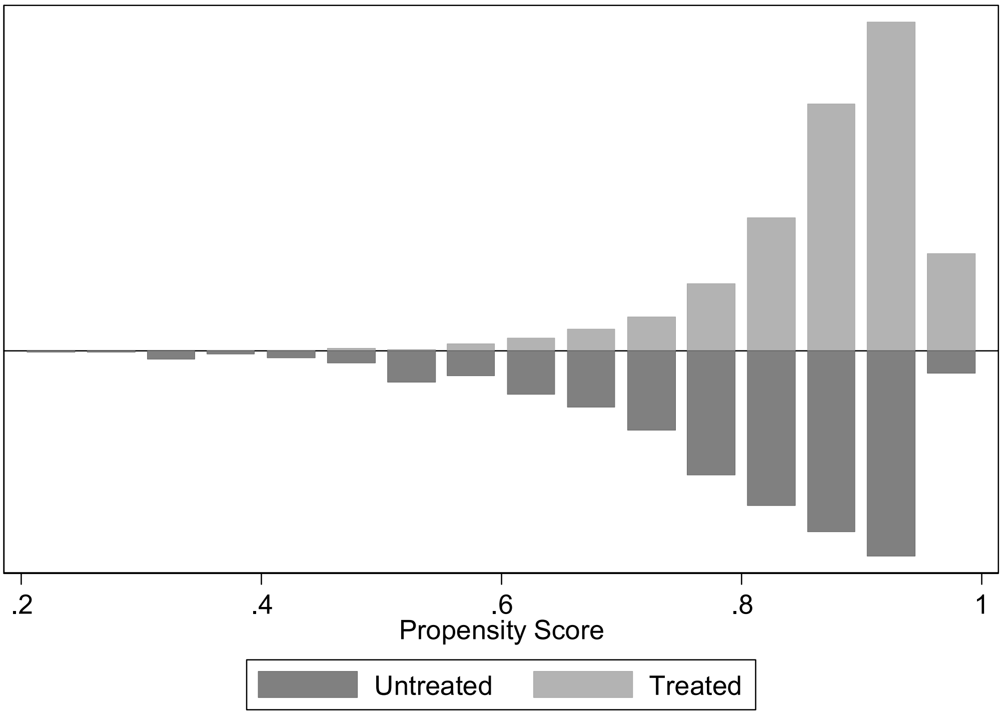

Figure 4: Region of common support for sample 3

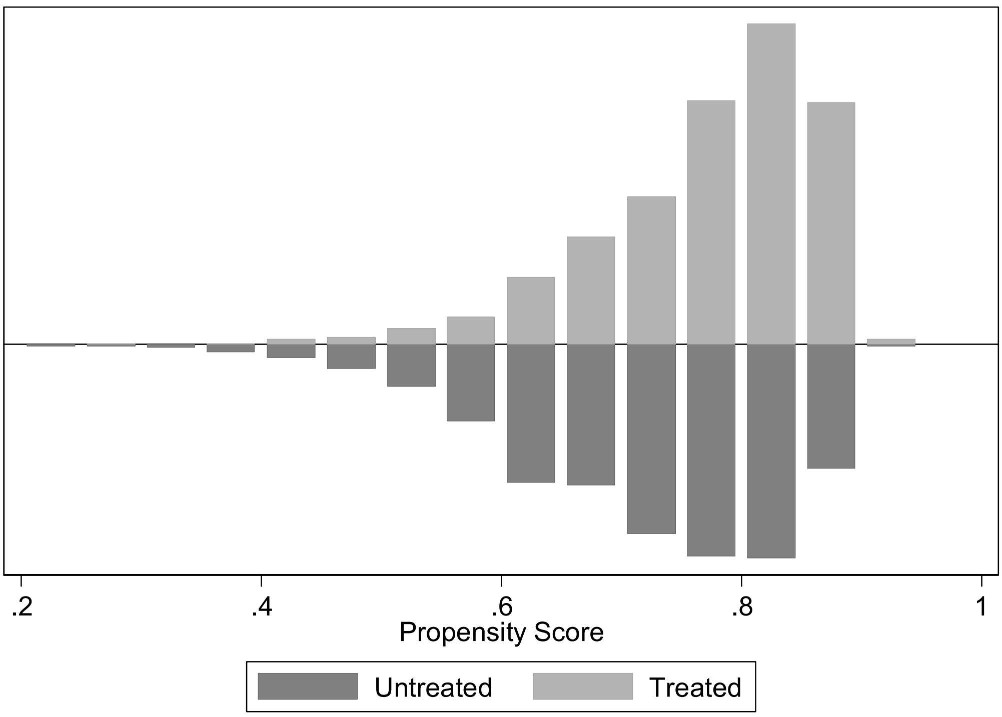

8.3. **Estimating the Average Treatment Effect using Different Matching Algorithms**

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 5.

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 is allows control units matched to each treated unit to be quite similar to the treated unit (in other words only units falling within the caliper are matched, while with nearest neighbor matching, the specified number of units are matched, while they could potentially be very far from the treated unit).[25] Again, the common support condition is imposed.

With the exception of the expenditure outcome variables, 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 household expenditure, however, seems to be significant and robust. All of the PSM techniques applied estimate a negative effect of receipt of the CSG on the *per capita household* expenditure on food and the 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

[25] Again, as discussed above, there is a trade-off between variance and bias, since calliper matching could potentially provide more unbiased results, however at the cost of increased variance since fewer matches are performed (Caliendo & Kopeinig, 2008: 42).

for these two outcome variables are very small, especially for expenditure on adult goods. As for the household expenditure on food, another concern is the limited number of control units used in the matching. The concern is therefore that PSM produces biased estimates of the treatment effect.

Although no clear rules for the use of the matching algorithms exist (i.e. which algorithm to use, how many neighbours to match to or what the size of the caliper should be), some general observations can be made regarding the results in Table 5. First, it is noticeable that extending the number observations to be included in matching (from the nearest neighbor to the 10 nearest neighbours) does not improve the results.

Table 5: Estimated average treatment effects (standard errors in parenthesis)

Matching Algorithm	Sample 1 HAZ	Sample 1 WHZ	Sample 2 Monthly *per capita* Household Food Expenditure	Sample 2 Monthly Adult Goods Expenditure per adult	Sample 3 School Enrolment	Sample 3 Repeat School Year
Nearest 1 Neighbour[#]	-0.04	-0.01	-48.28	-11.56	0.00	-0.02
	(0.07)	(0.12)	(5.50)	(4.23)	(0.01)	(0.03)
	N=3397	*N=1729*	*N=4899*	*N=922*	*N=2741*	*N=2777*
	T=2619	*T=1393*	*T=4201*	*T=756*	*T=2084*	*T=2115*
Nearest 10 Neighbours[#]	0.03	-0.01	-43.73	-9.89	0.01	-0.01
	(0.07)	(0.12)	(5.50)	(4.23)	(0.01)	(0.02)
	N=3397	*N=1729*	*N=4899*	*N=922*	*N=2741*	*N=2777*
	T=2619	*T=1393*	*T=4201*	*T=765*	*T=2084*	*T=2115*
Kernel	0.04	-0.02	-42.82	-9.66	0.01	-0.01
	(0.07)	(0.12)	(5.50)	(4.23)	(0.01)	(0.02)
	N=3398	*N=1736*	*N=4206*	*N=924*	*N=2742*	*N=2778*
	T=2620	*T=1400*	*T=4904*	*T=767*	*T=2085*	*T=2116*
Radius (caliper=0.01)[#]	0.03	-0.05	-44.37	-10.54	0.00	-0.01
	(0.07)	(0.12)	(5.50)	(4.23)	(0.01)	(0.02)
	N=3395	*N=1735*	*N=4903*	*N=923*	*N=2741*	*N=2777*
	T=2617	*T=1399*	*T=4205*	*T=766*	*T=2084*	*T=2115*
Radius (caliper=0.001)[#]	0.00	-0.01	-43.53	-9.14	0.01	-0.01
	(0.07)	(0.12)	(5.50)	(4.23)	(0.01)	(0.02)
	N=3284	*N=1482*	*N=4780*	*N=642*	*N=2632*	*N=2669*
	T=2506	*T=1146*	*T=4082*	*T=485*	*T=1975*	*T=2007*
Radius (caliper=0.0001)[#]	0.05	0.01	-48.75	-25.22	0.01	0.02
	(0.07)	(0.12)	(5.50)	(4.23)	(0.01)	(0.02)
	N=1673	*N=610*	*N=2672*	*N=231*	*N=1288*	*N=1304*
	T=895	*T=273*	*T=1974*	*T=74*	*T=631*	*T=642*

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.

Also, the effect of reducing the size of the caliper can be observed as the number of observations used in the matching process is reduced significantly with a corresponding reduction in the size of the caliper from 0.01, 0.001 and finally to 0.0001. As the caliper becomes smaller, fewer control units are matched to every treatment unit, also decreasing the quality (i.e. sign and size) of the estimates.

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 of the CSG is that the CSG is not implemented in a binary fashion. In addition, as already mentioned, the effect of the CSG on the outcome variables might take the form of a type of threshold effect, in other words the effect of the CSG transfers on the outcomes may only be observable subsequent to a certain monetary threshold being reached (after receiving the grants for a certain period). This criticism and a possible alternative (more appropriate) technique to correct for this flaw are discussed in more detail in the next section of the paper.[26]

Second technique: The use of Generalised Propensity Scores with a Continuous Treatment Variable

9. Methodology

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 that all individuals may be classified as falling into one of two categories, namely 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 similar and comparable and received equal dosages of treatment. In the case of the CSG, this assumption clearly does not hold. 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

[26] Theoretically, as a final step in conducting PSM, it is necessary to ascertain whether the balancing property holds. One option in this regard would be to make use of the method suggested by Leuven and Sianesi (2003), which tests whether, for each covariate, the difference before and after matching is statistically significant by conducting a t-test. This was done for all samples and matching techniques reported above. However, since none of the treatment effects estimated in Table 6 are statistically significant, these results are not reported here.

explain inconclusive results regarding the treatment effect of CSG (Agüero *et al*, 2006). 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. The authors use data collected from winners of the Megabucks lottery in Massachusetts during the mid-1980's to estimate the average labour-response of winners, more specifically their earnings six years after winning. Although the data were generated as a random process, and should, accordingly, satisfy the general unconfoundedness condition, 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 (as defined in the previous section); the level of treatment, D_i; and the outcome corresponding to the treatment received $Y_i = Y_i(D_i)$.

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

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

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, 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

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

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 emphasized 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 (2007, 2008) and Agüero *et al* (2009).

As a first step, the conditional distribution of treatment given the covariates is modeled, 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

$$\hat{R}_i = \frac{1}{\sqrt{2\pi\hat{\sigma}^2}} \exp\left(-\frac{1}{2\hat{\sigma}^2}\left(D_i - \hat{\beta}_0 - \hat{\beta}_1^{\prime x_i}\right)^2\right).$$

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

$$\lambda(D_i, \hat{R}_i) = E[Y_i | D_i, \hat{R}_i] = \alpha_0 + \alpha_1 D_i + \alpha_2 D_i^2 + \alpha_3 \hat{R}_i + \alpha_4 \hat{R}_i^2 + \alpha_5 D_i \hat{R}_i$$

These coefficients may be estimated using ordinary least squares or maximum likelihood (depending on the type of outcome variable). Hirano and Imbens (2004: 76) emphasize 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

$$E[\hat{Y}(d)] = \frac{1}{N}\sum_{i=1}^{N}\left(\hat{\alpha}_0 + \hat{\alpha}_1 d_i + \hat{\alpha}_2 d_i^2 + \hat{\alpha}_3 \hat{r}(d, x_i) + \hat{\alpha}_4 \hat{r}(d, x_i)^2 + \hat{\alpha}_5 (d_i) \hat{r}(d, x_i)\right)$$

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

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

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. **Modeling the Conditional Distribution of Treatment**

As a first step, a maximum likelihood estimation was conducted for all three samples and the results presented below. For the sake of consistency, the specification of the model remains the same as discussed under the section dealing with PSM. As with PSM, one of the assumptions to be satisfied 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. Again, as discussed in the section above, it is argued that compliance with this assumption is achieved 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, the caregiver motivation variable is also added to the model in order to control for unobserved caregiver characteristics. Since the relevant sample for this technique includes only children who have received the CSG, none of the *caveats* surrounding the use of the variable discussed in the previous section apply.

The results from the maximum likelihood estimation are set out in Table 6 below, and correspond to a large extent with the results from the binary treatment estimation, apart from the inclusion of the caregiver motivation variable. This coefficient is large and significant, indicating the expected positive relationship between more motivated caregivers and larger dosages of treatment.

Furthermore, children with employed and better educated caregivers are more likely to be covered for a larger proportion of their lives, as one would expect their caregivers to generally be more eager to apply for the CSG and to ensure that the grant is received as early as possible in the child's life. The effect of education is again non-linear.[27].

The age variables are also jointly significant[28] and confirm previous evidence that older children are more likely to be receiving the CSG for a smaller percentage of their lives, possible 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 or a flush toilet decreases the probability of the child having received the CSG for a large portion of his/her life.

Table 6: Maximum likelihood estimates with different samples

[27] The education coefficients are jointly significant at the 5% significance level for all three samples.
[28] At the 5% significance level for all three samples.

	Sample 1		Sample 2		Sample 3	
Variable	**Coefficient**	**Standard Error**	**Coefficient**	**Standard Error**	**Coefficient**	**Standard Error**
Caregiver Characteristics						
Motivation	13.18***	0.32	11.05***	0.33	3.38***	0.37
Education	-0.44	0.44	-0.35	0.46	-0.50	0.58
Education squared	0.11***	0.04	0.09**	0.04	0.09*	0.05
Employed	4.47***	1.51	3.81**	1.53	0.43	1.77
Female	-1.88	2.84	-0.49	2.83	-0.41	3.23
Married	-0.93	1.16	-0.78	1.20	-2.46	1.51
Child Characteristics						
Age	0.50	0.45	-0.11	0.54	-3.17	1.95
Age squared	-0.14***	0.04	-0.49***	0.48	-0.07*	0.11
Boy	1.19***	0.98	-0.14	0.04	0.69	1.30
Household Characteristics						
Electricity	1.01	1.19	1.62	1.25	2.21	1.60
Piped water	-0.71	1.29	-0.05	1.35	-1.36	1.72
Landline	-4.98***	1.81	-4.70**	1.91	-3.20	2.45
Flush toilet	-3.53**	1.37	-5.22***	1.42	-0.73	1.84
Female head	-1.67	1.05	0.34	1.10	0.17	1.44
Per capita expenditure	-1.12	0.72	-	-	-1.37	0.97
Constant	55.00***	4.95	55.17***	3.42	95.79***	10.27
Observations	3304		2888		1421	

*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 7 below. 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, education and electricity covariates) 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 5 above), 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 7: The balancing property: t-statistics for equality of means between treatment terciles

				Sample 1			Sample 2			Sample 3		
	Unadjusted for GPS Treatment Intervals			Adjusted for GPS Treatment Intervals			Adjusted for GPS Treatment Intervals			Adjusted for GPS Treatment Intervals		
Variable	**Low [0,35)**	**Medium [35,66)**	**High [66,100]**	**Low [0,35)**	**Medium [35,66)**	**High [66,100]**	**Low [0,35)**	**Medium [35,66)**	**High [66,100]**	**Low [0,35)**	**Medium [35,66)**	**High [66,100]**
Caregiver Characteristics												
Motivation	10.41	4.26	-12.40	3.01	-6.68	-8.08	-2.20	-3.40	-6.34	3.13	1.80	-5.33
Education	6.73	0.71	-6.04	-0.43	1.21	0.43	-0.41	-1.24	1.19	1.89	-0.43	0.28
Education squared	7.04	0.96	-6.51	-0.34	1.47	0.94	-0.28	-1.22	1.65	1.90	-0.29	0.31
Employed	-0.86	1.28	-0.48	-1.45	-0.30	2.08	-1.64	-0.23	2.23	-0.02	0.46	0.11
Female	0.68	-0.28	-0.29	-0.61	0.82	-0.21	-0.26	-0.60	0.17	0.71	-0.38	0.33
Married	-0.18	-0.47	0.57	1.06	-1.61	-1.25	1.02	-0.91	-1.32	0.59	-0.23	-0.30
Child Characteristics												
Age	-17.53	-3.35	17.24	0.74	-12.94	6.71	-2.61	1.19	5.02	-0.88	-6.13	4.85
Age squared	-18.39	-4.72	19.43	1.54	-15.96	7.50	-2.91	0.45	6.28	-1.34	-6.27	6.12
Boy	0.70	0.17	-0.72	0.88	-0.13	-0.03	0.51	-0.28	-0.27	-0.14	-0.38	0.32
Household Characteristics												
Electricity	2.52	0.80	-2.75	1.26	0.03	-0.42	0.25	0.13	-0.58	1.55	-1.20	0.35
Piped water	-0.19	0.03	0.12	-0.27	0.08	-0.52	0.56	0.66	-0.98	0.54	-0.33	-0.65
Landline	-0.42	-0.01	0.34	0.77	-0.16	-0.51	1.21	0.08	-0.83	0.80	-0.23	-0.39
Flush toilet	-0.60	1.77	-1.12	-0.43	1.71	-1.09	0.98	2.45	-2.51	0.07	0.72	-1.27
Female head	-0.79	-0.38	0.98	0.99	-0.21	1.56	-0.26	-0.97	1.02	-0.85	-0.27	0.90
Log per capita expenditure	0.35	-0.62	0.29	1.27	-1.50	0.60	-	-	-	1.52	-0.79	-0.40

10.2. **Estimating the Outcomes Conditioned on the GPS and Treatment**

The results of the estimates in the second stage are presented below. 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.

Table 8: Second stage estimates per outcome

Outcome variable (estimation technique)	**HAZ (OLS)**	**WHZ (OLS)**	***Per capita* Household Food Expenditure (OLS)**	**Adult Goods Expenditure per Adult (OLS)**	**School Enrolment (Maximum Likelihood)**	**Repeat School Year (Maximum Likelihood)**
Parameter	**Coefficient (Robust Standard Error)**					
α_0	-0.63	1.36	191.70	21.88	6.47	-2.47
	(0.21)	(0.51)	(12.68)	(12.65)	(3.19)	(0.46)
α_1	-0.01	0.01	-1.25	-0.41	0.07	0.05
	(0.00)	(0.01)	(0.33)	(0.31)	(0.07)	(0.02)
α_2	0.00	-0.00	0.00	0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
α_3	60.49	-177.25	-4202.28	2182.84	-410.24	149.58
	(43.67)	(106.70)	(2888.57)	(2787.85)	(523.81)	(86.85)
α_4	-3487.41	2253.73	162214.2	-25647.31	12446.69	-2693.45
	(2536.86)	(5702.60)	(176625.9)	(164037.1)	(23243.95)	(4264.11)
α_5	0.87	-1.36	31.10	-22.24	-0.01	-2.50
	(0.32)	(0.51)	(22.12)	(19.88)	(2.90)	(0.65)
Observations	2161	1308	2888	609	1267	1294
Prob>F	0.00	0.00	0.00	0.04		
Prob>Chi2					0.63	0.00

As indicated above, the coefficients from these estimations have no direct meaning, however, testing whether the coefficients in this estimation are statistically different to zero 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 8. The null-hypothesis can be rejected for all of the models except where school enrolment is specified as the outcome variable. The analysis is therefore taken to the next stage in order to estimate the treatment effect.

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 result are presented as Figures 5 to 7, with the treatment level indicated on the y-axis and the change in the outcome variable at each level of treatment is indicated on the x-axis.

The 95% confidence intervals (calculated using bootstrapped standard errors with 500 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 household expenditure variables (based on the confidence intervals calculated), these two figures are not included.

Figures 5 and 6 reflect positive gains from the CSG on both HAZ and WHZ. The estimated average treatment effects are, however, not very large. Positive and statistically significant gains in HAZ from receiving the CSG are recorded for children receiving the grant for 30-40% of their lives. The gains in HAZ reach a maximum where children have been receiving the CSG for 40% of their lives, compared to a children who have only been receiving it for 30%, where the increase in the z-score is approximately 0.03, translating into 3% of the standard deviation. This estimated gain in HAZ occur at a lower level of treatment and is somewhat less that 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.

As for the estimated effect on WHZ, this reaches a maximum at an increase in the z score of 0.06 (i.e. 6% of a standard deviation), which occurs where children have been receiving the CSG for 50% of their lives compared to children who have only received it for 40%. Neither of these estimates is large, and although they indicate some positive effect on the health and nutritional well-being of children receiving the grant, they do not provide convincing evidence of a substantial impact.

Both of these effects are estimated to be negative at higher dosages of treatment. However, these portions of the curves may be ignored since they fall outside the confidence interval band.

Figure 5: Estimated gains in HAZ

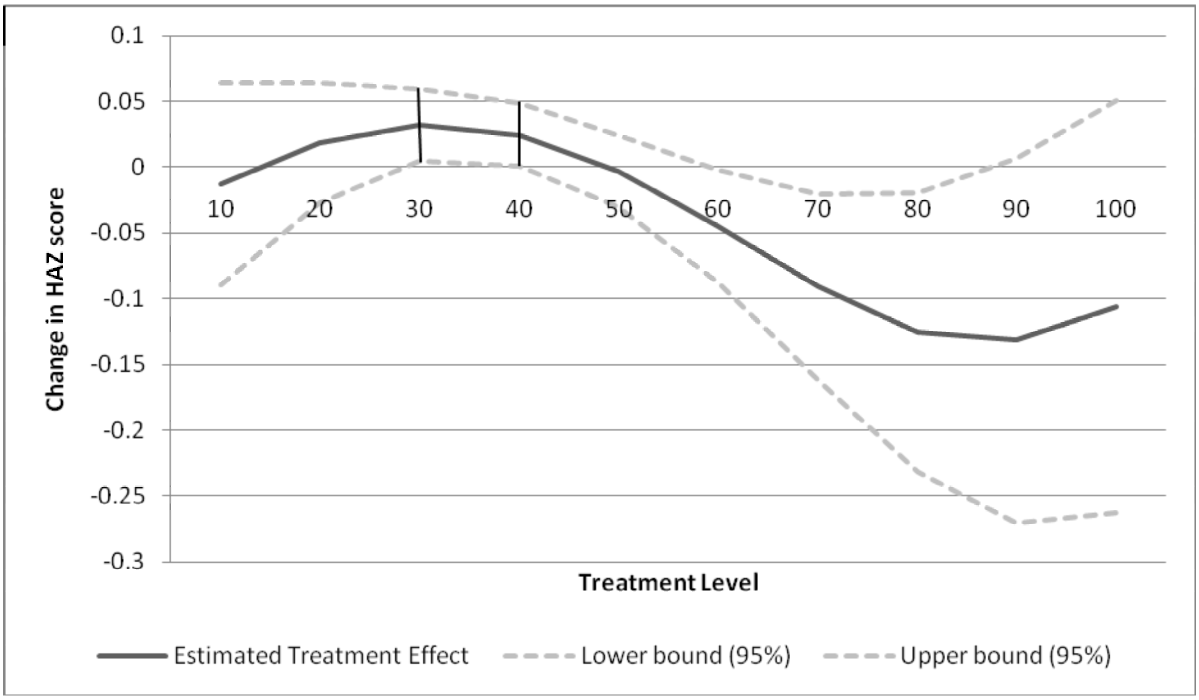

Figure 6: Estimated gains in WHZ

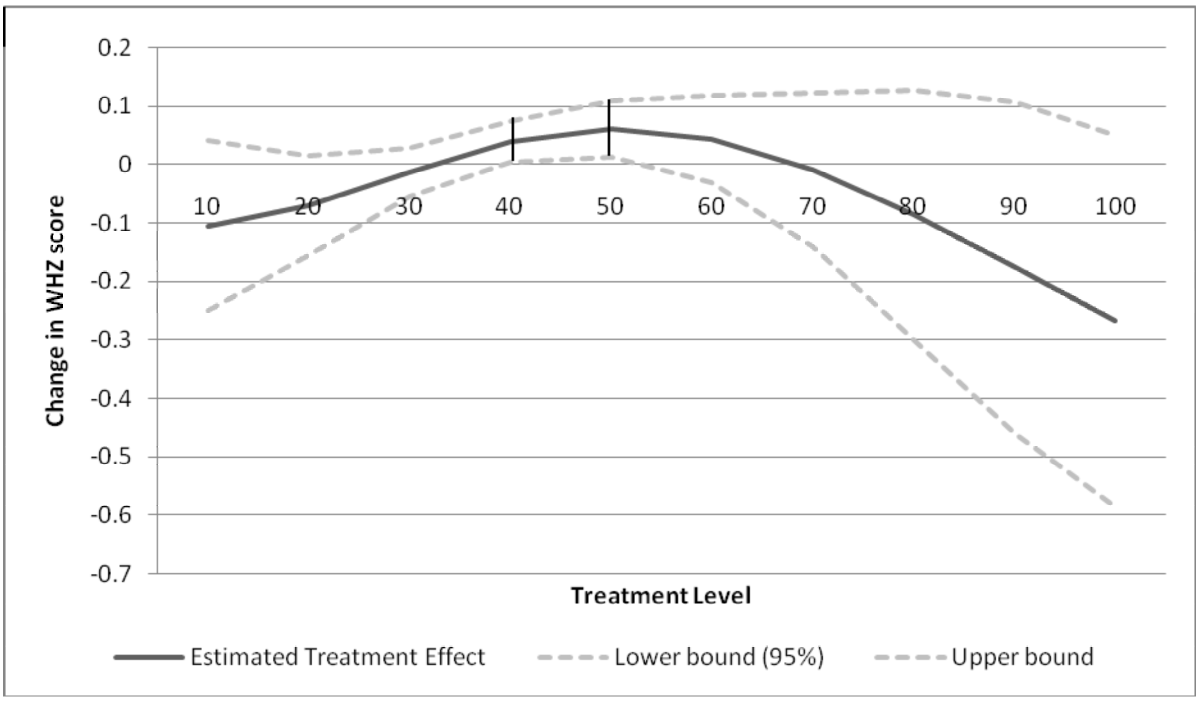

Figure 7: Estimated change in probability of repeating a school year (children older than 5 years)

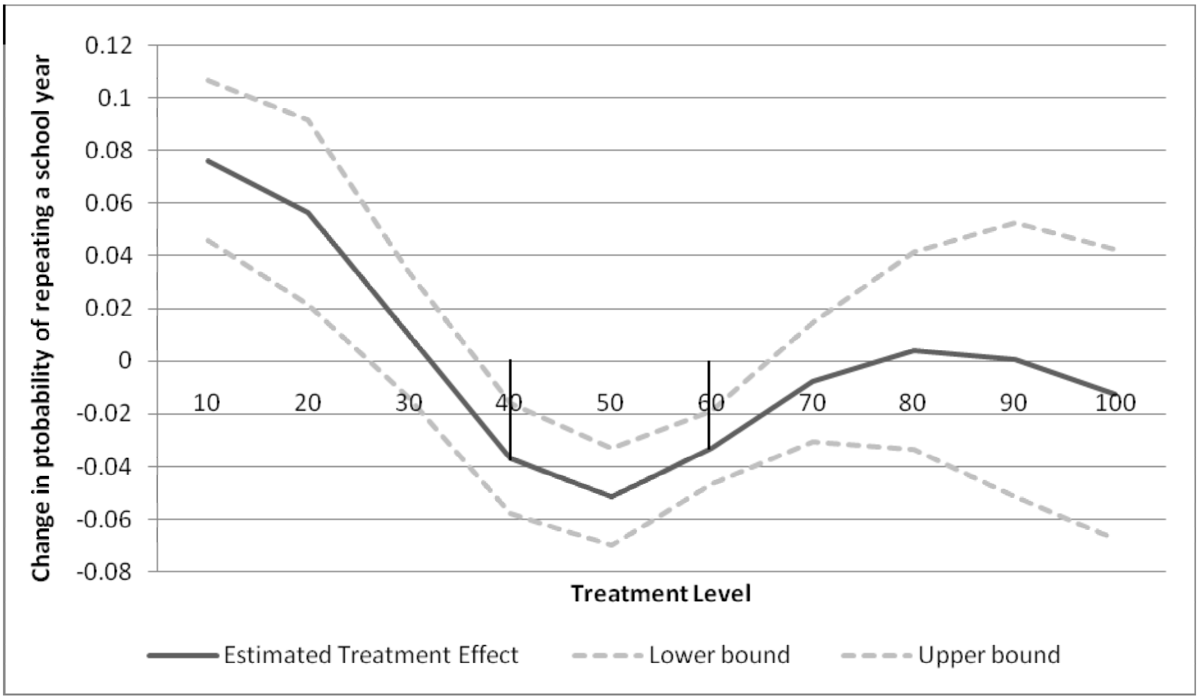

Looking at Figure 7, there appears to be some positive impact on the probability of repeating a school year. A significant decrease in this probability is estimated for children receiving the CSG between 40% and 60% of their lives. The maximum impact is a decrease in the probability of approximately 0.05 percentage points. 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 recipient children, they 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 channeled towards the purchase of other goods 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 the cash is spread across the entire household and that 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 recipient 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, outcomes, an evaluation which is warranted 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 randomized 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 technique 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 and education. Although a significant treatment effect is estimated for household expenditure, the absence of a sufficiently large overlap between treated and control units as well as a relatively small number of control observations indicate that these estimates may be biased and accordingly unreliable as an indication of the effect of the CSG. In addition, 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, WHZ and whether the child has repeated a school year. 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 Uncondtional 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. T*he 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 Microeconometric 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 Responjses to Anticipated Cash Transfers in South Africa. *National Bureau of Economic Research Working Paper No. 10265*.

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.

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.

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.

Samson, M., *et al*. 2001. The Impact of Social Security on School Enrolment in South Africa. *EPRI Research Paper number 25*.

Smith, J.A. & Todd, P.E. 2005. Does matching oversome 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).

Straus, 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.

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.

World Bank. 2006. *Repositioning Nutrition as Central to Development: A Strategy for Large-Scale Action*. Directions in Development, The World Bank.

Appendix A – Method 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).